

Online Supplement: Inference for Directional Population Sensitivity

Lucas D. Konrad and Nikolas Kuschnig

October 19, 2026

The supplement states and proves the primitive regression verifications, develops the unrestricted-refitting bridge and reports additional Monte Carlo, sample-audit and calibrated population-geometry detail. References to results outside this supplement refer to the main paper.

S1 Estimated conclusion margins

The paired procedure allows the conclusion margin and the directional curve to be estimated from the same observations $O_i = (x_i, y_i)$. The following assumption states the additional conditions on the margin estimator and its fitted influence score.

Assumption S1 (Asymptotic linearity of the conclusion margin). *For a population margin $c = c(P) > 0$, the estimator $\hat{c}$ and fitted influence scores $\hat{\varphi}_{c,i}$ satisfy*

$$\begin{aligned} \sqrt{N}(\hat{c} - c) &= \frac{1}{\sqrt{N}} \sum_{i=1}^N \varphi_c(O_i) + o_p(1), & \mathbb{E}\varphi_c &= 0, & \mathbb{E}\varphi_c^2 &< \infty, \\ \frac{1}{N} \sum_{i=1}^N \{\hat{\varphi}_{c,i} - \varphi_c(O_i)\}^2 &= o_p(1). \end{aligned} \tag{S1}$$

A fixed, known margin is the special case $\hat{c} = c$ and $\varphi_c = \hat{\varphi}_c = 0$.

Contrast margins and empirical orientation. For a prespecified contrast $a'_0\beta \neq b$, let $\iota = \text{sign}(a'_0\beta - b)$ orient the adverse direction as $a = \iota a_0$. The population margin, its

plug-in estimate, and its influence score are

$$c(P) = |a'_0\beta - b|, \quad \hat{c} = |a'_0\hat{\beta} - b|, \quad \varphi_c(x, y) = \iota(a'_0Q^{-1}x)u = (r'_a x)u.$$

In sample computation, use $\hat{\iota} = \text{sign}(a'_0\hat{\beta} - b)$, $\hat{a} = \hat{\iota}a_0$, and $\hat{\varphi}_{c,i} = (\hat{r}'_{\hat{a}}x_i)\hat{u}_i$, evaluating the fitted curve and its scores in direction $\hat{a}$. Because $c(P) > 0$ and OLS consistency imply $\mathbb{P}(\hat{\iota} = \iota) \rightarrow 1$, this empirical orientation shares the same asymptotic law as the fixed population direction a . The maintained moment and OLS consistency conditions deliver the empirical- L^2 consistency in (S1). A zero population margin ($a'_0\beta = b$) lies outside this regular regime: the absolute-value map is not differentiable there, and the estimated sign need not converge to a fixed orientation.

Other margin specifications. Conclusion margins can also measure the distance to a nominal significance boundary, either overturning or creating apparent evidence: with a fixed reference scale s_0 and critical value z , the downward and upward creation margins for a coefficient β_T inside $[-zs_0, zs_0]$ are $\beta_T + zs_0$ and $zs_0 - \beta_T$. These margins inherit the coefficient's influence score with the corresponding sign; an estimated reference scale requires its own joint expansion, and crossing a nominal boundary after outcome-dependent deletion does not establish selection-adjusted significance. Only strictly positive margins satisfying Assumption S1 are covered by the regular paired procedure.

Pointwise gap inference. For the gap and paired score in (17), Theorem 2 and assumption S1 give

$$\sqrt{N}\{\hat{g}(p) - g(p)\} = N^{-1/2} \sum_{i=1}^N h_i(p) + o_p(1)$$

uniformly on $\mathcal{P}$. At a single prespecified share p , the pointwise expansion instead requires only Assumption 1(a) and Assumption S1. With $\text{Var}(h(p)) > 0$, the pointwise gap standard error is $\hat{\sigma}_g(p)/\sqrt{N}$, where $\hat{\sigma}_g^2(p) = \widehat{\text{Var}}(\hat{h}(p))$ incorporates the covariance between the fragility curve and the estimated margin. This gives the Wald interval $\hat{g}(p) \pm z_{1-\alpha/2}\hat{\sigma}_g(p)/\sqrt{N}$ and a one-sided test of $g(p) \leq 0$, rejecting when $\sqrt{N}\hat{g}(p)/\hat{\sigma}_g(p) > z_{1-\alpha}$. Choosing a share after inspecting the path instead requires simultaneous inference, as in Proposition 2.

S2 Primitive Regression Verifications

S2.1 Continuous and mixed covariate laws

Proposition S1 (Continuous- and mixed-control primitive verification). *Let (x, y) be i.i.d. in fixed dimension, fix $a \neq 0$, and suppose $Q = \mathbb{E}[xx'] > 0$, $\mathbb{E}\|x\|^4 < \infty$, and $P(x'Q^{-1}a = 0) = 0$. Write $y = m(x) + \sigma_0(x)\varepsilon$ with measurable m, σ_0 satisfying*

$$\mathbb{E}[(1 + \|x\|^2)m(x)^2] < \infty, \quad 0 < \sigma_- \leq \sigma_0(x) \leq \sigma_+ < \infty.$$

Suppose $\mathbb{E}[\varepsilon \mid x] = 0$, and $\varepsilon \mid x$ has a jointly measurable regular conditional density $h_x(t)$, strictly positive and continuous in t for almost every x , with deterministic constants $M_h, V < \infty$ such that

$$\sup_{t \in \mathbb{R}} (1 + t^2)h_x(t) \leq M_h, \quad t^2 h_x(t) \longrightarrow 0 \quad (|t| \longrightarrow \infty), \quad \mathbb{E}[\varepsilon^2 \mid x] \leq V \quad \text{a.s.}$$

The tail limit is pointwise for almost every x , not necessarily uniform in x . With $\beta = Q^{-1}\mathbb{E}[xm(x)]$, the standing assumptions Assumptions 2 to 5 and Assumption 6 hold at every $p \in (0, 1)$, and Assumption 7 holds on every compact $\mathcal{P} \subset (0, 1)$. Consequently the pointwise intervals and functional multiplier bands apply.

