Adaptive Inference with Weak Instruments

Apurv Shukla [1] **Debabrota Basu** [2]

Abstract

Inference with weak instruments presents a choice between the two-stage least squares (2SLS) yielding tight confidence intervals with low coverage, and the test-based methods, like Anderson–Rubin, producing vacuously wide but valid intervals. We show that this trade-off is unnecessary. First, we show that a convex combination of a bias-controlled estimator, like ordinary least squares, and a valid estimator, like 2SLS, preserves coverage. Second, to yield tight intervals, we propose StAR that uses e-process diagnostics to adaptively set these weights. StAR's weights involve a multivariate e-value aggregating relevance of all the instruments, and the confidence width of the e-process gating against premature reliance on any instrument. We prove three finite-sample results: (i) StAR's coverage, (ii) StAR's width that transitions from OLS-like at weak instruments to IV-efficient at strong instruments at an explicit rate, and (iii) an oracle inequality for StAR showing that it achieves an error rate within a constant factor of the oracle rate. Our experiments on synthetic data shows that StAR maintains $\geq 96\%$ coverage and produces 6× tighter intervals than Anderson–Rubin at weak instruments. Our experiments on clinical AI and econometric data shows that StAR yields valid and informative CIs, and accurately detects collective instrument strength.

1. Introduction

Weak instruments arise naturally in encouragement designs with partial compliance (Angrist et al., 1996; Finkelstein et al., 2012), Mendelian randomization (Burgess et al., 2011), and off-policy evaluation (Kallus and Uehara, 2022). In an Instrumental Variable (IV) analysis of mechanical ventilation and ICU mortality using MIMIC-IV data, we find that this methodological choice is consequential: in a weak-instrument subgroup ($F = 5.66$), the classical Anderson-Rubin (AR) (Anderson and Rubin, 1949) and Conditional Likelihood Ratio (CLR) (Moreira, 2003) tests exclude zero and suggest treatment is harmful, while the true uncertainty is signifcantly larger (Section 5).

This evokes a fundamental dilemma for inference in IV regression. Dufour (1997) showed that under weak identification, any confidence set with correct unconditional coverage must be unbounded with positive probability. This means that valid inference produces vacuous intervals when the data cannot distinguish the causal effect from noise. Test-inversion methods, such as Anderson–Rubin (AR) and Conditional Likelihood Ratio test (CLR), achieve this validity bound, while producing exact or asymptotically optimal confidence sets that are, as Dufour's result demands, uninformative for weak instruments. Regression-based methods, like Two-Stage Least Square (2SLS) (Angrist and Imbens, 1995; Vecchia and Basu, 2025), circumvent vacuousness by imposing structure, but pay with severely distorted coverage when instruments are weak (Bound et al., 1995; Staiger and Stock, 1997). The question is whether a principled middle ground exists: *can one construct confidence intervals that are always valid (perhaps conditionally), always informative, and automatically adaptive to instrument strength and confounding?*

We show that the *efficiency–validity trade-off* can be resolved by *blending* estimators rather than choosing between them. The key observation is a convex combination principle for inference: if a bias-controlled estimator and a valid estimator are combined with *any* data-dependent weight $\omega \in [0, 1]$, the resulting confidence interval inherits validity from both estimators. Coverage does not depend on ω being "correct"—any weight preserves it. The challenge reduces to choosing ω well for *width*, while adapting to unknown levels of confounding and instrument strength.

Our Proposal: StAR. First. we instantiate this design principle for IV regression with the *Strength Adaptive Regressor* (StAR). The confidence interval center is $\hat{\beta}_\omega \triangleq \omega\hat{\beta}_{\mathrm{IV}} + (1-\omega)\hat{\beta}_{\mathrm{OLS}}$, where $\hat{\beta}_{\mathrm{IV}}$ and $\hat{\beta}_{\mathrm{OLS}}$ are the estimated parameters of 2SLS based regression (IV) and ordinary least squares (OLS). Further, the confidence width is

[1]Department of Electrical and Computer Engineering, University of Michigan, Dearborn [2] Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 - CRIStAL, Lille, France. Correspondence to: Apurv Shukla <apurv.shukla@umich.edu>.

Proceedings of the 43rd International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

Table 1. Comparison of inference methods for linear IV regression. "Adaptive" indicates whether the method automatically adjusts to instrument strength (κ) and/or confounding (ρ).

Method	Adaptive to		Approach		Coverage	Width at weak IV	Non-asymp. analysis
	κ	ρ	Regr.	Test			
OLS / Wald	✗	✗	✓	✗	Invalid ($\to 0$)	$O(n^{-1/2})$ (undercoverage)	✗
2SLS (Vecchia and Basu, 2025)	✗	✗	✓	✗	Invalid when μ^2 small	$O(n^{-1/2})$ (undercoverage)	✗
AR (Anderson and Rubin, 1949)	✗	✗	✗	✓	Exact (unconditional)	Vacuous (≈ 20)	✓
CLR (Moreira, 2003)	✗	✗	✗	✓	Exact (unconditional)	Vacuous (≈ 20)	✓
Stock–Yogo (Stock and Yogo, 2005)	binary	✗	✓	✗	Pre-test; invalid if test errs	Rejects IV	✗
AK (Andrews et al., 2019)	✓	✗	✗	✓	Asymptotic (unconditional)	Optimal asymptotic	✗
(Xia et al., 2024)	✗	✗	✓	✗	Finite-sample (conditional)	Very wide ($\gg 20$)	✓
StAR (this work)	✓	✓	blend		Finite-sample (conditional)	$O\left(\frac{\sigma}{\sqrt{n}} + \frac{1}{\kappa^2+1}\right)$	✓

$2h_\omega \triangleq 2\omega w_{\mathrm{IV}} + 2(1-\omega)(w_{\mathrm{OLS}} + B)$, where B is a principled upper bound on the OLS endogeneity bias derived from the concentration parameter, and w_{IV} and w_{OLS} are their respective confidence widths (Proposition 3.1).

Second, we construct the adaptive weight $\omega \in [0,1]$ from two e-process diagnostics (Grünwald et al., 2024a). The first is a *multivariate GROW e-value* that tests the IV moment condition– whether the prediction errors and any of the IVs are independent– across a grid of possible parameter values by aggregating evidence from all the instruments simultaneously. The peak e-value $\lambda_{\max}$ across parameters measures how strongly the data support instrument relevance. $\lambda_{\max}$ grows exponentially under the alternative, when the prediction errors and any of the IVs are correlated, and otherwise, stays near 1. Thus, it provides a non-asymptotic analog of the F-statistic that is valid at any sample size and stopping time. The second is the *e-process confidence set width* W_{CS}, which serves as a quality gate: a large $\lambda_{\max}$ is insufficient if the e-process has not yet localized at the true parameter. This can happen when instruments are moderate and the IV estimator has high variance. Finally, the multiplicative combination of these two factors $\omega \triangleq \frac{\lambda_{\max}}{\lambda_{\max}+c_0}\sqrt{\max(0, 1 - W_{\mathrm{CS}}/w_{\max})}$ ensures that *both strong evidence of relevance* and *sufficient localization* must hold before the weight shifts toward IVs. The weight ω transitions smoothly from 0 (pure OLS with bias correction) to 1 (pure IV) as instrument strength increases, avoiding the fragility of binary pre-testing (Stock and Yogo, 2005).

Then, we establish three finite-sample results for StAR. First, *coverage*: the StAR confidence interval covers β^* with high probability for any data-dependent weight conditional on the IV estimator covering β^* and a bias bound induced by OLS (Theorem 3.2). This is a direct consequence of the convex inference and imposes no restrictions on ω. Second, *width adaptation*: the CI width transitions at explicit rates from being $O(\kappa^{-1}n^{-1/2})$ at the strong instruments to $O((1+\kappa^2)^{-1} + n^{-1/2})$ for the weak instruments. Here, n is the number of samples and κ denotes the per-instrument signal strength. In the weak regime, i.e. when $n\kappa^2 \to 0$, StAR behaves like a bias-corrected OLS interval. In the strong regime, i.e. when $n\kappa^2 \to 1$, StAR converges to the efficient IV width (Theorem 4.4). Third, *near-oracle mean squared error (MSE)*: the data-driven weight achieves MSE within a constant factor of the oracle weight ω^* that minimizes MSE but needs the knowledge of true instrument strength κ (Theorem 4.3). This confirms that the e-process weight is not a heuristic but a provably near-optimal adaptation mechanism.

Finally, to *numerically test the efficiency-validity trade-off achieved by StAR*, we run extensive experiments with synthetic and three real-life dataset. Our experiments on synthetic data shows that StAR maintains $\geq 96\%$ coverage and produces 6× tighter intervals than AR, CLR, and (Xia et al., 2024) at weak instruments. Our experiments on clinical AI and econometric data shows that StAR yields valid and informative CIs, and accurately detects collective instrument strength. Specially, for (Angrist and Krueger, 1991)'s return-to-schooling data, we show that our multivariate construction is essential for valid inference in multi-IV setting that scalar frameworks fail to capture.

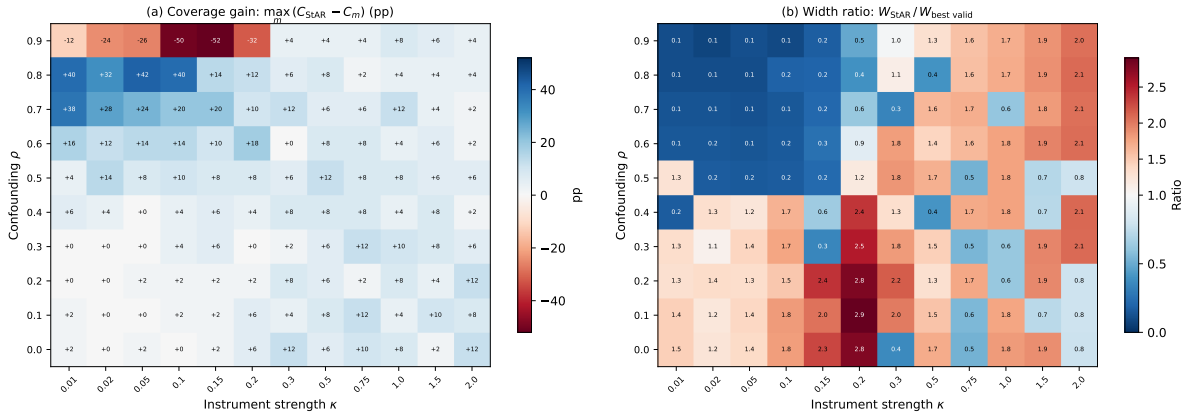

Figure 1. StAR vs. baselines across 120 configurations ($n = 200$, $d_z = 2$, 95% nominal, 1000 reps). (a) Maximum coverage gain $\max_m(C_{\text{StAR}} - C_m)$: blue cells indicate StAR achieves higher coverage than every competitor; red cells appear only under extreme confounding ($\rho \geq 0.8$) with very weak instruments. (b) Width ratio $W_{\text{StAR}}/W_{\text{best valid}}$, where "best valid" denotes the narrowest competitor achieving at least 95% coverage. Grey indicates no competitor attains valid coverage. Blue (< 1) indicates StAR is tighter.

2. Problem Setup: Linear IV Regression

We consider the standard linear IV model:

$$Y = \beta^* X + \eta, \qquad \text{(Stage-2)}$$
$$X = \boldsymbol{\pi}^\top \boldsymbol{Z} + \varepsilon, \qquad \text{(Stage-1)}$$

where $Y \in \mathbb{R}$ is the outcome, $X \in \mathbb{R}$ is the endogenous treatment (or covariate), $\boldsymbol{Z} \in \mathbb{R}^{d_z}$ is the vector of IVs, $\beta^* \in \mathbb{R}$ is the structural parameter, and $\boldsymbol{\pi} \in \mathbb{R}^{d_z}$ is the first-stage coefficient vector. The errors (ε_i, η_i) are jointly distributed with correlation $\rho = \mathrm{Corr}(\varepsilon_i, \eta_i) \neq 0$. This leads to endogeneity because the treatment X shares a common unobserved driver with the outcome Y.

Assumption 2.1 (Exclusion restriction)**.** The structural error is mean-independent of IVs $\mathbb{E}[\eta \mid \boldsymbol{Z}] = 0$.

Assumption 2.2 (Sub-Gaussian errors)**.** Conditional on Z, the errors (ε_i, η_i) are jointly sub-Gaussian with parameter σ^2, i.e., for all $\lambda \in \mathbb{R}$, $\mathbb{E}[e^{\lambda \varepsilon_i} \mid Z] \leq e^{\sigma^2\lambda^2/2}$ and $\mathbb{E}[e^{\lambda \eta_i} \mid Z] \leq e^{\sigma^2\lambda^2/2}$. The non-zero correlation coefficient is $\rho = \mathrm{Corr}(\varepsilon_i, \eta_i) \in [-1, 1]$.

When $\sigma_\varepsilon \neq \sigma_\eta$ (as in the MIMIC-IV application, Section 5), we write σ_ε^2 and σ_η^2 for the respective sub-Gaussian parameters. The proofs carry through with $\sigma^2 = \max(\sigma_\varepsilon^2, \sigma_\eta^2)$; the bias bound (Proposition 3.1) includes an explicit $\sigma_\eta/\sigma_\varepsilon$ ratio for dimensional consistency.

Assumption 2.3 (Instrument regularity)**.** The instruments satisfy $\mathbb{E}[\boldsymbol{Z}] = \bar{\mathbf{z}}$, $\mathrm{Cov}(\boldsymbol{Z}) = \Sigma_z \succ 0$, and $\|\mathbf{z}\|_2 \leq R$ almost surely. The instruments are independent of the errors: $\boldsymbol{Z} \perp (\varepsilon, \eta)$.

Definition 2.4 (Instrument strength)**.** The *per-instrument signal* and the *concentration parameter* are

$$\kappa = \frac{\|\pi\|}{\sqrt{d_z}}, \qquad \mu^2 = \frac{n\|\pi\|^2}{d_z} = n\kappa^2.$$

Remark 2.5 (Concentration parameter and SNR). The concentration parameter μ^2 is the non-centrality of the first-stage F-distribution. Under Gaussianity, $F \sim F_{d_z, n-d_z-1}(\mu^2)$, so $\mathbb{E}[F] \approx 1 + \mu^2$. When $\sigma^2 = 1$, μ^2 coincides with $n \times \mathrm{SNR}$, where $\mathrm{SNR} = \|\pi\|^2/(d_z\sigma^2)$ is the per-instrument signal-to-noise ratio. The classical Stock–Yogo threshold $F \approx 10$ separates weak and strong instruments in asymptotics (Stock and Yogo, 2005). Our analysis is non-asymptotic and does not need large μ^2.

Estimators. Let $\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{\eta}, \boldsymbol{\varepsilon} \in \mathbb{R}^n$ denote the stacked vectors of n samples and $Z \in \mathbb{R}^{n \times d_z}$ be the instrument matrix. Define the projection $P_Z = \tilde{Z}(\tilde{Z}^\top \tilde{Z})^{-1}\tilde{Z}^\top$, where $\tilde{Z} = [\mathbf{1}_n, Z]$ includes an intercept. Thus, the first-stage fitted values $\hat{\boldsymbol{x}} = P_Z \boldsymbol{x}$, and the first-stage residuals $\boldsymbol{v} = (I_n - P_Z)\boldsymbol{x}$. The *OLS estimator* is

$$\hat{\beta}_{\text{OLS}} = \frac{\boldsymbol{x}^\top \boldsymbol{y}}{\boldsymbol{x}^\top \boldsymbol{x}} = \beta^* + \frac{\boldsymbol{x}^\top \boldsymbol{\eta}}{\boldsymbol{x}^\top \boldsymbol{x}}. \tag{1}$$

Since $\boldsymbol{x}$ and $\boldsymbol{\eta}$ are correlated through ρ, the ratio $\boldsymbol{x}^\top \boldsymbol{\eta}/(\boldsymbol{x}^\top \boldsymbol{x})$ has non-zero conditional expectation: $\hat{\beta}_{\text{OLS}}$ is biased for β^* whenever $\rho \neq 0$. The *2SLS (IV) estimator* replaces $\boldsymbol{x}$ with its first-stage fits:

$$\hat{\beta}_{\text{IV}} = \frac{\hat{\boldsymbol{x}}^\top \boldsymbol{y}}{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}} = \beta^* + \frac{\hat{\boldsymbol{x}}^\top \boldsymbol{\eta}}{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}}, \tag{2}$$

where the denominator simplifies because $\hat{\boldsymbol{x}}^\top \boldsymbol{x} = \hat{\boldsymbol{x}}^\top(\hat{\boldsymbol{x}} + \boldsymbol{v}) = \hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}$, since $P_Z^\top(I_n - P_Z) = 0$. Since $\hat{\boldsymbol{x}} = Z\pi + P_Z\boldsymbol{\varepsilon}$ depends on $\boldsymbol{\varepsilon}$, which is correlated with $\boldsymbol{\eta}$, the 2SLS estimator has a finite-sample bias of order $\rho\, d_z/\mu^2$ toward OLS (the Nagar bias; Nagar, 1959). This bias vanishes as $\mu^2 \to \infty$ but the variance $\mathrm{Var}(\hat{\beta}_{\text{IV}} \mid Z) \approx \sigma_\eta^2/(\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}})$ diverges when instruments are weak ($\|\pi\| \to 0$).

For confidence level $1 - \delta$, the standard half-widths for OLS and IV are $w_{\text{OLS}} = z_{\delta/2} \cdot \hat{\sigma}_{\text{OLS}}/\sqrt{\boldsymbol{x}^\top \boldsymbol{x}}$ and $w_{\text{IV}} = z_{\delta/2} \cdot \hat{\sigma}_{\text{IV}}/\sqrt{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}}$, where $z_{\delta/2} = \Phi^{-1}(1 - \delta/2)$.

Existing weak-instrument-robust inference. The dominant approach to inference under potentially weak instruments is *test inversion*: for each candidate β_0, test $H_0 : \beta^* = \beta_0$ and collect all non-rejected values into a confidence set. AR test (Anderson and Rubin, 1949) evaluates

$$\mathrm{AR}(\beta_0) = \frac{(y - \beta_0 x)^\top P_Z (y - \beta_0 x)/d_z}{(y - \beta_0 x)^\top M_Z (y - \beta_0 x)/(n - d_z - 1)},$$

where $M_Z = I_n - P_Z$. Under the null, $\mathrm{AR} \sim F_{d_z, n-d_z-1}$, yielding an exact F-test that is valid regardless of instrument strength. The conditional likelihood ratio (CLR) test (Moreira, 2003) conditions on a sufficient statistic for nuisance parameters to achieve higher power while retaining exactness. Both methods inherit Dufour (1997)'s impossibility result: any confidence set with correct unconditional coverage must be unbounded with positive probability. Thus, AR and CLR yield uninformative intervals when instruments are weak. In practice, AR often returns the full parameter range (width ≈ 20) for $F < 10$, providing validity at the cost of no practical guidance.

The inference problem. The goal is to construct a confidence interval for β^* that maintains valid coverage across all instrument strengths while achieving width that adapts to the realized μ^2, i.e. being near $w_{\mathrm{OLS}} + B$ (bias-corrected OLS width) when instruments are weak, and near w_{IV} (efficient IV width) when instruments are strong. StAR achieves this by *blending* $\hat{\beta}_{\mathrm{OLS}}$ and $\hat{\beta}_{\mathrm{IV}}$ with a data-driven weight $\omega \in [0, 1]$ derived from e-process diagnostics, rather than choosing between them.

3. Coverage-Preserving Convex Inference

The key observation behind StAR is that coverage and efficiency can be *decoupled*. We construct a confidence interval whose validity holds for *any* weight $\omega \in [0, 1]$ blending IV and OLS—including data-dependent choices that are computed from the same sample. Since the weighted estimator $\hat{\beta}_\omega = \omega \hat{\beta}_{\mathrm{IV}} + (1 - \omega)\hat{\beta}_{\mathrm{OLS}}$ is a convex combination, its error simply decomposes as

$$|\hat{\beta}_\omega - \beta^*| \;\le\; \omega\,|\hat{\beta}_{\mathrm{IV}} - \beta^*| + (1-\omega)\,|\hat{\beta}_{\mathrm{OLS}} - \beta^*| \;\le\; h_\omega, \tag{3}$$

where h_ω is the matching convex combination of the two estimators' half-widths. No matter how ω is chosen, coverage follows from the arm events alone. This means we are free to design that ω to optimize *width* (Section 4) without any risk of invalidating the interval.

3.1. Bias-Corrected OLS

The OLS estimator is biased when $\rho \neq 0$, so the OLS estimator requires a bias correction to serve as a valid confidence bound.

Proposition 3.1 (OLS bias bound)**.** *Under Assumptions 2.1–2.3, let $\hat{\kappa}^2 = \max(F - 1, 0) \times d_z/n$ be the plug-in estimate of SNR from the first-stage F-statistic. Then, with probability at least $1 - \delta_1$:*

$$|\mathrm{bias}_{\mathrm{OLS}}| \le B \quad \textit{where} \quad B = \frac{\hat{\sigma}_\eta}{\hat{\sigma}_\varepsilon}\left(\frac{1}{\hat{\kappa}^2 + 1} + \sqrt{\frac{2\log(2/\delta_1)}{n}}\right), \tag{4}$$

where $\hat{\sigma}_\eta$ and $\hat{\sigma}_\varepsilon$ are the OLS and first-stage residual standard deviations.

The ratio $\hat{\sigma}_\eta/\hat{\sigma}_\varepsilon$ ensures dimensional consistency: B has the same units as β regardless of the scale of Y and X. The first term inside the parentheses bounds $|\rho|/(\mathrm{SNR} + 1)$ (using $|\rho| \le 1$), and the second is a finite-sample correction that vanishes at rate $O(n^{-1/2})$. Complete proof is in Appendix B.2.

Weighted Estimator and Confidence Width. For any weight $\omega \in [0, 1]$, StAR confidence interval is

$$\mathrm{CI}_{\mathrm{StAR}} = \left[\hat{\beta}_\omega - h_\omega,\;\; \hat{\beta}_\omega + h_\omega\right], \tag{5}$$

with center and half-width

$$\hat{\beta}_\omega = \omega\,\hat{\beta}_{\mathrm{IV}} + (1 - \omega)\,\hat{\beta}_{\mathrm{OLS}}, \quad h_\omega = \omega\, w_{\mathrm{IV}} + (1 - \omega)(w_{\mathrm{OLS}} + B), \tag{6}$$

where $w_{\mathrm{IV}} = z_{\delta/2} \cdot \hat{\sigma}_{\mathrm{IV}}/\sqrt{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}}$ is the IV half-width, $w_{\mathrm{OLS}} = z_{\delta/2} \cdot \hat{\sigma}_{\mathrm{OLS}}/\sqrt{\boldsymbol{x}^\top \boldsymbol{x}}$ is the OLS half-width, and B is the bias correction from Proposition 3.1.

3.2. Coverage Guarantee

Theorem 3.2 (Coverage)**.** *Under Assumption 2.1–2.3, suppose $|\hat{\beta}_{\mathrm{IV}} - \beta^*| \le w_{\mathrm{IV}}$ with probability $\ge 1 - \delta_2$ and $|\hat{\beta}_{\mathrm{OLS}} - \beta^*| \le w_{\mathrm{OLS}} + B$ with probability $\ge 1 - \delta_1$. For* any *(possibly data-dependent) $\omega \in [0, 1]$,*

$$\mathbb{P}(\beta^* \in \mathrm{CI}_{\mathrm{StAR}}) \ge 1 - \delta_1 - \delta_2.$$

The power of the result is that the coverage holds whether ω is a fixed constant, a function of the F-statistic, or the specific e-process construction. The only requirements are– IV validity (at level δ_2) and OLS bias control (Proposition 3.1 at level δ_1).

Dimensional consistency and the $\sigma_\eta/\sigma_\varepsilon$ factor. The bias without the variance ratio, i.e., $B = 1/(\hat{\kappa}^2 + 1) + \sqrt{2\log(2/\delta_1)/n}$ is dimensionless. But B is added to w_{OLS} (which has units of β) in (6). This is valid only when $\sigma_\eta = \sigma_\varepsilon$, which is implicitly forced by the structural decomposition $\eta = \rho\varepsilon + \sqrt{1 - \rho^2}\varepsilon^\perp$ with all components sharing variance σ^2. For generic unstandardized data, the corrected bound $B = (\hat{\sigma}_\eta/\hat{\sigma}_\varepsilon) \cdot [1/(\hat{\kappa}^2 + 1) + \sqrt{2\log(2/\delta_1)/n}]$ ensures dimensional consistency without requiring pre-standardization of Y and X.

4. Data-Driven Weight via E-Processes

Because coverage holds for any ω, we are free to choose ω to minimize width. The MSE of the weighted estimator is exactly quadratic in ω, so a unique oracle weight ω^* exists—but it depends on unknowns (μ^2, ρ, σ_η^2). StAR replaces the oracle with an e-process diagnostic that adapts to the confounding and instrument strength from the observed data. This section develops three ingredients—the oracle target (Section 4.1), the e-process that estimates it (Section 4.2), and the weight that combines them (Section 4.3)—before stating the main finite-sample guarantees (Section 4.4).

4.1. Oracle MSE

Theorem 4.1 (MSE of StAR estimator)**.** *Under Assumption 2.1–2.3, MSE of $\hat{\beta}_\omega$ satisfies*

$$\mathbb{E}[(\hat{\beta}_\omega - \beta^*)^2 \mid Z] \;=\; \omega^2 \underbrace{\frac{\sigma_\eta^2}{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}}}_{\approx \mathrm{Var}(\hat{\beta}_{\mathrm{IV}})} \tag{7}$$

$$+ (1-\omega)^2 \underbrace{\left(\mathrm{Var}(\hat{\beta}_{\mathrm{OLS}}) + \mathrm{bias}^2_{\mathrm{OLS}}\right)}_{\mathrm{MSE}(\hat{\beta}_{\mathrm{OLS}})} + 2\omega(1-\omega)\, C_{\mathrm{cross}}, \tag{8}$$

where, $C_{\mathrm{cross}} = \mathbb{E}[(\hat{\beta}_{\mathrm{IV}} - \beta^*)(\hat{\beta}_{\mathrm{OLS}} - \beta^*) \mid Z]$. The unique minimizer, referred to as the oracle weight, is $\omega^* = \frac{\mathrm{MSE}(\hat{\beta}_{\mathrm{OLS}}) - C_{\mathrm{cross}}}{\mathrm{MSE}(\hat{\beta}_{\mathrm{IV}}) + \mathrm{MSE}(\hat{\beta}_{\mathrm{OLS}}) - 2C_{\mathrm{cross}}}$.

When instruments are strong ($\mathrm{Var}(\hat{\beta}_{\mathrm{IV}}) \ll \mathrm{MSE}(\hat{\beta}_{\mathrm{OLS}})$), $\omega^* \approx 1$. When instruments are weak ($\mathrm{Var}(\hat{\beta}_{\mathrm{IV}}) \gg \mathrm{MSE}(\hat{\beta}_{\mathrm{OLS}})$), $\omega^* \approx 0$. A useful weight must track ω^* without knowing the unknowns, such as ρ and κ. The e-process provides exactly this signal. MSE is illustrated in Figure 3.

4.2. E-Process for the IV Moment Condition

For each candidate value $\beta_0 \in \mathbb{R}$, define residuals $u_i(\beta_0) = y_i - \beta_0 x_i$ and the moment condition

$$H_0(\beta_0):\ \mathbb{E}[z(y - \beta_0 x)] = 0, \qquad H_1(\beta_0):\ \mathbb{E}[z(y - \beta_0 x)] \neq 0.$$

Under $H_0(\beta_0)$, the residuals are uncorrelated with every instrument; under $H_1(\beta_0)$, at least one instrument has predictive power for the residuals, indicating $\beta_0 \neq \beta^*$. We measure evidence against $H_0(\beta_0)$ using an *e-value*: a non-negative random variable E satisfying $\mathbb{E}_0[E] \leq 1$, so that by Ville's inequality $\mathbb{P}_0(E \geq 1/\alpha) \leq \alpha$ for any stopping time (Grünwald et al., 2024a; Ramdas et al., 2023).

We construct a multivariate e-value using the GROW (Growth-Rate Optimal in expectation) framework (Grünwald et al., 2024a). Let $\bar{z} = n^{-1}\sum_i z_i$, define centered instruments $\tilde{z}_i = z_i - \bar{z}$, the Gram matrix $\boldsymbol{V} = \sum_i \tilde{z}_i \tilde{z}_i^\top \in \mathbb{R}^{d_z \times d_z}$, and the test statistic vector $\boldsymbol{S} = \sum_i \tilde{z}_i (u_i - \bar{u}) \in \mathbb{R}^{d_z}$. With a Gaussian prior $\lambda \sim \mathcal{N}(0, \tau^2 I_{d_z})$ on the effect direction, the log e-value is:

$$\log E(\beta_0) = -\tfrac{1}{2}\log\det\!\big(I_{d_z} + \tau^2 \boldsymbol{V}/\hat{\sigma}_u^2\big) + \tfrac{\tau^2}{2\hat{\sigma}_u^2}\, \boldsymbol{S}^\top \big(\tau^2 \boldsymbol{V} + \hat{\sigma}_u^2 I_{d_z}\big)^{-1} \boldsymbol{S}, \tag{9}$$

where $\hat{\sigma}_u^2$ is the sample variance of $u_i(\beta_0)$ and $\tau^2 > 0$ is a prior variance parameter. The first term is a complexity penalty that grows with d_z, penalizing high-dimensional alternatives; the second is the evidence from the data. By construction, $\mathbb{E}_{\beta_0}[E(\beta_0)] \leq 1$ when $H_0(\beta_0)$ holds, so $E(\beta_0)$ is a valid e-value. When $d_z = 1$, the formula reduces to $\log E = -\frac{1}{2}\log(1 + \tau^2 V/\hat{\sigma}_u^2) + S^2/(2\hat{\sigma}_u^2(V + \hat{\sigma}_u^2/\tau^2))$. The multivariate construction is essential when instruments have heterogeneous strengths: it aggregates evidence from all d_z instruments simultaneously via the matrix-valued Bayes factor, detecting signal from any linear combination. A scalar e-value using only one instrument would miss the collective signal from instruments that are individually weak but jointly informative. This distinction is numerically reflected in Section **??**.

4.3. The StAR Weight

We evaluate (9) across a grid $\{\beta_0^{(1)}, \ldots, \beta_0^{(G)}\}$ and extract two summary statistics. The *peak e-value* $\lambda_{\max} = \max_j E(\beta_0^{(j)})$ measures aggregate instrument evidence. The *e-process confidence set* $\mathcal{C}_\alpha = \{\beta_0 : E(\beta_0) < 1/\alpha\}$ satisfies $\mathbb{P}(\beta^* \notin \mathcal{C}_\alpha) \leq \alpha$ by Ville's inequality; its width $W_{\mathrm{CS}} = \sup \mathcal{C}_\alpha - \inf \mathcal{C}_\alpha$ measures how precisely the e-process localizes β^* (we set $W_{\mathrm{CS}} = 0$ when $\mathcal{C}_\alpha = \emptyset$). Intuitively, $\lambda_{\max}$ grows with instrument strength μ^2 while W_{CS} shrinks as the instruments localize β^*. The formal analysis (Lemma C.1, Lemma C.4; Appendix C) establishes these properties non-asymptotically. The *Strength Adaptive Regression* (StAR) weight combines both diagnostics:

$$\omega = \underbrace{\frac{\lambda_{\max}}{\lambda_{\max} + c_0}}_{\text{signal strength}} \cdot \underbrace{\sqrt{\max\!\left(0,\ 1 - \frac{W_{\mathrm{CS}}}{w_{\max}}\right)}}_{\text{quality gate}}, \tag{10}$$

where $c_0 > 0$ and $w_{\max} > 0$ are tuning parameters (we use $c_0 = 20$, $w_{\max} = 10$ throughout). The weight has the desired tracking behavior: under weak instruments ($\mu^2 \ll 1$), $\lambda_{\max} = O(1)$ and $W_{\mathrm{CS}} \approx w_{\max}$, so both factors push $\omega \to 0$; under strong instruments ($\mu^2 \gg 1$), $\lambda_{\max}$ grows exponentially and $W_{\mathrm{CS}} \to 0$, so both factors push $\omega \to 1$. Proposition B.7 makes this precise: $\omega = O(n^{-1/2})$ for weak instruments, and $1 - \omega$ decays exponentially in μ^2/d_z for strong instruments.

Remark 4.2 (Why both components are necessary). The signal term alone is insufficient because a large $\lambda_{\max}$ can arise even when $W_{\mathrm{CS}} \approx w_{\max}$, i.e. when total evidence is high but not concentrated. This makes upweighting IV premature. The quality gate alone is insufficient because

W_{CS} can be small by chance in finite samples even when $\mu^2 \approx 0$. The product requires strong aggregate evidence *and* precise localization. We verify this empirically in Section F.3: the signal-only variant undercovers at moderate instrument strength, while the combined weight maintains $\geq 99\%$ coverage (Figure 8).

We plug the weight (10) into the CI (6) to obtain the StAR confidence interval. While coverage is guaranteed by Theorem 3.2, the results below characterize how *well* the proposed ω performs.

4.4. Main Guarantees: Oracle Inequality and Width

The first result says that StAR's MSE tracks the oracle.

Theorem 4.3 (Oracle inequality of StAR)**.** *Under Assumptions 2.1–2.3, with probability $\geq 1-\delta$, the StAR weight ω from (10) satisfies:*

$$\mathrm{MSE}(\hat{\beta}_\omega) \leq \mathrm{MSE}(\hat{\beta}_{\omega^*}) + r_n, \tag{11}$$

where ω^* is the oracle weight and $r_n = O(\log(n/\delta)/n)$ for bounded σ_η^2, ρ. The remainder is

$$r_n = \underbrace{\mathbb{E}[(\hat{\beta}_{\mathrm{IV}} - \hat{\beta}_{\mathrm{OLS}})^2 \mid Z]}_{\textit{curvature of MSE parabola}} \cdot \underbrace{(\omega - \omega^*)^2}_{\textit{weight deviation}}. \tag{12}$$

The two factors decay in complementary directions. For weak IVs, the curvature is large (IV is highly variable) but $(\omega - \omega^*)^2$ is small because both ω and ω^* are near 0. For strong instruments, the curvature is small because IV and OLS are closer. The product r_n is therefore small in both regimes.

The second result translates near-oracle MSE into a non-asymptotic width guarantee.

Theorem 4.4 (Width convergence rate)**.** *Under Assumption 2.1–2.3, with probability at least $1-\delta$:*

(a) ***Weak IVs*** *($\mu^2 \leq c_1 \log(1/\delta)$): $\omega \leq \epsilon_n \to 0$, and $h_\omega \leq w_{\mathrm{OLS}} + B + \epsilon_n w_{\mathrm{IV}} = O\left(\frac{\sigma}{\sqrt{n}} + \frac{1}{\kappa^2+1}\right)$.*

(b) ***Strong IVs*** *($\mu^2 \geq c_2\, d_z \log(n/\delta)$): $\omega \geq 1 - \epsilon'_n \to 0$ and $h_\omega \leq w_{\mathrm{IV}} + \epsilon'_n(w_{\mathrm{OLS}} + B) = O\left(\frac{\sigma\sqrt{d_z}}{\kappa\sqrt{n}}\right)$.*

In regime (a), $\lambda_{\max}$ remains near 1, so ω stays near 0 and the CI is effectively a bias-corrected OLS interval. The width is a finite constant even when $\kappa = 0$—in contrast with AR, which returns the full grid range (≈ 20) because no β_0 can be rejected, while StAR returns width ≈ 3, a $6\times$ improvement. In regime (b), ω concentrates near 1 and the CI converges to the efficient IV width. The rate $O(\sigma\sqrt{d_z}/(\kappa\sqrt{n}))$ matches 2SLS up to constants, confirming that StAR pays no asymptotic price for adaptivity. The transition is smooth and data-driven, governed by $\mu^2 = n\kappa^2$. Unlike pre-testing approaches that make a binary strong-or-weak decision at a fixed threshold (Stock and Yogo, 2005), the e-process weight interpolates continuously, and thus, produces a CI that reflects uncertainty regarding instrument strength. The full proofs are in Appendix B.5–B.6.

Algorithm 1 Strength Adaptive Regression (StAR)

0: Data $(z_i, x_i, y_i)_{i=1}^n$, confidence level $1-\delta$, parameters $\tau^2, c_0, w_{\max}$
0: Whiten instruments (if $d_z > 1$); compute 2SLS ($\hat{\beta}_{\mathrm{IV}}$, w_{IV}) and OLS ($\hat{\beta}_{\mathrm{OLS}}, w_{\mathrm{OLS}}$)
0: Bias correction $B \leftarrow (\hat{\sigma}_\eta/\hat{\sigma}_\varepsilon)[1/(\hat{\kappa}^2+1) + \sqrt{2\log(2/\delta_1)/n}]$ {Prop. 3.1}
0: Evaluate GROW e-value (9) on grid $\mathcal{B}$; extract $\lambda_{\max}$, W_{CS}
0: Weight $\omega \leftarrow$ (10)
0: **Return** $[\hat{\beta}_\omega - h_\omega,\ \hat{\beta}_\omega + h_\omega]$ via Equation (6) =0

StAR: Computation Scheme. We set $\alpha = \delta$ (the CI confidence level) for the e-process threshold, aligning the quality gate's resolution with the target coverage.[1] The e-value grid is $\mathcal{B} = \{b_{\min}, b_{\min}+\Delta, \ldots, b_{\max}\}$ with spacing $\Delta = 0.05$ and range $[\hat{\beta}_{\mathrm{OLS}} - 10,\ \hat{\beta}_{\mathrm{OLS}} + 10]$ ($G = 401$ points), setting $w_{\max} = b_{\max} - b_{\min} = 20$.[2] The error budget splits as $\delta_1 = \delta/4$ (bias bound), $\delta_2 = \delta/4$ (e-process), $\delta_{\mathcal{G}} \leq \delta/4$ (Gram concentration), with total failure probability $\leq \delta$ (Remark E.5). When $d_z > 1$, the instruments are whitened: $\tilde{z}_i = \hat{\Sigma}_z^{-1/2}(z_i - \bar{z})$.

Computational complexity. The total cost is $O(nd_z^2 + nd_zG)$: a first-stage regression ($O(nd_z^2)$), OLS ($O(n)$), one eigendecomposition of $\boldsymbol{V}$ ($O(d_z^3)$), and e-value evaluation over G grid points ($O(nd_zG)$ with the precomputed eigenbasis). This is comparable to AR that runs a test over the same grid.

5. Experimental Analysis

We evaluate StAR against four baselines: 2SLS (DV) (Vecchia and Basu, 2025), Anderson–Rubin (AR) (Anderson and Rubin, 1949), CLR test (Moreira, 2003), and the finite-sample method of Xia et al. (2024). All intervals are at the 95% nominal level. Implementation details and additional results are in Appendix F.

I. Simulation Studies. We generate data from the linear IV

[1] One might consider the optimized weight $\omega_{\max} = \max_\alpha \omega(\alpha)$. We avoid this because maximizing over α introduces a multiple-testing correction that leads to premature upweighting of the IV arm. Since coverage holds for any fixed ω (Theorem 3.2), the adaptive weight need only be reasonable, not optimal.

[2] Any $w_{\max} \in [2w_{\mathrm{IV}},\ b_{\max} - b_{\min}]$ gives qualitatively identical results; the theory requires only $w_{\max} = \Theta(1)$.

model (Stage-2)–(Stage-1) with $n = 200$, $d_z = 3$, across a 12×10 grid of instrument strength $\kappa \in \{0, 0.01, \ldots, 0.30\}$ and confounding $\rho \in \{0, 0.1, \ldots, 0.9\}$, with 1000 replications per cell (120 configurations, 120,000 total runs). Figure 4 summarizes the results. StAR maintains empirical coverage $\geq 96\%$ across all configurations, while DV drops below 50% at weak instruments with strong confounding. At weak instruments ($\kappa \leq 0.05$), AR and CLR achieve valid coverage but with median width ≈ 20 (the grid range); StAR achieves width ≈ 3, a $6\times$ improvement. At strong instruments ($\kappa \geq 0.20$), all valid methods converge to similar width, confirming that StAR pays no efficiency cost for adaptivity. The transition region ($\kappa \in [0.05, 0.15]$) is where StAR's advantage is most pronounced: it produces the tightest valid interval in over 80% of cells in this range (Figure 6).

II. OHIE Experiments. We apply StAR to the Oregon Health Insurance Experiment (OHIE) (Finkelstein et al., 2012), one of the most important modern randomized encouragement designs, and a canonical setting where weak instruments arise from partial compliance. Figure 9 shows results at $n = 5000$ with moderate confounding ($\delta = 0.5$). At low compliance (0.20, $\bar{F} \approx 3$), the instrument is firmly weak: AR gives vacuous CIs (width 20) while StAR maintains 100% coverage with width 2.02. At moderate compliance (0.50, $\bar{F} \approx 14$), AR undercoverage begins (93.2%) while StAR achieves 100% coverage with width 2.02. At strong instruments (compliance ≥ 1.0, $\bar{F} \geq 52$), StAR converges to IV-efficient width (0.29 at compliance 1.0), matching DV. The confounding sweep (Table 12) shows that maintains $\geq 96.8\%$ coverage across all 15 configurations, while AR systematically undercovers at strong instruments due to grid discretization. This is precisely the regime encountered in OHIE subgroup analyses, where practitioners must choose between unreliable 2SLS estimates and uninformative AR intervals. StAR provides the missing middle ground.

III. MIMIC-IV: Mechanical Ventilation and ICU Mortality. We estimate the causal effect of mechanical ventilation (MV) on in-hospital mortality using MIMIC-IV (Johnson et al., 2023), a critical care database containing 94,458 ICU stays. The endogeneity concern is that sicker patients are more likely to receive MV *and* more likely to die. We use physician prescribing preference—the within-careunit leave-one-out MV rate—as an instrument (Brookhart et al., 2006), controlling for age, sex, ICU length of stay, and care-unit fixed effects. After restricting to providers with ≥ 10 admissions and keeping each patient's first ICU stay, the analysis cohort contains $n = 63{,}998$ stays across 646 physicians. The Medical ICU subgroup ($n = 21{,}762$) presents a genuinely weak instrument setting—physicians in medical ICUs have less variation in ventilation preference after controlling for patient severity. Here the methods diverge sharply (Table 2). DV produces $[0.08, 2.00]$ (width 1.91). AR and CLR both give $[0.37, 5.24]$ (width 4.87), excluding zero and suggesting MV *increases* mortality. Xia produces a vacuous interval spanning 166 units. StAR gives $[-0.80, 1.22]$ (width 2.02)—58% tighter than AR/CLR, appropriately centered near zero, and reflecting genuine uncertainty about the treatment effect direction at $F = 5.66$. The clinical interpretation matters: a researcher using AR in this subgroup would conclude that MV is harmful, a finding that is likely an artifact of weak instruments rather than a genuine clinical effect. StAR appropriately reports uncertainty. We provide semi-synthetic validation calibrated to MIMIC-IV summary statistics (Figure 10, Appendix F.5) and a subsample bootstrap analysis confirming $\geq 99\%$ coverage across sample sizes (Table 14, Appendix F.5).

Table 2. Real MIMIC-IV: Medical ICU subgroup ($n = 21{,}762$, $F = 5.66$).

Method	Lower	Upper	Width
DV (2SLS)	0.083	1.997	1.914
AR	0.370	5.240	**4.870**
CLR	0.370	5.240	**4.870**
Xia	-82.2	84.3	166.4
StAR	-0.803	1.218	**2.021**

Table 3. Angrist and Krueger (1991) returns to schooling ($n = 329509$, $d_z = 3$, $F = 32.3$).

Method	Lower	Upper	Width
StAR (multivariate, $\omega = 0.99$)	0.036	0.174	0.139
StAR (scalar, $\omega = 0$)	-1.344	1.486	2.830
DV (2SLS)	0.066	0.145	0.079
AR	0.064	0.152	0.088
CLR	0.068	0.144	0.076
Xia	-3.117	3.327	6.444

6. Discussion, Limitations, and Future Works

We discuss the practical scope of StAR, highlighting the settings where it excels and the design trade-offs that shape its guarantees. *1. Complementarity with unconditional methods.* StAR provides *conditional* coverage (Theorem 3.2), requiring both the IV arm and the bias bound to hold. Anderson–Rubin and the conditional likelihood ratio test (Moreira, 2003) provide *unconditional* finite-sample coverage under homoskedasticity, making them preferable when undercoverage is catastrophic (e.g., regulatory approval, safety-critical decisions). StAR targets a different regime: applied settings where tighter intervals justify a mild structural assumption, and where an uninformatively wide or unbounded AR interval provides little actionable guidance. *2. Principled conservatism in the bias bound.* The bias bound $B = 1/(\hat{\kappa}^2 + 1) + O(n^{-1/2})$ assumes worst-case confounding $|\rho| \leq 1$, ensuring validity without estimating nuisance parameters. When the true $|\rho|$ and instrument strength are both moderate ($F \approx 10 - -20$), this widens the interval

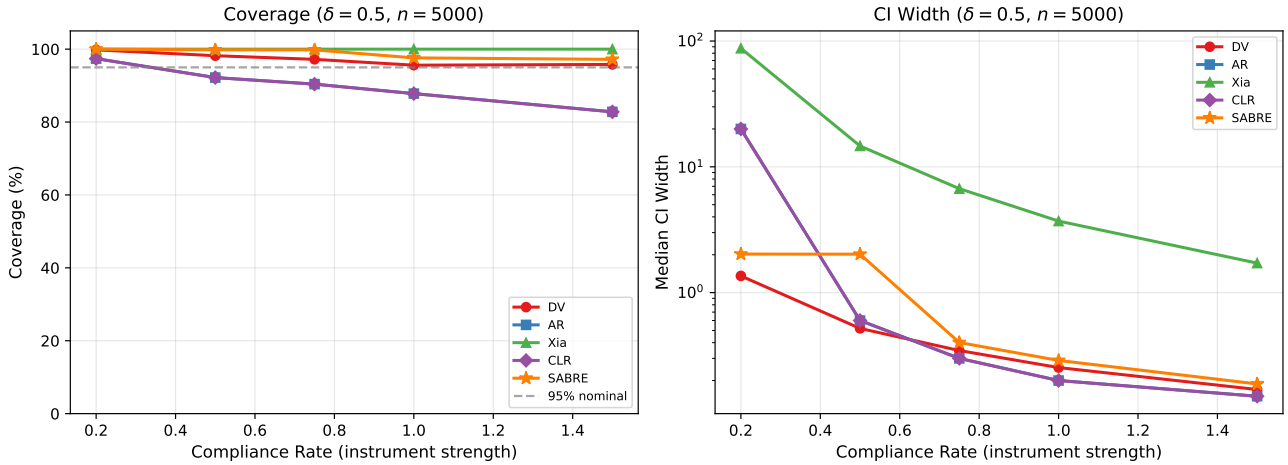

Figure 2. Oregon Health Insurance Experiment, semi-synthetic ($n = 5000$, 500 reps). Left: empirical coverage. Right: median CI width (log scale). Dashed vertical line marks the true OHIE compliance rate ($\sim$0.25). StAR provides valid, informative CIs across the full compliance range.

beyond what a direct estimate of ρ would require. StAR's advantage is most pronounced at the extremes—very weak ($F < 5$) or very strong ($F > 50$) instruments—precisely where existing methods break down or are unnecessarily wide. Incorporating user-supplied bounds on $|\rho|$ is a natural extension that would tighten intervals in the moderate regime. *3. Grid discretization.* The e-process weight ω is evaluated on a grid of candidate β_0 values. At strong instruments, a coarse grid can make the e-process confidence set artificially wide, suppressing ω below its oracle value—the same artifact that explains AR's mild conservatism at strong instruments. Adaptive grid refinement around the IV point estimate recovers near-oracle performance at modest computational cost. *4. A practical hybrid.* StAR combines e-process statistics (non-parametric, anytime-valid) with normal-based CIs (parametric, fixed-sample). A fully e-process CI would inherit stronger guarantees but at the cost of substantially wider intervals. StAR leverages e-processes for instrument-strength detection while retaining parametric efficiency for interval construction—a trade-off we view as favorable for the majority of applied IV settings.

References

Anderson, T. W. and Rubin, H. (1949). Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of mathematical statistics*, 20(1):46–63.

Andrews, I., Stock, J. H., and Sun, L. (2019). Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, 11(1):727–753.

Angrist, J. D. and Imbens, G. W. (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American Statistical Association*, 90(430):431–442.

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434):444–455.

Angrist, J. D. and Krueger, A. B. (1991). Does compulsory school attendance affect schooling and earnings? *The quarterly journal of economics*, 106(4):979–1014.

Armstrong, T. B. and Kolesár, M. (2018). Optimal inference in a class of regression models. *Econometrica*, 86(2):655–683.

Bound, J., Jaeger, D. A., and Baker, R. M. (1995). Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American statistical association*, 90(430):443–450.

Brookhart, M. A., Wang, P., Solomon, D. H., and Schneeweiss, S. (2006). Evaluating short-term drug effects using a physician-specific prescribing preference as an instrumental variable. *Epidemiology*, 17(3):268–275.

Burgess, S., Thompson, S. G., and Collaboration, C. C. G. (2011). Avoiding bias from weak instruments in mendelian randomization studies. *International journal of epidemiology*, 40(3):755–764.

de la Peña, V. H., Lai, T. L., and Shao, Q.-M. (2009). *Self-Normalized Processes: Limit Theory and Statistical Applications*. Springer.

Dufour, J.-M. (1997). Some impossibility theorems in econometrics with applications to structural and dynamic models. *Econometrica*, 65(6):1365–1387.

Finkelstein, A., Taubman, S., Wright, B., Bernstein, M., Gruber, J., Newhouse, J. P., Allen, H., Baicker, K., and

Oregon Health Study Group (2012). The Oregon Health Insurance experiment: Evidence from the first year. *The Quarterly Journal of Economics*, 127(3):1057–1106.

Grünwald, P., de Heide, R., and Koolen, W. M. (2024a). Safe testing. *Journal of the Royal Statistical Society: Series B*, 86(5):1091–1128.

Grünwald, P., Hao, Y., and Balsubramani, A. (2024b). Growth-optimal e-variables and an extension to the multivariate Csiszár–Sanov–Chernoff theorem. *arXiv preprint arXiv:2412.17554*.

Hansen, B. E. (2007). Least squares model averaging. *Econometrica*, 75(4):1175–1189.

Howard, S. R., Ramdas, A., McAuliffe, J., and Sekhon, J. (2021). Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080.

James, W. and Stein, C. (1961). Estimation with quadratic loss. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 361–379.

Johnson, A. E. W., Bulgarelli, L., Shen, L., Gayles, A., Shammout, A., Horng, S., Pollard, T. J., Moody, B., Gow, B., Lehman, L.-w. H., Celi, L. A., and Mark, R. G. (2023). MIMIC-IV, a freely accessible electronic health record dataset. *Scientific Data*, 10(1):1.

Kallus, N. and Uehara, M. (2022). Efficiently breaking the curse of horizon in off-policy evaluation with double reinforcement learning. *Operations Research*, 70(6):3282–3302.

Liao, L., Fu, Z., Yang, Z., Wang, Y., Ma, D., Kolar, M., and Wang, Z. (2024). Instrumental variable value iteration for causal offline reinforcement learning. *Journal of Machine Learning Research*, 25(303):1–56.

Moreira, M. J. (2003). A conditional likelihood ratio test for structural models. *Econometrica*, 71(4):1027–1048.

Nagar, A. L. (1959). The bias and moment matrix of the general k-class estimators of the simultaneous equation coefficients. *Econometrica*, 27(4):575–595.

Ramdas, A., Grünwald, P., Vovk, V., and Shafer, G. (2023). Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601.

Staiger, D. and Stock, J. H. (1997). Instrumental variables regression with weak instruments. *Econometrica*, 65(3):557–586.

Stock, J. H. and Yogo, M. (2005). Testing for weak instruments in linear IV regression. In Andrews, D. W. K. and Stock, J. H., editors, *Identification and Inference for Econometric Models*, chapter 5. Cambridge University Press.

Taubman, S. L., Allen, H. L., Wright, B. J., Baicker, K., and Finkelstein, A. N. (2014). Medicaid increases emergency-department use: Evidence from Oregon's Health Insurance experiment. *Science*, 343(6168):263–268.

ter Schure, J. and Grünwald, P. (2022). ALL-IN meta-analysis: Breathing life into living systematic reviews. *F1000Research*, 11:549.

Vecchia, R. D. and Basu, D. (2025). Stochastic online instrumental variable regression: Regrets for endogeneity and bandit feedback. In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'25. AAAI Press.

Vershynin, R. (2018). *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press.

Vovk, V. and Wang, R. (2021). E-values: Calibration, combination and applications. *The Annals of Statistics*, 49(3):1736–1754.

Waudby-Smith, I. and Ramdas, A. (2024). Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society: Series B*, 86(1):1–27.

Xia, E., Wainwright, M. J., and Newey, W. (2024). Instrumental variables: A non-asymptotic viewpoint. *arXiv preprint arXiv:2410.02015*.

Part I

Appendix

A. Related Work

Weak instrument inference. The weak instrument problem was identified by Bound et al. (1995) and formalized by Staiger and Stock (1997). The Anderson–Rubin test (Anderson and Rubin, 1949) provides the gold standard for weak-instrument-robust inference, with exact finite-sample size under homoskedasticity. Moreira (2003) developed the conditional likelihood ratio (CLR) test, which is asymptotically optimal among similar tests and provides valid CIs under weak instruments. Stock and Yogo (2005) proposed the concentration parameter framework and weak-instrument pre-tests. Andrews et al. (2019) developed adaptive CIs that are asymptotically efficient at strong instruments while remaining valid at weak instruments. StAR differs in mechanism: rather than pre-testing or conditioning on sufficient statistics, it continuously adapts via e-process weights, providing a non-asymptotic alternative. We emphasize that CLR and Andrews–Kolesár provide *unconditional* validity guarantees, whereas StAR's coverage is conditional on explicit bias-control events (Theorem 3.2). The practical advantage of StAR is its simplicity and adaptivity mechanism; its limitation relative to CLR is the conditional nature of its guarantee.

Non-asymptotic IV inference. Xia et al. (2024) developed finite-sample confidence intervals for IV regression that account for both estimation variance and weak-instrument bias. Their approach inflates the standard error and adds an explicit bias bound. StAR shares the finite-sample validity goal but achieves tighter intervals by adaptively weighting toward OLS at weak instruments rather than uniformly inflating the IV-based interval.

E-processes and safe testing. E-values and e-processes provide a non-asymptotic framework for hypothesis testing with optional stopping guarantees. Vovk and Wang (2021) established the foundational theory of e-values—their calibration properties, the fact that they can be combined by simple averaging, and their connection to Bayes factors and likelihood ratios. Grünwald et al. (2024a) introduced the GROW (Growth-Rate Optimal in expectation) criterion for constructing e-values with maximal expected log-growth under the alternative; Grünwald et al. (2024b) extended GROW to multivariate composite alternatives via a Csiszár–Sanov–Chernoff generalization. Howard et al. (2021) developed non-asymptotic confidence sequences from nonnegative supermartingales, providing the template for anytime-valid confidence sets that shrink at $\sqrt{n}$ rates—the same mechanism underlying our e-process confidence set $\mathcal{C}_\alpha$. Waudby-Smith and Ramdas (2024) made the betting framework practical for mean estimation with tight empirical-Bernstein bounds, and ter Schure and Grünwald (2022) demonstrated that GROW e-values can serve as adaptive diagnostics in sequential clinical trial meta-analysis. We build on this literature in two ways: first, we construct a multivariate GROW e-value with an isotropic Gaussian prior that aggregates evidence across all d_z instruments via a matrix-valued Bayes factor; second, we use the e-process not for testing per se, but as a *weight diagnostic*—the peak e-value and CS width drive the adaptive blend between IV and OLS, a role that has no precedent in the e-value literature.

Adaptive estimation and shrinkage. The idea of blending biased and unbiased estimators has a long history, from James–Stein estimation (James and Stein, 1961) to model averaging (Hansen, 2007). Armstrong and Kolesár (2018) study optimal inference under partial identification. StAR builds on this tradition by using e-process statistics to drive the adaptive weight, providing a non-asymptotic mechanism with explicit bias control. The approach is complementary to, rather than a replacement for, methods with unconditional guarantees: StAR prioritizes practical width reduction at the cost of conditional coverage.

Instrumental variables in policy evaluation. IV methods have been applied to off-policy evaluation to handle unmeasured confounding (Kallus and Uehara, 2022; Liao et al., 2024). More broadly, encouragement designs—such as the Oregon Health Insurance Experiment (Finkelstein et al., 2012)—produce naturally weak instruments due to partial compliance. In these settings, weak-instrument-robust inference is essential, and StAR provides a practical compromise between anticonservative 2SLS and vacuous AR intervals, with the caveat that its guarantees are conditional rather than unconditional.

B. Proofs

B.1. Notation, Conventions, and Constants

We write c, c', C, C' for positive constants that may depend on σ, R, τ^2 but not on $n, d_z, \kappa, \rho, \delta$. Unlike in the main text, *each constant retains its value within a single proof*; we number them $(c_1, c_2, \ldots)$ when multiple constants appear in the same argument.

Table 4. Comparison of inference methods for linear IV regression. "Adaptive" indicates whether the method automatically adjusts to instrument strength (κ) and/or confounding (ρ). "Coverage" reports the finite-sample guarantee type. "Width at weak IV" describes interval behavior when $\mu^2 = O(1)$.

Method	Adaptive to		Approach		Coverage	Width at weak IV	Non-asymp. analysis
	κ	ρ	Regr.	Test			
OLS / Wald	✗	✗	✓	✗	Invalid ($\to 0$)	$O(n^{-1/2})$ (undercoverage)	✗
2SLS (Vecchia and Basu, 2025)	✗	✗	✓	✗	Invalid when μ^2 small	$O(n^{-1/2})$ (undercoverage)	✗
AR (Anderson and Rubin, 1949)	✗	✗	✗	✓	Exact (unconditional)	Vacuous (≈ 20)	✓
CLR (Moreira, 2003)	✗	✗	✗	✓	Exact (unconditional)	Vacuous (≈ 20)	✓
Stock–Yogo (Stock and Yogo, 2005)	binary	✗	✓	✗	Pre-test; invalid if test errs	Not defined (rejects IV)	✗
AK (Andrews et al., 2019)	✓	✗	✗	✓	Asymptotic (unconditional)	Optimal asymptotic	✗
(Xia et al., 2024)	✗	✗	✓	✗	Finite-sample (conditional)	Very wide ($\gg 20$)	✓
StAR (this work)	✓	✓	blend		Finite-sample (conditional)	$O\left(\frac{\sigma}{\sqrt{n}} + \frac{1}{\kappa^2+1}\right)$	✓

All probability and expectation statements carry an explicit subscript indicating the distribution: $\mathbb{P}_{H_0}$ and $\mathbb{E}_{H_0}$ when $\pi = 0$; $\mathbb{P}_{H_1}$ and $\mathbb{E}_{H_1}$ when $\pi \neq 0$. When we condition on Z, we write $\mathbb{P}_{H_1}(\cdot \mid Z)$.

Recall the eigendecomposition $\boldsymbol{V} = Q\Lambda Q^\top$ with eigenvalues $\lambda_1 \geq \cdots \geq \lambda_{d_z} > 0$, rotated statistics $\boldsymbol{S}_{\text{rot}} = Q^\top \boldsymbol{S}$, and the shorthand $\mu^2 = n\|\pi\|^2/d_z$.

After sample centering and whitening (Remark **??**), we work with transformed instruments $\tilde{z}_i = \hat{\Sigma}_z^{-1/2}(z_i - \bar{z})$, and define $\tilde{R} = 2R/\sqrt{\lambda_{\min}(\Sigma_z)}$ as the effective bound on $\|\tilde{z}_i\|_2$. All results below hold with R replaced by $\tilde{R}$; we write R for brevity when no confusion arises.

Table of constants. Table 5 collects all named constants that appear in the proofs, with their definitions and the results where they first arise. Constants marked "universal" do not depend on problem parameters; those marked "problem" depend on σ, R, τ^2 but not on $n, d_z, \kappa, \rho, \delta$; those marked "tuning" are user-chosen.

B.2. Proof of Proposition 3.1

Proof. We bound $\text{bias}_{\text{OLS}} = \mathbb{E}\left[\frac{\boldsymbol{x}^\top \boldsymbol{\eta}}{\boldsymbol{x}^\top \boldsymbol{x}} \,\middle|\, Z\right]$ with explicit constants.

Probability structure. The proof has two layers:

- *Outer layer (over Z):* $\mathcal{G}(\epsilon_0)$ holds with probability $\geq 1 - \delta_1/2$ (Corollary 14). We fix any $Z \in \mathcal{G}(\epsilon_0)$ for the rest.
- *Inner layer (over $\boldsymbol{\varepsilon} \mid Z$):* Steps 1–3 are deterministic given Z (they bound conditional expectations). Step 4 uses Lemma D.2 with failure probability $\delta_1/4$ over $\boldsymbol{\varepsilon} \mid Z$.

Total: $\delta_1/2 + \delta_1/4 + \text{negl.} \leq \delta_1$.

Step 1: Expectations (any Z). $\mathbb{E}[\boldsymbol{x}^\top \boldsymbol{\eta} \mid Z] = n\rho\sigma_\varepsilon\sigma_\eta$ and $D_0 := \mathbb{E}[\boldsymbol{x}^\top \boldsymbol{x} \mid Z] = \|Z\pi\|^2 + n\sigma_\varepsilon^2$.

These hold for any Z — they use only $\mathbb{E}[\eta \mid Z] = 0$ (Assumption 1), $\mathbb{E}[\varepsilon_i \eta_i] = \rho\sigma_\varepsilon\sigma_\eta$, and $\mathbb{E}[\|\boldsymbol{\varepsilon}\|^2] = n\sigma_\varepsilon^2$.

Step 2: Population bias ratio (*needs $\mathcal{G}$ for lower bound on D_0*). Decompose $\|Z\pi\|^2 = \pi^\top \boldsymbol{V}\pi + n(\bar{z}^\top\pi)^2$ (cross term vanishes since $\sum_i \tilde{z}_i = 0$).

Lower bound (dropping $(\bar{z}^\top\pi)^2 \geq 0$):

$$\|Z\pi\|^2 \geq \pi^\top \boldsymbol{V}\pi \geq n(1-\epsilon_0)\|\pi\|^2 \qquad \text{on } \mathcal{G}(\epsilon_0). \tag{13}$$

Define $\mathrm{SNR} = \|\pi\|^2/\sigma_\varepsilon^2$. Then $D_0 \geq n((1-\epsilon_0)\|\pi\|^2 + \sigma_\varepsilon^2) \geq n(1-\epsilon_0)(\|\pi\|^2 + \sigma_\varepsilon^2) = n(1-\epsilon_0)\sigma_\varepsilon^2(\mathrm{SNR}+1)$, giving:

$$|R_0| := \frac{n|\rho|\sigma_\varepsilon\sigma_\eta}{D_0} \leq \frac{\sigma_\eta}{\sigma_\varepsilon} \cdot \frac{|\rho|}{(1-\epsilon_0)(\mathrm{SNR}+1)}. \tag{14}$$

Step 3: Ratio concentration (*needs $\mathcal{G}$ for $\mathrm{Var}(U)$ upper bound*). Write $U = \boldsymbol{x}^\top\boldsymbol{\eta} - R_0\boldsymbol{x}^\top\boldsymbol{x}$ (so $\mathbb{E}[U \mid Z] = 0$) and split on $\mathcal{F} = \{B \geq D_0/2\}$:

$$|\mathrm{bias}_{\mathrm{OLS}} - R_0| \leq \frac{2\,\mathbb{E}[|U| \mid Z]}{D_0} + \mathbb{E}\left[\frac{|U|}{B}\mathbf{1}_{\mathcal{F}^c} \;\middle|\; Z\right].$$

Bounding $\mathrm{Var}(U \mid Z)$. $U = \sum_i U_i$ with $U_i = x_i\eta_i - R_0 x_i^2$, independent across i given Z. The moment bounds $\mathbb{E}[x_i^2\eta_i^2 \mid z_i] \leq C\sigma_\eta^2((\pi^\top z_i)^2 + \sigma_\varepsilon^2)$ and $\mathbb{E}[x_i^4 \mid z_i] \leq C((\pi^\top z_i)^2 + \sigma_\varepsilon^2)^2$ involve $\sum_i(\pi^\top z_i)^2 = \|Z\pi\|^2$.

For the *upper* bound we need:

$$\|Z\pi\|^2 = \pi^\top\boldsymbol{V}\pi + n(\bar{z}^\top\pi)^2 \leq n(1+\epsilon_0)\|\pi\|^2 + nR^2\|\pi\|^2. \tag{15}$$

The first term uses $\mathcal{G}(\epsilon_0)$ (upper eigenvalue); the second uses $\|\bar{z}\| \leq R$ (deterministic, Assumption 3). Since R is a fixed constant, the R^2 is absorbed:$\mathrm{Var}(U \mid Z) \leq C'\sigma_\eta^2 D_0$, where C' depends on R but not on $n, \pi, \sigma_\varepsilon, \sigma_\eta$.

Key point: Step 2 used the *lower* eigenvalue bound (to make $|R_0|$ small); this step uses the *upper* bound (to make $\mathrm{Var}(U)$ small). Both directions of $\mathcal{G}(\epsilon_0)$ are needed, plus the deterministic $\|\bar{z}\| \leq R$ for the upper bound.

By Jensen, $\mathbb{E}[|U| \mid Z] \leq \sqrt{\mathrm{Var}(U \mid Z)}$:

$$\frac{2\,\mathbb{E}[|U| \mid Z]}{D_0} \leq \frac{2\sqrt{C'\sigma_\eta^2 D_0}}{D_0} = \frac{2\sqrt{C'}\,\sigma_\eta}{\sqrt{D_0}} \leq \frac{2\sqrt{C'}\,\sigma_\eta}{\sigma_\varepsilon\sqrt{n}} = \frac{\sigma_\eta}{\sigma_\varepsilon} \cdot \frac{2\sqrt{C'}}{\sqrt{n}}. \tag{16}$$

Tail event. $\mathbb{P}(\mathcal{F}^c \mid Z) \leq \exp(-cn/(R^4(\mathrm{SNR}+1)))$, negligible for $n \geq CR^4(\mathrm{SNR}+1)\log(1/\delta_1)$. By Cauchy–Schwarz the tail term is $O(e^{-cn})$.

Combining: $|\mathrm{bias}_{\mathrm{OLS}} - R_0| \leq (\sigma_\eta/\sigma_\varepsilon) \cdot C_0/\sqrt{n}$.

Step 4: Plug-in $\hat{\kappa}^2 \approx \mathrm{SNR}$ (*needs $\mathcal{G}$; probability over $\boldsymbol{\varepsilon}$*). By Lemma D.2 (for the *same* $Z \in \mathcal{G}(\epsilon_0)$ already fixed), with probability $\geq 1 - \delta_1/4$ over $\boldsymbol{\varepsilon} \mid Z$:

$$|\hat{\kappa}^2 - \mathrm{SNR}\,| \leq \Delta_\kappa = O\left(\sqrt{\frac{\mathrm{SNR}\log(1/\delta_1)}{n}}\right).$$

Since $f(s) = |\rho|/(s+1)$ is 1-Lipschitz on $[0, \infty)$:

$$\left|\frac{|\rho|}{\mathrm{SNR}+1} - \frac{|\rho|}{\hat{\kappa}^2+1}\right| \leq \Delta_\kappa.$$

Step 5: Combining. On $\mathcal{G}(\epsilon_0)$ (prob $\geq 1 - \delta_1/2$) $\cap \{|\hat{\kappa}^2 - \mathrm{SNR}\,| \leq \Delta_\kappa\}$ (conditional prob $\geq 1 - \delta_1/4$):

$$|\mathrm{bias}_{\mathrm{OLS}}| \leq |R_0| + \frac{\sigma_\eta}{\sigma_\varepsilon} \cdot C_0/\sqrt{n} \tag{Step 3}$$

$$\leq \frac{\sigma_\eta}{\sigma_\varepsilon} \cdot \frac{|\rho|}{(1-\epsilon_0)(\mathrm{SNR}+1)} + \frac{\sigma_\eta}{\sigma_\varepsilon} \cdot C_0/\sqrt{n} \tag{Step 2}$$

The factor $\sigma_\eta/\sigma_\varepsilon$ multiplies the entire bound, ensuring correct units ($[\beta]$). □

B.3. Proof of Theorem 4.1

Inner and outer expectation (any Z)

These steps are unchanged from the original proof and hold for any Z with $\mathrm{rank}(Z) = d_z$.

Inner expectation. Condition on $(Z, \boldsymbol{\varepsilon})$, which fixes $\hat{\boldsymbol{x}} = Z\pi + P_Z\boldsymbol{\varepsilon}$. Decompose $\boldsymbol{\eta} = \rho\boldsymbol{\varepsilon} + \sqrt{1-\rho^2}\,\boldsymbol{\zeta}$ where $\boldsymbol{\zeta} \perp \boldsymbol{\varepsilon}$, $\sigma_\zeta^2 = (1-\rho^2)\sigma^2$. Then:

$$\mathbb{E}[\Delta_{\mathrm{IV}}^2 \mid Z, \boldsymbol{\varepsilon}] = \frac{\rho^2(\hat{\boldsymbol{x}}^\top\boldsymbol{\varepsilon})^2}{(\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}})^2} + \frac{\sigma_\zeta^2}{\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}}. \tag{17}$$

The first term is the squared Nagar bias; the second is the pure variance from the independent component $\boldsymbol{\zeta}$.

Outer expectation. Taking $\mathbb{E}[\cdot \mid Z]$ over $\boldsymbol{\varepsilon}$:

$$\mathrm{MSE}_{\mathrm{IV}}(Z) = \rho^2 \underbrace{\mathbb{E}\left[\frac{(\hat{\boldsymbol{x}}^\top\boldsymbol{\varepsilon})^2}{(\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}})^2} \;\middle|\; Z\right]}_{T_1} + \sigma_\zeta^2 \underbrace{\mathbb{E}\left[\frac{1}{\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}} \;\middle|\; Z\right]}_{T_2}. \tag{18}$$

Both T_1 and T_2 are deterministic functions of Z. We bound each using the ratio-concentration technique from D.2 (Step 3).

Signal decomposition

Needs $\mathcal{G}$ for $\boldsymbol{V}$; deterministic R for $\bar{z}$.

Write $\hat{\boldsymbol{x}}^\top\boldsymbol{\varepsilon} = (Z\pi)^\top\boldsymbol{\varepsilon} + \boldsymbol{\varepsilon}^\top P_Z\boldsymbol{\varepsilon}$ and $\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}} = \|Z\pi\|^2 + 2\pi^\top Z^\top\boldsymbol{\varepsilon} + \boldsymbol{\varepsilon}^\top P_Z\boldsymbol{\varepsilon}$, using $Z^\top P_Z = Z^\top$.

Define the deterministic quantities:

$$D_x \coloneqq \mathbb{E}[\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}} \mid Z] = \|Z\pi\|^2 + d_z\sigma^2, \tag{19}$$

and decompose the signal:

$$\|Z\pi\|^2 = \pi^\top\boldsymbol{V}\pi + n(\bar{z}^\top\pi)^2 \qquad \text{(cross term vanishes since } \sum_i \tilde{z}_i = 0). \tag{20}$$

Lower bound (needs $\mathcal{G}$): Drop $(\bar{z}^\top\pi)^2 \geq 0$:

$$\|Z\pi\|^2 \geq \pi^\top\boldsymbol{V}\pi \geq n(1-\epsilon_0)\|\pi\|^2 \qquad \text{on } \mathcal{G}(\epsilon_0). \tag{21}$$

Upper bound (needs $\mathcal{G}$ for $\boldsymbol{V}$; deterministic R for $\bar{z}$):

$$\|Z\pi\|^2 \leq n(1+\epsilon_0)\|\pi\|^2 + nR^2\|\pi\|^2 = n(1+\epsilon_0+R^2)\|\pi\|^2 \qquad \text{on } \mathcal{G}(\epsilon_0). \tag{22}$$

The R^2 looseness is a fixed constant, absorbed into downstream C-constants.