Proof of Proposition S1. Write $r_0 = Q^{-1}a$, $\mu = a'r_0$, $v = r_0'x$, $Y = vu$, and $Z = v^2/\mu$. The bounded positive scales, fourth design moment and exact mixed moment $\mathbb{E}[(1 + \|x\|^2)m(x)^2] < \infty$ imply

$$\beta = Q^{-1}\mathbb{E}[xm(x)], \quad \mathbb{E}[xu] = 0, \quad 0 < \mathbb{E}u^2 < \infty, \quad \mathbb{E}[\|x\|^2 u^2] < \infty.$$

Thus $\mathbb{E}Y = 0$, $\mathbb{E}Z = 1$, $\mathbb{E}(Y^2 + Z^2) < \infty$, and $H = 1 + \|x\| |u| + \|x\|^2 \in L^2$, proving the standing moments and C1.

Leverage and roots before cutoff smoothness. Fix $\mathcal{P} = [\underline{p}, \bar{p}] \subset (0, 1)$ and put

$$b_0 = \int_0^{1-\bar{p}} F_Z^{-1}(t) dt > 0.$$

Here positivity follows from $P(v = 0) = 0$. Bottom-quantile allocation gives $\mathbb{E}[ZR] \geq b_0$ for every intensity $0 \leq R \leq 1$ with $\mathbb{E}R = 1 - p$, $p \leq \bar{p}$. This is all-rule scalar leverage, not retained-matrix rank. Conditional full support makes Y nonconstant with both tails; its mean zero implies $\mathcal{T}_p(Y) > 0$ for every proper upper tail. Hence Assumptions 3

and 4 and Proposition 4 give a unique positive relaxed root $\delta(p) \leq \mathbb{E}|Y|/b_0$. Quantile integration is continuous in p for integrable scores, and $|\mathcal{T}_p(W_\lambda) - \mathcal{T}_p(W_{\lambda'})| \leq |\lambda - \lambda'| \mathbb{E}Z$. Joint continuity and strict descent by b_0 therefore give root continuity without cutoff regularity. For every finite λ , the oracle score $W_\lambda = \nu u + \lambda Z$ mixes full-support absolutely continuous laws, since $P(\nu = 0) = 0$. Its CDF is continuous and strictly increasing, proving Assumption 5 for every $\lambda \geq 0$. At the root,

$$\mathbb{E}W_\delta = \delta = \mathcal{T}_p(W_\delta), \quad \int_0^{1-p} F_{W_\delta}^{-1}(t) dt = 0.$$

There is strictly negative score mass. A nonpositive $(1-p)$ -quantile would make the last integral strictly negative. Thus $\xi(p) = \xi_a^*(p) > 0$, and the root region has exactly mass p , identifying $\delta(p) = \delta(p)$.

Joint nuisance regularity through zero indices. Let $\theta = (b, r)$, $c = a'r$, $s = r'x$, $d = m(x) - x'b$, $k = d + \lambda s/c$, and

$$W_{\theta, \lambda} = s(y - x'b) + \lambda s^2/c, \quad t = (w/s - k)/\sigma_0(x) \quad (s \neq 0).$$

Work on compact parameter sets with $c \geq \mu/2$, bounded λ , and $w \in [\eta, M]$, $\eta > 0$. Holding x fixed, the conditional score CDF is $F_x(t)$ for $s > 0$, $1 - F_x(t)$ for $s < 0$, and 1 for $s = 0$. For $s \neq 0$ its derivatives are

$$\begin{aligned} \partial_w F &= h_x(t)/(s|\sigma_0|), & \partial_b F &= \text{sign}(s)h_x(t)x/\sigma_0, \\ \partial_\lambda F &= -|s|h_x(t)/(c\sigma_0), & \nabla_r F &= \frac{\text{sign}(s)h_x(t)}{\sigma_0} \left[\left(-\frac{w}{s^2} - \frac{\lambda}{c} \right) x + \frac{\lambda s}{c^2} a \right]. \end{aligned}$$

They extend continuously by zero through $s = 0$, using $t^2 h_x(t) \rightarrow 0$ pointwise in x . Actual differentiability there follows since, for fixed x and locally small s , the CDF differs from 1 by at most $P(|\varepsilon| > C_x/|s| \mid x) = O(s^2) = o(|s|)$. No density derivative is assumed. The identity $w/s = k + \sigma_0 t$ gives

$$\frac{h_x(t)}{|s|\sigma_0} \leq C(1 + |m(x)| + \|x\|)/\eta, \quad \frac{w}{s^2} h_x(t) \leq C(1 + m(x)^2 + \|x\|^2)/\eta.$$

All first CDF derivatives therefore have the integrable envelope $C\{1 + |m(x)| + \|x\|^3 + \|x\| m(x)^2\}$, even for unbounded controls. Dominated differentiation makes the unconditional CDF jointly C^1 . Its density $\mathbb{E}[h_x(t)/(s|\sigma_0|)]$, with the integrand extended

by zero at $s = 0$, is finite, continuous and positive on positive-score domains: $c > 0$ and $Q > 0$ imply $P(s \neq 0) > 0$. No global density bound at score zero is asserted; nearby nuisance directions may have zero-index atoms.

The quantile implicit function theorem gives locally C^1 nuisance quantiles about each positive root cutoff. Root continuity implies cutoff continuity, hence compactness gives $0 < \min_p \xi(p) \leq \max_p \xi(p) < \infty$. A sufficiently small common compact nuisance/root/cutoff tube lies entirely at positive score coordinates, with density bounded above and away from zero. Quantile bracketing and uniqueness join its local quantile maps, proving C2 and U1.

Leverage, differentiability and empirical expansions. Uniform L^1 continuity of $Z_\theta = (r'x)^2/(a'r)$ preserves retained leverage at least $b_0/2$ on a small nuisance neighborhood, proving C3 and the slope clause of U2. The moments and cutoff regularity already established give C1–C3; Lemma 5 then supplies the criterion equicontinuity and direction differentiability conclusions C4–C5, with the full gradient in (8).