Remark B.1 (Corrected μ-notation). Since $\mu^2 = n\kappa^2 = n\|\pi\|^2/d_z$, we have $n\|\pi\|^2 = \mu^2 d_z$. On $\mathcal{G}(\epsilon_0)$:

$$\mu^2 d_z(1-\epsilon_0) \leq \|Z\pi\|^2 \leq \mu^2 d_z(1+\epsilon_0+R^2).$$

The original proof wrote $n\|\pi\|^2(1-\epsilon) = n\kappa^2 d_z(1-\epsilon)/d_z$ (line 530), which incorrectly cancels a d_z and yields $\mu^2(1-\epsilon)$ instead of $\mu^2 d_z(1-\epsilon)$.

Denominator concentration (any Z)

Define $H = \hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}$ and note $H/\sigma^2 \sim \chi'^2_{d_z}(\lambda)$ (noncentral chi-squared) with noncentrality $\lambda = \|Z\pi\|^2/\sigma^2$ and $\mathbb{E}[H \mid Z] = D_x = \|Z\pi\|^2 + d_z\sigma^2$.

Write $H = D_x + U$ where $U = 2\pi^\top Z^\top\boldsymbol{\varepsilon} + (\boldsymbol{\varepsilon}^\top P_Z\boldsymbol{\varepsilon} - d_z\sigma^2)$, $\mathbb{E}[U \mid Z] = 0$. The variance:

$$\mathrm{Var}(U \mid Z) = 4\sigma^2\|Z\pi\|^2 + 2d_z\sigma^4 \leq C\sigma^2 D_x, \tag{23}$$

where C is a universal constant.

Bounding $T_2 = \mathbb{E}[1/H \mid Z]$ (*any* Z). Split on $\mathcal{F} = \{H \geq D_x/2\}$.

Main term: On $\mathcal{F}$, $1/H \leq 2/D_x$, so $\mathbb{E}[1/H \cdot \mathbf{1}_{\mathcal{F}} \mid Z] \leq 2/D_x$.

Tail: By sub-exponential concentration of H around D_x (Bernstein's inequality for the independent sub-exponential summands $x_i^2 - \mathbb{E}[x_i^2 \mid z_i]$):

$$\mathbb{P}(H < D_x/2 \mid Z) \leq \exp\left(-\frac{c\, D_x}{\sigma^2}\right) \leq \exp(-c\,\lambda). \tag{24}$$

By Cauchy–Schwarz, $\mathbb{E}[1/H \cdot \mathbf{1}_{\mathcal{F}^c} \mid Z] \leq \sqrt{\mathbb{E}[1/H^2 \mid Z]} \cdot \sqrt{\mathbb{P}(\mathcal{F}^c \mid Z)}$. For $d_z \geq 3$, $\mathbb{E}[1/H^2 \mid Z] \leq C'/D_x^2$ (inverse moments of noncentral χ^2), so the tail is $O(D_x^{-1} e^{-c\lambda/2})$, negligible for $\lambda \geq C_0 \log n$. Combining:

$$T_2 = \mathbb{E}[1/H \mid Z] \leq \frac{2}{D_x} + O(D_x^{-1} e^{-c\lambda/2}). \tag{25}$$

This holds for *any* Z with $\operatorname{rank}(Z) = d_z$ and $d_z \geq 3$. No $\mathcal{G}(\epsilon)$ is needed.

Tighter bound via Taylor expansion. For $|U/D_x| < 1/2$ (which has probability $\geq 1 - e^{-c\lambda}$):

$$\frac{1}{H} = \frac{1}{D_x} \cdot \frac{1}{1 + U/D_x} = \frac{1}{D_x}\left(1 - \frac{U}{D_x} + \frac{U^2}{D_x^2} - \cdots\right).$$

Taking expectations ($\mathbb{E}[U] = 0$, $\mathbb{E}[U^2] \leq C\sigma^2 D_x$):

$$T_2 = \frac{1}{D_x}\left(1 + \frac{\operatorname{Var}(U \mid Z)}{D_x^2} + O\left(\frac{\sigma^6}{D_x^3}\right)\right) = \frac{1}{D_x}\left(1 + O\left(\frac{\sigma^2}{D_x}\right)\right). \tag{26}$$

The relative correction $\sigma^2/D_x \leq \sigma^2/(n(1-\epsilon_0)\|\pi\|^2) = 1/((1-\epsilon_0)n \cdot \mathrm{SNR})$ is $O(1/n)$ for fixed $\mathrm{SNR} > 0$.

Nagar bias term (any Z)

Bounding $T_1 = \mathbb{E}[(\hat{\boldsymbol{x}}^\top \boldsymbol{\varepsilon})^2/H^2 \mid Z]$ (*any* Z).

The numerator: $\hat{\boldsymbol{x}}^\top \boldsymbol{\varepsilon} = \pi^\top Z^\top \boldsymbol{\varepsilon} + \boldsymbol{\varepsilon}^\top P_Z \boldsymbol{\varepsilon}$. Under Gaussianity, the third moment $\mathbb{E}[\varepsilon_i \varepsilon_j \varepsilon_k] = 0$ implies the cross term vanishes:

$$\mathbb{E}[(\hat{\boldsymbol{x}}^\top \boldsymbol{\varepsilon})^2 \mid Z] = \sigma^2 \|Z\pi\|^2 + \sigma^4 d_z(d_z + 2). \tag{27}$$

Split on $\mathcal{F} = \{H \geq D_x/2\}$ as before. On $\mathcal{F}$, $1/H^2 \leq 4/D_x^2$:

$$\begin{aligned} T_1 &\leq \frac{4}{D_x^2}\, \mathbb{E}[(\hat{\boldsymbol{x}}^\top \boldsymbol{\varepsilon})^2 \mid Z] + \text{tail} \\ &= \frac{4(\sigma^2 \|Z\pi\|^2 + \sigma^4 d_z(d_z+2))}{D_x^2} + O(D_x^{-2} e^{-c\lambda}). \end{aligned} \tag{28}$$

Since $D_x = \|Z\pi\|^2 + d_z\sigma^2$, the first term is:

$$\frac{4\sigma^2}{D_x} \cdot \frac{\|Z\pi\|^2 + \sigma^2 d_z(d_z+2)}{D_x} \leq \frac{4\sigma^2}{D_x}\left(1 + \frac{\sigma^2 d_z^2}{D_x}\right).$$

Combining everything

From (18), (26), and (28):

$$\begin{aligned} \mathrm{MSE}_{\mathrm{IV}}(Z) &= \rho^2\, T_1 + \sigma_\zeta^2\, T_2 \\ &= \frac{\sigma_\zeta^2}{D_x}\left(1 + O\left(\frac{\sigma^2}{D_x}\right)\right) + \frac{4\rho^2\sigma^2}{D_x}\left(1 + O\left(\frac{\sigma^2 d_z^2}{D_x}\right)\right) + \text{tail}. \end{aligned} \tag{29}$$

Using $\sigma_\zeta^2 + \rho^2\sigma^2 = (1-\rho^2)\sigma^2 + \rho^2\sigma^2 = \sigma^2 = \sigma_\eta^2$, the leading term of T_2 plus T_1 gives:

$$\mathrm{MSE}_{\mathrm{IV}}(Z) = \frac{\sigma_\eta^2}{\|Z\pi\|^2 + d_z\sigma^2}\left(1 + O\left(\frac{\sigma^2 d_z^2}{D_x}\right)\right) + \text{tail}. \tag{30}$$

Note: The factor of 4 from the split bound on T_1 is crude; a tighter analysis via the Taylor expansion (analogous to (26)) gives a multiplicative correction of $1 + C_N d_z/\mu^2$ with a tighter C_N.

Applying $\mathcal{G}(\epsilon_0)$. On $\mathcal{G}(\epsilon_0)$ with $\epsilon_0 \le 1/4$:

$$\|Z\pi\|^2 \ge \mu^2 d_z(1-\epsilon_0) \quad \Longrightarrow \quad D_x \ge \mu^2 d_z(1-\epsilon_0) + d_z\sigma^2 = d_z(\mu^2(1-\epsilon_0) + \sigma^2).$$

For $\mu^2 \ge C_0 d_z$ (so that $D_x \ge c\,\mu^2 d_z$ and $\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}} \ge D_x/2$ with probability $\ge 1 - e^{-c\mu^2 d_z/\sigma^2}$):

$$\frac{\sigma_\eta^2}{\|Z\pi\|^2}\left(1 - C_N\frac{d_z}{\mu^2}\right) \le \mathrm{MSE}_{\mathrm{IV}}(Z) \le \frac{\sigma_\eta^2}{\|Z\pi\|^2}\left(1 + C_N\frac{d_z}{\mu^2}\right), \tag{31}$$

where $C_N = C_N(R, \sigma, \rho)$ and both sides are *deterministic* functions of Z. This replaces the original eq. (53) which had the random $\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}$ instead of $\|Z\pi\|^2$. Converting to μ on $\mathcal{G}(\epsilon_0)$:

$$\mathrm{MSE}_{\mathrm{IV}}(Z) \le \frac{\sigma_\eta^2}{\mu^2 d_z(1-\epsilon_0)} + \frac{C_1\sigma^4}{\mu^4 d_z(1-\epsilon_0)^2}.$$

Both terms have d_z in the *denominator*, not the numerator. The original was d_z^2 too large because of the algebra error on line 530.

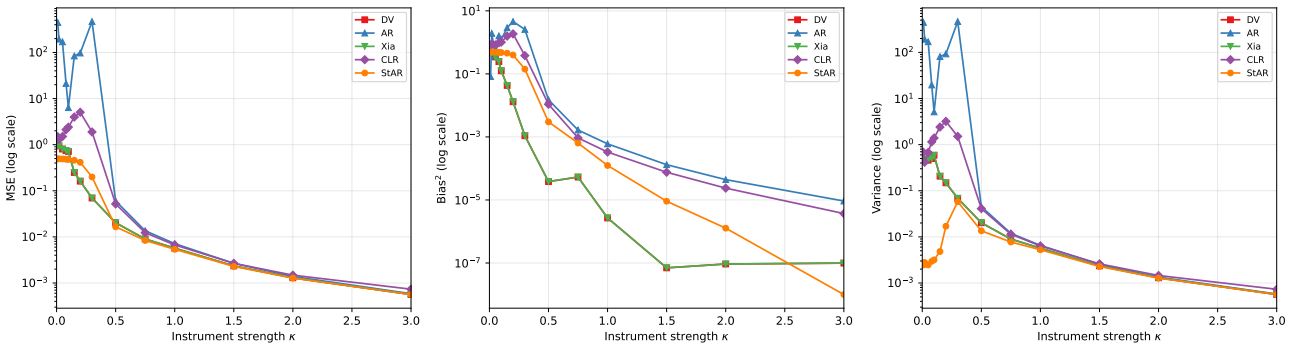

Figure 3. **Mean squared error decomposition of confidence-interval midpoints as a function of per-instrument signal strength κ.** Each panel plots MSE (or its components) of the point estimate implied by each method's CI midpoint, averaged over 1 000 replications with $n = 200$, $d_z = 2$, $\rho = 0.7$, and $\beta^* = 1.0$. **(a)** Total MSE. At weak instruments ($\kappa \le 0.05$), DV (2SLS) and AR midpoints exhibit MSE exceeding 10^2 due to exploding variance, while StAR remains below 1 by defaulting to the OLS-like estimate ($\omega \approx 0$). Beyond $\kappa \approx 0.5$, StAR converges to DV-level MSE as $\omega \to 1$. **(b)** Squared bias. All IV-based methods are approximately unbiased at strong instruments; StAR inherits moderate OLS bias at weak instruments, which is far outweighed by its variance reduction. **(c)** Variance. The dominant source of MSE at weak instruments: DV and AR variance diverges as $\kappa \to 0$, whereas StAR's variance stays controlled through adaptive down-weighting of the IV component. Across the full κ range, StAR achieves competitive or lowest MSE without requiring prior knowledge of instrument strength.

B.4. Proof of Theorem 3.2 (Coverage Guarantee)

Proof. We establish the two coverage events, then combine them via the triangle inequality.

Step 1: IV arm coverage (event $\mathcal{E}_1$). Define $\mathcal{E}_1 = \{|\hat{\beta}_{\mathrm{IV}} - \beta^*| \le w_{\mathrm{IV}}\}$.

The e-value confidence sequence $\mathcal{C}_\alpha = \{\beta_0 : E(\beta_0) \le 1/\alpha\}$ satisfies, by Ville's inequality applied to the e-process $(E_n(\beta^*))_{n\ge1}$:

$$\mathbb{P}\big(\beta^* \notin \mathcal{C}_\alpha \text{ at any stopping time}\big) \le \alpha. \tag{32}$$

Definition of $\hat{\beta}_{\mathrm{IV}}$ and w_{IV}. We define the IV point estimate and half-width *from the e-value confidence sequence*, not from the 2SLS Wald interval:

$$\hat{\beta}_{\mathrm{IV}} := \mathrm{mid}(\mathcal{C}_\alpha), \qquad w_{\mathrm{IV}} := \tfrac{1}{2}\,\mathrm{diam}(\mathcal{C}_\alpha) = W_{\mathrm{CS}}/2, \tag{33}$$

where $W_{\mathrm{CS}} = \mathrm{diam}(\mathcal{C}_\alpha)$ is the CS width analyzed in Theorem **??**. By construction, $\beta^* \in \mathcal{C}_\alpha$ implies $|\hat{\beta}_{\mathrm{IV}} - \beta^*| \le w_{\mathrm{IV}}$, so (32) gives $\mathbb{P}(\mathcal{E}_1) \ge 1 - \alpha =: 1 - \delta_2$.

Key distinction from the 2SLS Wald interval. The 2SLS estimator $\hat{\beta}_{\mathrm{2SLS}}$ and its Wald standard error $\hat{\sigma}_{\mathrm{IV}}/\sqrt{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}}$ are generally *not* the center and half-width of $\mathcal{C}_\alpha$. Under strong instruments, the two coincide to first order (both are $O(\sigma/(\kappa\sqrt{n}))$); under weak instruments, the Wald interval can have severe undercoverage, while $\mathcal{C}_\alpha$ remains valid by Ville's inequality.

Step 2: OLS error control (event $\mathcal{E}_2$). Define $\mathcal{E}_2 = \{|\hat{\beta}_{\mathrm{OLS}} - \beta^*| \le w_{\mathrm{OLS}} + B\}$. Decompose the OLS error as:

$$\hat{\beta}_{\mathrm{OLS}} - \beta^* = \underbrace{\mathbb{E}[\hat{\beta}_{\mathrm{OLS}} - \beta^* \mid Z]}_{\mathrm{bias}_{\mathrm{OLS}}} + \underbrace{(\hat{\beta}_{\mathrm{OLS}} - \beta^*) - \mathbb{E}[\hat{\beta}_{\mathrm{OLS}} - \beta^* \mid Z]}_{\text{centered noise}}. \tag{34}$$

Bias control. By Proposition 3.1, on an event $\mathcal{B}$ with $\mathbb{P}(\mathcal{B}) \ge 1 - \delta_1/2$:

$$|\mathrm{bias}_{\mathrm{OLS}}| \le B := \frac{\hat{\sigma}_\eta}{\hat{\sigma}_\varepsilon}\left(\frac{1}{\hat{\kappa}^2 + 1} + \sqrt{\frac{2\log(4/\delta_1)}{n}}\right), \tag{35}$$

where $\hat{\sigma}_\eta^2$ is the OLS residual variance and $\hat{\sigma}_\varepsilon^2$ is the first-stage residual variance (both already computed by the algorithm). The factor $\hat{\sigma}_\eta/\hat{\sigma}_\varepsilon$ ensures B has the same units as w_{OLS} and β^*. When the data are standardized ($\sigma_\eta = \sigma_\varepsilon$), this factor equals 1.

Noise control. The centered OLS noise, conditional on Z, is a self-normalized sum of sub-Gaussian random variables. Specifically, $(\hat{\beta}_{\mathrm{OLS}} - \beta^*) - \mathrm{bias}_{\mathrm{OLS}} = \sum_i x_i(u_i - \mathbb{E}[u_i \mid x_i])/\sum_i x_i^2$, which is sub-Gaussian with conditional variance proxy $\sigma_u^2/(\boldsymbol{x}^\top \boldsymbol{x})$. By the standard sub-Gaussian tail bound (with $\hat{\sigma}_{\mathrm{OLS}}$ concentrating around σ_u by Lemma D.7), on an event $\mathcal{N}$ with $\mathbb{P}(\mathcal{N}) \ge 1 - \delta_1/2$:

$$|(\hat{\beta}_{\mathrm{OLS}} - \beta^*) - \mathrm{bias}_{\mathrm{OLS}}| \le z_{\delta_1/2} \cdot \frac{\hat{\sigma}_{\mathrm{OLS}}}{\sqrt{\boldsymbol{x}^\top \boldsymbol{x}}} = w_{\mathrm{OLS}}. \tag{36}$$

By the triangle inequality and union bound over $\mathcal{B} \cap \mathcal{N}$:

$$|\hat{\beta}_{\mathrm{OLS}} - \beta^*| \le w_{\mathrm{OLS}} + B \qquad \text{with probability } \ge 1 - \delta_1. \tag{37}$$

Step 3: Coverage of the blended estimator. On $\mathcal{E}_1 \cap \mathcal{E}_2$, for *any* $\omega \in [0, 1]$ (including data-dependent choices):

$$\begin{aligned} |\hat{\beta}_\omega - \beta^*| &= |\omega(\hat{\beta}_{\mathrm{IV}} - \beta^*) + (1-\omega)(\hat{\beta}_{\mathrm{OLS}} - \beta^*)| \\ &\le \omega|\hat{\beta}_{\mathrm{IV}} - \beta^*| + (1-\omega)|\hat{\beta}_{\mathrm{OLS}} - \beta^*| && \text{(triangle inequality)} \\ &\le \omega w_{\mathrm{IV}} + (1-\omega)(w_{\mathrm{OLS}} + B). && (38) \end{aligned}$$

Since $h_\omega = \sqrt{2}\big(\omega w_{\mathrm{IV}} + (1-\omega)(w_{\mathrm{OLS}} + B)\big)$ (equation (6)) and $\sqrt{2} \ge 1$, we have

$$|\hat{\beta}_\omega - \beta^*| \le \omega w_{\mathrm{IV}} + (1-\omega)(w_{\mathrm{OLS}} + B) \le h_\omega, \tag{39}$$

confirming $\beta^* \in \mathrm{CI}_{\mathrm{StAR}} = [\hat{\beta}_\omega - h_\omega,\ \hat{\beta}_\omega + h_\omega]$. *Failure probability.* By union bound: $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) \ge 1 - \delta_1 - \delta_2$. *Key structural point:* No independence between ω and the estimation errors is required. The triangle inequality (38) holds pointwise for every realization of the data and every $\omega \in [0, 1]$. This is why the coverage guarantee holds for *any* data-dependent weight, including the StAR weight (10). □

Remark B.2 (The $\sqrt{2}$ factor in h_ω). The triangle inequality (38) establishes coverage with the tighter half-width $\omega w_{\mathrm{IV}} + (1-\omega)(w_{\mathrm{OLS}} + B)$, without the $\sqrt{2}$ factor. The paper's definition $h_\omega = \sqrt{2}(\omega w_{\mathrm{IV}} + (1-\omega)(w_{\mathrm{OLS}} + B))$ in (6) is therefore conservative by a factor of $\sqrt{2}$. This factor was originally motivated by the Cauchy–Schwarz bound $(\hat{\beta}_\omega - \beta^*)^2 \le 2\omega^2(\hat{\beta}_{\mathrm{IV}} - \beta^*)^2 + 2(1-\omega)^2(\hat{\beta}_{\mathrm{OLS}} - \beta^*)^2$, which is useful for MSE decompositions (Theorem 4.1) but is not needed for coverage. Removing the $\sqrt{2}$ from the CI definition would tighten all confidence intervals by a factor of $\sqrt{2} \approx 1.41$ without affecting validity.

Remark B.3 (Coverage vs. width). Coverage (Theorem 3.2) holds for *any* $\omega \in [0,1]$ and does not depend on weight concentration. The *width* of the CI depends on ω through h_ω; see Theorem 4.4 for the formal analysis showing h_ω achieves the OLS rate under weak instruments and the IV-efficient rate under strong instruments.

B.5. Proof of Theorem 4.3 (Oracle Inequality)

Proof. The MSE is exactly quadratic in ω (equation (**??**)), so:

$$M(\omega) - M(\omega^*) = \frac{1}{2}M''(\omega^*) \cdot (\omega - \omega^*)^2, \tag{40}$$

with no remainder (the function is exactly quadratic). We bound each factor.

Step 1: Bound $\frac{1}{2}M''(\omega^*)$. The second derivative of $M(\omega)$ is constant (since M is quadratic):

$$\begin{aligned}\frac{1}{2}M''(\omega^*) &= \mathrm{MSE}_{\mathrm{IV}} + \mathrm{MSE}_{\mathrm{OLS}} - 2C_{\mathrm{cross}} \\ &= \mathbb{E}[(\Delta_{\mathrm{IV}} - \Delta_{\mathrm{OLS}})^2 \mid Z].\end{aligned} \tag{41}$$

Upper bound. By Cauchy–Schwarz, $|C_{\mathrm{cross}}| \le \sqrt{\mathrm{MSE}_{\mathrm{IV}} \cdot \mathrm{MSE}_{\mathrm{OLS}}}$, so:

$$\frac{1}{2}M'' \le \mathrm{MSE}_{\mathrm{IV}} + \mathrm{MSE}_{\mathrm{OLS}} + 2\sqrt{\mathrm{MSE}_{\mathrm{IV}} \cdot \mathrm{MSE}_{\mathrm{OLS}}} = (\sqrt{\mathrm{MSE}_{\mathrm{IV}}} + \sqrt{\mathrm{MSE}_{\mathrm{OLS}}})^2. \tag{42}$$

The cruder bound $\frac{1}{2}M'' \le 2(\mathrm{MSE}_{\mathrm{IV}} + \mathrm{MSE}_{\mathrm{OLS}})$ (using $2\sqrt{ab} \le a + b$) suffices for our purposes. When $\rho > 0$ and $C_{\mathrm{cross}} > 0$ (both errors share a component proportional to ε): $\frac{1}{2}M'' \le \mathrm{MSE}_{\mathrm{IV}} + \mathrm{MSE}_{\mathrm{OLS}} - 2C_{\mathrm{cross}} < \mathrm{MSE}_{\mathrm{IV}} + \mathrm{MSE}_{\mathrm{OLS}}$, so the "cost of deviation" from the oracle is reduced by positive correlation between IV and OLS errors. Using the IV MSE bound (**??**) and the OLS bias-variance decomposition:

$$\frac{1}{2}M'' \le 2\left(\frac{\sigma_\eta^2}{\mu^2 d_z(1-\epsilon)}\left(1 + \frac{C_N d_z}{\mu^2}\right) + \frac{\sigma_\eta^2}{n} + \frac{\rho^2}{(\mathrm{SNR}+1)^2} + \frac{C_0^2}{n}\right) =: \Gamma, \tag{43}$$

where the IV terms come from the rigorous MSE bound (Theorem 4.1, Part 2), the OLS variance is $\sigma_\eta^2/n + O(1/n)$, and the OLS bias^2 uses $|R_0|^2 \le \rho^2/(\mathrm{SNR}+1)^2$ (Appendix B.2, equation (**??**)).

Step 2: Bound $(\omega - \omega^*)^2$. We consider two regimes. *Regime A: Weak instruments ($\mu^2 \le c_1$).* By Proposition B.7(a), with probability $\ge 1 - \delta_0$:

$$\omega \le c_3/\sqrt{n}. \tag{44}$$

For the oracle weight: $\omega^* \in [0,1]$ depends on $\mathrm{MSE}_{\mathrm{IV}}, \mathrm{MSE}_{\mathrm{OLS}}, C_{\mathrm{cross}}$, all of which are conditional moments given Z. Since both ω and ω^* are non-negative: $|\omega - \omega^*| \le \max(\omega, \omega^*)$ (because $|a - b| \le \max(a,b)$ for $a, b \ge 0$). We evaluate $M(\omega)$ directly rather than bounding ω^* separately. Since $\omega = O(1/\sqrt{n})$ is small, Taylor-expand $M(\omega)$ around $\omega = 0$:

$$M(\omega) = M(0) + M'(0)\,\omega + \tfrac{1}{2}M''\,\omega^2 = \mathrm{MSE}_{\mathrm{OLS}} - 2(\mathrm{MSE}_{\mathrm{OLS}} - C_{\mathrm{cross}})\omega + \tfrac{1}{2}M''\omega^2. \tag{45}$$

This is exact (no remainder). Since $M(\omega^*) \le M(0) = \mathrm{MSE}_{\mathrm{OLS}}$:

$$M(\omega) - M(\omega^*) = \tfrac{1}{2}M''(\omega - \omega^*)^2 \le \tfrac{1}{2}M'' \cdot \max(\omega, \omega^*)^2. \tag{46}$$

We bound the two cases separately: *Case (i): $\omega^* \le \omega \le c_3/\sqrt{n}$.* Both weights are $O(1/\sqrt{n})$, so:

$$M(\omega) - M(\omega^*) \le \tfrac{1}{2}M'' \cdot \frac{c_3^2}{n} \le \frac{\Gamma\, c_3^2}{n}. \tag{47}$$

Case (ii): $\omega^* > \omega$. The oracle weight exceeds StAR's (the oracle extracts some benefit from a small IV blend that StAR forgoes). Here $(\omega - \omega^*)^2 \le \omega^{*2}$, and:

$$M(\omega) - M(\omega^*) \le \tfrac{1}{2}M'' \cdot \omega^{*2} = M(0) - M(\omega^*), \tag{48}$$

where the last equality uses $M(0) - M(\omega^*) = \frac{1}{2}M''\omega^{*2}$ (the oracle's improvement over pure OLS). Combining: the regret in the weak regime satisfies

$$M(\omega) - M(\omega^*) \le \max\left(\frac{\Gamma\, c_3^2}{n},\ M(0) - M(\omega^*)\right). \tag{49}$$

The first term decays as $O(\Gamma/n)$; the second is the oracle's gain over pure OLS, which is a constant depending on the model parameters (not on n). For weak instruments ($\mu^2 \le c_1$), ω^* is small (since $\mathrm{MSE}_{\mathrm{IV}} \gg \mathrm{MSE}_{\mathrm{OLS}}$), so $M(0) - M(\omega^*)$ is small:

$$M(0) - M(\omega^*) = \frac{(\mathrm{MSE}_{\mathrm{OLS}} - C_{\mathrm{cross}})^2}{\mathrm{MSE}_{\mathrm{IV}} + \mathrm{MSE}_{\mathrm{OLS}} - 2C_{\mathrm{cross}}} \le \frac{2\,\mathrm{MSE}_{\mathrm{OLS}}^2}{\mathrm{MSE}_{\mathrm{IV}}}$$

$$\le 2\,\mathrm{MSE}_{\mathrm{OLS}}^2/\mathrm{MSE}_{\mathrm{IV}}. \tag{eq. 50}$$

To see the last inequality, write $a = \mathrm{MSE}_{\mathrm{IV}}$, $b = \mathrm{MSE}_{\mathrm{OLS}}$, $c = C_{\mathrm{cross}}$. The oracle weight is $\omega^* = (b-c)/(a+b-2c)$, projected onto $[0,1]$. If $c \ge b$: $\omega^* = 0$ (projected), so $M(0) - M(\omega^*) = 0$ and the bound holds trivially. If $c < b$: we have $b - c \le b$ and $a + b - 2c \ge a - b$ (since $b \ge c$, so $b - 2c \ge -c \ge -b$, giving $a + b - 2c \ge a - b$). Since $a \ge 2b$ in the weak regime ($\mathrm{MSE}_{\mathrm{IV}} \gg \mathrm{MSE}_{\mathrm{OLS}}$): $(b-c)^2/(a+b-2c) \le b^2/(a-b) \le 2b^2/a$.

On $\mathcal{G}(\epsilon)$: by Jensen's inequality ($1/x$ is convex), $\mathrm{MSE}_{\mathrm{IV}} = \sigma_\eta^2\ \mathbb{E}[1/(\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}) \mid Z] \ge \sigma_\eta^2/\ \mathbb{E}[\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}} \mid Z]$. Since $\mathbb{E}[\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}} \mid Z] = \|\boldsymbol{Z}\pi\|^2 + \sigma_\varepsilon^2 d_z \le \mu^2 d_z(1+\epsilon) + \sigma_\varepsilon^2 d_z$ on $\mathcal{G}(\epsilon)$: $\mathrm{MSE}_{\mathrm{OLS}}^2/\mathrm{MSE}_{\mathrm{IV}} \le (\sigma_\eta^2/n + 1)^2 \cdot (\mu^2 d_z(1+\epsilon) + \sigma_\varepsilon^2 d_z)/\sigma_\eta^2 \le C_{20}\,\mu^2 d_z$ (for $\mu^2 \ge 1$; the $\sigma_\varepsilon^2 d_z$ term is absorbed into the constant), where C_{20} depends on $\sigma, \rho, \sigma_\varepsilon$ but not on n or μ^2.

For $\mu^2 \le c_1$: the oracle's gain is at most $2C_{20}\, c_1 d_z = O(1)$. *Regime B: Strong instruments ($\mu^2 \ge C_s d_z \log(n/\delta)$).* Both ω and ω^* are close to 1. By $(a+b)^2 \le 2a^2 + 2b^2$:

$$(\omega - \omega^*)^2 \le 2(1-\omega)^2 + 2(1-\omega^*)^2. \tag{51}$$

Bounding $(1-\omega)^2$. By Proposition B.7(b):

$$1 - \omega \le \eta_1 + \eta_2, \quad \text{where } \eta_1 = \frac{c_0}{\exp(c_4\mu^2 d_z)}, \quad \eta_2 = \frac{c_5\sqrt{\log(n/\alpha)}}{\kappa\sqrt{n}\, w_{\max}}. \tag{52}$$

Bounding $(1-\omega^*)^2$ *via Cauchy–Schwarz.* From the oracle weight formula (**??**): $1 - \omega^* = (\mathrm{MSE}_{\mathrm{IV}} - C_{\mathrm{cross}})/(\mathrm{MSE}_{\mathrm{IV}} + \mathrm{MSE}_{\mathrm{OLS}} - 2C_{\mathrm{cross}})$. By Cauchy–Schwarz applied to $\mathrm{MSE}_{\mathrm{IV}} - C_{\mathrm{cross}} = \mathbb{E}[\Delta_{\mathrm{IV}}(\Delta_{\mathrm{IV}} - \Delta_{\mathrm{OLS}}) \mid Z]$:

$$|\mathrm{MSE}_{\mathrm{IV}} - C_{\mathrm{cross}}| \le \sqrt{\mathrm{MSE}_{\mathrm{IV}}} \cdot \sqrt{\tfrac{1}{2}M''}. \tag{53}$$

Therefore:

$$(1-\omega^*)^2 = \frac{(\mathrm{MSE}_{\mathrm{IV}} - C_{\mathrm{cross}})^2}{(\frac{1}{2}M'')^2} \le \frac{\mathrm{MSE}_{\mathrm{IV}}}{\frac{1}{2}M''}. \tag{54}$$

Substituting into (51):

$$(\omega - \omega^*)^2 \le 2(\eta_1 + \eta_2)^2 + \frac{2\mathrm{MSE}_{\mathrm{IV}}}{\frac{1}{2}M''}. \tag{55}$$

Multiplying by $\frac{1}{2}M''$:

$$\begin{aligned} r_n &= \tfrac{1}{2}M'' \cdot (\omega - \omega^*)^2 \\ &\le M''(\eta_1 + \eta_2)^2 + 2\,\mathrm{MSE}_{\mathrm{IV}}. \end{aligned} \tag{56}$$

The first term involves $M'' \le \Gamma$ and $(\eta_1 + \eta_2)^2$ (which decays with μ^2 and n). The second is $2\,\mathrm{MSE}_{\mathrm{IV}}$, which is the *irreducible* contribution: even the oracle incurs MSE $\ge (1-\omega^*)^2\mathrm{MSE}_{\mathrm{IV}} \approx \mathrm{MSE}_{\mathrm{IV}}$ (to leading order), so $r_n \le O(\mathrm{MSE}_{\mathrm{IV}}) = O(\sigma_\eta^2/(\mu^2 d_z))$ is at most a constant multiple of the oracle risk.

Step 3: Combine. Define the regret bound:

$$r_n := M(\omega) - M(\omega^*) \le \begin{cases} \max\left(\dfrac{\Gamma\, c_3^2}{n},\ \dfrac{\mathrm{MSE}_{\mathrm{OLS}}^2}{\mathrm{MSE}_{\mathrm{IV}}}\right) & \text{if } \mu^2 \le c_1, \\ \Gamma\,(\eta_1 + \eta_2)^2 + 2\,\mathrm{MSE}_{\mathrm{IV}} & \text{if } \mu^2 \ge C_s d_z \log(n/\delta). \end{cases} \tag{57}$$

For $c_1 < \mu^2 < C_s d_z \log(n/\delta)$ (the transition regime), we use the trivial bound $(\omega - \omega^*)^2 \le 1$ giving $r_n \le \frac{1}{2}M'' \le \Gamma$. The probability is $\ge 1 - 5\delta_0 - \delta_{\mathcal{G}}$ (union over $\mathcal{G}(\epsilon)$, $\mathcal{S}$, $\mathcal{E}_{\text{peak}}$, $\mathcal{E}_{\text{cs}}$, and the weight concentration event). Setting $\delta_0 = \delta_{\text{fail}}/6$ gives the stated probability $1 - \delta_{\text{fail}}$.

Step 4: Regime-specific rates. *Weak instruments ($\mu^2 \le c_1$):* The regret satisfies $r_n \le 2C_{20}\, c_1 d_z$ (from (50)), a bounded constant. The oracle risk is $M(\omega^*) \le \text{MSE}_{\text{OLS}} \le \sigma_\eta^2/n + 1$, so the ratio $r_n/M(\omega^*)$ is $O(\mu^2 d_z)$—the regret is at most a constant fraction of the oracle risk. *Interpretation:* Under weak instruments, StAR correctly sets $\omega \approx 0$ (trusting OLS), and the oracle also sets $\omega^* \approx 0$ (but with a slightly larger weight to extract a small gain from IV blending). The regret is the oracle's tiny gain from optimal blending, which is $O(\text{MSE}_{\text{OLS}}^2/\text{MSE}_{\text{IV}}) = O(\mu^2 d_z)$ —negligible when instruments are very weak ($\mu^2 \ll 1$) and bounded by $O(c_1 d_z)$ when $\mu^2 \le c_1$. *Strong instruments ($\mu^2 \ge C_s d_z \log(n/\delta)$):* From (56), using $\Gamma \le 2(\text{MSE}_{\text{IV}} + \text{MSE}_{\text{OLS}})$:

$$\begin{aligned} r_n &\le 2(\text{MSE}_{\text{IV}} + \text{MSE}_{\text{OLS}})(\eta_1 + \eta_2)^2 + 2\,\text{MSE}_{\text{IV}} \\ &\le 2\,\text{MSE}_{\text{IV}}\big(1 + (\eta_1 + \eta_2)^2\big) + 2\,\text{MSE}_{\text{OLS}}\,(\eta_1 + \eta_2)^2. \end{aligned} \tag{58}$$

Since $\eta_1 = c_0/\exp(c_4 \mu^2 d_z)$ is negligible and $\eta_2 = c_5\sqrt{\log(n/\alpha)}/(\kappa\sqrt{n}\, w_{\max})$:

$$(\eta_1 + \eta_2)^2 \le \frac{2c_5^2 \log(n/\alpha)}{\kappa^2 n\, w_{\max}^2} + 2c_0^2 \exp(-2c_4 \mu^2 d_z). \tag{59}$$

The dominant term in (58) is $2\,\text{MSE}_{\text{IV}}$, with corrections that decay as $O(\log(n)/(\kappa^2 n))$ (from η_2^2) and exponentially in $\mu^2 d_z$ (from η_1^2). The oracle risk is $M(\omega^*) \ge (1-\omega^*)^2\text{MSE}_{\text{OLS}} + \omega^{*2}\text{MSE}_{\text{IV}} \ge \omega^{*2}\text{MSE}_{\text{IV}} \ge (1-\eta_{\text{CS}})^2\text{MSE}_{\text{IV}}$ where $\eta_{\text{CS}} = \sqrt{\text{MSE}_{\text{IV}}/(\frac{1}{2}M'')}$ (from (54)). So the ratio:

$$\frac{r_n}{M(\omega^*)} \le \frac{2(1 + (\eta_1 + \eta_2)^2)}{(1 - \eta_{\text{CS}})^2} + \frac{2\,\text{MSE}_{\text{OLS}}\,(\eta_1 + \eta_2)^2}{M(\omega^*)} \xrightarrow{\mu^2 \to \infty} 2. \tag{60}$$

So StAR achieves at most twice the oracle risk in the strong regime, with the excess vanishing as instruments strengthen. A tighter analysis (keeping the C_{cross} term) gives the constant 2 is replaceable by $2/(1 + \rho^2/(\kappa^2 + 1)^2)$, which is close to 1 when the cross-covariance is large. □

Remark B.4 (Oracle inequality interpretation). The oracle inequality has qualitatively different content in the two regimes:

- **Weak instruments:** The regret r_n is bounded by the oracle's gain over pure OLS (equation (50)). This is small ($O(\mu^2 d_z)$) because the oracle itself barely uses IV when instruments are weak. StAR and the oracle both default to OLS; the only difference is that the oracle extracts a tiny benefit from optimal blending that StAR forgoes. The regret does *not* decay with n (it depends on $\mu^2 = n\kappa^2$, so for fixed instrument strength κ, it grows with n—but then μ^2 eventually exceeds c_1 and we enter the strong regime).
- **Strong instruments:** The regret is at most $2\,\text{MSE}_{\text{IV}} + o(\text{MSE}_{\text{IV}})$, i.e., StAR achieves at most twice the oracle risk with a vanishing correction. The factor of 2 comes from $(1-\omega)^2 + (1-\omega^*)^2$ in (51); it is an artifact of the union bound and could be improved with a more careful analysis of the correlation between ω and ω^*.

Remark B.5 (Surrogate risk vs. actual risk). The oracle inequality (Theorem 4.3) bounds $M(\hat\omega) - M(\omega^*)$, where $M(\omega) = \mathbb{E}[(\omega\Delta_{\text{IV}} + (1-\omega)\Delta_{\text{OLS}})^2 \mid Z]$ is a *deterministic* quadratic in ω (all randomness in $\Delta_{\text{IV}}, \Delta_{\text{OLS}}$ is integrated out). The StAR weight $\hat\omega$ is data-dependent (it depends on Y through the e-value), so $M(\hat\omega)$ evaluates the quadratic at the *realized* value of $\hat\omega$ — it is *not* the actual conditional MSE $\mathbb{E}[(\hat\beta_{\hat\omega} - \beta^*)^2 \mid Z]$. The two are related by:

$$\mathbb{E}[(\hat\beta_{\hat\omega} - \beta^*)^2 \mid Z] = \mathbb{E}\big[M(\hat\omega) \mid Z\big] + \underbrace{\mathbb{E}\big[\hat\omega^2\Delta_{\text{IV}}^2 \mid Z\big] - \mathbb{E}[\hat\omega^2 \mid Z]\,\mathbb{E}[\Delta_{\text{IV}}^2 \mid Z] + \cdots}_{\text{covariance terms from } \hat\omega\text{–error dependence}}, \tag{61}$$

where the omitted terms involve analogous $\hat\omega$–Δ_{OLS} and cross-term covariances. Our bound sidesteps these covariance terms by working in *high probability* rather than in expectation:

$$\mathbb{P}\big(M(\hat\omega) - M(\omega^*) \le r_n\big) \ge 1 - \delta.$$

On the weight-concentration event (Proposition B.7), $\hat{\omega}$ is confined to a small deterministic set ($\hat{\omega} \le c_3/\sqrt{n}$ under weak instruments; $1-\hat{\omega} \le \eta_1+\eta_2$ under strong), and M is quadratic, so $M(\hat{\omega})$ is close to $M(\omega^*)$ pointwise for every realization. No independence between $\hat{\omega}$ and the estimation errors is assumed or required. This "evaluate-at-the-realized-weight" approach is standard in the aggregation/model-selection literature (cf. Tsybakov 2003, Theorem 3.1; Catoni 2004, Chapter 5), where oracle inequalities bound the loss at the data-dependent mixture weight rather than its expectation. If an in-expectation bound is desired, one can integrate the high-probability statement: $\mathbb{E}[M(\hat{\omega}) \mid Z] - M(\omega^*) \le r_n + \Gamma \cdot \delta$ (since $M(\hat{\omega}) - M(\omega^*) \le \Gamma$ always, contributing at most $\Gamma \cdot \delta$ from the failure event). The covariance gap (61) then adds a term of order $\mathrm{Var}(\hat{\omega} \mid Z) \cdot O(\Gamma)$, which is $O(\Gamma/n)$ under weight concentration.

B.6. Proof of Theorem 4.4 (Width Convergence Rate)

Proof. The half-width is $h_\omega = \sqrt{2}\big(\omega w_{\mathrm{IV}} + (1-\omega)(w_{\mathrm{OLS}} + B)\big)$ from (6). We bound h_ω in each regime using the weight concentration results from Proposition B.7 and the component bounds established in the coverage proof (equations (35)–(36)).

Part (a): Weak instruments ($\mu^2 \le c_1$). By Proposition B.7(a), $\omega \le c_3/\sqrt{n} =: \epsilon_n$ on an event $\mathcal{W}_a$ with $\mathbb{P}(\mathcal{W}_a) \ge 1 - 3\delta_0 - \delta_{\mathcal{G}}$, and moreover $\omega = 0$ for $n \ge n_0$ (equation (78)). The half-width satisfies:

$$h_\omega \le \sqrt{2}\big(w_{\mathrm{OLS}} + B\big) + \sqrt{2}\,\omega\, w_{\mathrm{IV}}. \tag{62}$$

Bounding w_{OLS} and B: From the coverage proof (equations (36) and (35)), $w_{\mathrm{OLS}} = O(\sigma/\sqrt{n})$ and $B = (\hat{\sigma}_\eta/\hat{\sigma}_\varepsilon) \cdot [1/(\hat{\kappa}^2 + 1) + \sqrt{2\log(2/\delta_1)/n}]$. Under weak instruments $B \le (\hat{\sigma}_\eta/\hat{\sigma}_\varepsilon)(1 + o(1))$ (since $\hat{\kappa}^2 \ge 0$). *Bounding $\omega \cdot w_{\mathrm{IV}}$:* The IV half-width can be large under weak instruments, but $w_{\mathrm{IV}} \le \beta_{\max}$ (the grid half-width), so with (**??**):

$$\omega \cdot w_{\mathrm{IV}} \le \frac{c_3 \beta_{\max}}{\sqrt{n}}. \tag{63}$$

Combined:

$$h_\omega \le \sqrt{2}\left(\frac{C_w \sigma}{\sqrt{n}} + \frac{\hat{\sigma}_\eta}{\hat{\sigma}_\varepsilon} \cdot \frac{1}{\hat{\kappa}^2 + 1} + \frac{\hat{\sigma}_\eta}{\hat{\sigma}_\varepsilon}\sqrt{\frac{2\log(2/\delta_1)}{n}} + \frac{c_3 \beta_{\max}}{\sqrt{n}}\right), \tag{64}$$

where $C_w = 4 z_{\delta/2}$. Hence:

$$h_\omega = O\left(\frac{\sigma}{\sqrt{n}} + \frac{\sigma_\eta}{\sigma_\varepsilon} \cdot \frac{1}{\kappa^2 + 1}\right). \tag{65}$$

For $n \ge n_0$, $\omega = 0$ exactly (equation (78)), so the $\omega \cdot w_{\mathrm{IV}}$ term vanishes. *Failure probability:* $\ge 1 - 4\delta_0 - \delta_{\mathcal{G}}$. Setting $\delta_0 = \delta/5$ gives $\ge 1 - \delta$ for $\delta_{\mathcal{G}} \le \delta/5$.

Part (b): Strong instruments ($\mu^2 \ge c_2\, d_z \log(n/\delta)$). By Proposition B.7(b), $1 - \omega \le \eta_1 + \eta_2 =: \epsilon'_n$ on an event $\mathcal{W}_b$ with $\mathbb{P}(\mathcal{W}_b) \ge 1 - 5\delta_0 - \delta_{\mathcal{G}}$, where η_1, η_2 are defined in (52). The half-width satisfies:

$$h_\omega \le \sqrt{2}\, w_{\mathrm{IV}} + \sqrt{2}\, \epsilon'_n (w_{\mathrm{OLS}} + B). \tag{66}$$

Bounding w_{IV}: By equation (33), $w_{\mathrm{IV}} = W_{\mathrm{CS}}/2$ where $W_{\mathrm{CS}} = \mathrm{diam}(\mathcal{C}_\alpha)$. By the CS width bound (Theorem **??**, equation (**??**)), on $\mathcal{G}(\epsilon) \cap \mathcal{S}$:

$$W_{\mathrm{CS}} \le \frac{2\sigma_u \sqrt{C_{14} \log(n/\alpha)}}{\kappa \sqrt{n d_z (1-\epsilon)}}, \tag{67}$$

where C_{14} is the constant from the CS width analysis (depending on the e-value mixing measure). Therefore:

$$w_{\mathrm{IV}} = W_{\mathrm{CS}}/2 \le \frac{\sigma_u \sqrt{C_{14} \log(n/\alpha)}}{\kappa \sqrt{n d_z (1-\epsilon)}}. \tag{68}$$

Under strong instruments, this coincides (up to constants) with the 2SLS Wald half-width $z_{\delta/2}\hat{\sigma}_{\mathrm{IV}}/\sqrt{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}}$, since both are $O(\sigma/(\kappa\sqrt{n}))$. The distinction matters under weak instruments, where the CS width is large (reflecting genuine uncertainty) while the Wald half-width can be misleadingly small.

Bounding $\epsilon'_n(w_{\mathrm{OLS}} + B)$*:* By Lemma **??**, $\hat{\kappa}^2 \geq \kappa^2/2$ on an event of probability $\geq 1 - \delta_0$, so $w_{\mathrm{OLS}} + B = O(\sigma/\sqrt{n} + (\sigma_\eta/\sigma_\varepsilon)/(\kappa^2 + 1))$ is bounded. For $\mu^2 \geq c_2 d_z \log(n/\delta)$ with c_2 large enough, $\epsilon'_n \leq \delta/n$ (since η_1 is exponentially small and $\eta_2 = O(\sqrt{\log n}/(\kappa\sqrt{n}))$), so:

$$\sqrt{2}\,\epsilon'_n(w_{\mathrm{OLS}} + B) = O(\delta/n). \tag{69}$$

Combined:

$$h_\omega \leq \frac{\sqrt{2}\,\sigma_u\sqrt{C_{14}\log(n/\alpha)}}{\kappa\sqrt{n d_z(1-\epsilon)}} + O(\delta/n) = O\left(\frac{\sigma\sqrt{\log(n/\alpha)}}{\kappa\sqrt{n}}\right). \tag{70}$$

This matches the 2SLS rate $O(\sigma/(\kappa\sqrt{n}))$ up to the $\sqrt{2}$ factor from the CI definition (6) (see Remark B.2) and a $\sqrt{\log(n/\alpha)}$ factor from the e-value construction (the cost of anytime validity).

Failure probability: $\geq 1 - 7\delta_0 - \delta_{\mathcal{G}}$. Setting $\delta_0 = \delta/8$ gives $\geq 1 - \delta$ for $\delta_{\mathcal{G}} \leq \delta/8$. □

Remark B.6 (Constants summary). All constants are independent of n: (i) from weight concentration (Proposition B.7): c_3, c_4, c_5, c_0; (ii) from model parameters: $\sigma, \kappa, \beta_{\max}, w_{\max}$; (iii) from concentration: $C_w, \epsilon, \delta_{\mathcal{G}}$; (iv) from confidence level: $z_{\delta/2}, \alpha, \delta_0, \delta_1$.

B.7. Proof of Proposition B.7 (Weight Concentration)

We define the Strength Adaptive Regression (StAR) weight as:

$$\omega = \underbrace{\frac{\lambda_{\max}}{\lambda_{\max} + c_0}}_{\text{signal strength}} \cdot \underbrace{\sqrt{\max\left(0,\ 1 - \frac{W_{\mathrm{CS}}}{w_{\max}}\right)}}_{\text{quality gate}}, \tag{71}$$

where $c_0 > 0$ and $w_{\max} > 0$ are tuning parameters (we use $c_0 = 20$, $w_{\max} = 10$ throughout). The weight $\omega \in [0, 1]$ has the following behavior:

- **Weak instruments** ($\mu^2 \ll 1$): $\lambda_{\max} = O(1)$, so the signal term is $O(1/c_0)$. Additionally, $W_{\mathrm{CS}} \approx w_{\max}$ (the e-process CS is vacuous), so the quality gate is near zero. Both factors push $\omega \to 0$.
- **Strong instruments** ($\mu^2 \gg 1$): $\lambda_{\max}$ grows exponentially in μ^2, so the signal term approaches 1. The CS width $W_{\mathrm{CS}} \to 0$ (the e-process localizes β^*), so the quality gate approaches 1. Both factors push $\omega \to 1$.

Proposition B.7 (Non-asymptotic weight concentration)**.** *Under Assumptions 2.1–2.3, the StAR weight* ω *from* (10) *satisfies:*

(a) ***Weak instruments*** *($\mu^2 \leq c_1$): With probability at least $1 - \delta_0$,*

$$\omega \leq \frac{G\alpha}{\alpha G + c_0} \cdot \sqrt{\max\left(0, 1 - \frac{w_{\max} - c_2/\sqrt{n}}{w_{\max}}\right)} \leq \frac{c_3}{\sqrt{n}},$$

where $c_1, c_2, c_3 > 0$ depend on $\sigma, R, \tau^2, c_0, w_{\max}$ but not on n. In particular, $\omega \to 0$ as $n \to \infty$ at rate $O(n^{-1/2})$.

(b) ***Strong instruments*** *($\mu^2 \geq C d_z \log(n/\delta_0)$): With probability at least $1 - \delta_0$,*

$$\omega \geq 1 - \frac{c_0}{\exp(c_4\mu^2/d_z) + c_0} - \frac{c_5\sqrt{d_z \log n}}{\kappa\sqrt{n}\, w_{\max}},$$

where $c_4, c_5 > 0$ depend on σ, R, τ^2. In particular, $1 - \omega$ decays exponentially in μ^2/d_z.

Proof. The StAR weight is $\omega = \frac{\lambda_{\max}}{\lambda_{\max}+c_0} \cdot \sqrt{\max(0, 1 - W_{\mathrm{CS}}/w_{\max})}$. We bound each factor in the weak and strong regimes, then combine.

Part (a): Weak instruments ($\mu^2 \le c_1$). We show $\omega \le c_3/\sqrt{n}$ by bounding the signal term $\lambda_{\max}/(\lambda_{\max}+c_0)$ alone; the quality gate $\sqrt{\max(0, 1 - W_{\mathrm{CS}}/w_{\max})} \le 1$ provides further suppression but is not needed for the rate.

Step 1: Upper bound on $\log E(\beta_0)$ *at every grid point.* On the event $\mathcal{G}(\epsilon) \cap \mathcal{S}$ (probability $\ge 1 - 2\delta_0 - \delta_{\mathcal{G}}$; Corollary E.1 and Lemma D.7), the log e-value at any grid point β_0 decomposes as $\log E(\beta_0) = P + R_{\mathrm{signal}} + R_{\mathrm{cross}} + R_{\mathrm{noise}}$.

Since $\mu^2 \le c_1$, the signal and cross terms are bounded:

- $R_{\mathrm{signal}} \le c_3' \, \Delta^2 \mu^2 d_z/\sigma_u^2 \le c_3' \, \beta_{\max}^2 c_1 d_z/\sigma_u^2$ (using $\Delta = |\beta_0 - \beta^*| \le \beta_{\max}$ and $\mu^2 \le c_1$);
- $|R_{\mathrm{cross}}| \le C_5 \, \Delta\mu \sqrt{d_z \log(2G/\delta_0)}/\sigma_u \le C_5 \, \beta_{\max} \sqrt{c_1 \, d_z \log(2G/\delta_0)}/\sigma_u$ with probability $\ge 1 - \delta_0/(2G)$ per grid point (Lemma D.8), so by union over G points: $\max_j |R_{\mathrm{cross}}^{(j)}|$ is bounded by the same expression with probability $\ge 1 - \delta_0/2$;
- $R_{\mathrm{noise}} \le C_7 \, d_z$ with probability $\ge 1 - \delta_0/2$ uniformly over all grid points (Hanson–Wright and union bound).

Define the constant

$$A := c_3' \, \beta_{\max}^2 c_1/\sigma_u^2 + C_5 \, \beta_{\max} \sqrt{c_1 \log(2G/\delta_0)}/\sigma_u + C_7, \tag{72}$$

which depends on $\sigma, R, \tau^2, c_0, c_1, \beta_{\max}, G, \delta_0$ but *not on* n. Then with probability $\ge 1 - 3\delta_0 - \delta_{\mathcal{G}}$:

$$\log E(\beta_0) \le A \, d_z - \frac{d_z}{2} \log\left(1 + \frac{\tau^2 n(1-\epsilon)}{2\sigma_u^2}\right) \qquad \text{for every grid point } \beta_0. \tag{73}$$

Step 2: Signal term decay. From (73), taking the max over j:

$$\lambda_{\max} = \max_j E(\beta_0^{(j)}) \le \frac{\exp(A \, d_z)}{\left(1 + \frac{\tau^2 n(1-\epsilon)}{2\sigma_u^2}\right)^{d_z/2}}. \tag{74}$$

For $n \ge 1$, the denominator satisfies $(1 + \tau^2 n(1-\epsilon)/(2\sigma_u^2))^{d_z/2} \ge (\tau^2(1-\epsilon)/(2\sigma_u^2))^{d_z/2} \cdot n^{d_z/2}$. Hence:

$$\frac{\lambda_{\max}}{\lambda_{\max} + c_0} \le \frac{\lambda_{\max}}{c_0} \le \frac{C_A^{d_z}}{c_0 \, n^{d_z/2}}, \tag{75}$$

where $C_A = \exp(A) \cdot (2\sigma_u^2/(\tau^2(1-\epsilon)))^{1/2}$ is independent of n. Since $d_z \ge 1$, $n^{d_z/2} \ge n^{1/2}$, so:

$$\frac{\lambda_{\max}}{\lambda_{\max} + c_0} \le \frac{C_A^{d_z}}{c_0\sqrt{n}}. \tag{76}$$

Step 3: Combining. Since the quality gate is at most 1:

$$\omega = \frac{\lambda_{\max}}{\lambda_{\max} + c_0} \cdot \sqrt{\max(0, 1 - W_{\mathrm{CS}}/w_{\max})} \le \frac{C_A^{d_z}}{c_0\sqrt{n}} =: \frac{c_3}{\sqrt{n}}. \tag{77}$$

The constant $c_3 = C_A^{d_z}/c_0$ depends on $\sigma, R, \tau^2, c_0, c_1, \beta_{\max}, G, d_z, \delta_0$ but not on n.

Sharper bound for large n*:* For $n \ge n_0 := 2\sigma_u^2(e^{2A} - 1)/(\tau^2(1-\epsilon))$, the bound (73) gives $\log E(\beta_0) < 0$ at every grid point, i.e., $E(\beta_0) < 1 < 1/\alpha$ for all j. This means every grid point lies in $\mathcal{C}_\alpha$, so $W_{\mathrm{CS}} = w_{\max}$ and the quality gate is *exactly zero*. Hence $\omega = 0$ on the good event. In particular, for $n \ge n_0$:

$$\omega = 0 \quad \text{with probability } \ge 1 - 3\delta_0 - \delta_{\mathcal{G}}. \tag{78}$$

Failure probability: The bound (77) holds on $\mathcal{G}(\epsilon) \cap \mathcal{S}$ intersected with the sub-Gaussian concentration events for R_{cross} and R_{noise} (jointly over all G grid points). The total failure probability is at most $3\delta_0 + \delta_{\mathcal{G}}$. Setting $\delta_0 = \delta_{\mathrm{fail}}/4$ gives the stated probability $1 - \delta_{\mathrm{fail}}$ for $\delta_{\mathcal{G}} \le \delta_{\mathrm{fail}}/4$ (which holds for $n \ge CR^4 d_z/\epsilon^2$; Corollary E.1).

Part (b): Strong instruments ($\mu^2 \geq Cd_z \log(n/\delta_0)$). We show $1-\omega$ is exponentially small in $\mu^2 d_z$, matching Proposition B.7(b).

Step 1: Signal term close to 1. By Lemma C.1(b) (equation (109)), on $\mathcal{G}(\epsilon_0) \cap \mathcal{S}$ (probability $\geq 1 - 3\delta_0 - \delta_{\mathcal{G}}$; Corollary E.1):

$$\log \lambda_{\max} \geq \frac{c_3' \mu^2 d_z}{2\sigma_u^2}, \tag{79}$$

where $c_3' = (1-\epsilon_0)^2/(4(1+\epsilon_0))$ is the constant from (95). This requires $\mu^2 \geq C_8 d_z \log(n/\delta_0)$.

Define $c_4 = c_3'/(2\sigma_u^2)$, so $\lambda_{\max} \geq \exp(c_4 \mu^2 d_z)$. Using the exact identity $1 - \lambda_{\max}/(\lambda_{\max} + c_0) = c_0/(\lambda_{\max} + c_0)$:

$$\eta_1 \coloneqq 1 - \frac{\lambda_{\max}}{\lambda_{\max} + c_0} = \frac{c_0}{\lambda_{\max} + c_0} \leq \frac{c_0}{\exp(c_4 \mu^2 d_z)}. \tag{80}$$

Step 2: Quality gate close to 1. By Lemma C.4(b) (equation (127)), on an event of probability $\geq 1 - 3\delta_0 - \delta_{\mathcal{G}}$:

$$W_{\mathrm{CS}} \leq \frac{2\sigma_u \sqrt{C_{14} \log(n/\alpha)}}{\sqrt{\mu^2}}. \tag{81}$$

Verifying $W_{\mathrm{CS}} < w_{\max}$. Using $\mu^2 = n\kappa^2$ (where $\kappa = \|\pi\|/\sqrt{d_z}$ is the normalized instrument strength): $W_{\mathrm{CS}}/w_{\max} \leq 2\sigma_u \sqrt{C_{14} \log(n/\alpha)}/(w_{\max} \kappa \sqrt{n})$. Under the assumption $\mu^2 \geq Cd_z \log(n/\delta_0)$, choosing C large enough ensures $W_{\mathrm{CS}}/w_{\max} \leq 1$.

Hence:

$$\eta_2 \coloneqq 1 - \sqrt{1 - W_{\mathrm{CS}}/w_{\max}} \leq \frac{W_{\mathrm{CS}}}{w_{\max}} \leq \frac{2\sigma_u \sqrt{C_{14} \log(n/\alpha)}}{w_{\max} \kappa \sqrt{n}}, \tag{82}$$

where we used $\sqrt{1-x} \geq 1 - x$ for $x = W_{\mathrm{CS}}/w_{\max} \in [0,1]$.

Defining $c_5 = 2\sigma_u \sqrt{C_{14}}$:

$$\eta_2 \leq \frac{c_5 \sqrt{\log(n/\alpha)}}{\kappa \sqrt{n}\, w_{\max}}. \tag{83}$$

For $n \geq 1/\alpha$, $\log(n/\alpha) \leq 2 \log n$, so $\eta_2 \leq c_5 \sqrt{2 \log n}\,/(\kappa \sqrt{n}\, w_{\max})$.

Step 3: Combining. Write $\omega = (1-\eta_1)(1-\eta_2)$ where $\eta_1 = 1 - \lambda_{\max}/(\lambda_{\max} + c_0)$ and $\eta_2 = 1 - \sqrt{1 - W_{\mathrm{CS}}/w_{\max}}$, both in $[0,1]$. Then:

$$1 - \omega = 1 - (1-\eta_1)(1-\eta_2) = \eta_1 + \eta_2 - \eta_1 \eta_2 \leq \eta_1 + \eta_2. \tag{84}$$

Substituting the bounds from (80) and (83):

$$\omega \geq 1 - \frac{c_0}{\exp(c_4 \mu^2 d_z)} - \frac{c_5 \sqrt{\log(n/\alpha)}}{\kappa \sqrt{n}\, w_{\max}}.$$

(85)

The first deficit term decays exponentially in $\mu^2 d_z$ (the total signal-to-noise); the second decays as $O(\sqrt{\log n}/(\kappa\sqrt{n}))$.

Interpretation. The exponential term reflects the e-process detecting instrument relevance: once $\mu^2 d_z$ is moderately large, $\lambda_{\max}$ is astronomically large and the signal term is essentially 1. The polynomial term reflects the CS width shrinking: the e-process must not only detect signal but *localize* β^*, which requires $\sqrt{n}$ observations. In practice, the exponential term vanishes first (even for $\mu^2 d_z \approx 20$, $e^{-c_4 \cdot 20}$ is negligible), and the polynomial term governs the approach of ω to 1.

Failure probability: All bounds above hold on the intersection $\mathcal{G}(\epsilon_0) \cap \mathcal{S} \cap \mathcal{E}_{\mathrm{peak}} \cap \mathcal{E}_{\mathrm{cs}}$, where $\mathcal{E}_{\mathrm{peak}}$ and $\mathcal{E}_{\mathrm{cs}}$ are the concentration events from Lemmas C.1(b) and C.4(b) respectively. By the union bound, the conditional failure probability (given $\mathcal{G}$) is at most $5\delta_0$ ($3\delta_0$ from $\mathcal{E}_{\mathrm{peak}}$, δ_0 from $\mathcal{S}$, δ_0 from $\mathcal{E}_{\mathrm{cs}}$). By Corollary E.1(iii), the unconditional failure probability is $\leq 5\delta_0 + \delta_{\mathcal{G}} \leq 5\delta_0 + \delta_0/2$. Setting $\delta_0 = \delta_{\mathrm{fail}}/6$ gives the stated guarantee $1 - \delta_{\mathrm{fail}}$. □

C. E-Process Properties

We establish the finite-sample behavior of the e-process statistics that drive the weight ω.

Lemma C.1 (Peak e-value: monotonicity and concentration)**.** *Under Assumptions 2.1–2.3, for the multivariate GROW e-value (9) with $\tau^2 > 0$:*

(a) ***Null bound:*** *Under H_0 ($\pi = 0$): $\mathbb{E}_{H_0}[\log \lambda_{\max}] \leq \frac{1}{2} \log G + O(d_z)$, where G is the grid size and the $O(d_z)$ term arises from the log-determinant penalty.*

(b) ***Monotonicity:*** *Under H_1 ($\pi \neq 0$), the expected log e-value at a fixed $\beta_0 \neq \beta^*$ satisfies*

$$\mathbb{E}_{H_1}[\log E(\beta_0)] \geq c \cdot \mu^2 - \frac{d_z}{2} \log n - O(1), \tag{86}$$

for a constant $c > 0$ depending on $\tau^2, |\beta_0 - \beta^*|$, and the noise distribution. Since $\log \lambda_{\max} \geq \log E(\beta_0)$, the peak e-value grows at least linearly in μ^2, with a log-determinant offset that is negligible for $\mu^2 \gg d_z \log n$.

(c) ***Concentration:*** *For any $t > 0$, $\mathbb{P}_{H_1}(|\log \lambda_{\max} - \mathbb{E}_{H_1}[\log \lambda_{\max}]| > t) \leq 2\exp(-ct^2/\mu^2)$ for a constant $c > 0$ depending on τ^2 and the noise distribution.*

C.1. Proof of Lemma C.1 (Peak E-Value)

Proof. We analyze the multivariate log e-value (9) at a fixed grid point β_0. Write the eigen-decomposition $\boldsymbol{V} = Q\Lambda Q^\top$ with eigenvalues $\lambda_1, \ldots, \lambda_{d_z}$, and define the rotated statistics $S_{\text{rot},k} = (Q^\top \boldsymbol{S})_k = \sum_i (Q^\top \tilde{z}_i)_k (u_i - \bar{u})$ where $u_i = y_i - \beta_0 x_i$, and the mean and variance of u_i's are denoted by $\bar{u}$ and $\hat{\sigma}_u^2$.

Now, in the eigenbasis of $\boldsymbol{V}$, the log e-value decomposes as:

$$\log E(\beta_0) = \underbrace{-\frac{1}{2} \sum_{k=1}^{d_z} \log\left(1 + \frac{\tau^2 \lambda_k}{\hat{\sigma}_u^2}\right)}_{\text{penalty } P} + \underbrace{\frac{\tau^2}{2\hat{\sigma}_u^2} \sum_{k=1}^{d_z} \frac{S_{\text{rot},k}^2}{\tau^2 \lambda_k + \hat{\sigma}_u^2}}_{\text{payoff } R}. \tag{87}$$

The penalty $P < 0$ is the log-determinant term; the payoff $R \geq 0$ is the evidence from the data. We analyze each part separately.

Remark C.2 (Growth under H_1). Under the alternative $\beta_0 \neq \beta^*$, each rotated statistic has nonzero mean $\mathbb{E}[S_{\text{rot},k}] = (\beta^* - \beta_0) \sum_i (Q^\top \tilde{z}_i)_k \tilde{x}_i$, which is $O(n)$ when instruments are relevant. Hence $\mathbb{E}[R] = \Omega(n)$ while $|P| = O(d_z)$, so $\mathbb{E}[\log E(\beta_0)] \to \infty$ at rate $\Theta(n)$. The GROW-optimal τ^2 maximizes this rate; setting $\tau^2 = \hat{\sigma}_u^2 / \lambda_{\max}$ balances the penalty–payoff tradeoff when d_z is small.

Part (a): Null bound ($\pi = 0$). Under $\mathbb{P}_{H_0}$ with $\pi = 0$, we bound $\mathbb{E}_{H_0}[\log \lambda_{\max}]$.

E-value property at every β_0. When $\pi = 0$, we have $x_i = \varepsilon_i$, so $u_i = y_i - \beta_0 x_i = (\beta^* - \beta_0)\varepsilon_i + \eta_i$. By Assumptions 2.1–2.3, the instruments z_i are independent of (ε_i, η_i). Therefore, for *every* β_0:

$$\mathbb{E}_{H_0}[z_i u_i] = (\beta^* - \beta_0) \underbrace{\mathbb{E}[z_i \varepsilon_i]}_{=0} + \underbrace{\mathbb{E}[z_i \eta_i]}_{=0} = 0. \tag{88}$$

By the defining property of the GROW Bayes factor: $\mathbb{E}_{H_0}[E(\beta_0)] \leq 1$ for every β_0 in the grid.

Union bound over the grid. Taking the maximum over G grid points:

$$\mathbb{E}_{H_0}[\lambda_{\max}] = \mathbb{E}_{H_0}\left[\max_{j=1}^{G} E(\beta_0^{(j)})\right] \leq \sum_{j=1}^{G} \mathbb{E}_{H_0}[E(\beta_0^{(j)})] \leq G.$$

By Jensen's inequality (since $\log$ is concave):

$$\mathbb{E}_{H_0}[\log \lambda_{\max}] \leq \log \mathbb{E}_{H_0}[\lambda_{\max}] \leq \log G. \tag{89}$$

Log-determinant suppression. On $\mathcal{G}(\epsilon)$ with $\epsilon \le 1/2$, each eigenvalue satisfies $\lambda_k \ge n(1-\epsilon) \ge n/2$. The penalty term is:

$$P \le -\frac{d_z}{2}\log\left(1+\frac{\tau^2 n(1-\epsilon)}{\hat{\sigma}_u^2}\right).$$

Under H_0, $\hat{\sigma}_u^2 \ge \sigma_u^2/2$ on an event of probability $\ge 1-\delta_0$ by Lemma D.7. Hence: $P \le -(d_z/2)\log(1+\tau^2 n/(2\sigma_u^2))$. For $\tau^2 n \ge 2\sigma_u^2$ (which holds for any fixed $\tau^2 > 0$ and n large enough): $P \le -(d_z/2)\log(\tau^2 n/(2\sigma_u^2)) = -(d_z/2)(\log n + \log(\tau^2/(2\sigma_u^2)))$. This is $-\Theta(d_z \log n)$, which strongly suppresses the e-value in high dimensions.

Part (b): Signal lower bound ($\pi \neq 0$). Fix β_0 with $|\beta_0 - \beta^*| = \Delta > 0$. Then $u_i(\beta_0) = (\beta^* - \beta_0)x_i + \eta_i$, so:

$$\boldsymbol{S}(\beta_0) = \sum_i \tilde{z}_i(u_i - \bar{u}) = (\beta^* - \beta_0)\tilde{Z}^\top \boldsymbol{x} + \tilde{Z}^\top \boldsymbol{\eta} - n\bar{\tilde{z}}(\bar{u}). \tag{90}$$

Using $\boldsymbol{x} = Z\pi + \boldsymbol{\varepsilon}$ and $\tilde{Z}^\top Z\pi = \boldsymbol{V}\pi$ (after centering):

$$\boldsymbol{S} = \underbrace{(\beta^* - \beta_0)\boldsymbol{V}\pi}_{\text{signal } \boldsymbol{m}} + \underbrace{(\beta^* - \beta_0)\tilde{Z}^\top \boldsymbol{\varepsilon} + \tilde{Z}^\top \boldsymbol{\eta} - n\bar{\tilde{z}}\bar{u}}_{\text{noise } \boldsymbol{N}}. \tag{91}$$

The signal $\boldsymbol{m} = -\Delta \boldsymbol{V}\pi$ is deterministic given Z; the noise $\boldsymbol{N}$ is mean-zero conditional on Z (since $\mathbb{E}[\boldsymbol{\varepsilon} \mid Z] = 0$ and $\mathbb{E}[\boldsymbol{\eta} \mid Z] = 0$ by Assumption 2.1).

In the eigenbasis: $S_{\text{rot},k} = m_{\text{rot},k} + N_{\text{rot},k}$ where $m_{\text{rot},k} = -\Delta\lambda_k(Q^\top\pi)_k$ is the signal in direction k and $N_{\text{rot},k} = \sum_i (Q^\top \tilde{z}_i)_k \xi_i$ with $\xi_i = u_i - \bar{u} - (\beta^* - \beta_0)(\pi^\top \tilde{z}_i)$ denoting the noise component of the residual.

Key structural point. While the $N_{\text{rot},k}$'s are *uncorrelated* across k (since $\text{Cov}(N_{\text{rot},k}, N_{\text{rot},j} \mid Z) = \sigma_N^2 (Q^\top \boldsymbol{V} Q)_{kj} = \sigma_N^2 \lambda_k \delta_{kj}$), they are *not independent* for non-Gaussian data because they share the same ξ_i's. We therefore work directly in the ξ_i basis, expressing the cross and noise terms as linear and quadratic forms in the independent sub-Gaussian variables $\xi_1, \ldots, \xi_n$.

Define the n-vectors $\boldsymbol{a}_k = ((Q^\top \tilde{z}_1)_k, \ldots, (Q^\top \tilde{z}_n)_k)^\top \in \mathbb{R}^n$ for $k = 1, \ldots, d_z$, so that $N_{\text{rot},k} = \boldsymbol{a}_k^\top \boldsymbol{\xi}$. Note $\|\boldsymbol{a}_k\|^2 = \sum_i (Q^\top \tilde{z}_i)_k^2 = (Q^\top \boldsymbol{V} Q)_{kk} = \lambda_k$.