Positive variances and feasible inference. For fixed p , put $\alpha(x) = m_a(p)'Q^{-1}x$. Conditional on x with $v \neq 0$, (9) is continuous and piecewise affine in u , with retained and active slopes $-\alpha(x)$ and $v - \alpha(x)$; the design correction is constant given x . On $\{v \neq 0\}$ both branches have positive conditional probability and their slopes differ by $v \neq 0$, so conditional variance is positive there. Since this event has positive probability, the unconditional variance is positive. The oracle score has active slope v , proving its positivity too. Continuous root/cutoff paths, boundary atomlessness and the common L^2 envelope give continuous selected moments and influence variances. Compactness supplies their positive variance floors, proving C6 and U2. The common cutoff tube, leverage and variance floors and $H \in L^2$ meet the hypotheses of Lemmas 12 and 13. Now Lemma 14 supplies U3, and Lemmas 9, 15 and 16 supply fitted empirical- L^2 score consistency, uniform variance consistency, studentization and multipliers. $\square$

Proposition S2 (Localized primitive verification with zero-index atoms). *Maintain the moment, conditional-mean, scale and conditional-error-density conditions of Proposition S1, except replace $P(x'Q^{-1}a = 0) = 0$ by the following condition. For a compact budget interval $\mathcal{P} = [\underline{p}, \bar{p}] \subset (0, 1)$, let $\rho_0 = P(x'Q^{-1}a = 0)$ and assume $\rho_0 < 1 - \bar{p}$. In addition, assume the threshold is unique and has exact upper-tail mass on a neighborhood of the population root path; on a common neighborhood of its cutoff path, the score laws have jointly continuous densities bounded above and bounded away from zero. These are local conditions only and need not hold at other multiplier values or score levels.*

Then the directional root is attained by a region of mass p for each $p \in \mathcal{P}$. The pointwise CLT and studentization, and the curve functional CLT and multiplier band, hold on $\mathcal{P}$ with these local conditions replacing the global threshold and cutoff-density requirements in their proofs. The band-inversion crossing interval also applies; the crossing Wald result additionally requires an interior crossing. For paired gap-band conclusions, retain the fixed nonzero-margin conditions of Corollary S1 and the stated conditions of Proposition 2. Inference remains directional: no unrestricted-refit agreement, exact-refit certificate or finite-sample coverage guarantee is implied. If $p \geq 1 - \rho_0$, the conclusion is not asserted.

Proof of Proposition S2. Write $r_0 = Q^{-1}a$, $v = r'_0x$, $\mu = a'r_0$, $Y = vu$, and $Z = v^2/\mu$. The moment calculation at the start of the preceding proof is unchanged. For $\tau_1 = 1 - \bar{p}$, define

$$b_0 = \int_0^{\tau_1} F_Z^{-1}(t) dt.$$

Since $Z = 0$ exactly on the zero-index event, $\rho_0 < \tau_1$ implies $b_0 > 0$. The bottom-quantile allocation bound gives $\mathbb{E}[ZR] \geq b_0$ for every retained intensity of mass $1 - p$, uniformly over $p \in \mathcal{P}$. Also $P(v \neq 0) > 0$ because $Q > 0$ and $a \neq 0$; conditional full support on this event makes Y nonconstant with both tails. The population root therefore exists, is unique and positive by Proposition 4, and its path is continuous by strict descent with slope bounded away from zero.

At the root, $\mathbb{E}W_\delta = \mathcal{T}_p(W_\delta) = \delta$, so the lower-tail integral of W_δ through its $1 - p$ quantile is zero. A nonpositive quantile would make this integral strictly negative: the score has positive probability of being negative, while the zero-index atom is at zero. Hence the root cutoff is strictly positive. Conditional full support for $v \neq 0$ implies a continuous, strictly increasing score distribution at every positive score; consequently the root cutoff is unique and has exact upper-tail mass p . The region target is thus attained and equals the relaxed root.

The positive-score nuisance-CDF argument in the preceding proof already includes $s = r'x = 0$: its derivatives extend continuously by zero through this event and have the same integrable envelope. It therefore applies unchanged near the positive root cutoffs here, despite the oracle zero-index atom.

The assumed local exact-mass thresholds and density bounds supply a common positive-cutoff tube about the compact root path, and the retained-leverage floor persists under small nuisance perturbations. The proofs of Lemmas 5, 12 and 13

use precisely these cutoff-local properties, the L^2 envelope CH and weighted laws for score half-spaces on this tube. Their profiling, value-sandwich and root-transfer arguments therefore apply with the local replacements, without imposing the global threshold condition at multipliers outside the tube. After the variance floors below are established, the uniform first-stage and fitted-multiplier arguments yield the stated inference conclusions on $\mathcal{P}$.

Finally, conditional on x with $v \neq 0$, the feasible influence is piecewise affine in u with branch slopes $-\alpha(x)$ and $v - \alpha(x)$, where $\alpha(x) = m_a(p)'Q^{-1}x$. Their difference is $v \neq 0$, and both branches have positive conditional probability, so the conditional variance is positive on a set of positive probability. For the paired gap the slope difference is v/B ; it is nonzero on $\{v \neq 0\}$, not almost surely when $\rho_0 > 0$. The root and cutoff paths are continuous, and the local boundary atomlessness and common L^2 envelope give continuity of selected moments and influence variances, including the paired gap. Hence the pointwise variances have positive minima on the compact budget interval. This supplies the variance floors when the paired-gap result's other stated margin conditions hold; the curve inference conclusions do not require a margin. No claim is made at or above the budget cap. $\square$

Corollary S1 (Paired sign-margin inference). *Under Proposition S1 for a fixed $a_0 \neq 0$ and threshold b , suppose $c = |a'_0\beta - b| > 0$ and orient $a = \iota a_0$, where $\iota = \text{sign}(a'_0\beta - b)$. Then (S1) holds with $\varphi_c = (r'_a x)u$ and its fitted score, and*

$$\inf_{p \in \mathcal{P}} \text{Var} \left\{ \frac{\Psi_a(p)}{B_a^*(p)} - (r'_a x)u \right\} > 0.$$

Thus Proposition 2 and corollary 1 apply; Theorem 3 additionally requires an actual interior crossing. Validity is pointwise in the population, not uniform as c approaches zero.