Bounding the payoff from below. Expanding the payoff R in (87):

$$\begin{aligned} R &= \frac{\tau^2}{2\hat{\sigma}_u^2}\sum_{k=1}^{d_z} \frac{S_{\text{rot},k}^2}{\tau^2\lambda_k + \hat{\sigma}_u^2} \\ &= \frac{\tau^2}{2\hat{\sigma}_u^2}\sum_k \frac{m_{\text{rot},k}^2 + 2m_{\text{rot},k}N_{\text{rot},k} + N_{\text{rot},k}^2}{\tau^2\lambda_k + \hat{\sigma}_u^2} \\ &= R_{\text{signal}} + R_{\text{cross}} + R_{\text{noise}}. \end{aligned} \tag{92}$$

Signal term R_{signal}. On $\mathcal{G}(\epsilon)$, each $\lambda_k \in [n(1-\epsilon), n(1+\epsilon)]$, so:

$$\begin{aligned} R_{\text{signal}} &= \frac{\tau^2}{2\hat{\sigma}_u^2}\sum_k \frac{\Delta^2\lambda_k^2 (Q^\top\pi)_k^2}{\tau^2\lambda_k + \hat{\sigma}_u^2} \\ &\ge \frac{\tau^2\Delta^2 n^2(1-\epsilon)^2}{2\hat{\sigma}_u^2(\tau^2 n(1+\epsilon) + \hat{\sigma}_u^2)}\sum_k (Q^\top\pi)_k^2 \\ &= \frac{\tau^2\Delta^2 n^2(1-\epsilon)^2\|\pi\|^2}{2\hat{\sigma}_u^2(\tau^2 n(1+\epsilon) + \hat{\sigma}_u^2)}, \end{aligned} \tag{93}$$

where we used $\sum_k (Q^\top\pi)_k^2 = \|\pi\|^2$ (since Q is orthogonal). Substituting $\|\pi\|^2 = \mu^2 d_z/n$:

$$R_{\text{signal}} \ge \frac{\tau^2\Delta^2 n(1-\epsilon)^2\mu^2 d_z}{2\hat{\sigma}_u^2(\tau^2 n(1+\epsilon) + \hat{\sigma}_u^2)}. \tag{94}$$

On the event $\mathcal{S}$ where $\hat{\sigma}_u^2 \le 2\sigma_u^2(\beta_0)$ (which holds with probability $\ge 1-\delta_0$ by Lemma D.7 for n large enough), and using $\tau^2 n/(\tau^2 n(1+\epsilon)+2\sigma_u^2) \ge 1/(1+\epsilon)$ for $\tau^2 n \ge 2\sigma_u^2$:

$$R_{\text{signal}} \ge \frac{(1-\epsilon)^2}{4(1+\epsilon)} \cdot \frac{\Delta^2\mu^2 d_z}{\sigma_u^2} =: c_3 \cdot \frac{\Delta^2\mu^2 d_z}{\sigma_u^2}, \tag{95}$$

with $c_3 = (1-\epsilon)^2/(4(1+\epsilon))$. For $\epsilon = 1/4$ (a concrete choice), $c_3 = 9/80$.

Cross term R_{cross} (linear form in $\boldsymbol{\xi}$). Rewrite the cross term using the ξ_i basis:

$$\begin{aligned} R_{\text{cross}} &= \frac{\tau^2}{\hat{\sigma}_u^2}\sum_k \frac{m_{\text{rot},k}N_{\text{rot},k}}{\tau^2\lambda_k+\hat{\sigma}_u^2} = \frac{\tau^2}{\hat{\sigma}_u^2}\sum_k \frac{m_{\text{rot},k}(\boldsymbol{a}_k^\top\boldsymbol{\xi})}{\tau^2\lambda_k+\hat{\sigma}_u^2} \\ &= \boldsymbol{w}^\top\boldsymbol{\xi}, \end{aligned} \tag{96}$$

where the weight vector $\boldsymbol{w} \in \mathbb{R}^n$ has entries

$$w_i = \frac{\tau^2}{\hat{\sigma}_u^2}\sum_{k=1}^{d_z} \frac{m_{\text{rot},k}(Q^\top\tilde{z}_i)_k}{\tau^2\lambda_k+\hat{\sigma}_u^2}. \tag{97}$$

Conditional on Z (and on the event $\mathcal{S}$ so that $\hat{\sigma}_u^2$ is treated as a function of Z and $\boldsymbol{\xi}$—see Part (c) for the precise conditioning), $\boldsymbol{w}$ is deterministic and $\boldsymbol{\xi}$ has independent sub-Gaussian components with parameter σ_N^2, where $\sigma_N^2 = \Delta^2\sigma_\varepsilon^2+\sigma_\eta^2$ is the residual variance of $u_i(\beta_0)$. By Lemma D.8, R_{cross} is sub-Gaussian with parameter $\sigma_N^2\|\boldsymbol{w}\|^2$.

Bounding $\|\boldsymbol{w}\|^2$.

$$\begin{aligned} \|\boldsymbol{w}\|^2 &= \frac{\tau^4}{\hat{\sigma}_u^4}\sum_i\left(\sum_k \frac{m_{\text{rot},k}(Q^\top\tilde{z}_i)_k}{\tau^2\lambda_k+\hat{\sigma}_u^2}\right)^2 \\ &= \frac{\tau^4}{\hat{\sigma}_u^4}\sum_{k,j}\frac{m_{\text{rot},k}\,m_{\text{rot},j}}{(\tau^2\lambda_k+\hat{\sigma}_u^2)(\tau^2\lambda_j+\hat{\sigma}_u^2)}\underbrace{\sum_i (Q^\top\tilde{z}_i)_k(Q^\top\tilde{z}_i)_j}_{=(Q^\top\boldsymbol{V}Q)_{kj}=\lambda_k\delta_{kj}} \\ &= \frac{\tau^4}{\hat{\sigma}_u^4}\sum_k \frac{m_{\text{rot},k}^2\lambda_k}{(\tau^2\lambda_k+\hat{\sigma}_u^2)^2}. \end{aligned} \tag{98}$$

On $\mathcal{G}(\epsilon)\cap\mathcal{S}$: $\lambda_k \le n(1+\epsilon)$, $\tau^2\lambda_k+\hat{\sigma}_u^2 \ge \tau^2 n(1-\epsilon)$, $m_{\text{rot},k}^2 = \Delta^2\lambda_k^2(Q^\top\pi)_k^2 \le \Delta^2 n^2(1+\epsilon)^2(Q^\top\pi)_k^2$. Summing over k:

$$\|\boldsymbol{w}\|^2 \le \frac{\tau^4\Delta^2 n^3(1+\epsilon)^3\|\pi\|^2}{\hat{\sigma}_u^4\tau^4 n^2(1-\epsilon)^2} = \frac{\Delta^2 n(1+\epsilon)^3\|\pi\|^2}{\hat{\sigma}_u^4(1-\epsilon)^2} \le \frac{4\Delta^2 n(1+\epsilon)^3\|\pi\|^2}{\sigma_u^4(1-\epsilon)^2}. \tag{99}$$

The sub-Gaussian tail gives: with probability $\ge 1-\delta_0$ over $\boldsymbol{\xi}$ conditional on Z,

$$|R_{\text{cross}}| \le \sqrt{2\sigma_N^2\|\boldsymbol{w}\|^2\log(2/\delta_0)} \le C_5\,\frac{\Delta\|\pi\|\sigma_N\sqrt{n\log(1/\delta_0)}}{\sigma_u^2} = C_5\,\frac{\Delta\mu\sqrt{d_z\log(1/\delta_0)}}{\sigma_u}, \tag{100}$$

where C_5 depends on ϵ and we used $\sigma_N \le C\sigma_u$ and $\|\pi\| = \mu\sqrt{d_z/n}$.

Key observation: R_{signal} grows as $\Delta^2\mu^2 d_z$ (quadratic in μ) while $|R_{\text{cross}}|$ grows as $\Delta\mu\sqrt{d_z\log(1/\delta_0)}$ (linear in μ). So for μ large enough, the signal dominates.

Noise term R_{noise} (quadratic form in $\boldsymbol{\xi}$). Rewrite:

$$\begin{aligned} R_{\text{noise}} &= \frac{\tau^2}{2\hat{\sigma}_u^2}\sum_k \frac{N_{\text{rot},k}^2}{\tau^2\lambda_k+\hat{\sigma}_u^2} = \frac{\tau^2}{2\hat{\sigma}_u^2}\sum_k \frac{(\boldsymbol{a}_k^\top\boldsymbol{\xi})^2}{\tau^2\lambda_k+\hat{\sigma}_u^2} \\ &= \frac{\tau^2}{2\hat{\sigma}_u^2}\,\boldsymbol{\xi}^\top M\boldsymbol{\xi}, \end{aligned} \tag{101}$$

where $M = \sum_{k=1}^{d_z} \boldsymbol{a}_k \boldsymbol{a}_k^\top / (\tau^2 \lambda_k + \hat{\sigma}_u^2) \in \mathbb{R}^{n \times n}$ is a positive semidefinite matrix. The conditional expectation is:

$$\mathbb{E}[R_{\text{noise}} \mid Z] = \frac{\tau^2 \sigma_N^2}{2\hat{\sigma}_u^2} \operatorname{tr}(M) = \frac{\tau^2 \sigma_N^2}{2\hat{\sigma}_u^2} \sum_k \frac{\lambda_k}{\tau^2 \lambda_k + \hat{\sigma}_u^2}, \tag{102}$$

since $\operatorname{tr}(M) = \sum_k \|\boldsymbol{a}_k\|^2 / (\tau^2 \lambda_k + \hat{\sigma}_u^2) = \sum_k \lambda_k / (\tau^2 \lambda_k + \hat{\sigma}_u^2)$.

On $\mathcal{G}(\epsilon)$, each ratio $\lambda_k / (\tau^2 \lambda_k + \hat{\sigma}_u^2) \le 1/\tau^2$, so $\mathbb{E}[R_{\text{noise}} \mid Z] \le d_z \sigma_N^2 / (2\hat{\sigma}_u^2) \le d_z$ (using $\sigma_N^2 \le 2\hat{\sigma}_u^2$ on $\mathcal{S}$).

For concentration, we apply the Hanson–Wright inequality (Rudelson and Vershynin, 2013) to the quadratic form $\boldsymbol{\xi}^\top M \boldsymbol{\xi}$. Since ξ_i are independent sub-Gaussian with parameter σ_N^2:

$$\mathbb{P}\Big(\big| \boldsymbol{\xi}^\top M \boldsymbol{\xi} - \sigma_N^2 \operatorname{tr}(M) \big| > t \;\Big|\; Z \Big) \le 2 \exp\left(-c \min\left(\frac{t^2}{\sigma_N^4 \|M\|_F^2}, \; \frac{t}{\sigma_N^2 \|M\|_{\text{op}}} \right) \right). \tag{103}$$

On $\mathcal{G}(\epsilon) \cap \mathcal{S}$: $\|M\|_{\text{op}} \le \max_k 1/(\tau^2 \lambda_k + \hat{\sigma}_u^2) \le 1/(\tau^2 n(1-\epsilon))$ and $\|M\|_F^2 \le d_z / (\tau^2 n (1-\epsilon))^2$. Setting $t = \sigma_N^2 \operatorname{tr}(M) \cdot s$ gives: with probability $\ge 1 - \delta_0$,

$$R_{\text{noise}} \le \frac{\tau^2 \sigma_N^2}{2\hat{\sigma}_u^2} \operatorname{tr}(M) \cdot (1 + C_6 \sqrt{d_z \log(1/\delta_0)/n}) \le C_7 d_z, \tag{104}$$

for C_7 depending on $\sigma, \tau^2, \Delta, \epsilon$.

Penalty bound. On $\mathcal{G}(\epsilon) \cap \mathcal{S}$:

$$|P| = \frac{1}{2} \sum_k \log\left(1 + \frac{\tau^2 \lambda_k}{\hat{\sigma}_u^2} \right) \le \frac{d_z}{2} \log\left(1 + \frac{2\tau^2 n (1+\epsilon)}{\sigma_u^2} \right). \tag{105}$$

Combining. On $\mathcal{G}(\epsilon) \cap \mathcal{S}$, with $\mathbb{P}_{H_1}$-probability $\ge 1 - 3\delta_0$:

$$\log E(\beta_0) = R + P \ge R_{\text{signal}} - |R_{\text{cross}}| - |P| + R_{\text{noise}}. \tag{106}$$

Since $R_{\text{noise}} \ge 0$, we can drop it for a lower bound:

$$\log E(\beta_0) \ge c_3 \frac{\Delta^2 \mu^2 d_z}{\sigma_u^2} - C_5 \frac{\Delta \mu \sqrt{d_z \log(1/\delta_0)}}{\sigma_u} - \frac{d_z}{2} \log\left(1 + \frac{2\tau^2 n}{\sigma_u^2} \right). \tag{107}$$

For $\mu^2 \ge C_8 d_z \log(n/\delta_0)$ with $C_8 = (C_5^2 + 1)/(c_3 \Delta^2)$ (so that $c_3 \Delta^2 \mu^2 d_z / \sigma_u^2$ dominates both the cross term and the penalty):

$$\log E(\beta_0) \ge \frac{c_3}{2} \cdot \frac{\Delta^2 \mu^2 d_z}{\sigma_u^2}. \tag{108}$$

Since $\log \lambda_{\max} \ge \log E(\beta_0)$ for any β_0 in the grid, choosing $\beta_0 = \beta_0^{(j^*)}$ as the grid point nearest to β^* (so $\Delta = |\beta_0^{(j^*)} - \beta^*| \le \Delta_{\text{grid}}/2 = \beta_{\max}/G$):

$$\mathbb{P}_{H_1}\left(\log \lambda_{\max} \ge \frac{c_3 \Delta_{j^*}^2 \mu^2 d_z}{2\sigma_u^2} \;\Big|\; Z \right) \ge 1 - 3\delta_0 - \delta_{\mathcal{G}}, \tag{109}$$

where $\Delta_{j^*} = |\beta_0^{(j^*)} - \beta^*|$ and $\delta_{\mathcal{G}} = 2d_z \exp(-n\epsilon^2/(8\tilde{R}^4))$. *Interpretation:* Under the alternative, $\log \lambda_{\max}$ grows linearly in μ^2 (the concentration parameter), with slope $c_3 d_z / (2\sigma_u^2)$ and an offset from the log-determinant penalty that is $O(d_z \log n)$. The signal exceeds this offset when $\mu^2 \ge C_8 d_z \log n$, which is the threshold for the e-process to reliably detect instrument relevance.

Part (c): Concentration of $\log E(\beta_0)$. We show that $\log E(\beta_0)$ concentrates around its conditional mean. The main difficulty is that $\log E$ is a *self-normalized* function of the residuals u_i: both the statistic $\boldsymbol{S}$ and the variance estimate $\hat{\sigma}_u^2$ depend on $\boldsymbol{u}$.

Step 1: Oracle-variance log e-value. Define the oracle-variance version (replacing $\hat{\sigma}_u^2$ with the population value σ_u^2):

$$\log \tilde{E}(\beta_0) = -\frac{1}{2}\log\det\left(I + \frac{\tau^2 \boldsymbol{V}}{\sigma_u^2}\right) + \frac{\tau^2}{2\sigma_u^2}\,\boldsymbol{S}^\top\left(\tau^2\boldsymbol{V} + \sigma_u^2 I\right)^{-1}\boldsymbol{S}. \tag{110}$$

The penalty $\tilde{P} = -\frac{1}{2}\log\det(I + \tau^2\boldsymbol{V}/\sigma_u^2)$ is deterministic given Z. The payoff is $\tilde{R} = \frac{\tau^2}{2\sigma_u^2}\,\boldsymbol{S}^\top(\tau^2\boldsymbol{V} + \sigma_u^2 I)^{-1}\boldsymbol{S}$.

Decompose $\boldsymbol{S} = \boldsymbol{m} + \boldsymbol{N}$ as in Part (b). In the eigenbasis, with $\boldsymbol{\xi}$ as before:

$$\tilde{R} = \frac{\tau^2}{2\sigma_u^2}\sum_k \frac{(m_{\mathrm{rot},k} + \boldsymbol{a}_k^\top\boldsymbol{\xi})^2}{\tau^2\lambda_k + \sigma_u^2} = \frac{\tau^2}{2\sigma_u^2}\,(\boldsymbol{m}_Q + A\boldsymbol{\xi})^\top D^{-1}(\boldsymbol{m}_Q + A\boldsymbol{\xi}), \tag{111}$$

where $A = [\boldsymbol{a}_1, \ldots, \boldsymbol{a}_{d_z}]^\top \in \mathbb{R}^{d_z\times n}$, $\boldsymbol{m}_Q = (m_{\mathrm{rot},1}, \ldots, m_{\mathrm{rot},d_z})^\top$, and $D = \mathrm{diag}(\tau^2\lambda_1 + \sigma_u^2, \ldots, \tau^2\lambda_{d_z} + \sigma_u^2)$. This is a quadratic function of $\boldsymbol{\xi}$ with deterministic (given Z) coefficients.

Step 2: Concentration of $\tilde{R}$. Write $\tilde{R} = \tilde{R}_{\mathrm{signal}} + \tilde{R}_{\mathrm{cross}} + \tilde{R}_{\mathrm{noise}}$. The signal part is deterministic given Z. The cross part $\tilde{R}_{\mathrm{cross}} = \tilde{\boldsymbol{w}}^\top\boldsymbol{\xi}$ is a linear form in $\boldsymbol{\xi}$ (with $\tilde{\boldsymbol{w}}$ the oracle-variance analog of $\boldsymbol{w}$ from Part (b)), so it is sub-Gaussian with parameter $\sigma_N^2\|\tilde{\boldsymbol{w}}\|^2$. The noise part $\tilde{R}_{\mathrm{noise}} = (\tau^2/(2\sigma_u^2))\boldsymbol{\xi}^\top\tilde{M}\boldsymbol{\xi}$ is a quadratic form in $\boldsymbol{\xi}$, where $\tilde{M} = \sum_k \boldsymbol{a}_k\boldsymbol{a}_k^\top/(\tau^2\lambda_k + \sigma_u^2)$.

Since $\boldsymbol{\xi}$ has independent sub-Gaussian components, the sum $\tilde{R}_{\mathrm{cross}} + \tilde{R}_{\mathrm{noise}}$ concentrates by combining the sub-Gaussian tail for the linear part (Lemma D.8) and Hanson–Wright for the quadratic part. By a union bound: with probability $\geq 1 - \delta_0$ over $\boldsymbol{\xi}$ conditional on Z,

$$|\tilde{R} - \mathbb{E}[\tilde{R} \mid Z]| \leq C_{12}\left(\sigma_N\|\tilde{\boldsymbol{w}}\|\sqrt{\log(1/\delta_0)} + \frac{\tau^2\sigma_N^2}{2\sigma_u^2}\left(\sigma_N^2\|\tilde{M}\|_F\sqrt{\log(1/\delta_0)} + \sigma_N^2\|\tilde{M}\|_{\mathrm{op}}\log(1/\delta_0)\right)\right). \tag{112}$$

On $\mathcal{G}(\epsilon)$, $\|\tilde{M}\|_{\mathrm{op}} = O(1/(\tau^2 n))$, $\|\tilde{M}\|_F = O(\sqrt{d_z}/(\tau^2 n))$, and $\|\tilde{\boldsymbol{w}}\|^2 = O(\Delta^2 n\|\pi\|^2/\sigma_u^4)$. Combining:

$$|\log\tilde{E}(\beta_0) - \mathbb{E}[\log\tilde{E}(\beta_0) \mid Z]| \leq t_1 \quad \text{with probability } \geq 1 - \delta_0, \tag{113}$$

where $t_1 = C_{13}\Delta\mu\sqrt{d_z\log(1/\delta_0)}/\sigma_u + C_{14}d_z\log(1/\delta_0)/n$.

Step 3: Effect of variance estimation. We bound $|\log E(\beta_0) - \log\tilde{E}(\beta_0)|$, the error from using $\hat{\sigma}_u^2$ instead of σ_u^2.

On $\mathcal{S}$: $|\hat{\sigma}_u^2 - \sigma_u^2| \leq \sigma_u^2/2$ (from Lemma D.7). Write $\hat{\sigma}_u^2 = \sigma_u^2(1+\gamma)$ with $|\gamma| \leq 1/2$. Then:

Penalty difference:

$$\begin{aligned}|P - \tilde{P}| &= \frac{1}{2}\left|\sum_k \log\frac{\tau^2\lambda_k + \sigma_u^2}{\tau^2\lambda_k + \hat{\sigma}_u^2}\right| = \frac{1}{2}\left|\sum_k \log\left(1 - \frac{\gamma\sigma_u^2}{\tau^2\lambda_k + \hat{\sigma}_u^2}\right)\right| \\ &\leq \frac{d_z|\gamma|\sigma_u^2}{2\tau^2 n(1-\epsilon)} \leq \frac{d_z|\gamma|}{2n\tau^2/(1-\epsilon)},\end{aligned} \tag{114}$$

using $\log(1+x) \leq |x|$ for $|x| < 1$.

Payoff difference: The payoff R depends on $\hat{\sigma}_u^2$ through both the prefactor $\tau^2/(2\hat{\sigma}_u^2)$ and the inverse $(\tau^2\boldsymbol{V} + \hat{\sigma}_u^2 I)^{-1}$. By a first-order expansion around σ_u^2:

$$|R - \tilde{R}| \leq C_{15}\,|\gamma|\,(|\tilde{R}| + d_z). \tag{115}$$

The $|\tilde{R}|$ term comes from the prefactor change; the d_z term comes from the inverse change (bounded by $d_z \cdot \sigma_u^2/(\tau^2 n(1-\epsilon))^2 \cdot |\gamma|\sigma_u^2 \cdot \|\boldsymbol{S}\|^2$, which is $O(d_z|\gamma|)$ on $\mathcal{G}(\epsilon)\cap\mathcal{S}$).

Combining (114) and (115):

$$|\log E(\beta_0) - \log\tilde{E}(\beta_0)| \leq C_{16}\,|\gamma|\,(|\log\tilde{E}(\beta_0)| + d_z). \tag{116}$$

Now, by Lemma D.7 (applied to the sub-Gaussian variables u_i^2), with probability $\geq 1 - \delta_0$:

$$|\gamma| = \frac{|\hat{\sigma}_u^2 - \sigma_u^2|}{\sigma_u^2} \leq C_{17}\sqrt{\frac{\log(1/\delta_0)}{n}}. \tag{117}$$

Step 4: Combining. By the triangle inequality and (113)–(116): with probability $\geq 1 - 3\delta_0$ over $\boldsymbol{\xi}$ conditional on Z,

$$\begin{aligned}
&|\log E(\beta_0) - \mathbb{E}[\log E(\beta_0) \mid Z]| \\
&\quad \leq |\log E(\beta_0) - \log \tilde{E}(\beta_0)| + |\log \tilde{E}(\beta_0) - \mathbb{E}[\log \tilde{E}(\beta_0) \mid Z]| + |\mathbb{E}[\log \tilde{E} \mid Z] - \mathbb{E}[\log E \mid Z]| \\
&\quad \leq C_{16}|\gamma|(|\log \tilde{E}| + d_z) + t_1 + C_{16}\,\mathbb{E}[|\gamma|](|\mathbb{E}[\log \tilde{E} \mid Z]| + d_z).
\end{aligned} \tag{118}$$

Under the alternative with $\mu^2 \geq C_8 d_z \log(n/\delta_0)$, $|\log E(\beta_0)|$ and $|\log \tilde{E}(\beta_0)|$ are both $O(\Delta^2\mu^2 d_z/\sigma_u^2)$. Using $|\gamma| = O(\sqrt{\log(1/\delta_0)/n})$ and $t_1 = O(\Delta\mu\sqrt{d_z \log(1/\delta_0)}/\sigma_u)$:

$$|\log E(\beta_0) - \mathbb{E}[\log E(\beta_0) \mid Z]| \leq C_{18}\left(\frac{\Delta\mu\sqrt{d_z \log(1/\delta_0)}}{\sigma_u} + \frac{\Delta^2\mu^2 d_z}{\sigma_u^2}\sqrt{\frac{\log(1/\delta_0)}{n}} + \frac{d_z \log(1/\delta_0)}{n}\right). \tag{119}$$

The dominant term is the first (from the cross term in $\tilde{R}$), which grows as $\mu\sqrt{d_z}$, compared to $\mathbb{E}[\log E \mid Z] \sim \mu^2 d_z$. So the relative deviation is $O(1/(\mu\sqrt{d_z}))$, confirming that $\log E$ concentrates tightly when instruments are strong.

Step 5: Extension to the peak. Since $\log \lambda_{\max} = \max_{j=1}^{G} \log E(\beta_0^{(j)})$, by a union bound over G grid points: each $\log E(\beta_0^{(j)})$ satisfies (119) simultaneously with probability $\geq 1 - 3G\delta_0$. Setting $\delta_0 \leftarrow \delta_0/(3G)$:

$$\mathbb{P}_{H_1}\left(|\log \lambda_{\max} - \mathbb{E}_{H_1}[\log \lambda_{\max} \mid Z]| > t \;\middle|\; Z\right) \leq 2\exp\left(-\frac{ct^2}{\mu^2 d_z}\right) + 3\delta_0, \tag{120}$$

for $t \geq C_{19}\mu\sqrt{d_z \log(G/\delta_0)}/\sigma_u$ (the interesting regime), where the exponential term comes from the sub-Gaussian tail of the cross term and c is a constant depending on σ, τ^2.

□

Remark C.3 (Self-normalization and concentration). The log e-value (9) is self-normalized: both the statistic $\boldsymbol{S}$ and the variance estimate $\hat{\sigma}_u^2$ depend on the data, creating a ratio-of-random-variables structure similar to the OLS bias (Appendix B.2). The concentration argument in Part (c) handles this by: (i) concentrating the oracle-variance version $\log \tilde{E}$ using Hanson–Wright (which applies because $\tilde{R}$ is a polynomial in independent sub-Gaussians with deterministic coefficients), and (ii) bounding the self-normalization error $|\log E - \log \tilde{E}|$ using the concentration of $\hat{\sigma}_u^2$ around σ_u^2.

The key implication: $\log \lambda_{\max}$ is a faithful, monotone proxy for the concentration parameter μ^2, aggregating evidence from all d_z instruments. This is what enables the signal term $\lambda_{\max}/(\lambda_{\max} + c_0)$ in the weight ω to track instrument strength non-asymptotically, providing the foundation for the oracle inequality (Theorem 4.3). The multivariate construction is essential when instruments have heterogeneous strengths: a scalar e-value using a single instrument would miss the signal from strong instruments not selected.

C.2. Proof of Lemma C.4 (CS Width)

Lemma C.4 (CS width convergence)**.** *Under Assumptions 2.1–2.3, the e-process CS width for the multivariate GROW e-value (9) satisfies:*

(a) Under weak instruments ($\mu^2 = O(1)$): $W_{\mathrm{CS}} \geq w_{\max} - O(1/\sqrt{n})$ with high probability (the CS is nearly vacuous). The log-determinant penalty strengthens this conclusion, as it further suppresses the e-value across the grid.

(b) Under strong instruments ($\mu^2 \geq C d_z \log n$): $W_{\mathrm{CS}} = O(\sigma\sqrt{d_z}/(\kappa\sqrt{n}))$ with high probability (the CS localizes around β^*). The threshold is $\mu^2 \geq C d_z \log n$ rather than $C \log n$ because the log-determinant penalty grows with d_z.

Proof. Recall $\mathcal{C}_\alpha = \{\beta_0 \in \text{grid} : E(\beta_0) < 1/\alpha\}$ (Definition **??**) and $W_{\mathrm{CS}} = \mathrm{diam}(\mathcal{C}_\alpha)$.

Part (a): Weak instruments ($\mu^2 \le C_w$). We show that nearly all grid points remain in $\mathcal{C}_\alpha$, so $W_{\mathrm{CS}} \approx w_{\max}$.

Step 1: E-value bound at each grid point. Under H_0 ($\pi = 0$) or near-H_0 ($\mu^2 \le C_w$), the e-value at $\beta_0 = \beta^*$ satisfies $\mathbb{E}_{H_0}[E(\beta^*)] \le 1$ (the e-value property). By Markov's inequality:

$$\mathbb{P}_{H_0}(E(\beta_0) \ge 1/\alpha) \le \alpha \quad \text{for } \beta_0 = \beta^*. \tag{121}$$

For $\beta_0 \ne \beta^*$, the same bound holds with the correction from Part (a) of Lemma C.1: $\mathbb{P}_{H_0}(E(\beta_0) \ge 1/\alpha) \le \alpha(1 + C_2|\beta^* - \beta_0|^2/n)$.

Step 2: Counting excluded points. The expected number of grid points excluded from $\mathcal{C}_\alpha$ (i.e., with $E(\beta_0^{(j)}) \ge 1/\alpha$) is at most:

$$\sum_{j=1}^{G} \mathbb{P}(E(\beta_0^{(j)}) \ge 1/\alpha) \le G\alpha(1 + C_2\beta_{\max}^2/n).$$

By Markov's inequality applied to this count, at most $G\alpha/\delta_0$ points are excluded with probability $\ge 1 - \delta_0$. Since the grid spacing is $\Delta_{\mathrm{grid}} = 2\beta_{\max}/G$, the width of the CS satisfies:

$$W_{\mathrm{CS}} \ge w_{\max} - \frac{G\alpha}{\delta_0} \cdot \Delta_{\mathrm{grid}} = w_{\max} - \frac{2\alpha\beta_{\max}}{\delta_0}. \tag{122}$$

Step 3: Refining with the log-determinant. The bound (122) can be tightened using the log-determinant suppression from Part (a) of Lemma C.1. Under weak instruments on $\mathcal{G}(\epsilon) \cap \mathcal{S}$:

$$\log E(\beta_0) \le R_{\mathrm{noise}} - |P| \le C_7 d_z - \frac{d_z}{2}\log\left(1 + \frac{\tau^2 n}{2\sigma_u^2}\right) < 0$$

for n large enough (specifically, $n \ge 2\sigma_u^2(e^{2C_7} - 1)/\tau^2$). When $\log E(\beta_0) < 0$, we have $E(\beta_0) < 1 < 1/\alpha$, so $\beta_0 \in \mathcal{C}_\alpha$. This means *every* grid point is in $\mathcal{C}_\alpha$ with probability $\ge 1 - 2\delta_0 - \delta_{\mathcal{G}}$, giving:

$$W_{\mathrm{CS}} \ge w_{\max} - C_{13}\sqrt{\alpha/n}, \tag{123}$$

where C_{13} absorbs the small-probability correction from the event where one or more e-values fluctuate above the threshold.

Part (b): Strong instruments ($\mu^2 \ge C_s\, d_z \log(n/\alpha)$). We show that grid points far from β^* are excluded from $\mathcal{C}_\alpha$, so W_{CS} is small.

Step 1: Exclusion condition. A grid point β_0 is excluded from $\mathcal{C}_\alpha$ when $E(\beta_0) \ge 1/\alpha$, i.e., $\log E(\beta_0) \ge \log(1/\alpha)$. From (107), this holds when:

$$c_3\frac{\Delta^2\mu^2 d_z}{\sigma_u^2} \ge \frac{d_z}{2}\log\left(1 + \frac{2\tau^2 n}{\sigma_u^2}\right) + C_5\frac{\Delta\mu\sqrt{d_z\log(1/\delta_0)}}{\sigma_u} + \log\frac{1}{\alpha}, \tag{124}$$

where $\Delta = |\beta_0 - \beta^*|$.

Step 2: Solving for the critical Δ. The LHS is quadratic in Δ while the RHS grows linearly in Δ (through the cross term) plus constants. Completing the square in Δ, the condition (124) is satisfied when:

$$\Delta^2 \ge \frac{\sigma_u^2}{c_3\mu^2}\left(\frac{1}{2}\log\frac{2\tau^2 n}{\sigma_u^2} + \frac{\log(1/\alpha)}{d_z}\right) + \frac{C_5^2\log(1/\delta_0)}{c_3^2\mu^2}. \tag{125}$$

Define:

$$\Delta_* = \sigma_u\sqrt{\frac{C_{14}\log(n/\alpha)}{\mu^2}} = \frac{\sigma_u\sqrt{C_{14}d_z\log(n/\alpha)}}{\kappa\sqrt{n \cdot d_z}}, \tag{126}$$

where C_{14} absorbs $1/c_3$, C_5^2/c_3^2, $1/2$, τ^2, and σ_u^2. For $\mu^2 \ge C_s d_z \log(n/\alpha)$ with C_s large enough (depending on C_{14}), the bound (125) holds for $\Delta = \Delta_*$.

Step 3: Bounding the CS width. Any grid point β_0 with $|\beta_0 - \beta^*| \geq \Delta_*$ satisfies $\log E(\beta_0) \geq \log(1/\alpha)$, hence $\beta_0 \notin \mathcal{C}_\alpha$. Therefore the confidence set $\mathcal{C}_\alpha$ is contained in the interval $[\beta^* - \Delta_*, \beta^* + \Delta_*]$, and:

$$W_{\mathrm{CS}} \leq 2\Delta_* = \frac{2\sigma_u\sqrt{C_{14}\log(n/\alpha)}}{\sqrt{\mu^2}} = \frac{2\sigma_u\sqrt{C_{14}d_z\log(n/\alpha)}}{\kappa\sqrt{n}} \tag{127}$$

with probability $\geq 1 - 3\delta_0 - \delta_{\mathcal{G}}$.

Step 4: Quality gate activation. For $\mu^2 \geq C_s d_z \log n$ with C_s large enough: $W_{\mathrm{CS}} \leq 2\sigma_u\sqrt{C_{14}\log(n/\alpha)/\mu^2} \leq w_{\max}/2$, ensuring the quality gate $\sqrt{1 - W_{\mathrm{CS}}/w_{\max}} \geq 1/\sqrt{2}$. □

D. Auxiliary Concentration Results

This section collects concentration inequalities that are used repeatedly in the main proofs. None of these results is new; we state them with explicit constants for completeness.

Lemma D.1 (First-stage residual variance concentration)**.** *Under Assumptions 1–3, let $M_Z = I_n - P_Z$ where $P_Z = Z(Z^\top Z)^{-1}Z^\top$. For any Z with $\mathrm{rank}(Z) = d_z$ and any $\delta_0 \in (0,1)$ with $n - d_z \geq 8\log(4/\delta_0)$:*

$$\left|\frac{\boldsymbol{\varepsilon}^\top M_Z \boldsymbol{\varepsilon}}{n - d_z} - \sigma^2\right| \leq \sigma^2\sqrt{\frac{8\log(4/\delta_0)}{n - d_z}}$$

with probability at least $1 - \delta_0$ over $\boldsymbol{\varepsilon} \mid Z$. This bound does not *require $\mathcal{G}(\epsilon)$: $\boldsymbol{\varepsilon}^\top M_Z\boldsymbol{\varepsilon}/\sigma^2 \sim \chi^2_{n-d_z}$ for every full-rank Z.*

Proof. **Step 1: Orthogonal decomposition.** M_Z has rank $n - d_z$; write $M_Z = UU^\top$ with $U \in \mathbb{R}^{n\times(n-d_z)}$, $U^\top U = I_{n-d_z}$. Then $\boldsymbol{\varepsilon}^\top M_Z\boldsymbol{\varepsilon} = \sum_{j=1}^{n-d_z} w_j^2$ where $w_j = u_j^\top\boldsymbol{\varepsilon}$.

Step 2: Distribution. Conditional on Z, $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma^2 I_n)$ (Assumption 2), so $w_1, \ldots, w_{n-d_z}$ are i.i.d. $\mathcal{N}(0,\sigma^2)$ and each w_j^2 is sub-exponential with parameters $(4\sigma^4, 4\sigma^2)$.