Proof of Corollary S1. Let $c = |a'_0\beta - b| > 0$, $\iota = \text{sign}(a'_0\beta - b)$, and $a = \iota a_0$. OLS asymptotic linearity and sign consistency give (S1) with $\varphi_c = vu$ and $\hat{\varphi}_{c,i} = (\hat{r}'_a x_i)\hat{u}_i$. Indeed $\mathbb{P}(\hat{\iota} = \iota) \rightarrow 1$; on this event the fitted-minus-population margin score is bounded by $o_p(1)H$, proving empirical- L^2 consistency since $P_N H^2 = O_p(1)$. The paired gap is $h(p) = \Psi_a(p)/B_a^*(p) - vu$, as in (17). Put $B = B_a^*(p)$ and $\alpha(x) = m_a(p)'Q^{-1}x$. Conditional on x , its retained and active slopes are

$$-\alpha(x)/B - v, \quad (v - \alpha(x))/B - v.$$

On $\{v \neq 0\}$ both branches have positive probability and their slopes differ by $v/B \neq 0$; this event has positive probability, though it need not have probability one. Thus $\text{Var}(h(p)) > 0$; continuity and compactness give $\inf_{p \in \mathcal{P}} \text{Var}(h(p)) > 0$. The margin-score bound, Lemma 15, and the slope floor give $\sup_p P_N[\{\hat{h}(p) - h(p)\}^2] = o_p(1)$. The fixed compact score family has an integrable squared envelope; its uniform law and Cauchy–Schwarz for the fitted replacement give uniform gap-variance consistency. Hence Proposition 2 and corollary 1 apply. Crossing Wald inference additionally requires an actual interior crossing; band inversion does not. This is a fixed nonzero-margin result, not uniform validity as margins approach zero. $\square$

Remark S1. *The independent symmetric density proportional to $\{(1 + t^2)^{3/2}[\log(e + t^2)]^2\}^{-1}$ satisfies these density and variance bounds but has no moment above order two and no polynomial density envelope of exponent greater than three. The fourth-design and mixed-mean moments remain required.*

S2.2 Finite covariate support

Proposition S3 (Finite-covariate primitive verification). *Let x take fixed values $x_1, \dots, x_J \in \mathbb{R}^d$ with probabilities $\pi_j > 0$, where J and d are fixed and $Q = \sum_j \pi_j x_j x_j' > 0$. For each j , suppose that $y \mid x = x_j$ has a strictly positive continuous density h_j on $\mathbb{R}$ and a finite second moment. Fix $a \neq 0$ and suppose $v_j = x_j' Q^{-1} a \neq 0$ for every j . Under i.i.d. sampling of (x, y) , with β the linear projection coefficient of (1), Assumption 6 holds at every $p \in (0, 1)$, and Assumption 7 holds on every compact $\mathcal{P} \subset (0, 1)$. In particular, with $\mu = a' Q^{-1} a$,*

$$B_a^*(p) \geq (1 - p) \min_{1 \leq j \leq J} \frac{v_j^2}{\mu} > 0.$$

Proof of Proposition S3. Write $c_j = \mathbb{E}[y \mid x = x_j]$, $\mu = a' Q^{-1} a$, and $r_a = Q^{-1} a$. The projection coefficient is $\beta = Q^{-1} \sum_j \pi_j x_j c_j$, so $\mathbb{E}[xu] = 0$; the finite support and finite conditional second moments give C1 and the standing model moments. Put $z_j = v_j^2 / \mu > 0$. Every retained region has leverage at least $(1 - p) \min_j z_j$, so Assumption 3 holds. The conditional densities h_j imply that $Y(a) = v_j u$ has a continuous, nondegenerate full-support distribution and mean zero, hence its upper- p tail integral is positive. Thus Assumption 4 and Proposition 4 give a positive finite root.

Joint cutoff regularity. For $\theta = (b, r)$ near (β, r_a) , let $s_j = r'x_j$, $c(r) = a'r$, and

$$W_{\theta, \lambda} = (r'x)(y - x'b) + \lambda \frac{(r'x)^2}{a'r}, \quad \gamma_j = -s_j x_j' b + \lambda s_j^2 / c(r).$$

Choose a compact neighborhood where $c(r) \geq \mu/2$ and each s_j is bounded away from zero with the sign of v_j . In cell j , the score is $s_j y + \gamma_j$, with mixture density

$$f_{\theta, \lambda}(w) = \sum_{j=1}^J \frac{\pi_j}{|s_j|} h_j \left(\frac{w - \gamma_j}{s_j} \right). \quad (\text{S2})$$

It is positive and jointly continuous in (θ, λ, w) . On compact parameter and cutoff sets it is bounded above and away from zero. The cellwise CDFs are $F_j((w - \gamma_j)/s_j)$ for $s_j > 0$ and $1 - F_j((w - \gamma_j)/s_j)$ for $s_j < 0$; since $F_j' = h_j$ is continuous, the mixture CDF is C^1 and strictly increasing. This verifies Assumption 5 and the local density regularity.

Fix $\mathcal{P} = [p, \bar{p}] \subset (0, 1)$. The retained-slope floor is $b_0 = (1 - \bar{p}) \min_j z_j > 0$, so the root is bounded by $\mathbb{E}|Y(a)|/b_0$. Joint continuity of the tail criterion on bounded root sets and this uniform descent give continuity of the root path; strict quantile monotonicity gives continuity of its cutoff path. On a common band about that compact path, the density in (S2) has uniform positive lower and finite upper bounds. Quantile bracketing and the implicit function theorem therefore give a C^1 cutoff map locally, verifying C2, U1, and the slope clauses of C3 and U2.

Derived process regularity. The established moments, leverage and cutoff conditions give C1–C3, so Lemma 5 supplies C4–C5.