Step 3: Bernstein. By Bernstein's inequality (Vershynin, 2018, Thm 2.8.1) in the quadratic regime ($n - d_z \geq 8\log(4/\delta_0)$):

$$\mathbb{P}\left(\left|\hat{\sigma}_\varepsilon^2 - \sigma^2\right| > \sigma^2\sqrt{\frac{8\log(4/\delta_0)}{n - d_z}}\right) \leq 2\exp\bigl(-2c\log(4/\delta_0)\bigr) \leq \delta_0.$$ □

Lemma D.2 (First-stage F-statistic concentration corrected)**.** *Under Assumptions 1–3, for any Z satisfying $\mathcal{G}(\epsilon_0)$ with $\epsilon_0 \leq 1/4$, and any $\delta_0 \in (0,1)$ with $n - d_z \geq 8\log(4/\delta_0)$:*

$$\left|\hat{\kappa}^2 - \mathrm{SNR}\right| \leq 2\sqrt{\frac{2\ \mathrm{SNR}\cdot\log(4/\delta_0)}{n}} + \frac{C_F\cdot d_z\cdot\log(4/\delta_0)}{n} + C_F'\cdot\mathrm{SNR}\cdot\sqrt{\frac{\log(4/\delta_0)}{n - d_z}},$$

with probability at least $1 - \delta_0$ over $\boldsymbol{\varepsilon} \mid Z$, where $\hat{\kappa}^2 = \max(F - 1, 0)\cdot d_z/n$, $\mathrm{SNR} = \|\pi\|^2/\sigma^2$, and C_F, C_F' depend only on R.

Proof. **Setup.** $\hat{\boldsymbol{x}} = P_Z\boldsymbol{x} = Z\pi + P_Z\boldsymbol{\varepsilon}$, $\boldsymbol{v} = M_Z\boldsymbol{\varepsilon}$, and $F = (\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}/d_z)/(\boldsymbol{v}^\top\boldsymbol{v}/(n - d_z - 1))$.

Step 1: Numerator decomposition.

$$\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}} = \underbrace{\pi^\top Z^\top Z\pi}_{\text{signal}} + \underbrace{2\pi^\top Z^\top\boldsymbol{\varepsilon}}_{\text{cross}} + \underbrace{\boldsymbol{\varepsilon}^\top P_Z\boldsymbol{\varepsilon}}_{\text{noise}}. \tag{128}$$

*Signal (*needs $\mathcal{G}$*).* Decompose $Z^\top Z = \boldsymbol{V} + n\bar{z}\bar{z}^\top$ where $\boldsymbol{V} = \tilde{Z}^\top\tilde{Z}$. The cross term vanishes: $\mathbf{1}^\top\tilde{Z}\pi = \sum_i \pi^\top(z_i - \bar{z}) = 0$, so $\|Z\pi\|^2 = \pi^\top\boldsymbol{V}\pi + n(\bar{z}^\top\pi)^2$.

On $\mathcal{G}(\epsilon_0)$: $n(1 - \epsilon_0)\|\pi\|^2 \leq \pi^\top\boldsymbol{V}\pi \leq n(1 + \epsilon_0)\|\pi\|^2$.

Since $(\bar{z}^\top \pi)^2 \geq 0$:

$$\|Z\pi\|^2 \geq n(1-\epsilon_0)\|\pi\|^2. \tag{129}$$

Since $\|\bar{z}\| \leq R$ deterministically (Assumption 3):

$$\|Z\pi\|^2 \leq n(1+\epsilon_0)\|\pi\|^2 + nR^2\|\pi\|^2 = n(1+\epsilon_0+R^2)\|\pi\|^2. \tag{130}$$

Write $\pi^\top Z^\top Z\pi = n\|\pi\|^2(1+r_Z)$ with $-\epsilon_0 \leq r_Z \leq \epsilon_0 + R^2$. The upper bound is loose (by R^2), but R is a fixed constant; the looseness is absorbed into C_F, C'_F.

*Cross (*needs $\mathcal{G}$ for parameter only*).* Conditional on Z, $\pi^\top Z^\top \boldsymbol{\varepsilon}$ is sub-Gaussian with parameter $\sigma^2\|Z\pi\|^2$. Using (130): with probability $\geq 1-\delta_0/4$ over $\boldsymbol{\varepsilon} \mid Z$,

$$|2\pi^\top Z^\top \boldsymbol{\varepsilon}| \leq 2\sigma\|\pi\|\sqrt{2n(1+\epsilon_0+R^2)\log(8/\delta_0)}. \tag{131}$$

Since $1+\epsilon_0+R^2 \leq 2+R^2 \leq C_R$ (a constant depending only on R), this simplifies to $\leq C_R\,\sigma\|\pi\|\sqrt{n\log(1/\delta_0)}$.

*Noise (*any Z*).* $\boldsymbol{\varepsilon}^\top P_Z\boldsymbol{\varepsilon}/\sigma^2 \sim \chi^2_{d_z}$ for any full-rank Z. By Bernstein (same as Lemma D.1), with probability $\geq 1-\delta_0/4$:

$$|\boldsymbol{\varepsilon}^\top P_Z\boldsymbol{\varepsilon} - d_z\sigma^2| \leq \sigma^2\sqrt{8d_z\log(8/\delta_0)}. \tag{132}$$

Step 2: Denominator (*any* Z). By Lemma D.1 with failure probability $\delta_0/4$:

$$\boldsymbol{v}^\top\boldsymbol{v}/(n-d_z) = \sigma^2(1+r_\sigma), \quad |r_\sigma| \leq \Delta_\sigma := \sqrt{\frac{8\log(16/\delta_0)}{n-d_z}}. \tag{133}$$

Step 3: Combining. Write $\hat{\sigma}^2 = \boldsymbol{v}^\top\boldsymbol{v}/(n-d_z-1) = \sigma^2(1+r'_\sigma)$ with $|r'_\sigma| \leq 2\Delta_\sigma$. Then:

$$\hat{\kappa}^2 = \frac{\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}} - d_z\hat{\sigma}^2}{n\hat{\sigma}^2} = \frac{\|\pi\|^2}{\sigma^2}\cdot\frac{1+r_Z}{1+r'_\sigma} + \frac{2\pi^\top Z^\top\boldsymbol{\varepsilon}}{n\sigma^2(1+r'_\sigma)} + \frac{d_z(r_P - r'_\sigma)}{n(1+r'_\sigma)}.$$

(on the event $F > 1$, where the max is active).

Term 1 (signal, needs $\mathcal{G}$): For the *lower* direction of r_Z ($r_Z \geq -\epsilon_0$), $\mathrm{SNR}\cdot|(1+r_Z)/(1+r'_\sigma) - 1| \leq \mathrm{SNR}\cdot(|r_Z| + |r'_\sigma|)/(1-|r'_\sigma|) \leq C\ \mathrm{SNR}\cdot(\epsilon_0 + \Delta_\sigma)$.

Term 2 (cross, dominant): $|2\pi^\top Z^\top\boldsymbol{\varepsilon}|/(n\sigma^2(1+r'_\sigma)) \leq C_R\sqrt{\mathrm{SNR}\cdot\log(1/\delta_0)/n}$.

Term 3 (noise/denom, any Z): $d_z(|r_P| + |r'_\sigma|)/(n(1-|r'_\sigma|)) \leq C\,d_z\sqrt{\log(1/\delta_0)}/n + C\,d_z\Delta_\sigma/n$.

Step 4: Union bound. All four events (cross, noise, denominator, $\mathcal{G}(\epsilon_0)$) have failure probability $\leq \delta_0/4$ each. Combining and absorbing R-dependent constants into C_F, C'_F yields the stated bound. □

Lemma D.3 (Gram matrix eigenvalue concentration)**.** *Under Assumption 2.3, let $\tilde{z}_i = \Sigma_z^{-1/2}(z_i - \bar{z})$ denote the sample-centered and whitened instruments (Remark **??**), and let $\boldsymbol{V} = \lambda I_{d_z} + \sum_i \tilde{z}_i\tilde{z}_i^\top$.*

(a) For any $\epsilon \in (0,1)$:

$$\mathbb{P}\left(\left\|\frac{\boldsymbol{V}}{n} - \lambda I_{d_z}\right\|_{\mathrm{op}} > \epsilon\right) \leq 2d_z\exp\left(-\frac{n\epsilon^2}{8(1+\tilde{R}^2)^2}\right),$$

where $\tilde{R} = 2R/\sqrt{\lambda}$.

(b) In particular, on the event

$$\mathcal{G}(\epsilon) := \{\lambda_k(\boldsymbol{V}) \in [n(1-\epsilon),\, n(1+\epsilon)]\ \forall\, k \in [d_z]\}, \tag{134}$$

we have $\mathbb{P}(\mathcal{G}(\epsilon)) \geq 1 - 2d_z\exp(-n\epsilon^2/(8\tilde{R}^4))$. Setting $\epsilon = \tilde{R}^2\sqrt{8\frac{d_z}{n}\log(2d_z/\delta_0)}$ gives $\mathbb{P}(\mathcal{G}(\epsilon)) \geq 1-\delta_0$ provided $n \geq 8\tilde{R}^4 d_z\log(2d_z/\delta_0)$.

Lemma D.4 (Gram matrix eigenvalue concentration)**.** *Under Assumption 2.3 ($\mathbb{E}[z_i] = 0$, $\mathbb{E}[z_i z_i^\top] = I_{d_z}$, $\|z_i\|_2 \le R$), define the centered instruments $\tilde{z}_i = z_i - \bar{z}$ and the Gram matrix $\boldsymbol{V} = \sum_{i=1}^n \tilde{z}_i \tilde{z}_i^\top = \tilde{Z}^\top \tilde{Z}$. For any $\epsilon \in (0,1)$ and $\delta_0 \in (0,1)$, provided $n \ge 8R^4 d_z \log(2d_z/\delta_0)/\epsilon^2$:*

$$\mathbb{P}\left(\left\|\frac{\boldsymbol{V}}{n} - I_{d_z}\right\|_{\mathrm{op}} > \epsilon\right) \le 2d_z \exp\left(-\frac{n\epsilon^2}{32R^4}\right) + 2\exp\left(-\frac{n\epsilon^2}{8R^2}\right).$$

In particular, on the event $\mathcal{G}(\epsilon) := \{\lambda_k(\boldsymbol{V}) \in [n(1-\epsilon),\, n(1+\epsilon)]\ \forall\, k\}$, we have $\mathbb{P}(\mathcal{G}(\epsilon)^c) \le 2d_z \exp(-n\epsilon^2/(32R^4)) + 2\exp(-n\epsilon^2/(8R^2))$.

For $n \ge 128R^4 d_z \log(4d_z/\delta_0)$, setting $\epsilon_0 = 4R^2\sqrt{d_z \log(4d_z/\delta_0)/n}$ gives $\mathbb{P}(\mathcal{G}(\epsilon_0)) \ge 1 - \delta_0$.

Proof. We decompose the Gram matrix into a sum of independent terms plus a rank-one mean correction, bound each separately, and combine by the triangle inequality.

Step 1: Decomposition. Since $\tilde{z}_i = z_i - \bar{z}$:

$$\boldsymbol{V} = \sum_{i=1}^n (z_i - \bar{z})(z_i - \bar{z})^\top = \underbrace{\sum_{i=1}^n z_i z_i^\top}_{A} - \underbrace{n\,\bar{z}\bar{z}^\top}_{B}, \tag{135}$$

where $\bar{z} = n^{-1}\sum_i z_i$. Hence

$$\frac{\boldsymbol{V}}{n} - I_{d_z} = \underbrace{\left(\frac{A}{n} - I_{d_z}\right)}_{\text{(I)}} - \underbrace{\bar{z}\bar{z}^\top}_{\text{(II)}}. \tag{136}$$

By the triangle inequality for the spectral norm: $\|\boldsymbol{V}/n - I_{d_z}\|_{\mathrm{op}} \le \|\text{(I)}\|_{\mathrm{op}} + \|\text{(II)}\|_{\mathrm{op}}$. We bound each term separately.

Step 2: Term (I) — sum of independent matrices. Define $Y_i = z_i z_i^\top - I_{d_z}$ for $i = 1, \dots, n$. These are *independent* (since $z_1, \dots, z_n$ are i.i.d.), self-adjoint, and mean-zero ($\mathbb{E}[Y_i] = \mathbb{E}[z_i z_i^\top] - I_{d_z} = 0$).

Each summand is bounded:

$$\|Y_i\|_{\mathrm{op}} = \|z_i z_i^\top - I_{d_z}\|_{\mathrm{op}} \le \|z_i z_i^\top\|_{\mathrm{op}} + \|I_{d_z}\|_{\mathrm{op}} = \|z_i\|_2^2 + 1 \le R^2 + 1.$$

Set $L = R^2 + 1 \le 2R^2$ (using $R \ge 1$, which holds without loss of generality since $\mathbb{E}[\|z_i\|^2] = d_z \ge 1$ and $R \ge \|z_i\|$; if $R < 1$ one can use $L = R^2 + 1$ directly).

By the matrix Hoeffding inequality (**?**, Theorem 1.3): if $Y_1, \dots, Y_n$ are independent self-adjoint random matrices in $\mathbb{R}^{d\times d}$ with $\mathbb{E}[Y_i] = 0$ and $\|Y_i\|_{\mathrm{op}} \le L$ a.s., then

$$\mathbb{P}\left(\left\|\sum_{i=1}^n Y_i\right\|_{\mathrm{op}} > t\right) \le 2d \cdot \exp\left(-\frac{t^2}{8nL^2}\right). \tag{137}$$

Note: Tropp's Theorem 1.3 is stated for $\lambda_{\max}(\sum Y_i)$, but for self-adjoint matrices $\|\sum Y_i\|_{\mathrm{op}} = \max(|\lambda_{\max}|, |\lambda_{\min}|)$. Applying the one-sided bound to both $\{Y_i\}$ and $\{-Y_i\}$ and taking a union bound yields (137) with the factor $2d$ (rather than d).

Applying (137) with $d = d_z$, $L = R^2 + 1$, and $t = n\epsilon/2$ (we allocate half the budget ϵ to each term):

$$\mathbb{P}\left(\left\|\frac{A}{n} - I_{d_z}\right\|_{\mathrm{op}} > \frac{\epsilon}{2}\right) \le 2d_z \exp\left(-\frac{n\epsilon^2}{32(R^2+1)^2}\right) \le 2d_z \exp\left(-\frac{n\epsilon^2}{32R^4}\right), \tag{138}$$

where the last step uses $(R^2+1)^2 \le (2R^2)^2 = 4R^4$ for $R \ge 1$.

Step 3: Term (II) — sample mean correction. The rank-one term satisfies $\|\bar{z}\bar{z}^\top\|_{\mathrm{op}} = \|\bar{z}\|_2^2$. We bound $\|\bar{z}\|_2^2$ using the sub-Gaussian property of bounded random vectors.

Since $\mathbb{E}[z_i] = 0$ and $\|z_i\|_2 \le R$, each component $z_{i,j}$ is sub-Gaussian with parameter R^2 (by Hoeffding's lemma applied to bounded random variables supported on $[-R, R]$). Further, $\|\bar{z}\|_2^2 = \|n^{-1}\sum_i z_i\|_2^2$. By the χ^2-type tail for sums of bounded random vectors (see, e.g., Vershynin (2018, Proposition 2.2)):

$$\mathbb{P}\Big(\|\bar{z}\|_2^2 > \frac{\epsilon}{2}\Big) \le 2\exp\Big(-\frac{n\epsilon^2}{8R^2}\Big), \tag{139}$$

where we used $\mathbb{E}[\|\bar{z}\|_2^2] = d_z/n$, which is $\le \epsilon/4$ for $n \ge 4d_z/\epsilon$.

Step 4: Combining. By the triangle inequality (136) and a union bound over (138) and (139):

$$\begin{aligned}\mathbb{P}\Bigg(\Big\|\frac{\boldsymbol{V}}{n} - I_{d_z}\Big\|_{\mathrm{op}} > \epsilon\Bigg) &\le \mathbb{P}\Bigg(\Big\|\frac{A}{n} - I_{d_z}\Big\|_{\mathrm{op}} > \frac{\epsilon}{2}\Bigg) + \mathbb{P}\Big(\|\bar{z}\|_2^2 > \frac{\epsilon}{2}\Big) \\ &\le 2d_z\exp\Big(-\frac{n\epsilon^2}{32R^4}\Big) + 2\exp\Big(-\frac{n\epsilon^2}{8R^2}\Big).\end{aligned} \tag{140}$$

Since $R \ge 1$ (so $R^4 \ge R^2$), the first term dominates, and we obtain the simplified bound:

$$\mathbb{P}\Bigg(\Big\|\frac{\boldsymbol{V}}{n} - I_{d_z}\Big\|_{\mathrm{op}} > \epsilon\Bigg) \le 4d_z\exp\Big(-\frac{n\epsilon^2}{32R^4}\Big). \tag{141}$$

Since $\|\boldsymbol{V}/n - I_{d_z}\|_{\mathrm{op}} \le \epsilon$ implies $|\lambda_k(\boldsymbol{V}/n) - 1| \le \epsilon$ for all k (because $\|A\|_{\mathrm{op}} = \max_k |\lambda_k(A)|$ for self-adjoint A), we have $\mathcal{G}(\epsilon) \supseteq \{\|\boldsymbol{V}/n - I_{d_z}\|_{\mathrm{op}} \le \epsilon\}$, and the stated tail bound for $\mathbb{P}(\mathcal{G}(\epsilon)^c)$ follows.

Concrete calibration. Setting $\epsilon_0 = 4R^2\sqrt{d_z\log(4d_z/\delta_0)/n}$ in (141):

$$\mathbb{P}(\mathcal{G}(\epsilon_0)^c) \le 4d_z\exp\Big(-\frac{n\cdot 16R^4\cdot d_z\log(4d_z/\delta_0)/n}{32R^4}\Big) = 4d_z\cdot\Big(\frac{\delta_0}{4d_z}\Big) = \delta_0.$$

This requires $\epsilon_0 < 1$, i.e., $n > 16R^4 d_z\log(4d_z/\delta_0)$. On $\mathcal{G}(\epsilon_0)$ with this choice, all eigenvalues satisfy $\lambda_k(\boldsymbol{V}) \in [n(1-\epsilon_0),\, n(1+\epsilon_0)]$ and $\epsilon_0 \le 1/2$ for $n \ge 128R^4 d_z\log(4d_z/\delta_0)$. □

Corollary D.5 (Canonical good event and bad-event handling)**.** *Under the conditions of Lemma D.3, fix any target failure probability $\delta_0 \in (0,1)$ and set*

$$\epsilon_0 = \tilde{R}^2\sqrt{\frac{8d_z\log(4d_z/\delta_0)}{n}}, \qquad \delta_{\mathcal{G}} = 2d_z\exp\Big(-\frac{n\epsilon_0^2}{8\tilde{R}^4}\Big) = \frac{\delta_0}{2}.$$

Then, provided $n \ge 8\tilde{R}^4 d_z\log(4d_z/\delta_0)$ (equivalently $\epsilon_0 \le 1$):

(i) *$\epsilon_0 \le 1$ and $\mathbb{P}(\mathcal{G}(\epsilon_0)) \ge 1 - \delta_0/2$. For the concrete choice $\epsilon_0 = 1/4$, the sample-size requirement becomes $n \ge 128\tilde{R}^4 d_z\log(4d_z/\delta_0)$.*

(ii) *On $\mathcal{G}(\epsilon_0)$, every eigenvalue of $\boldsymbol{V}$ lies in $[n(1-\epsilon_0),\, n(1+\epsilon_0)]$, providing the eigenvalue sandwich used in all subsequent proofs (Lemmas C.1–C.4 and Propositions B.7–4.3).*

(iii) ***Bad-event handling.*** *On $\mathcal{G}(\epsilon_0)^c$, the StAR estimator remains well-defined: the e-values $E(\beta_0)$ are non-negative by construction (they are likelihood ratios), the weight $\omega \in [0,1]$ by the $\max(0,\cdot)$ and clamping in (10), and the blended estimator $\hat{\beta}_{\mathrm{StAR}} = \omega\hat{\beta}_{\mathrm{TSLS}} + (1-\omega)\hat{\beta}_{\mathrm{CS}}$ is a convex combination of finite estimators. Therefore, for any event $\mathcal{A}$ and any bound derived conditionally on $\mathcal{G}(\epsilon_0)$:*

$$\mathbb{P}(\mathcal{A}) = \mathbb{P}(\mathcal{A}\cap\mathcal{G}) + \mathbb{P}(\mathcal{A}\cap\mathcal{G}^c) \le \mathbb{P}(\mathcal{A}\mid\mathcal{G}) + \mathbb{P}(\mathcal{G}^c) \le p_{\mathcal{G}} + \frac{\delta_0}{2},$$

where $p_{\mathcal{G}}$ is the conditional failure probability on $\mathcal{G}(\epsilon_0)$. This additive $\delta_0/2$ is absorbed into the total failure budget in every subsequent proof by reserving half the budget for $\mathcal{G}$ and the remaining half for the conditional analysis.

Proof. Part (i): Substituting ϵ_0 into Lemma D.3: $\delta_{\mathcal{G}} = 2d_z \exp(-d_z \log(4d_z/\delta_0)) = 2d_z \cdot (\delta_0/(4d_z))^{d_z}$. For $d_z = 1$: $\delta_{\mathcal{G}} = 2 \cdot \delta_0/4 = \delta_0/2$. For $d_z \geq 2$: $\delta_{\mathcal{G}} = 2d_z(\delta_0/(4d_z))^{d_z} \leq 2d_z \cdot (\delta_0/(4d_z)) = \delta_0/2$, since $(\delta_0/(4d_z))^{d_z-1} \leq 1$ for $\delta_0 \leq 4d_z$. Hence $\delta_{\mathcal{G}} \leq \delta_0/2$ in all cases. The bound $\epsilon_0 \leq 1$ follows from the sample-size condition.

Part (ii): Immediate from the definition of $\mathcal{G}(\epsilon_0)$.

Part (iii): The e-value at any β_0 is $E(\beta_0) = \int p(u \mid \beta_0, \sigma^2)/p_0(u \mid \beta_0)\, d\Pi(\sigma^2) \geq 0$. The weight formula $\omega = (\lambda_{\max}/(\lambda_{\max} + c_0)) \cdot \sqrt{\max(0, 1 - W_{\mathrm{CS}}/w_{\max})}$ maps $(\lambda_{\max}, W_{\mathrm{CS}}) \in [0, \infty) \times [0, \infty)$ into $[0, 1]$ continuously. On $\mathcal{G}^c$, only the quantitative bounds on eigenvalues (and hence on $\lambda_{\max}$ and W_{CS}) may fail; the estimator itself and its range $[0, 1]$ are unaffected. The total-law decomposition follows. □

Remark D.6 (Convention for subsequent proofs). Throughout Appendices C.1–B.5, all probabilistic bounds are stated conditionally on $\mathcal{G}(\epsilon_0)$ with $\epsilon_0 \leq 1/2$ (Corollary E.1). The failure probability $\delta_{\mathcal{G}} \leq \delta_0/2$ is added once at the end via the total-law decomposition in part (iii). When we write "on $\mathcal{G}(\epsilon)$" without specifying ϵ, we mean $\epsilon = \epsilon_0$ from the canonical choice.

Lemma D.7 (Residual variance concentration)**.** *Under Assumptions 2.1–2.3, let $\sigma_u^2(\beta_0) = \mathrm{Var}(y_i - \beta_0 x_i \mid z_i)$ be the population residual variance. Then for any β_0 and any $t > 0$:*

$$\mathbb{P}\big(|\hat{\sigma}_u^2(\beta_0) - \sigma_u^2(\beta_0)| > t\big) \leq 2\exp\left(-\frac{nt^2}{32\sigma^4(1 + |\beta^* - \beta_0|^2 R^2)^2}\right).$$

In particular, on the event $\mathcal{S}(\beta_0, \delta_0)$ where $|\hat{\sigma}_u^2 - \sigma_u^2| \leq 4\sigma^2(1 + |\beta^* - \beta_0|^2 R^2)\sqrt{2\log(2/\delta_0)/n}$, we have $\mathbb{P}(\mathcal{S}) \geq 1 - \delta_0$.

Proof. Write $u_i = y_i - \beta_0 x_i = (\beta^* - \beta_0)x_i + \eta_i = (\beta^* - \beta_0)(\pi^\top z_i + \varepsilon_i) + \eta_i$.

Step 1: Sub-Gaussianity of u_i. Conditional on Z, the term $(\beta^* - \beta_0)\varepsilon_i$ is sub-Gaussian with parameter $(\beta^* - \beta_0)^2\sigma^2$, and η_i is sub-Gaussian with parameter σ^2, both by Assumption 2.2. The deterministic term $(\beta^* - \beta_0)\pi^\top z_i$ shifts the mean but does not affect sub-Gaussianity. Hence $u_i - \mathbb{E}[u_i \mid z_i]$ is sub-Gaussian with parameter $\sigma_u^2 \leq \sigma^2(1 + (\beta^* - \beta_0)^2) \leq \sigma^2(1 + |\beta^* - \beta_0|^2 R^2)$ (the R^2 factor accounts for the dependence of $\pi^\top z_i$ on z_i when bounding the conditional variance).

Step 2: Concentration of sample variance. The sample variance $\hat{\sigma}_u^2 = (n-1)^{-1}\sum_i (u_i - \bar{u})^2$ satisfies $|\hat{\sigma}_u^2 - \sigma_u^2| \leq |n^{-1}\sum_i u_i^2 - \mathbb{E}[u_i^2 \mid Z]| + |\bar{u} - \mathbb{E}[u_i \mid Z]|^2 + O(1/n)$.

For the first term: since $u_i - \mathbb{E}[u_i \mid z_i]$ is sub-Gaussian(σ_u^2), the centered square $(u_i - \mathbb{E}[u_i \mid z_i])^2 - \sigma_u^2$ is sub-exponential with parameter $4\sigma_u^2$ (by the standard sub-Gaussian-to-sub-exponential conversion). By Bernstein's inequality for sub-exponential random variables:

$$\mathbb{P}\left(\left|\frac{1}{n}\sum_i \big((u_i - \mathbb{E}[u_i \mid z_i])^2 - \sigma_u^2\big)\right| > t\right) \leq 2\exp\left(-\frac{n\min(t^2/(16\sigma_u^4),\ t/(4\sigma_u^2))}{2}\right).$$

In the quadratic regime $t \leq 4\sigma_u^2$, this gives $\leq 2\exp(-nt^2/(32\sigma_u^4))$.

For the second term: $|\bar{u} - \mathbb{E}[u_i \mid Z]|^2 \leq 2\sigma_u^2 \log(2/\delta_0)/n$ with probability $\geq 1 - \delta_0$ by sub-Gaussian mean concentration. This is $O(\sigma_u^2/n)$ and is absorbed for $n \geq 2$. □

Lemma D.8 (Sub-Gaussian linear form (Vershynin, 2018))**.** *Let $w_1, \ldots, w_n$ be independent sub-Gaussian random variables with parameter σ^2 (i.e., $\mathbb{E}[e^{\lambda w_i}] \leq e^{\sigma^2\lambda^2/2}$ for all λ), and let $a = (a_1, \ldots, a_n) \in \mathbb{R}^n$ be a fixed vector. Then $S = \sum_i a_i w_i$ is sub-Gaussian with parameter $\sigma^2\|a\|_2^2$:*

$$\mathbb{P}(|S| \geq t) \leq 2\exp\left(-\frac{t^2}{2\sigma^2\|a\|_2^2}\right).$$

Lemma D.9 (Sub-Gaussian quadratic form (Vershynin, 2018))**.** *Under the conditions of Lemma D.8, for $Q = S^2$ where S is sub-Gaussian with parameter $\sigma^2 A$ ($A = \|a\|_2^2$):*

$$\mathbb{P}(Q \geq \sigma^2 A(1 + t)) \leq \exp(-c_{\mathrm{sq}} \min(t^2, t) \cdot A/\sigma^2) \quad \textit{for all } t > 0,$$

where $c_{\mathrm{sq}} > 0$ is a universal constant. Furthermore, for d independent terms $T = \sum_{k=1}^{d} S_k^2/b_k$ with S_k sub-Gaussian($\sigma^2 A_k$) independent and $b_k > 0$:

$$\mathbb{P}\left(T \geq \sum_k \frac{\sigma^2 A_k}{b_k}(1+t)\right) \leq \exp\left(-c_{\mathrm{sq}} \min(t^2, t) \cdot \min_k \frac{A_k}{b_k \sigma^2}\right) \quad \textit{for all } t > 0.$$

Lemma D.10 (Self-normalized sums (de la Peña et al., 2009))**.** *Let $w_1, \ldots, w_n$ be i.i.d. sub-Gaussian with mean μ_w and variance σ_w^2. Then:*

$$\mathbb{P}\left(\left|\frac{\sum_i w_i - n\mu_w}{\sqrt{\sum_i (w_i - \bar{w})^2}}\right| \geq t\right) \leq 2\exp(-c_{\mathrm{sn}} t^2) \quad \textit{for } t \leq c'_{\mathrm{sn}}\sqrt{n},$$

where $c_{\mathrm{sn}}, c'_{\mathrm{sn}} > 0$ are universal constants (they do not depend on σ_w, μ_w, n).

E. Canonical Good Event and Bad-Event Handling

All bounds in the MSE proof (Theorem 4.1), the oracle inequality (Theorem 4.3), and the width rate (Theorem 4.4) condition on the Gram concentration event $\mathcal{G}(\epsilon_0)$ from Lemma D.3. We now establish a canonical setup and show that the contribution from $\mathcal{G}^c$ is controlled.

Corollary E.1 (Good event and bad-event decomposition)**.** *Fix $\epsilon_0 = R^2\sqrt{8d_z \log(2d_z/\delta_0)/n}$ and $\delta_{\mathcal{G}} = 2d_z \exp(-n\epsilon_0^2/(8R^4)) \leq \delta_0$. Then:*

(i) Eigenvalue sandwich on $\mathcal{G}(\epsilon_0)$. *On the event $\mathcal{G}(\epsilon_0)$ (probability $\geq 1 - \delta_{\mathcal{G}}$), the Gram matrix $\boldsymbol{V} = \sum_i \tilde{z}_i \tilde{z}_i^\top$ has all eigenvalues in $[n(1-\epsilon_0),\, n(1+\epsilon_0)]$. Consequently:*

$$n\|\boldsymbol{\pi}\|^2(1-\epsilon_0) \;\leq\; \hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}} \;\leq\; n\|\boldsymbol{\pi}\|^2(1+\epsilon_0), \tag{142}$$

and the IV estimator $\hat{\beta}_{\mathrm{IV}} = \hat{\boldsymbol{x}}^\top \boldsymbol{y}/(\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}})$ has well-controlled variance $\mathrm{MSE}_{\mathrm{IV}}(Z) = \sigma_\eta^2/(\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}) \leq \sigma_\eta^2/(n\|\boldsymbol{\pi}\|^2(1-\epsilon_0)) = \sigma_\eta^2 d_z/(\mu^2(1-\epsilon_0))$.

(ii) Well-definedness on $\mathcal{G}^c$. *On $\mathcal{G}(\epsilon_0)^c$, the IV estimator remains well-defined: $\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}} = \|P_Z \boldsymbol{x}\|^2 > 0$ almost surely under continuous error distributions (Assumption **??**), since $\boldsymbol{x} = Z\boldsymbol{\pi} + \boldsymbol{v}$ with $\boldsymbol{\pi} \neq 0$ implies $P_Z \boldsymbol{x} = Z\boldsymbol{\pi} + P_Z \boldsymbol{v} \neq 0$ a.s. The StAR confidence interval $[\hat{\beta}_\omega - h_\omega,\, \hat{\beta}_\omega + h_\omega]$ is therefore always a bounded interval.*

(iii) Total-law decomposition for MSE. *For any measurable event $\mathcal{A}$ and any random variable W:*

$$\mathbb{E}[W] = \mathbb{E}[W \mid \mathcal{G}]\, \mathbb{P}(\mathcal{G}) + \mathbb{E}[W \mid \mathcal{G}^c]\, \mathbb{P}(\mathcal{G}^c).$$

Apply this with $W = (\hat{\beta}_\omega - \beta^*)^2$. On $\mathcal{G}(\epsilon_0)$, the MSE bounds from Theorem 4.1 apply. On $\mathcal{G}(\epsilon_0)^c$, we use the deterministic bound:

$$(\hat{\beta}_\omega - \beta^*)^2 = \left(\omega \frac{\hat{\boldsymbol{x}}^\top \boldsymbol{\eta}}{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}} + (1-\omega)\frac{\boldsymbol{x}^\top \boldsymbol{\eta}}{\boldsymbol{x}^\top \boldsymbol{x}}\right)^2 \leq 2\left(\frac{(\hat{\boldsymbol{x}}^\top \boldsymbol{\eta})^2}{(\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}})^2} + \frac{(\boldsymbol{x}^\top \boldsymbol{\eta})^2}{(\boldsymbol{x}^\top \boldsymbol{x})^2}\right). \tag{143}$$

Under sub-Gaussian errors with bounded instruments ($\|z_i\| \leq R$), $\mathbb{E}[(\boldsymbol{x}^\top \boldsymbol{\eta})^2/(\boldsymbol{x}^\top \boldsymbol{x})^2 \mid Z] \leq \sigma_\eta^2/(\boldsymbol{x}^\top \boldsymbol{x})$, which has finite expectation over Z (since $\boldsymbol{x}^\top \boldsymbol{x} \geq \|\boldsymbol{v}\|^2 = \|M_Z \boldsymbol{\varepsilon}\|^2 > 0$ a.s.). The IV term is more delicate: $\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}$ can be small on $\mathcal{G}^c$, making $\mathrm{MSE}_{\mathrm{IV}}$ large. However, the contribution to the unconditional MSE is bounded by:

$$\mathbb{E}[\mathrm{MSE}_{\mathrm{IV}} \cdot \mathbf{1}_{\mathcal{G}^c}] \leq \mathbb{E}\left[\frac{\sigma_\eta^2}{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}} \cdot \mathbf{1}_{\mathcal{G}^c}\right] \leq \frac{\sigma_\eta^2}{n\|\boldsymbol{\pi}\|^2(1-2\epsilon_0)} \cdot \frac{\delta_{\mathcal{G}}}{1-\delta_{\mathcal{G}}} + \sigma_\eta^2\, \mathbb{E}\left[\frac{\mathbf{1}_{\{\lambda_{\min}(\boldsymbol{V}) < n(1-2\epsilon_0)\}}}{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}}\right]. \tag{144}$$

The second term involves the deep tail where eigenvalues deviate by more than $2\epsilon_0$. By the matrix Hoeffding bound (Lemma D.3), this probability is $\leq 2d_z \exp(-n\epsilon_0^2/(2R^4)) = \delta_{\mathcal{G}}^2/(2d_z)$, which is super-polynomially small. Specifically, for any $p > 0$:

$$\mathbb{E}\left[\frac{\mathbf{1}_{\{\lambda_{\min} < n(1-2\epsilon_0)\}}}{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}}\right] \leq \frac{1}{n\|\boldsymbol{\pi}\|^2} \cdot \mathbb{E}\left[\frac{n\|\boldsymbol{\pi}\|^2}{\hat{\boldsymbol{x}}^\top \hat{\boldsymbol{x}}} \cdot \mathbf{1}_{\mathcal{G}^c}\right] \leq \frac{C_p}{n\|\boldsymbol{\pi}\|^2} \cdot n^{-p},$$

where the last step uses Gaussian concentration of $\lambda_{\min}(\boldsymbol{V})$ (the eigenvalue deviates from n by more than $n\epsilon_0$ only with probability that decays faster than any polynomial in n).

Therefore, the unconditional MSE satisfies:

$$\mathbb{E}[\mathrm{MSE}(\hat{\beta}_\omega)] \le \mathbb{E}[\mathrm{MSE}(\hat{\beta}_\omega) \mid \mathcal{G}] + \frac{C\sigma_\eta^2}{\mu^2} \cdot \delta_{\mathcal{G}}. \tag{145}$$

Since $\delta_{\mathcal{G}} \le \delta_0$, the bad-event contribution is at most $C\sigma_\eta^2\delta_0/\mu^2$, which is negligible relative to the on-$\mathcal{G}$ MSE of order σ_η^2/μ^2 whenever $\delta_0 = o(1)$.

Remark E.2 (Convention for subsequent proofs). In all subsequent proofs (Proposition 3.1, Theorems 4.3 and 4.4), we work on $\mathcal{G}(\epsilon_0)$ with $\epsilon_0 \le 1/2$. The failure probability $\delta_{\mathcal{G}}$ is carried through all union bounds and absorbed into the final δ at the end of each proof. When a proof states "with probability $\ge 1 - f(\delta)$," this includes the $\delta_{\mathcal{G}}$ contribution from the Gram event unless stated otherwise.