Influence variance. For fixed p , set $\alpha_j = m_a(p)'Q^{-1}x_j$. Conditional on cell j , (9) is continuous and piecewise affine in u , with slopes $-\alpha_j$ on the retained half-line and $v_j - \alpha_j$ on the active half-line. Both half-lines have positive probability under $h_j > 0$, and the slopes cannot both vanish since $v_j \neq 0$. Thus the conditional variance, and hence C6, is positive; the oracle score also has positive variance because its active-region slope is $v_j \neq 0$. The root and cutoff paths and selected moments vary continuously in p ; a common $C(1 + |y|)$ square-integrable envelope gives continuity of the influence variance. Its minimum on $\mathcal{P}$ is positive, establishing U2. The finite-support envelope $C(1 + |y|) \in L^2$, common cutoff tube and these floors give Lemmas 12 and 13; Lemma 14 then supplies U3. $\square$

S2.3 Gaussian regressors through the origin

Proposition S4 (Primitive Gaussian verification). *Let $x \sim \mathcal{N}(0, I_d)$ and $u \sim \mathcal{N}(0, 1)$ independently, let $y = x'\beta + u$, and fix $a = e_1$. Then Assumption 6 holds at every $p \in (0, 1)$ and Assumption 7 on every compact $\mathcal{P} \subset (0, 1)$. The design gradient is zero; with $T = x_1$ and $\nu(p) = \mathbb{E}[T^2 \mathbf{1}\{A_p^*\}]$,*

$$\Psi_a(p) = (Tu + \delta(p)T^2 - \xi_a^*(p))_+ + p\xi_a^*(p) - \delta(p)T^2 - \nu(p)Tu.$$

Proof of Proposition S4. Apply Proposition S1 with $m(x) = x'\beta$, $\sigma_0(x) = 1$, and $\varepsilon = u$. The Gaussian design has $Q = I_d$ and all fourth and mixed-mean moments; the independent standard normal conditional density is jointly measurable, positive and continuous, has variance one, obeys $\sup_t (1 + t^2)\phi(t) < \infty$ and $t^2\phi(t) \rightarrow 0$, and has conditional mean zero. Finally $r_a = e_1$ and $P(x_1 = 0) = 0$. Thus all standing assumptions, Assumption 6, and Assumption 7 follow from that proposition.

For $r \neq 0$, rotational invariance and independence give

$$((r'x)u, (r'x)^2) \stackrel{d}{=} (\|r\| Tu, \|r\|^2 T^2).$$

Positive homogeneity of the tail criterion therefore yields

$$\delta^{\text{sc}}(r; p) = \delta(p) / \|r\|, \quad \Phi_a(r; p) = \frac{a'r}{\|r\|} \delta(p).$$

The latter has zero gradient at $r_a = a = e_1$. Since A_p^* depends only on (T, u) , each coordinate x_j , $j > 1$, is independent of that event and of T and has mean zero. Consequently $\mathbb{E}[x_j T \mathbf{1}\{A_p^*\}] = 0$ for $j > 1$, whereas its first coordinate is $\mathbb{E}[T^2 \mathbf{1}\{A_p^*\}] = \nu(p)$: hence $m_a(p) = \nu(p)e_1$. In (9), the design layer vanishes and the residual layer is $-m_a' Q^{-1} x u = -\nu(p)Tu$. Substituting $Y = Tu$, $Z = T^2$ gives the stated score. $\square$

S3 Unrestricted-refitting bridge

The exact-refitting target provides the unrestricted comparison for the directional sensitivity developed in the main text. We define its population value and fixed-point characterization here; exact-statistic sampling theory is developed separately in the companion manuscript named in Section 1. The remaining results focus on the geometry, bounds, and conditions relating the exact and directional targets.

S3.1 Sample definitions and scalar equivalence

The exact maximum-influence statistic. The unrestricted sample benchmark maximizes the decrease in the estimated contrast $a'\hat{\beta}$ over all size- k deletion sets $S \in \mathcal{S}$. By the full-sample normal equations $\sum_{i=1}^N x_i \hat{u}_i = 0$, the retained residual sum satisfies $-\sum_{i \notin S} x_i \hat{u}_i = \sum_{i \in S} x_i \hat{u}_i$, which yields the standard leave- S -out identity.

Lemma S1 (Leave- S -out identity). *For every $S \in \mathcal{S}$, define the retained design matrix and deleted score sum*

$$\hat{Q}_{-S} := \frac{1}{N} \sum_{i \notin S} x_i x_i', \quad \hat{g}_S := \frac{1}{N} \sum_{i \in S} x_i \hat{u}_i.$$

Both use full- N normalization; the retained-sample design average is $\hat{Q}_{-S}/(1 - p_N)$. Whenever $\hat{Q}_N$ and $\hat{Q}_{-S}$ are invertible,

$$\Delta_a(S) := a'(\hat{\beta} - \hat{\beta}_{-S}) = a' \left(\sum_{i \notin S} x_i x_i' \right)^{-1} \sum_{i \in S} x_i \hat{u}_i = a' \hat{Q}_{-S}^{-1} \hat{g}_S,$$

where $\hat{\beta}_{-S}$ is the OLS estimator computed on the retained sample $[N] \setminus S$.

At the population level, removing a measurable region D of mass $P(D) = p$ obeys the analogous identity $\vartheta(P) - \vartheta(P_{-D}) = a' \Sigma(D)^{-1} g(D)$, where $\Sigma(D) := \mathbb{E}[xx' \mathbf{1}\{D^c\}]$ is the retained design moment and $g(D) := \mathbb{E}[xu \mathbf{1}\{D\}]$ is the deleted score moment (Section S3). Evaluating the exact sample maximum over all $\binom{N}{k}$ subsets requires every retained design matrix to be invertible.

Assumption S2 (Sample non-degeneracy). *At the reported deletion budget,*

$$\kappa_N := \min_{S \in \mathcal{S}} \lambda_{\min}(\hat{Q}_{-S}) > 0.$$

Definition S1. *The exact maximum-influence statistic in direction a , defined on the event of Assumption S2, is*

$$\hat{\delta}^{\text{ex}}(p) := \max_{S \in \mathcal{S}} a' \hat{Q}_{-S}^{-1} \hat{g}_S. \quad (\text{S3})$$

Assumption S4(E1) states the population counterpart of Assumption S2. Because the retained inverse $\hat{Q}_{-S}^{-1}$ varies with the deletion set S , subset deletion both rescales leverage and rotates the effective contrast direction $\hat{Q}_{-S}^{-1}a$, coupling all d coordinates of the design. The directional statistic in (4) holds the full-sample contrast direction fixed while adjusting for retained leverage along that direction, requiring a much

weaker scalar denominator condition; Section 3.3 compares the resulting population targets.

Lemma S2 (Scalar reduction when both statistics are defined). *Suppose $\hat{Q}_N$ is invertible and the sample satisfies Assumption S2, so that both statistics are defined. Let $\mathcal{D}_N$ denote the event of Proposition 3 that every retained denominator is strictly positive, read at the feasible pair $(\hat{s}_i(a), \hat{Z}_i(a))$. If $d = 1$ and $a = 1$, then on $\mathcal{D}_N$ the two statistics coincide exactly,*

$$\hat{\delta}(p) = \hat{\delta}^{\text{ex}}(p).$$

For $d = 1$ and $a = 1$, $\hat{s}_i(1) = x_i \hat{u}_i / \hat{Q}_N$ and $\hat{Z}_i(1) = x_i^2 / \hat{Q}_N$, so for each size- k set S ,

$$\frac{\sum_{i \in S} \hat{s}_i(1)}{\sum_{i \notin S} \hat{Z}_i(1)} = \frac{\sum_{i \in S} x_i \hat{u}_i}{\sum_{i \notin S} x_i^2}.$$

Maximizing over S gives $\hat{\delta}^{\text{ex}}(p)$, which establishes the scalar reduction directly. In one dimension, the retained inverse $\hat{Q}_S^{-1}$ changes only by a scalar factor, so the directional and unrestricted statistics coincide whenever defined. Under the conditions of Lemma 8(d), $\mathbb{P}(\mathcal{D}_N) \rightarrow 1$. For $d \geq 2$, unrestricted refitting may also rotate the contrast direction; Section 3.3 analyzes the resulting population target difference.

S3.2 Population comparison

The directional target optimizes exact retained least-squares movements on a fixed coefficient line; unrestricted refitting allows the displacement to leave that line. This restriction does not by itself order adverse contrast movements, because least squares minimizes prediction error rather than the target contrast. To compare its population target $\delta^{\text{ex}}(p)$ with the directional target $\delta(p)$ at a fixed deletion share $p \in (0, 1)$, the population normal equations imply that removing a measurable region D with $P(D) = p$ shifts the projection contrast by

$$\Phi(D) := \vartheta(P) - \vartheta(P_{-D}) = a' \Sigma(D)^{-1} g(D), \quad \Sigma(D) := \mathbb{E}[xx' \mathbf{1}\{D^c\}], \quad g(D) := \mathbb{E}[xu \mathbf{1}\{D\}],$$

so that $\delta^{\text{ex}}(p) = \sup_{P(D)=p} \Phi(D)$. For any D with $\Sigma(D) > 0$, $\Phi(D)$ is an attained movement; with $B_D = r_a' \Sigma(D) r_a / \mu_Z(a) > 0$, its difference from the line-restricted movement is

$$\Phi(D) - \frac{r_a' g(D)}{B_D} = g(D)' \{ \Sigma(D)^{-1} a - B_D^{-1} r_a \}. \quad (\text{S4})$$

At $D = A^*$ the line term is $\delta(p)$; this same-set rotation increment can have either sign. An admissible actual refit supplies a lower witness for the unrestricted supremum even without alignment.

Evaluating the retained design and deleted score moments at the directional active set $A^* = \{W_a^* \geq \xi_a^*\}$ defines the directional *retained-design metric* Q_{-A^*} and moment pair (Σ_{-A^*}, g_{A^*}) :

$$Q_{-A^*} := (1 - p)^{-1} \mathbb{E}[xx' \mathbf{1}\{A^{*,c}\}], \quad \Sigma_{-A^*} := (1 - p)Q_{-A^*}, \quad g_{A^*} := \mathbb{E}[xu \mathbf{1}\{A^*\}]. \quad (\text{S5})$$

Two alignments govern comparison with the directional refit. The first requires the retained design to rescale, rather than rotate, the contrast direction:

$$Q_{-A^*}^{-1}a \propto Q^{-1}a. \quad (\text{SN1})$$

The second requires the refitted shift $b = \Sigma_{-A^*}^{-1}g_{A^*}$ to align with the contrast direction:

$$\Sigma_{-A^*}^{-1}g_{A^*} \propto Q^{-1}a \quad (\text{equivalently, } g_{A^*} \propto a \text{ when Equation (SN1) holds}). \quad (\text{SN2})$$

Each condition imposes $d - 1$ directional restrictions rather than full matrix equality $Q_{-A^*} = Q$. The first restricts a retained second moment of x ; the second restricts an active-set cross-moment of x and u , and by Lemma S4 coincides under Equation (SN1) with the vanishing of the first-stage design correction ($g_a = 0$). No contrast rotation alone supplies the one-sided comparison:

Proposition S5 (No contrast rotation orders the population targets). *Under Assumption S4(E1)–(E2), Assumption 6, and Equation (SN1),*

$$\Sigma_{-A^*}^{-1}a = B_a^{*-1}r_a \quad \text{and} \quad \Phi(A^*) = \delta(p) \leq \delta^{\text{ex}}(p),$$

with strict inequality $\delta(p) < \delta^{\text{ex}}(p)$, under Assumption S4(E3), whenever Equation (SN2) fails and the design is boundary-regular.

Here boundary-regularity means positive Lebesgue density of (x, u) on an open ball in $\mathbb{R}^{d+1}$ intersecting $\{W_a^* = \xi_a^*\}$. It excludes intercept and discrete-covariate specifications, but is needed only for the strict-gap clause, not the sufficient lower bound.

Because $\delta(p) \leq \delta^{\text{ex}}(p)$ under Equation (SN1), directional fragility directly implies unrestricted fragility.