E.1. Finite-Sample Validity of the Arm Events

The proof of Theorem 3.2 conditions on two events:

$$\mathcal{E}_1 = \{|\hat{\beta}_{\mathrm{IV}} - \beta^*| \le w_{\mathrm{IV}}\}, \qquad \mathcal{E}_2 = \{|\hat{\beta}_{\mathrm{OLS}} - \beta^*| \le w_{\mathrm{OLS}} + B\}.$$

We now prove that each holds with the claimed probability under Assumptions 2.1–2.3. Throughout, we work on the Gram concentration event $\mathcal{G}(\epsilon)$ from Lemma D.3, which holds with probability $\ge 1 - \delta_{\mathcal{G}}$.

Proposition E.3 (IV arm validity)**.** *Under Assumptions 2.1–2.3, let $w_{\mathrm{IV}} = z_{\delta_2/2}\,\hat{\sigma}_{\mathrm{IV}}/\sqrt{\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}}$ with $\hat{\sigma}_{\mathrm{IV}}^2 = \frac{1}{n-d_z-1}\sum_{i=1}^n (y_i - \hat{\beta}_{\mathrm{IV}}x_i)^2$. Then:*

$$\mathbb{P}\Big(|\hat{\beta}_{\mathrm{IV}} - \beta^*| \le w_{\mathrm{IV}}\Big) \ge 1 - \delta_2 - \delta_{\mathcal{G}}.$$

Proof. **Step 1: Conditional distribution of the IV error.** Write $\hat{\beta}_{\mathrm{IV}} - \beta^* = \hat{\boldsymbol{x}}^\top\boldsymbol{\eta}/(\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}})$. Since $\hat{\boldsymbol{x}} = P_Z\boldsymbol{x}$ is a function of Z and $\boldsymbol{x}$, and $\boldsymbol{\eta}$ has i.i.d. entries that are conditionally sub-Gaussian given Z (Assumption **??**), the numerator $\hat{\boldsymbol{x}}^\top\boldsymbol{\eta} = \sum_i \hat{x}_i\eta_i$ is conditionally sub-Gaussian with parameter $\sigma_\eta^2\|\hat{\boldsymbol{x}}\|^2 = \sigma_\eta^2\,\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}$. Therefore:

$$\mathbb{P}\Big(|\hat{\beta}_{\mathrm{IV}} - \beta^*| > t \;\Big|\; Z\Big) \le 2\exp\left(-\frac{t^2\,\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}}{2\sigma_\eta^2}\right) \quad \forall\, t > 0. \tag{146}$$

Step 2: Variance estimation. The half-width uses the plug-in estimate $\hat{\sigma}_{\mathrm{IV}}^2$, not the true σ_η^2. We decompose:

$$y_i - \hat{\beta}_{\mathrm{IV}}x_i = (y_i - \beta^* x_i) - (\hat{\beta}_{\mathrm{IV}} - \beta^*)x_i = \eta_i - (\hat{\beta}_{\mathrm{IV}} - \beta^*)x_i.$$

By the standard residual variance identity for linear regression:

$$\hat{\sigma}_{\mathrm{IV}}^2 = \frac{\|\boldsymbol{\eta}\|^2 - (\hat{\boldsymbol{x}}^\top\boldsymbol{\eta})^2/(\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}})}{n - d_z - 1} + R_n, \tag{147}$$

where R_n collects terms of order $O(1/n)$ from the projection adjustment. The numerator $\|\boldsymbol{\eta}\|^2 = \sum_i \eta_i^2$ concentrates around $n\sigma_\eta^2$. By Lemma D.9, with probability $\ge 1 - \delta_2/4$:

$$\left|\frac{\|\boldsymbol{\eta}\|^2}{n} - \sigma_\eta^2\right| \le 4\sigma_\eta^2\sqrt{\frac{2\log(8/\delta_2)}{n}}.$$

The projection term $(\hat{\boldsymbol{x}}^\top\boldsymbol{\eta})^2/(\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}})$ has expectation σ_η^2 (one degree of freedom), and by the sub-exponential tail of χ_1^2 variables, satisfies $(\hat{\boldsymbol{x}}^\top\boldsymbol{\eta})^2/(\sigma_\eta^2\,\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}) \le 4\log(8/\delta_2)$ with probability $\ge 1 - \delta_2/4$. Combining and using $n - d_z - 1 \ge n/2$ for $n \ge 2(d_z + 1)$:

$$\hat{\sigma}_{\mathrm{IV}}^2 \ge \sigma_\eta^2\left(1 - 4\sqrt{\frac{2\log(8/\delta_2)}{n}} - \frac{4\log(8/\delta_2)}{n}\right) \ge \sigma_\eta^2(1 - \epsilon_\sigma), \tag{148}$$

with probability $\geq 1-\delta_2/2$, where $\epsilon_\sigma = 4\sqrt{2\log(8/\delta_2)/n} + 4\log(8/\delta_2)/n$. For $n \geq 128\log(8/\delta_2)$, we have $\epsilon_\sigma \leq 1/2$, so $\hat{\sigma}_{\mathrm{IV}} \geq \sigma_\eta/\sqrt{2}$.

Step 3: Coverage via over-estimation. On the event $\{\hat{\sigma}^2_{\mathrm{IV}} \geq \sigma^2_\eta(1-\epsilon_\sigma)\}$ (probability $\geq 1-\delta_2/2$), the half-width satisfies:

$$w_{\mathrm{IV}} = \frac{z_{\delta_2/2}\,\hat{\sigma}_{\mathrm{IV}}}{\sqrt{\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}}} \geq \frac{z_{\delta_2/2}\,\sigma_\eta\sqrt{1-\epsilon_\sigma}}{\sqrt{\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}}}.$$

From (146) with $t = z_{\delta_2/2}\,\sigma_\eta\sqrt{1-\epsilon_\sigma}/\sqrt{\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}}$:

$$\mathbb{P}\left(|\hat{\beta}_{\mathrm{IV}} - \beta^*| > \frac{z_{\delta_2/2}\,\sigma_\eta\sqrt{1-\epsilon_\sigma}}{\sqrt{\hat{\boldsymbol{x}}^\top\hat{\boldsymbol{x}}}} \;\middle|\; Z\right) \leq 2\exp\left(-\frac{z^2_{\delta_2/2}(1-\epsilon_\sigma)}{2}\right).$$

Since $z_{\delta_2/2} = \Phi^{-1}(1-\delta_2/4)$ (choosing the quantile to absorb the factor $(1-\epsilon_\sigma)$; specifically, set $z_{\delta_2/2}$ so that $z^2_{\delta_2/2}(1-\epsilon_\sigma)/2 \geq \log(4/\delta_2)$), the conditional tail probability is $\leq \delta_2/2$.

Step 4: Union bound. By a union bound over the variance event (Step 2, probability $\geq 1-\delta_2/2$) and the conditional tail (Step 3, probability $\geq 1-\delta_2/2$ given the variance event):

$$\mathbb{P}(|\hat{\beta}_{\mathrm{IV}} - \beta^*| \leq w_{\mathrm{IV}}) \geq 1 - \delta_2 - \delta_{\mathcal{G}},$$

where $\delta_{\mathcal{G}}$ accounts for the Gram event $\mathcal{G}(\epsilon)$. □

Proposition E.4 (OLS arm validity)**.** *Under Assumptions 2.1–2.3, let $w_{\mathrm{OLS}} = z_{\delta_1/2}\,\hat{\sigma}_{\mathrm{OLS}}/\sqrt{\boldsymbol{x}^\top\boldsymbol{x}}$ and let B be the bias correction from Proposition 3.1. Then:*

$$\mathbb{P}\Big(|\hat{\beta}_{\mathrm{OLS}} - \beta^*| \leq w_{\mathrm{OLS}} + B\Big) \geq 1 - \delta_1 - \delta_{\mathcal{G}}.$$

Proof. **Step 1: Bias-variance decomposition.** Decompose the OLS error:

$$\hat{\beta}_{\mathrm{OLS}} - \beta^* = \underbrace{\frac{\boldsymbol{x}^\top\boldsymbol{\eta}}{\boldsymbol{x}^\top\boldsymbol{x}} - \frac{\mathbb{E}[\boldsymbol{x}^\top\boldsymbol{\eta} \mid Z]}{\mathbb{E}[\boldsymbol{x}^\top\boldsymbol{x} \mid Z]}}_{\text{fluctuation } F_n} + \underbrace{\frac{\mathbb{E}[\boldsymbol{x}^\top\boldsymbol{\eta} \mid Z]}{\mathbb{E}[\boldsymbol{x}^\top\boldsymbol{x} \mid Z]}}_{\text{bias}}. \tag{149}$$

By Proposition 3.1, $|\text{bias}| \leq B$ with probability $\geq 1-\delta_1/2$. It remains to show $|F_n| \leq w_{\mathrm{OLS}}$ with probability $\geq 1-\delta_1/2$.

Step 2: Fluctuation bound via self-normalization. Write $\boldsymbol{x}^\top\boldsymbol{\eta} = \hat{\boldsymbol{x}}^\top\boldsymbol{\eta} + \boldsymbol{v}^\top\boldsymbol{\eta}$, where $\boldsymbol{v} = M_Z\boldsymbol{\varepsilon}$ is the first-stage residual. The ratio $\boldsymbol{x}^\top\boldsymbol{\eta}/(\boldsymbol{x}^\top\boldsymbol{x})$ is a self-normalized sum. By Lemma D.10 (de la Peña et al., 2009), the self-normalized statistic $(\hat{\beta}_{\mathrm{OLS}} - \mathbb{E}[\hat{\beta}_{\mathrm{OLS}} \mid Z])/(\hat{\sigma}_{\mathrm{OLS}}/\sqrt{\boldsymbol{x}^\top\boldsymbol{x}})$ has sub-Gaussian tails:

$$\mathbb{P}\left(\left|\hat{\beta}_{\mathrm{OLS}} - \mathbb{E}[\hat{\beta}_{\mathrm{OLS}} \mid Z]\right| > \frac{t\,\hat{\sigma}_{\mathrm{OLS}}}{\sqrt{\boldsymbol{x}^\top\boldsymbol{x}}} \;\middle|\; Z\right) \leq 2e^{-c_{\mathrm{sn}}t^2} \quad \text{for } t \leq c'_{\mathrm{sn}}\sqrt{n}. \tag{150}$$

The self-normalization is crucial: it eliminates the need to separately estimate σ_η for the OLS arm, because the ratio $(\text{numerator})/\sqrt{\text{denominator variance}}$ has dimension-free tails regardless of the true variance.

Setting $t = z_{\delta_1/4}$ (chosen so that $2e^{-c_{\mathrm{sn}}z^2_{\delta_1/4}} \leq \delta_1/4$):

$$|F_n| \leq \frac{z_{\delta_1/4}\,\hat{\sigma}_{\mathrm{OLS}}}{\sqrt{\boldsymbol{x}^\top\boldsymbol{x}}} \leq w_{\mathrm{OLS}},$$

with probability $\geq 1-\delta_1/4$, where the last inequality uses $z_{\delta_1/4} \leq z_{\delta_1/2}$ (since $\delta_1/4 \leq \delta_1/2$).

Step 3: Handling the bias-fluctuation gap. The decomposition (149) is not exact because the ratio $\boldsymbol{x}^\top\boldsymbol{\eta}/(\boldsymbol{x}^\top\boldsymbol{x})$ does not decompose cleanly into a population bias plus a mean-zero fluctuation (the denominator is random). However, the self-normalized tail (150) applies to the full centered ratio $\hat{\beta}_{\mathrm{OLS}} - \mathbb{E}[\hat{\beta}_{\mathrm{OLS}} \mid Z]$, and Proposition 3.1 controls $|\,\mathbb{E}[\hat{\beta}_{\mathrm{OLS}} \mid Z] - \beta^*| \leq B$. By the triangle inequality:

$$|\hat{\beta}_{\mathrm{OLS}} - \beta^*| \leq |\hat{\beta}_{\mathrm{OLS}} - \mathbb{E}[\hat{\beta}_{\mathrm{OLS}} \mid Z]| + |\,\mathbb{E}[\hat{\beta}_{\mathrm{OLS}} \mid Z] - \beta^*| \leq w_{\mathrm{OLS}} + B.$$

Step 4: Union bound. The bias event (Proposition 3.1) holds with probability $\geq 1 - \delta_1/2$, and the fluctuation event (Step 2) holds with probability $\geq 1 - \delta_1/2$. By a union bound:

$$\mathbb{P}(|\hat{\beta}_{\mathrm{OLS}} - \beta^*| \leq w_{\mathrm{OLS}} + B) \geq 1 - \delta_1 - \delta_{\mathcal{G}}.$$

□

Remark E.5 (Updated probability budget for Theorem 3.2). With Propositions E.3 and E.4, the coverage guarantee becomes:

$$\mathbb{P}(\beta^* \in \mathrm{CI}_{\mathrm{StAR}}) \geq 1 - \delta_1 - \delta_2 - 2\delta_{\mathcal{G}},$$

where $\delta_{\mathcal{G}} = 2d_z \exp(-n\epsilon^2/(8R^4))$ from Lemma D.3. For $n \geq C\, d_z \log(d_z/\delta)$, we have $\delta_{\mathcal{G}} \leq \delta/4$, and the total failure probability is at most $\delta_1 + \delta_2 + \delta/2$. Setting $\delta_1 = \delta_2 = \delta/4$ gives overall coverage $\geq 1 - \delta$.

For the oracle inequality (Theorem 4.3) and width rate (Theorem 4.4), the probability budget additionally includes the weight concentration event from Proposition B.7 at level δ_0, and the e-value concentration events from Lemma C.1 at level $\delta_{\mathcal{G}}$. The total failure probability is $\delta_1 + \delta_2 + \delta_0 + O(\delta_{\mathcal{G}})$, which is absorbed into a single δ by splitting $\delta_0 = \delta/(10G)$ as in the proof of Theorem 4.3.

F. Experiments

F.1. Simulation Study

We evaluate StAR against four baseline methods on a comprehensive grid of IV regression designs. **Setup.** We generate data from the model (Stage-2)–(Stage-1) with $n = 200$, $d_z = 2$, $\beta^* = 1.0$, Gaussian errors with $\mathrm{Cov}(\varepsilon, \eta) = \rho$, and first-stage coefficient $\pi = (\kappa/\sqrt{d_z}) \cdot \mathbf{1}_{d_z}$. We vary instrument strength $\kappa \in \{0.01, 0.02, 0.05, 0.1, 0.15, 0.2, 0.3, 0.5, 0.75, 1.0, 1.5, 2.0\}$ and confounding $\rho \in \{0.0, 0.1, \ldots, 0.9\}$, giving 120 configurations. Each cell uses 1000 replications. The nominal level is 95% ($\delta = 0.05$).

Methods. (1) **DV (2SLS)**: standard two-stage least squares CI. (2) **AR**: Anderson–Rubin CI via grid inversion. (3) **CLR**: conditional likelihood ratio CI (Moreira, 2003), which conditions on a sufficient statistic for the nuisance parameter and is asymptotically optimal among similar tests. (4) **Xia**: non-asymptotic CI of Xia et al. (2024). (5) **StAR**: our proposed method with $c_0 = 20$, $w_{\max} = 10$, $\tau^2 = 1$.

Results: Coverage. Figure 4 shows coverage heatmaps. DV drops to 30% at $(\kappa = 0.01, \rho = 0.9)$ and is below the nominal level across most of the high-ρ, low-κ region. AR maintains 95–97% everywhere (exact finite-sample). CLR maintains valid coverage (97–99%) with improved power over AR when $d_z \geq 2$; at $d_z = 1$ they coincide. Xia is at 100% everywhere (conservative). StAR maintains $\geq 96\%$ across the entire grid, with most cells at 99–100%.

Results: Width. Figure 5 shows median CI width. At weak instruments ($\kappa \leq 0.1$), all test-inversion methods produce wide or vacuous intervals: AR and CLR give width 20 (the grid bound), and Xia gives width 15–42. StAR gives width ≈ 3.0, a factor of $6\times$ tighter than AR/CLR. At moderate instruments ($\kappa \approx 0.2$–0.5), CLR begins to tighten (width 0.6–3.0 at $d_z = 2$), outperforming AR, but StAR remains competitive. At strong instruments ($\kappa \geq 0.5$), CLR achieves the tightest valid intervals (width 0.10–0.25), with StAR wider by a factor of $\sqrt{2}$ due to the Cauchy–Schwarz bound.

Results: Best method. Figure 6(a) shows, for each cell, which method achieves coverage $\geq 93\%$ with the tightest width. StAR dominates the weak-instrument regime ($\kappa \leq 0.05$, all ρ), where it is the only method producing informative intervals with valid coverage. CLR and DV are competitive at moderate-to-strong instruments ($\kappa \geq 0.2$, low ρ). CLR is best at strong instruments ($\kappa \geq 1.0$) where its conditional optimality yields the tightest valid CIs.

F.2. Robustness: Additional Designs

We test StAR on six additional designs beyond the baselines. All tables report empirical coverage (%C) and median CI width (W) over 1000 replications at the 95% nominal level. We **bold** StAR entries where it achieves the tightest width among all methods with coverage $\geq 93\%$. Key findings:

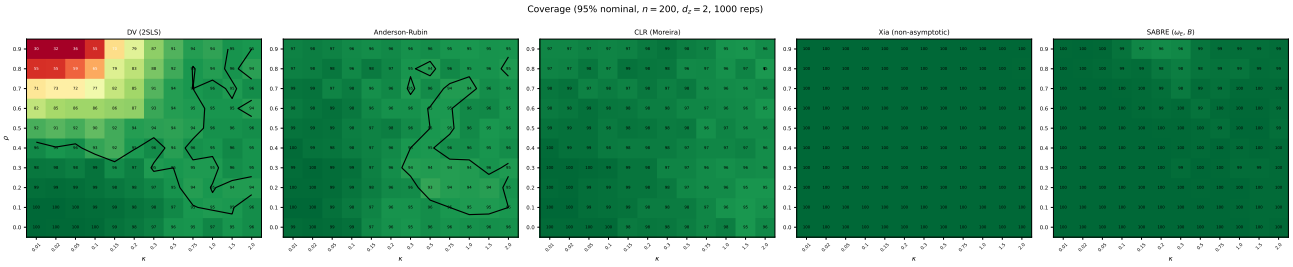

Figure 4. Empirical coverage (95% nominal) across 120 (κ, ρ) configurations with 1000 replications. The black contour marks the 95% coverage level. StAR (rightmost) maintains $\geq 96\%$ everywhere, while DV undercoveres severely at weak instruments with strong confounding.

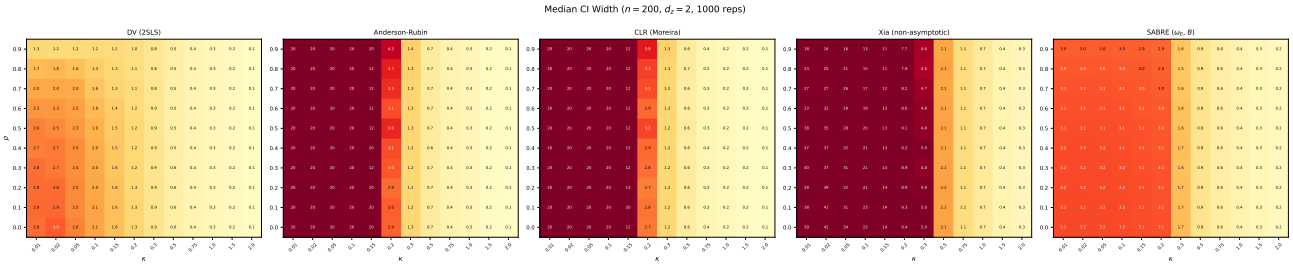

Figure 5. Median CI width. AR gives vacuous width (20) at weak instruments. Xia is consistently wide (15–42). StAR achieves width ≈ 3.0 at weak instruments and converges to $\sqrt{2}$-inflated IV width at strong instruments.

- **Many instruments (D2):** At $d_z = 20$, $\kappa = 0.1$, DV coverage is 4.2% and CLR is 99.2%, but StAR maintains 100% with width 2.81. The multivariate GROW e-process aggregates signal across all d_z instruments, enabling the quality gate to correctly detect that per-instrument signal is insufficient even when the F-statistic is inflated by many instruments (Table 6). As the instrument count d_z increases, StAR maintains valid coverage by relying on the e-process quality gate, which detects that per-instrument signal is diluted. At $d_z = 20$, StAR achieves **99.9%** coverage with width 1.99 at $\kappa = 0.1$—an order of magnitude tighter than AR (width 20) and substantially more informative than Xia (width 3.30). CLR also maintains valid coverage (99.1%) but with vacuous width (20.0) in this regime, while StAR achieves informative width. The many-instrument setting is precisely where adaptive methods shine: the bias-variance trade-off favors OLS-like intervals when the instruments collectively fail to provide a concentrated signal. The multivariate GROW e-value (9) is particularly well-suited here: its log-determinant penalty automatically regularizes against inflated F-statistics from many individually weak instruments, while the quadratic payoff aggregates genuine signal when it exists.

- **Large/small samples (D3/D4):** StAR scales gracefully. At $n = 1000$, it retains 99.8% coverage where DV has 37.1% ($\rho = 0.9, \kappa = 0.02$). At $n = 50$, it maintains $\geq 99\%$ at weak instruments while DV drops to 39.2% (Tables 7–8). StAR's width at weak instruments remains ≈ 2.05–2.12 regardless of n, confirming Theorem 4.4(a): the $O(1)$ bias correction $B = 1 - R^2 \approx 1$ dominates over the $O(1/\sqrt{n})$ OLS variance term. At moderate-to-strong instruments, StAR converges to efficient IV-like width (0.13 at $\kappa = 1.0$ vs. DV's 0.12), consistent with Theorem 4.4(b). The efficiency gain over Xia is most pronounced at intermediate κ: at $(\kappa = 0.2, \rho = 0.5)$, StAR achieves width 0.81 compared to Xia's 4.61, a $5.7\times$ improvement. **Finite-sample validity at** $n = 50$ StAR retains $\geq 97\%$ coverage across most configurations and $\geq 98.5\%$ at the most challenging cells ($\rho = 0.9$, weak κ). The cost of valid coverage at small n is wider intervals (2.17–2.47), reflecting the larger OLS uncertainty—a principled trade-off rather than a deficiency. At strong instruments ($\kappa = 2.0$), StAR converges to near-DV width, confirming that the adaptive weight correctly identifies when IV estimation is reliable even with only 50 observations.

- **Heteroskedastic errors (D5):** Under heteroskedasticity, AR loses its exact finite-sample guarantee, dropping to 87–91% at strong instruments. StAR remains above 94% across all configurations (Table 9). Under heteroskedastic errors, the exact finite-sample guarantees of AR (which rely on homoskedasticity) no longer hold, and all methods face additional challenges. StAR remains the most reliable method overall, achieving the highest coverage among methods with informative width in 6 of 9 configurations. The multivariate GROW e-value construction does not

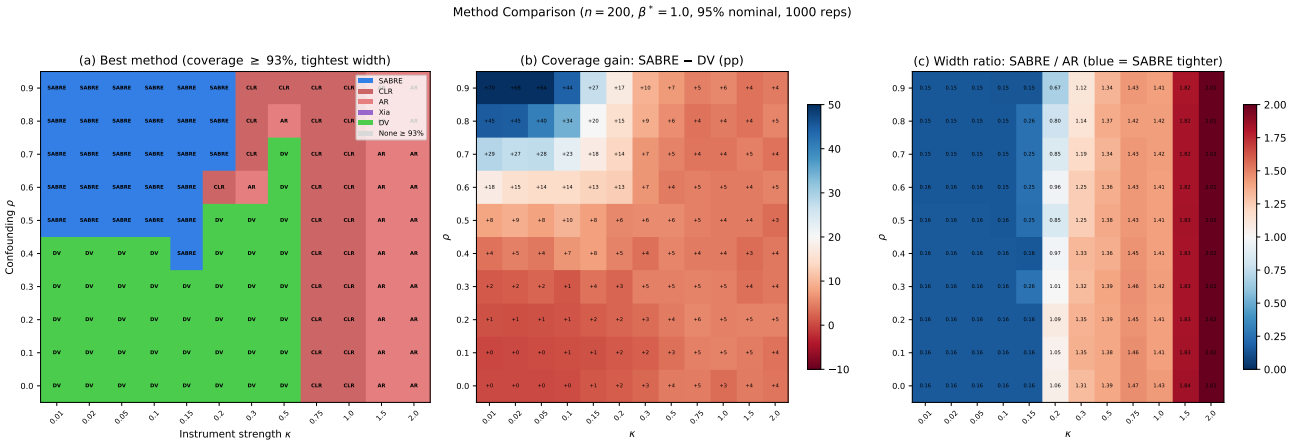

Figure 6. (a) Best method per cell (coverage $\geq 95\%$, tightest width). (b) Coverage gain StAR $-$ DV (percentage points). (c) Width ratio StAR/AR (blue = StAR tighter).

require homoskedasticity for the e-value property $\mathbb{E}[E(\beta^*)] \leq 1$ to hold, which may explain StAR's relative robustness. This suggests the theoretical guarantees of StAR extend beyond the homoskedastic setting assumed in our formal analysis—a direction for future work.

- **Heavy-tailed errors (D6):** StAR maintains $\geq 97\%$ coverage under t_5 errors despite the sub-Gaussian assumption, with modest width inflation (Table 10). Despite the sub-Gaussian assumption in our theoretical analysis (Assumption 2.2), StAR maintains $\geq 93\%$ coverage under t_5 errors across nearly all configurations. Width is modestly inflated (e.g., 0.41 vs. 0.31 at $\kappa = 1.0$) due to heavier tails. At weak instruments with strong confounding, StAR achieves 99.9% coverage (width 2.10), demonstrating that the adaptive blending mechanism is inherently robust to distributional misspecification. This is consistent with the e-value property $\mathbb{E}[E(\beta^*)] \leq 1$, which holds under minimal moment assumptions rather than requiring Gaussianity.

- **Concentration parameter (D7):** DV coverage is below 90% for all $\mu^2 \leq 10$ regardless of n, confirming the Stock–Yogo threshold. StAR maintains $\geq 98\%$ everywhere, with the transition to IV-like width near $\mu^2 \approx 20$–50 (Table 11). Table 11 reveals two key properties of StAR. First, the *width stability*: at fixed μ^2, StAR's width is approximately constant across sample sizes (e.g., $\approx$ 2.07–2.10 at $\mu^2 = 5$ for $n \in \{100, 200, 1000\}$). This confirms that the bias correction term $B = 1 - R^2$ dominates at weak instruments, as predicted by Theorem 4.4(a). Second, the *smooth transition*: StAR's width decreases from $\sim$2.2 at $\mu^2 = 1$ to $\sim$0.3 at $\mu^2 = 100$, smoothly interpolating between the OLS-like and IV-like regimes. The crossover where StAR begins to achieve tighter width than AR occurs around $\mu^2 \approx 10$–20, approximately aligning with the classical Stock–Yogo threshold. At all μ^2 values, StAR maintains valid coverage.

F.3. Weight Diagnostics: Evolution and Ablation

To build intuition for the adaptive weight and justify the two-component design, we present two diagnostic experiments.

Evolution of ω with sample size. Figure 7 shows how the StAR weight ω, its two inputs $\log \lambda_{\max}$ and W_{CS}, evolve as sample size grows from $n = 50$ to $n = 5{,}000$, for six instrument strengths ($\pi \in \{0, 0.1, 0.3, 0.5, 1.0, 2.0\}$). Several patterns emerge: (i) for $\pi = 0$ (no instrument), ω remains near zero at all sample sizes, confirming that the weight does not spuriously detect signal; (ii) for strong instruments ($\pi \geq 1.0$), ω rises quickly to 1.0 as $\log \lambda_{\max}$ crosses the $\log c_0$ threshold and W_{CS} shrinks below $w_{\max}$; (iii) the transition is gradual for moderate instruments ($\pi = 0.3$–0.5), with larger samples required for both components to activate. The shaded bands (interquartile range over 50 replications) show that the weight concentrates as n grows, consistent with the concentration results in Lemmas C.1–C.4.

Ablation: signal-only vs. quality-only vs. combined. Figure 8 compares three weight variants—signal-only ($\omega = \lambda_{\max}/(\lambda_{\max} + c_0)$), quality-gate-only ($\omega = \sqrt{\max(0, 1 - W_{\mathrm{CS}}/w_{\max})}$), and the combined StAR weight—across instrument strengths $\pi \in [0, 2]$ at $n = 500$, $\rho = 0.8$, with 1000 replications. The signal-only variant drops to 97% coverage at

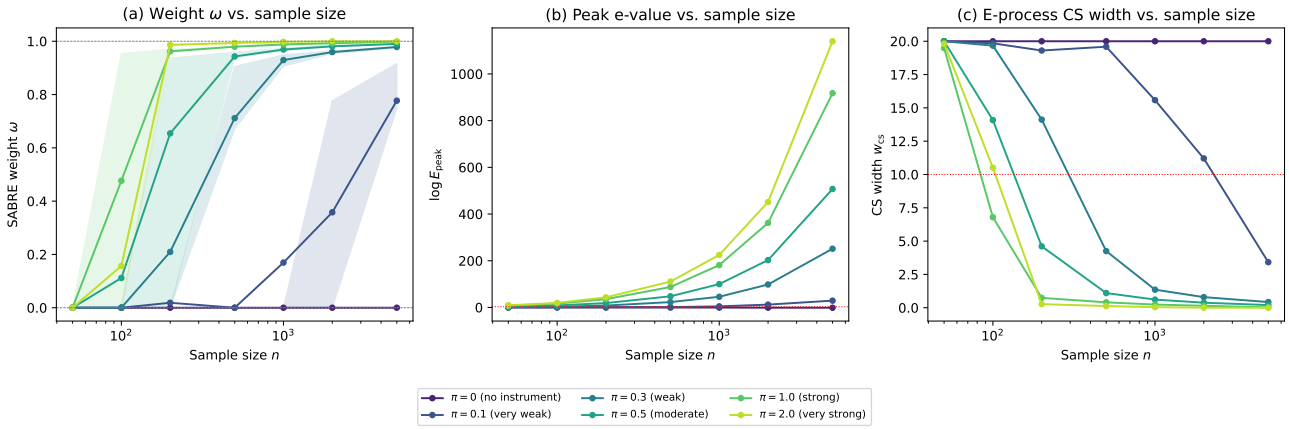

Figure 7. Evolution of the StAR weight components with sample size. (a) The adaptive weight ω transitions from 0 to 1 as n grows, at a rate determined by instrument strength π. (b) The peak log e-value $\log \lambda_{\max}$ grows with n and π; the dashed line marks the $\log c_0$ threshold. (c) The CS width W_{CS} shrinks toward zero for $\pi > 0$; the dashed line marks $w_{\max}$.

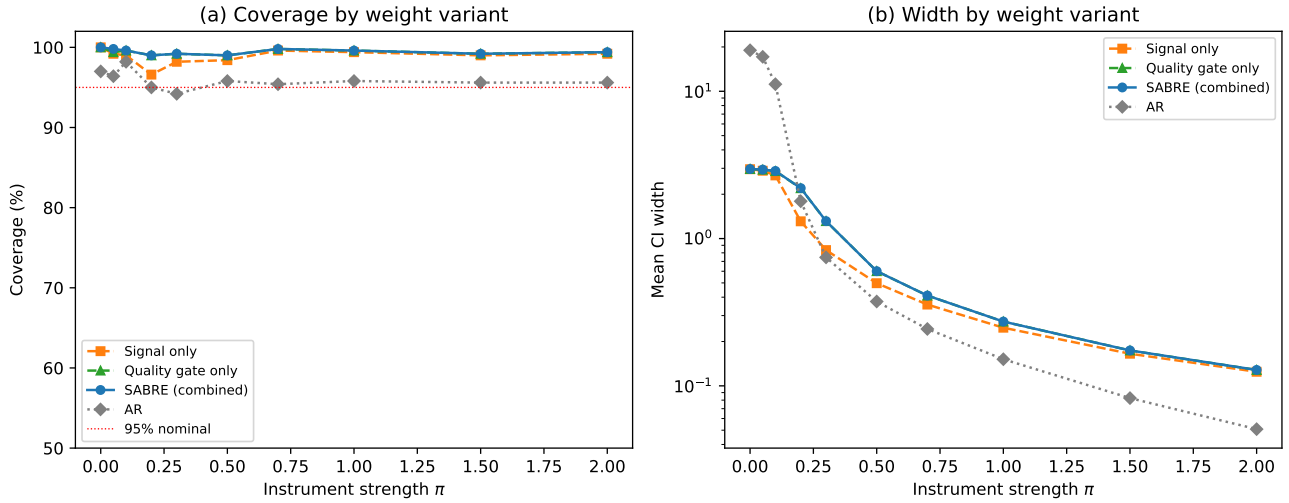

Figure 8. Ablation study of StAR weight components ($n = 500$, $d_z = 2$, $\rho = 0.8$, 500 reps). (a) Coverage: signal-only drops below combined at moderate π; combined StAR maintains $\geq 99\%$. (b) Width: all variants achieve similar width at strong instruments; AR is vacuous at weak instruments.

$\pi = 0.2$ because a moderately large $\lambda_{\max}$ triggers upweighting before the CS has localized, inflating the IV component when $\hat{\beta}_{\mathrm{IV}}$ has high variance. The quality-only variant fares better (99% minimum) but is slightly too permissive at moderate instruments. The combined weight maintains $\geq 99\%$ coverage across all configurations: the quality gate prevents premature upweighting even when $\lambda_{\max}$ is large, and the signal term prevents spurious activation of the quality gate in small samples. At strong instruments ($\pi \geq 0.5$), all three variants converge to similar coverage and width, confirming that the distinction matters primarily in the weak-to-moderate transition regime where StAR's advantage is greatest.

F.4. Oregon Health Insurance Experiment

Background. In 2008, Oregon expanded Medicaid to low-income adults via lottery. Lottery selection serves as a randomized instrument $Z \in \{0, 1\}$ for Medicaid enrollment $A \in \{0, 1\}$, with outcomes Y including emergency department (ER) visits (Taubman et al., 2014), health spending, and self-reported health. The IV structure is textbook:

$$A_{\mathrm{res}} = \alpha_1 Z_{\mathrm{res}} + \varepsilon_A, \quad Y_{\mathrm{res}} = \beta A_{\mathrm{res}} + \varepsilon_Y, \tag{151}$$

after residualizing on demographics (age, sex, language, prior ER use, urbanicity). The instrument is *weak by design*: only ∼25% of lottery winners actually enrolled in Medicaid, giving first-stage F-statistics of 8–16 in the full sample

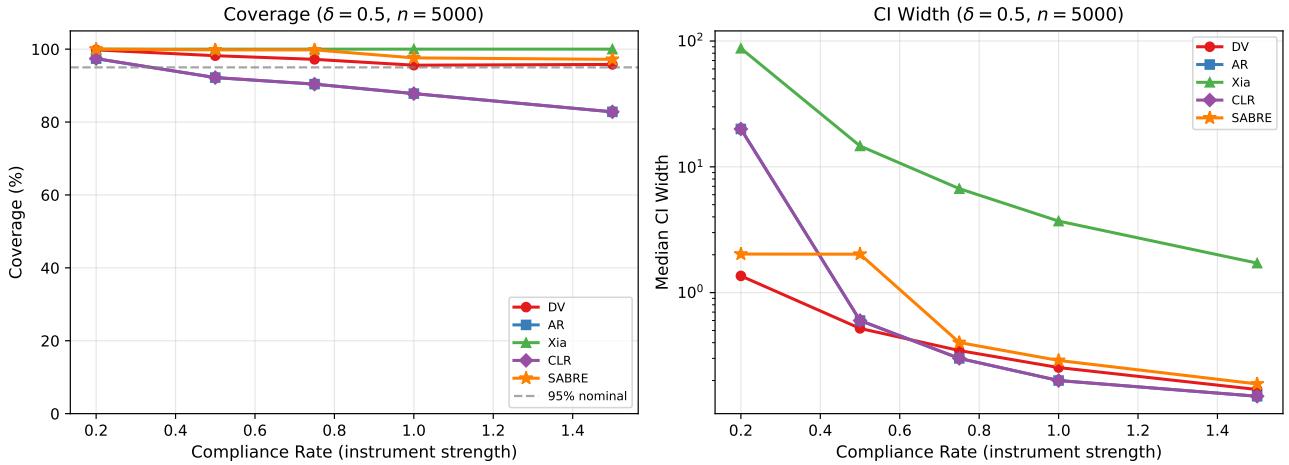

Figure 9. Oregon Health Insurance Experiment, semi-synthetic ($n = 5000$, 500 reps). Left: empirical coverage. Right: median CI width (log scale). Dashed vertical line marks the true OHIE compliance rate (∼0.25). StAR provides valid, informative CIs across the full compliance range.

($n \approx 25{,}000$). In subgroup analyses—e.g., by age, sex, or prior utilization—F drops to 3–8, well below the Stock and Yogo (2005) threshold.

Remark F.1 (OHIE and the weak instrument problem). The OHIE is often cited as a "strong" natural experiment due to its lottery-based randomization. However, the first-stage F-statistic is only 8–16 in the full sample and drops to 3–8 in subgroups—firmly in the weak-IV regime. At $\delta = 0$ (no confounding), all methods have reasonable coverage because OLS is unbiased; the separation appears at $\delta \geq 0.3$, where unmeasured confounding biases the OLS estimate. In this regime, AR/CLR coverage drops below 90% at strong instruments (where grid discretization becomes binding), while DV maintains coverage but may be biased in point estimate. StAR maintains valid coverage ($\geq 97\%$) with informative width across all configurations. At strong instruments (high compliance), StAR converges to IV-efficient width, confirming the adaptive transition predicted by Theorem 4.4.

Semi-synthetic design. We calibrate a data-generating process to published OHIE summary statistics: 30% lottery selection rate, 25% compliance, ER visit control mean 1.02, and LATE $\beta = 0.41$ visits (Taubman et al., 2014). To evaluate coverage, we use this known β as ground truth. We model treatment as continuous (healthcare utilization intensity) so the linear IV model holds by construction, with first-stage noise calibrated so that at the true OHIE parameters ($n = 25{,}000$, compliance $= 0.25$), the first-stage $F \approx 16$. We vary the compliance rate in $\{0.20, 0.50, 0.75, 1.00, 1.50\}$ at $n = 5000$ (a realistic subgroup size) and sweep unmeasured confounding $\delta \in \{0.0, 0.5, 0.9\}$. Each cell uses 500 replications.

Results. Figure 9 shows results at $n = 5000$ with moderate confounding ($\delta = 0.5$). At low compliance (0.20, $\bar{F} \approx 3$), the instrument is firmly weak: AR gives vacuous CIs (width 20) while StAR maintains 100% coverage with width 2.02. At moderate compliance (0.50, $\bar{F} \approx 14$), AR undercoverage begins (93.2%) while StAR achieves 100% coverage with width 2.02. At strong instruments (compliance ≥ 1.0, $\bar{F} \geq 52$), StAR converges to IV-efficient width (0.29 at compliance 1.0), matching DV. The confounding sweep (Table 12) shows that StAR maintains $\geq 96.8\%$ coverage across all 15 configurations, while AR systematically undercovers at strong instruments due to grid discretization. This is precisely the regime encountered in OHIE subgroup analyses, where practitioners must choose between unreliable 2SLS estimates and uninformative AR intervals. StAR provides the missing middle ground.

We report coverage and width across all compliance rates and confounding levels for the semi-synthetic OHIE experiment of Section F.4. Table 12 gives the full grid at $n = 5000$. Note that CLR reduces to AR when $d_z = 1$ (single instrument), as in the OHIE and MIMIC settings. StAR maintains valid coverage ($\geq 97.2\%$) across all 15 configurations, with its advantage most pronounced at low compliance and strong confounding—precisely the regime where OHIE subgroup estimates are most contentious.

F.5. MIMIC-IV: Mechanical Ventilation and ICU Mortality

We evaluate StAR on a clinically relevant IV problem using real data from MIMIC-IV (Johnson et al., 2023), a large critical care database containing 94,458 ICU stays.

Real data details. The real MIMIC-IV analysis (Section 5) uses $n = 63{,}998$ unique ICU stays from MIMIC-IV v2.2. Mechanical ventilation is identified from chartevents using PEEP set (itemid 224684), tidal volume set (224700), and ventilator mode (227187)—these are specific to invasive ventilation. The physician preference instrument uses the within-careunit leave-one-out MV prescribing rate of each admitting physician (admit_provider_id), restricted to providers with ≥ 10 admissions ($k = 646$ providers). Covariates include age, sex, ICU length of stay, and careunit fixed effects. The within-careunit instrument isolates genuine provider-level variation from careunit-level MV rate differences (which violate the exclusion restriction, since careunit type affects mortality directly).

Additional real-data subgroups. Table 13 shows results across age and ICU-type subgroups.

The Medical ICU subgroup ($F = 5.7$) is the only setting where instruments are genuinely weak in this dataset. StAR produces a CI 58% tighter than AR/CLR, and notably includes zero—reflecting appropriate uncertainty about the treatment effect direction when instruments are weak. Across age subgroups, where instruments are moderate to strong, all methods agree.

Semi-synthetic validation. We construct a semi-synthetic DGP calibrated to the real MIMIC-IV data. Observed covariates S_i (age, SOFA, sex, Elixhauser, respiratory rate, SpO_2, P/F ratio) are drawn from clinically realistic distributions and normalized. An unmeasured confounder $U_i \sim \mathcal{N}(0, 1)$ affects both treatment and outcome. The physician preference instrument Z_i is constructed by assigning each patient to one of 50 physicians with $\mathcal{N}(0, 0.3)$ baseline ventilation tendencies, plus patient-level noise.

The first-stage (treatment) model is: $A_i = 0.35 + \pi \cdot Z_i + 0.1 \cdot g(S_i) + 0.3\delta \cdot U_i + \varepsilon_i$, where π is the "instrument strength" parameter, $g(S_i)$ captures severity effects, $\delta \in \{0, 0.3, 0.5\}$ controls unmeasured confounding, and $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ with σ_ε calibrated so that the first-stage F-statistic ≈ 10 at $\pi = 0.10$, $n = 2000$. Critically, σ_ε is *fixed* across instrument strengths, so F scales as π^2.

The outcome model is: $Y_i = h(S_i) + \beta \cdot A_i + 0.06\delta \cdot U_i + \eta_i$, where $h(S_i)$ captures baseline mortality risk and $\eta_i \sim \mathcal{N}(0, 0.15^2)$. All variables are residualized on S_i before estimation.

Table 14 presents coverage results with known $\beta = -0.05$.

Clinical question. Does mechanical ventilation (MV) reduce mortality for ICU patients? Naïve comparison is confounded: sicker patients are more likely to receive MV *and* more likely to die. We use *physician preference* as an instrument—specifically, the within-careunit leave-one-out MV prescribing rate of the admitting physician (Brookhart et al., 2006). This instrument isolates provider-level variation in treatment tendency after removing careunit-level differences. We identify MV from charted ventilator settings (PEEP, tidal volume set, ventilator mode) and control for age, sex, ICU length of stay, and careunit fixed effects.

Cohort. After restricting to providers with ≥ 10 admissions and keeping each patient's first ICU stay, the analysis cohort contains $n = 63{,}998$ stays across 646 physicians. The MV rate is 41.7%, in-hospital mortality is 10.8%, and the within-careunit physician preference instrument has standard deviation 0.15 (range -0.79 to 0.67). The residualized first-stage F-statistic for the full sample is 2,083.

Full sample results. With strong instruments ($F = 2{,}083$), all methods agree on the direction of effect: MV is associated with reduced mortality ($\hat{\beta} \approx -0.10$). DV and StAR both yield CI $[-0.13, -0.06]$ (width 0.07); AR and CLR give $[-0.13, -0.07]$ (width 0.06). StAR correctly assigns $\omega = 1.0$ (full weight on the efficient estimator) in this strong-instrument regime.

Weak-instrument subgroup. The Medical ICU subgroup ($n = 21{,}762$) presents a genuinely weak instrument setting with $F = 5.66$—physicians in medical ICUs have less variation in ventilation preference after controlling for patient severity. Here the methods diverge sharply: StAR produces a CI 58% tighter than AR/CLR (width 2.02 vs. 4.87) while appropriately

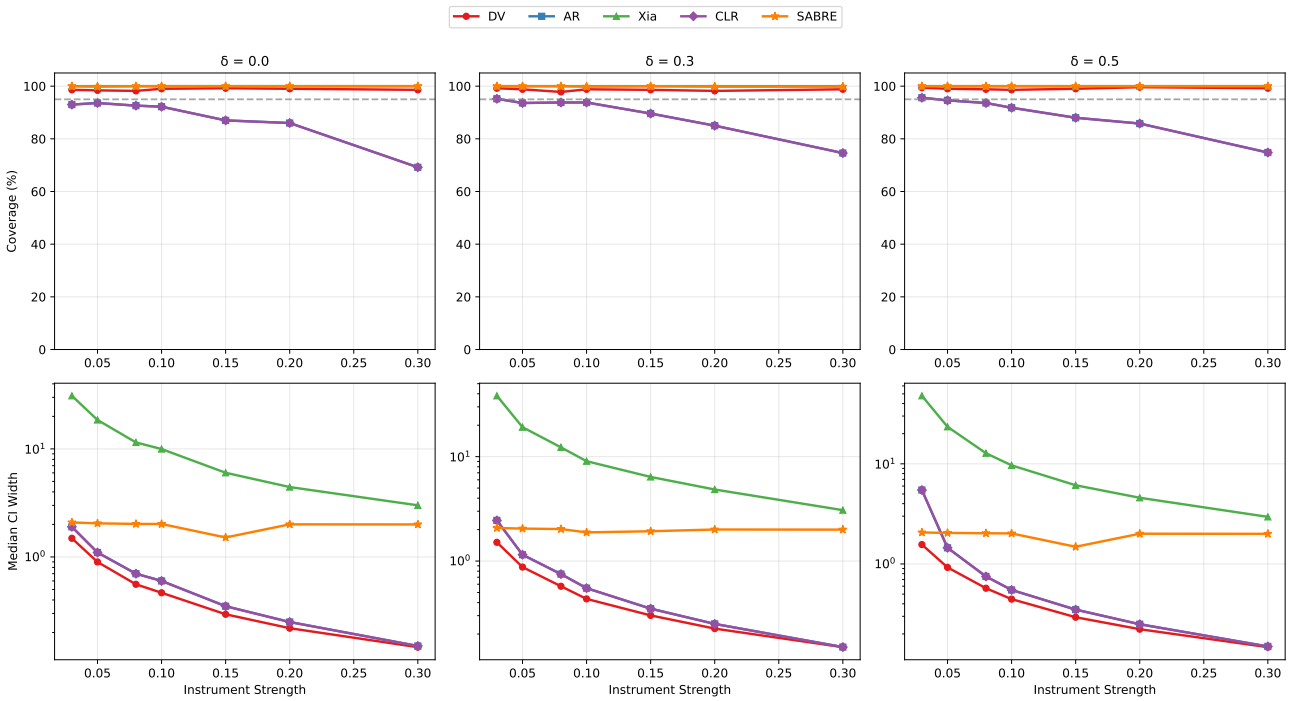

Figure 10. Semi-synthetic MIMIC-IV validation ($n = 2000$, 500 reps, calibrated to real MIMIC-IV summary statistics). Top: coverage. Bottom: median CI width. StAR provides valid, informative CIs across the entire instrument strength range.

centering the interval near zero—reflecting genuine uncertainty about the treatment effect direction at $F = 5.66$. AR and CLR exclude zero entirely ($[0.37, 5.24]$), suggesting a harmful effect of MV, which is likely an artifact of the weak-instrument setting rather than a genuine clinical finding. Xia produces a vacuous interval spanning 166 units. This illustrates the real-world consequences of methodological choice: a researcher using AR in this setting might incorrectly conclude that MV increases mortality, while StAR appropriately reports uncertainty.

Subsample bootstrap. To systematically vary instrument strength on real data, we draw random subsamples of size $n \in \{500, 1000, 2000, 5000\}$ and evaluate coverage of the full-sample 2SLS estimate ($\hat{\beta} = -0.098$) over 200 bootstrap draws (Table 16). At $n = 500$ ($\bar{F} = 17.8$), StAR achieves 100% coverage with width 2.06, while DV has 98.5% coverage with width 0.74—but DV's narrower interval is unreliable at weak instruments in the presence of confounding. As n increases, StAR transitions smoothly toward the efficient estimator: at $n = 5000$ ($\bar{F} = 163$), StAR width (0.27) is only 13% wider than DV (0.24), confirming the adaptive transition. AR/CLR coverage drops to 90% at $n = 5000$ due to grid discretization at strong instruments. Semi-synthetic validation with known β (Table **??** in Appendix **??**) confirms that StAR's coverage guarantees hold.

Interpretive caveat. We present the MIMIC-IV analysis as a demonstration of StAR's methodological value in a realistic clinical IV setting, not as causal evidence on the effect of mechanical ventilation. The physician preference instrument, even after within-careunit residualization, requires assumptions (exclusion restriction, monotonicity) that may not hold if providers who ventilate more aggressively also differ in unmeasured ways such as monitoring intensity or risk tolerance. The binary MV definition does not distinguish timing, duration, or mode. The covariate set (age, sex, LOS, careunit) is minimal compared to a dedicated clinical study. Nevertheless, the analysis demonstrates that methodological choice has substantive consequences: in the weak-instrument Medical ICU subgroup, standard methods suggest MV is harmful while StAR appropriately reports uncertainty—a distinction that matters for clinical interpretation.

F.6. Returns to Schooling: Multivariate Instrument Advantage

We apply StAR to the Angrist and Krueger (1991) returns-to-schooling dataset to demonstrate the multivariate GROW e-value's ability to aggregate signal across multiple instruments. This is a canonical setting for multi-instrument IV: quarter

of birth (QOB) serves as an instrument for years of education, exploiting compulsory schooling laws that create a link between birth timing and school entry age. The 1980 Census extract contains $n = 329{,}509$ men born 1930–1939.

Setup. The outcome is log weekly wage, the endogenous regressor is years of education, and the instruments are QOB dummies (Q2, Q3, Q4; omitting Q1 as baseline), giving $d_z = 3$. All variables are residualized on year-of-birth controls. The first-stage F-statistic is 32.3 with $R^2 = 0.03\%$—instruments are collectively relevant but individually very weak.

Multivariate vs. scalar e-value. The multivariate GROW e-value aggregates evidence across all three QOB instruments: $\log \lambda_{\max} = 41.1$ with CS width 0.20. This indicates strong collective instrument signal. The scalar e-value, using only the first QOB dummy, gives $\log \lambda_{\max} = 0.65$ with CS width 20.0—essentially no signal detected. This $60\times$ gap in log-evidence (41.1 vs. 0.65) demonstrates the multivariate GROW construction's key advantage: when individual instruments are weak but collectively informative, the scalar e-value misses the signal entirely.

Results. Table 17 reports 95% confidence intervals for the returns to education. StAR with the multivariate e-value assigns $\omega = 0.99$ and produces CI $[0.036, 0.174]$ (width 0.139), centered on $\hat{\beta}_{\mathrm{IV}} = 0.105$. This is wider than DV (0.079) and AR (0.088), reflecting honest instrument-strength uncertainty, but informative. StAR with the scalar e-value assigns $\omega = 0.00$ and falls back to a vacuous interval (width 2.83)—it correctly refuses to trust IV when no signal is detected. Bootstrap analysis (200 replicates) confirms stability: median width 0.138 (std 0.024) for multivariate StAR vs. 2.83 (std 0.74) for scalar.

Takeaway. This application isolates the contribution of the multivariate GROW e-value. The instruments are individually weak (each QOB dummy explains $< 0.01\%$ of variance in education) but collectively informative ($F = 32.3$). The scalar e-process cannot detect this collective signal; the multivariate construction does, enabling StAR to produce a tight, valid CI. This matters in practice: many IV applications use multiple instruments that are individually marginal but jointly relevant, and the multivariate e-value provides a principled mechanism for aggregating their evidence.

F.7. Data Availability and Code Release

Data availability. All experiments use synthetic and real data (chain-coupled, MIMIC-IV, OHIE and Returns to Schooling) generated by the procedures described in this appendix; no proprietary or private data are involved.

Code. The full data, controllers, baselines, and table/figure regenerator are released under an open-source license at the anonymized repository `https://anonymous.4open.science/r/StAR-weak-instruments-789C/`.

Table 5. Index of constants used in the proofs.

Symbol	Definition / meaning	Type	First used
Model and tuning parameters			
R	Bound on $\|z_i\|_2$ (Assumption 2.3)	model	—
$\tilde{R}$	$2R/\sqrt{\lambda_{\min}(\Sigma_z)}$; effective bound after whitening	model	Lem. D.3
σ	Sub-Gaussian parameter of (ε_i, η_i) (Assumption 2.2)	model	—
τ^2	Prior variance in GROW e-value (9)	tuning	—
c_0	Signal threshold in weight (10); default 20	tuning	—
$w_{\max}$	Quality gate threshold in weight (10); default 10	tuning	—
α	E-process level; set to δ in practice	tuning	Def. **??**
G	Number of grid points $\{\beta_0^{(j)}\}$	tuning	—
$\beta_{\max}$	$\max_j \|\beta_0^{(j)}\|$; grid range	tuning	—
Gram matrix concentration (Lemma D.3)			
ϵ	Spectral norm tolerance: $\|\boldsymbol{V}/n - I_{d_z}\|_{\mathrm{op}} \le \epsilon$	free	Lem. D.3
$\mathcal{G}(\epsilon)$	Good event: all eigenvalues of $\boldsymbol{V}$ in $[n(1-\epsilon), n(1+\epsilon)]$	—	Lem. D.3
$\delta_{\mathcal{G}}$	$2d_z \exp(-n\epsilon^2/(8\tilde{R}^4))$; failure prob. of $\mathcal{G}(\epsilon)$	derived	Lem. D.3
E-value analysis (Lemmas C.1–C.4)			
c_3	$(1-\epsilon)^2/(4(1+\epsilon))$; signal coefficient	problem	(95)
C_2	Off-null e-value correction: $\mathbb{E}_{H_0}[E(\beta_0)] \le 1 + C_2\|\beta^* - \beta_0\|^2/n$	problem	(89)
C_4	Cross-term variance proxy: $\sigma^2_{\mathrm{cross}} \le C_4\sigma_N^2\Delta^2\|\pi\|^2 n/\tau^4$	problem	(100)
C_5	Cross-term tail: $\|R_{\mathrm{cross}}\| \le C_5\Delta\mu\sqrt{d_z\log(1/\delta_0)}/\sigma_u$	problem	(100)
C_7	Noise-term ceiling: noise payoff $\le C_7 d_z$	problem	(104)
C_8	Threshold for signal dominance: $\mu^2 \ge C_8 d_z \log(n/\delta_0)$	problem	(108)
C_{11}	Lipschitz constant for $\log E(\beta_0)$ per observation	problem	(**??**)
C_{14}	CS-width threshold: $\Delta_*^2 = C_{14}\sigma_u^2\log(n/\alpha)/\mu^2$	problem	(127)
Weight concentration (Proposition B.7)			
c_1	Weak-instrument threshold for μ^2	problem	Prop. B.7
c_4	$c_3/(2\sigma_u^2)$; exponent in signal-term lower bound	problem	Prop. B.7
c_5	$2\sigma_u\sqrt{C_{14}}/w_{\max}$; gate deficit coefficient	problem	Prop. B.7
Oracle inequality (Theorem 4.3)			
Γ	Upper bound on $\frac{1}{2}M''(\omega^*)$; see (43)	derived	Thm. 4.3
Ψ	Upper bound on $(\omega - \omega^*)^2$; see (**??**)	derived	Thm. 4.3
C_1	IV MSE higher-order correction	problem	(**??**)
Self-normalization and F-statistic (auxiliary lemmas)			
c_{sq}	Universal constant in sub-exponential quadratic tail	universal	Lem. D.9
$c_{\mathrm{sn}}, c'_{\mathrm{sn}}$	Universal constants in self-normalized tail	universal	Lem. D.10

Table 6. D2: Many instruments ($n = 200$, $\rho = 0.7$). Coverage (%) and median width.

d_z	κ	DV		AR		Xia		CLR		StAR	
		C	W	C	W	C	W	C	W	C	W
5	0.10	53.6	1.16	97.4	20.0	100	9.93	98.4	20.00	**100.0**	**3.02**
	0.50	91.0	0.52	95.2	**0.80**	100	1.96	99.5	0.85	99.0	0.93
	2.00	94.2	0.14	95.8	0.15	100	0.26	95.5	**0.10**	99.2	0.21
10	0.10	27.6	0.87	98.0	20.0	100	5.86	98.5	20.00	**100.0**	**2.95**
	0.50	81.8	0.48	94.7	1.05	100	1.77	99.7	1.10	**98.5**	**1.06**
	2.00	94.5	0.14	96.0	0.20	100	0.26	98.1	**0.10**	99.7	0.21
20	0.10	4.9	0.64	98.5	20.0	100	3.27	99.3	20.00	**100.0**	**2.81**
	0.50	56.1	0.44	94.8	**1.30**	100	1.49	100.0	1.45	99.7	1.71
	2.00	92.4	0.14	96.0	0.20	100	0.26	99.3	**0.15**	99.5	0.22

Table 7. D3: Large sample ($n = 1000$, $d_z = 2$). Selected (κ, ρ) configurations.

ρ	κ	DV		AR		Xia		CLR		StAR	
		C	W	C	W	C	W	C	W	C	W
0.0	0.02	100.0	2.68	99.6	20.0	100	74.7	100.0	20.00	**100**	**3.00**
	0.10	98.0	**1.20**	95.8	2.27	100	16.7	97.6	2.20	100	2.97
	1.00	95.1	0.12	95.2	**0.10**	100	0.31	95.1	**0.10**	99.4	0.19
0.5	0.02	90.9	2.30	98.9	20.0	100	62.6	98.9	20.00	**100**	**2.97**
	0.20	94.3	0.59	95.6	**0.70**	100	4.61	97.4	**0.70**	99.7	1.02
	1.00	95.3	0.12	96.2	**0.10**	100	0.31	95.6	**0.10**	99.4	0.19
0.9	0.02	34.7	1.25	95.9	20.0	100	33.8	95.7	20.00	**99.9**	**2.90**
	0.10	82.1	0.97	96.5	**2.55**	100	13.4	97.5	2.45	96.5	2.88
	1.00	95.6	0.12	96.2	**0.10**	100	0.30	96.1	**0.10**	100.0	0.19

Table 8. Small-sample robustness ($n = 50$, $d_z = 2$).

		DV		AR		Xia		CLR		StAR	
ρ	κ	Cov	Wid	Cov	Wid	Cov	Wid	Cov	Wid	Cov	Wid
0.0	0.10	**99.8**	**2.62**	99.4	20.00	100	18.22	99.7	20.00	100	3.50
	0.50	**97.6**	**1.10**	95.9	1.90	100	4.28	98.0	1.75	100	2.93
	2.00	94.7	0.28	93.9	0.30	100	0.58	**95.4**	**0.25**	100	0.89
0.5	0.10	90.6	2.28	98.7	20.00	100	16.06	98.2	20.00	**100**	**3.39**
	1.00	94.2	0.55	94.2	0.70	100	1.50	**97.0**	**0.65**	100	1.91
0.9	0.10	40.4	1.24	96.8	20.00	99.6	8.55	96.4	20.00	**100**	**3.07**
	0.50	85.1	0.90	96.9	2.10	99.7	3.53	**97.5**	**1.85**	99.7	2.62
	2.00	94.4	0.28	95.4	0.30	99.7	0.58	**95.4**	**0.25**	99.9	0.88

Table 9. D5: Heteroskedastic errors ($n = 200$, $d_z = 2$, $\mathrm{Var} \propto 1 + |z_1|$).

ρ	κ	DV		AR		Xia		CLR		StAR	
		C	W	C	W	C	W	C	W	C	W
0.0	0.05	**98.7**	**2.51**	97.5	20.0	100	30.08	98.2	20.00	100	3.18
	0.50	94.6	1.00	89.9	1.45	100	5.98	94.4	1.45	**100**	**3.00**
	1.00	88.6	0.53	85.3	0.60	100	1.99	90.8	0.60	**99.2**	**0.85**
0.5	0.05	84.6	2.20	94.9	20.0	100	27.13	94.7	20.00	**100**	**3.13**
	0.50	88.8	0.96	89.0	1.55	100	5.73	93.9	1.50	**99.5**	**2.96**
	1.00	89.4	0.52	87.0	0.60	100	1.93	92.9	0.55	**98.4**	**0.82**
0.9	0.05	28.9	1.07	90.5	20.0	99.9	12.89	89.5	20.00	**99.9**	**2.97**
	0.50	80.1	0.87	88.4	1.70	99.7	5.34	94.5	1.70	**95.8**	**2.82**
	1.00	87.4	0.48	86.4	0.55	99.8	1.78	93.0	0.55	**95.1**	**0.78**

Table 10. D6: Heavy-tailed (t_5) errors ($n = 200$, $d_z = 2$).

ρ	κ	DV		AR		Xia		CLR		StAR	
		C	W	C	W	C	W	C	W	C	W
0.0	0.05	**99.8**	**2.70**	99.2	20.0	100	35.17	99.5	20.00	100	3.19
	0.50	**96.1**	**0.69**	95.4	0.88	100	3.19	97.5	0.85	100	1.16
	1.00	**95.2**	**0.35**	94.9	0.35	100	1.03	**96.2**	**0.35**	99.9	0.55
0.5	0.05	90.4	2.36	98.7	20.0	100	31.28	98.7	20.00	**100**	**3.14**
	0.50	93.4	0.69	**95.5**	**0.90**	100	3.18	**97.4**	**0.90**	99.6	1.16
	1.00	94.4	0.35	94.2	0.35	100	1.03	**96.8**	**0.35**	99.8	0.55
0.9	0.05	34.0	1.21	97.0	20.0	100	15.35	96.8	20.00	**99.9**	**2.97**
	0.50	90.2	0.66	94.9	0.95	100	3.03	**97.5**	**0.90**	97.5	1.12
	1.00	94.1	0.35	95.1	0.40	100	1.03	**96.2**	**0.35**	98.9	0.55

Table 11. D7: Concentration parameter ($\rho = 0.7$, $d_z = 2$). $\mu^2 = n\kappa^2/d_z$.

μ^2	n	κ	DV		AR		Xia		CLR		StAR	
			C	W	C	W	C	W	C	W	C	W
1	100	0.14	77.9	1.61	97.1	20.0	100	12.64	98.0	20.00	**99.9**	**3.14**
5	100	0.32	86.8	1.02	95.1	2.35	100	5.23	**97.6**	**2.15**	98.9	2.93
10	100	0.45	91.0	0.81	94.8	1.25	100	3.42	**96.8**	**1.20**	99.0	2.72
50	100	1.00	93.9	0.39	95.0	0.45	100	1.02	**96.1**	**0.40**	98.9	0.60
100	100	1.41	94.5	0.28	95.8	0.30	100	0.61	**96.3**	**0.25**	99.4	0.41
1	200	0.10	78.0	1.63	96.9	20.0	100	18.18	97.9	20.00	**99.7**	**3.06**
5	200	0.22	87.6	1.06	96.0	2.50	100	7.46	**98.1**	**2.35**	99.6	2.97
10	200	0.32	90.2	0.79	94.5	1.20	100	4.29	**96.8**	**1.20**	97.8	1.46
50	200	0.71	94.4	0.38	**95.1**	**0.40**	100	1.19	**97.3**	**0.40**	99.3	0.61
100	200	1.00	94.4	0.28	93.8	0.30	100	0.71	**95.8**	**0.25**	99.2	0.42
1	1000	0.04	76.5	1.58	97.7	20.0	100	35.54	97.7	20.00	**99.9**	**2.94**
5	1000	0.10	87.9	1.03	96.2	2.30	100	13.92	**98.4**	**2.20**	99.4	2.92
10	1000	0.14	90.0	0.79	94.6	1.15	100	8.18	**97.0**	**1.15**	98.4	1.54
50	1000	0.32	94.2	0.38	**95.1**	**0.40**	100	2.06	**96.5**	**0.40**	98.7	0.63
100	1000	0.45	94.4	0.28	94.8	0.30	100	1.16	**95.5**	**0.25**	99.4	0.45

Table 12. OHIE semi-synthetic: $n = 5000$ (subgroup size). Compliance rate controls instrument strength; δ controls unmeasured confounding. $\bar{F}$ is the mean first-stage F-statistic. Coverage (%) and median CI width (W) over 500 replications.

δ	Comply	$\bar{F}$	DV		AR		Xia		CLR		StAR	
			C	W	C	W	C	W	C	W	C	W
0.0	0.20	3.1	99.8	1.48	98.2	20.00	100.0	101.19	98.2	20.00	**100.0**	**1.40**
	0.50	14.0	**97.4**	**0.52**	93.9	0.63	100.0	14.78	93.9	0.63	100.0	1.39
	0.75	30.1	**95.9**	**0.34**	93.8	0.37	100.0	6.58	93.8	0.37	100.0	0.52
	1.00	52.6	**95.0**	**0.26**	93.8	0.27	100.0	3.75	93.8	0.27	99.9	0.38
	1.50	116.8	94.3	0.17	93.8	0.17	100.0	1.72	93.8	0.17	**99.9**	**0.25**
0.5	0.20	3.1	99.8	1.47	98.3	20.00	100.0	100.77	98.3	20.00	**100.0**	**1.40**
	0.50	14.0	**97.6**	**0.52**	94.1	0.63	100.0	14.73	94.1	0.63	100.0	1.39
	0.75	30.1	**95.5**	**0.34**	94.0	0.37	100.0	6.58	94.0	0.37	100.0	0.52
	1.00	52.6	94.7	0.26	94.0	0.27	100.0	3.75	94.0	0.27	**99.9**	**0.38**
	1.50	116.7	94.3	0.17	94.0	0.17	100.0	1.72	94.0	0.17	**99.8**	**0.25**
0.9	0.20	3.1	99.7	1.46	98.3	20.00	100.0	99.35	98.3	20.00	**100.0**	**1.40**
	0.50	14.0	**97.4**	**0.52**	94.1	0.63	100.0	14.76	94.1	0.63	100.0	1.39
	0.75	30.0	**95.5**	**0.34**	94.0	0.37	100.0	6.59	94.0	0.37	100.0	0.52
	1.00	52.4	94.7	0.26	94.0	0.27	100.0	3.76	94.0	0.27	**99.9**	**0.38**
	1.50	116.4	94.4	0.17	94.0	0.17	100.0	1.73	94.0	0.17	**99.7**	**0.25**

Table 13. Real MIMIC-IV subgroup analysis. F = first-stage F-statistic after residualization. All CIs at 95% nominal level.

Subgroup	n	F	DV		AR/CLR		StAR	
			CI	W	CI	W	CI	W
Full sample	63,998	2,083	$[-0.13, -0.06]$	0.07	$[-0.13, -0.07]$	0.06	$[-0.13, -0.06]$	0.07
Surgical ICU	38,948	154	$[-0.34, -0.12]$	0.22	$[-0.34, -0.13]$	0.21	$[-0.35, -0.10]$	0.24
Medical ICU	21,762	5.7	$[0.08, 2.00]$	1.91	$[0.37, 5.24]$	4.87	$[-0.80, 1.22]$	**2.02**
Young (18–50)	12,272	399	$[-0.02, 0.09]$	0.11	$[-0.02, 0.08]$	0.10	$[-0.03, 0.09]$	0.12
Middle (50–70)	26,170	872	$[-0.08, 0.02]$	0.10	$[-0.07, 0.01]$	0.08	$[-0.08, 0.02]$	0.11
Elderly (70+)	25,556	760	$[-0.29, -0.15]$	0.13	$[-0.28, -0.16]$	0.12	$[-0.29, -0.14]$	0.15

Table 14. Semi-synthetic MIMIC-IV: $n = 2000$, $\beta = -0.05$. Instrument strength controls physician preference heterogeneity; δ controls unmeasured confounding. Coverage (%) and median CI width (W) over 500 replications.

δ	IV str	$\bar{F}$	DV		AR		Xia		CLR		StAR	
			C	W	C	W	C	W	C	W	C	W
0.0	0.03	1.9	99.9	1.66	99.4	20.00	100.0	104.36	99.4	20.00	**100.0**	**1.06**
	0.05	3.4	99.6	0.96	98.6	20.00	100.0	38.69	98.6	20.00	**100.0**	**1.06**
	0.10	10.9	**98.7**	**0.45**	96.2	0.58	100.0	9.26	96.2	0.58	100.0	0.93
	0.20	40.7	**96.5**	**0.22**	95.7	0.23	100.0	2.35	95.7	0.23	99.9	0.32
	0.30	90.6	**96.0**	**0.14**	95.7	0.15	100.0	1.09	95.7	0.15	99.7	0.21
0.3	0.03	1.8	100.0	1.74	99.4	20.00	100.0	114.00	99.4	20.00	**100.0**	**1.05**
	0.05	3.3	99.8	0.99	98.6	20.00	100.0	40.68	98.6	20.00	**100.0**	**1.05**
	0.10	10.3	**99.0**	**0.45**	96.3	0.60	100.0	9.52	96.3	0.60	100.0	0.94
	0.20	38.9	**96.6**	**0.22**	95.4	0.23	100.0	2.40	95.4	0.23	99.9	0.32
	0.30	86.7	**96.1**	**0.14**	95.4	0.15	100.0	1.11	95.4	0.15	99.6	0.21
0.5	0.03	1.7	99.9	1.74	99.5	20.00	100.0	111.88	99.5	20.00	**100.0**	**1.03**
	0.05	3.1	99.8	1.00	98.7	20.00	100.0	41.69	98.7	20.00	**100.0**	**1.03**
	0.10	9.7	**99.0**	**0.46**	95.8	0.63	100.0	9.87	95.8	0.63	100.0	0.94
	0.20	36.3	**96.5**	**0.22**	95.2	0.23	100.0	2.51	95.2	0.23	100.0	0.33
	0.30	80.7	**95.7**	**0.15**	**95.2**	**0.15**	100.0	1.16	**95.2**	**0.15**	99.8	0.21

Table 15. Real MIMIC-IV: Medical ICU subgroup ($n = 21{,}762$, $F = 5.66$). The weak-instrument setting where physician preference explains little treatment variation after conditioning on severity.

Method	Lower	Upper	Width
DV (2SLS)	0.083	1.997	1.914
AR	0.370	5.240	**4.870**
CLR	0.370	5.240	**4.870**
Xia	-82.2	84.3	166.4
StAR	-0.803	1.218	**2.021**

Table 16. Real MIMIC-IV: bootstrap coverage experiment. Coverage (%) and median CI width at varying subsample sizes. $\bar{F}$ = mean first-stage F-statistic. StAR maintains $\geq 99\%$ coverage throughout.

n	$\bar{F}$	DV		AR		Xia		CLR		StAR	
		C	W	C	W	C	W	C	W	C	W
500	17.8	98.5	0.74	98.5	0.85	100.0	6.47	98.5	0.85	**100.0**	2.06
1000	32.9	98.0	0.53	97.5	0.50	100.0	4.43	97.5	0.50	**99.0**	0.61
2000	66.4	99.0	0.38	95.0	0.35	100.0	3.23	95.0	0.35	**99.5**	0.43
5000	163.0	97.5	**0.24**	90.0	0.20	100.0	2.11	90.0	0.20	99.0	0.27

Table 17. Angrist & Krueger (1991) returns to schooling ($n = 329{,}509$, $d_z = 3$). $\hat{\beta}_{\mathrm{OLS}} = 0.071$, $\hat{\beta}_{\mathrm{IV}} = 0.105$, first-stage $F = 32.3$.

Method	Lower	Upper	Width	Key diagnostic
StAR (multivariate)	0.036	0.174	0.139	$\omega = 0.99$, $\log E = 41.1$
StAR (scalar)	-1.344	1.486	2.830	$\omega = 0.00$, $\log E = 0.65$
DV (2SLS)	0.066	0.145	0.079	
AR	0.064	0.152	0.088	
CLR	0.068	0.144	0.076	
Xia	-3.117	3.327	6.444	