Corollary S2 (Transfer of evidence of exact fragility). *Let the conditions of Proposition 2 hold, and suppose Assumption S4(E1)–(E2) and Equation (SN1) hold at every $p \in \mathcal{P}$. For the lower gap band $L_{g,N}$ of Proposition 2 and margin $c = c(P) > 0$,*

$$\liminf_{N \rightarrow \infty} \mathbb{P}\{L_{g,N}(p) \geq 0 \implies \delta^{\text{ex}}(p) \geq c \text{ for every } p \in \mathcal{P}\} \geq 1 - \alpha.$$

Proof. On simultaneous gap coverage, $L_{g,N}(p) \leq \delta(p) - c \leq \delta^{\text{ex}}(p) - c$ at every budget, by Proposition S5. The implication therefore holds on an event whose probability tends to $1 - \alpha$. $\square$

The simultaneous band and its inverse concern the directional population target. Under no contrast rotation throughout $\mathcal{P}$, a nonnegative lower gap band transfers evidence of fragility to unrestricted refitting; a negative upper gap band certifies directional stability, not unrestricted stability. The latter requires target agreement or a valid upper confidence bound for $\delta^{\text{ex}}(p)$.

Remark S2 (The qualification is sharp: an audit where the ordering reverses). *Equation (SN1) restricts the retained design at the directional active set, and when it fails the two targets are ordered in neither direction. In the intercept-and-slope regression with covariate values $w = (0, 1, 2, 3, 4)$ and outcomes $y = (1, 2, -1, 4, -4)$, so that $N = 5$ and the full-sample fit is $\hat{\beta} = (2, -0.8)'$, the slope contrast $a = (0, 1)'$ gives $\hat{r}_a = (-1, 0.5)'$ and $\hat{\mu}_Z(a) = 0.5$. At $k = 2$ ($p = 2/5$) all ten retained designs are non-singular ($\min_s \det \hat{Q}_{-s} = 0.24$) and every retained directional denominator is positive ($h_Z = T_{N,p}(\hat{Z}) = 0.8 < 1$). The directional statistic, read either as the ratio in (4) or as the unique root of the threshold dual $f_N(\lambda) = T_{N,p}(\hat{s} + \lambda \hat{Z}) - \lambda$, equals 1.2800 and is attained uniquely at $S^* = \{1, 4\}$, which deletes the observations with $w = 0$ and $w = 3$. The exact maximum-influence statistic over the same ten sets equals 1.1286, likewise attained at S^* , and the first-order benchmark is $T_{N,p}(\hat{s}) = 0.6400$ on that same set, so all three rules agree on the deletion and the reversal lies in the criterion value rather than the selection. The mechanism is the increment above: at S^* the retained design rotates the contrast, $\hat{Q}_{-S^*}^{-1}a = (-2.5, 1.0714)'$ is not proportional to $\hat{r}_a$ (angle 3.37°), and the attained refit falls short of the line-restricted movement by $\hat{g}_{S^*}'\{\hat{Q}_{-S^*}^{-1}a - \hat{B}_{S^*}^{-1}\hat{r}_a\} = -0.1514$, exactly the observed gap (Table S1). Proposition S5 is illustrated rather than contradicted, since*

its sufficient condition fails at S^* ; reporting $\hat{\delta}(p)$ alone overstates the largest realised refit by 13.4% at this budget, whereas the same-set check step of Section 4.3 reveals the rotation directly. The audit is finite-sample and discrete: boundary-regularity and Assumption S4(E2) fail here, and the reversal is specific to this contrast orientation and budget, the ordinary ordering holding at $k = 1$ and $k = 3$.

TABLE S1: Directional and unrestricted refits in a five-observation audit where the directional statistic exceeds the unrestricted optimum.

Statistic	Value	Exact	Deleted set
Directional refitting, $\hat{\delta}(2/5)$	1.2800	32/25	{1, 4}
Unrestricted refitting, $\hat{\delta}^{\text{ex}}(2/5)$	1.1286	79/70	{1, 4}
Attained refit on the $\hat{\delta}$ set	1.1286	79/70	{1, 4}
First-order benchmark, $T_{N,p}(\hat{s})$	0.6400	16/25	{1, 4}
Gap, $\hat{\delta}(2/5) - \hat{\delta}^{\text{ex}}(2/5)$	0.1514	53/350	—
Same-set rotation increment	−0.1514	−53/350	—
Contrast rotation at S^*	3.37°	—	—

Notes. Intercept-and-slope ordinary least squares on $N = 5$ observations with $w = (0, 1, 2, 3, 4)$ and $y = (1, 2, -1, 4, -4)$, slope contrast $a = (0, 1)'$, budget $k = 2$, so $p = 2/5$. Directional refitting is the maximum over the ten size-two deletion sets of the line-restricted movement $\sum_{i \in S} \hat{s}_i(a) / \sum_{i \in S} \hat{Z}_i(a)$ (4); unrestricted refitting is the exact maximum-influence statistic $\max_S a' \hat{Q}_S^{-1} \hat{g}_S$ (S3), that is, the largest realised decrease in the fitted slope, whose refitted value is -1.9286 against a full-sample slope of -0.8 ; the first-order benchmark is $T_{N,p}(\hat{s})$. All ten retained designs are nonsingular and all ten retained directional denominators are positive. This discrete-design sample calculation illustrates finite-sample behavior; it is neither a population nor coverage result, and does not furnish a smooth-law counterexample to Proposition S5, whose sufficient condition fails at the selected set.

Both alignments make the directional region an exact fixed point, but target agreement additionally requires global optimality (Proposition S7). Index linearity supplies both alignments; mean-zero ellipticity with independent errors also yields global agreement (Corollary S3), a through-origin benchmark rather than a result for ordinary intercept specifications. The remaining subsections develop these distinctions and the population envelope bound, which has no sampling theory here.

S3.3 The exact statistic

The exact population target is defined directly, as the best that deletion of a p -mass can do. Identify a deletion set with a measurable region $D \subset \mathbb{R}^d \times \mathbb{R}$ and write

$$\Phi(D) := a' \Sigma(D)^{-1} g(D), \quad \delta^{\text{ex}}(p) := \sup_{P(D)=p} \Phi(D), \quad (\text{S6})$$

with $\Sigma(D) = \mathbb{E}[xx' \mathbf{1}\{D^c\}]$ the design second moment over the survivors and $g(D) = \mathbb{E}[xu \mathbf{1}\{D\}]$ the score mass the deleted region carries, using the population normal equations. Taking the supremum over *all* regions of mass p matches the unrestricted scope of the sample statistic, which maximizes over all $\binom{N}{k}$ subsets. Optimality yields the fixed-point characterization below, although finding a global maximizer remains difficult.

Proposition S6 (Attainment and the fixed-point characterisation). *Under the maintained moments $\mathbb{E} \|x\|^2 < \infty$ and $\mathbb{E}[\|x\| |u|] < \infty$ of Section 2, and Assumption S4 (E1) and (E2), the supremum in (S6) is attained, and every maximiser is almost surely a threshold set of its own deletion score, hence an exact fixed point in the sense of (S7).*

Every population maximizer satisfies the following self-consistency system, where $r_{a,-p} = \Sigma(D)^{-1}a$, $b = \Sigma(D)^{-1}g(D)$, and $s = (r'_{a,-p}x)(u + x'b)$:

$$\begin{aligned} D &= \{s \geq \xi\}, & P(D) &= p, \\ \Sigma(D) &= \mathbb{E}[xx' \mathbf{1}\{D^c\}] = (1-p)Q_{-D}, & g(D) &= \mathbb{E}[xu \mathbf{1}\{D\}]. \end{aligned} \quad (\text{S7})$$

Definition S2 (Exact fixed point). *An exact fixed point is a region and its moment pair $(D, \Sigma(D), g(D))$ satisfying (S7), with attained value $\Phi(D) = a' \Sigma(D)^{-1}g(D)$. This value equals $\delta^{\text{ex}}(p)$ only if D is globally maximizing; the fixed-point equations alone do not ensure this. The retained metric $Q_{-D} = \Sigma(D)/(1-p)$ differs in general from the directional retained-design metric Q_{-A^*} of (S5).*

S3.4 The exact target and index linearity

The exact statistic generally has a different population target because $\hat{Q}_s^{-1}$ is evaluated on the complement of the active set. The retained-design second moment at the directional optimum is the metric Q_{-A^*} of (S5), the population analogue of $\tilde{Q}_s :=$

$\hat{Q}_S/(1 - p_N)$. The exact statistic instead inverts the retained moment $\Sigma(D)$ at its maximizing region D , so replacing Q by Q_{-A^*} alone need not make the targets coincide.

Full design balance, $Q_{-A^*} = Q$, is strictly stronger than needed: it implies condition [SN1](#) but neither implies condition [SN2](#) nor holds in the independent-error elliptical case under the support conditions of Lemma [S6\(ii\)](#), so we work with the pair of directional conditions throughout.

The primitive single-index condition delivering both alignments is the following.

Assumption S3 (Index linearity). *The regressors are mean-linear in the contrast index, conditionally on the error,*

$$\mathbb{E}[x \mid r'_a x, u] = \frac{a}{\mu_Z(a)} r'_a x \quad a.s. \quad (\text{S8})$$

Assumption [S3](#) restricts the design in the *single* direction r_a and delivers both alignments in one line, $\Sigma_{-A^*} r_a = B_a^* a$ and $g_{A^*} = \mu_Z(a)^{-1} \delta B_a^* a$, hence condition [SN1](#) and condition [SN2](#); it is implied by, but strictly weaker than, mean-zero ellipticity with $x \perp\!\!\!\perp u$.

The exact population criterion depends on a region through the retained-design inverse. Its first variation gives the deletion score; combined with attainment, this yields the fixed-point characterization in Definition [S2](#).

Lemma S3 (Deletion score). *Adding mass ε at (x, u) to D changes $\Phi(D)$ by $\varepsilon s_D(x, u) + O(\varepsilon^2)$, with*

$$s_D(x, u) = (r'_{a,-D} x)(u + x' b_D), \quad r_{a,-D} := \Sigma(D)^{-1} a, \quad b_D := \Sigma(D)^{-1} g(D). \quad (\text{S9})$$

If a region D satisfies the relaxed first-order variational inequality $\mathbb{E}[(\pi - \mathbf{1}\{D\})s_D] \leq 0$ for every intensity $0 \leq \pi \leq 1$ with $\mathbb{E}\pi = p$, then $\{s_D > \xi\} \subseteq D \subseteq \{s_D \geq \xi\}$ up to null sets for a cutoff ξ . Under cutoff atomlessness this gives $D = \{s_D \geq \xi\}$ almost surely.

Proof. Adding mass ε at (x, u) sends $g \mapsto g + \varepsilon x u$ and $\Sigma \mapsto \Sigma - \varepsilon x x'$, so by Sherman–Morrison $\Sigma^{-1} \mapsto \Sigma^{-1} + \varepsilon \Sigma^{-1} x x' \Sigma^{-1} + O(\varepsilon^2)$. Substituting into [\(S6\)](#) and collecting the ε term gives $\dot{\Phi} = \varepsilon[(a' \Sigma^{-1} x)u + (a' \Sigma^{-1} x)(x' \Sigma^{-1} g)] = \varepsilon(r'_{a,-D} x)(u + x' b_D)$. The stated variational inequality makes $\mathbf{1}\{D\}$ maximize the linear functional $\pi \mapsto \mathbb{E}[\pi s_D]$ over the mass- p intensity class. Upper-tail allocation gives the threshold inclusions, allowing allocation on a tied cutoff; cutoff atomlessness removes that allocation. $\square$

Proof of Proposition S6. Attainment. Relax the region to a deletion intensity $\pi : \mathbb{R}^d \times \mathbb{R} \rightarrow [0, 1]$ with $\mathbb{E}\pi = p$, and write $\Sigma(\pi) = \mathbb{E}[xx'(1 - \pi)]$, $g(\pi) = \mathbb{E}[xu\pi]$, $\Phi(\pi) = a'\Sigma(\pi)^{-1}g(\pi)$; every region D induces a feasible indicator $\pi = \mathbf{1}\{D\}$. The feasible set is weak-* compact in $L^\infty(P) = (L^1(P))^*$ by Banach–Alaoglu, since it is a weak-* closed subset of the unit ball. The moment bounds make $\Sigma(\pi)$ and $g(\pi)$ weak-* continuous, while E1 gives $\Sigma(\pi) \succeq \kappa_p I$ for every feasible intensity. Thus inversion and Φ are continuous on this compact set, so a maximizer π^* exists.

Bang-bang. Let η be any feasible direction at π^* ($\mathbb{E}\eta = 0$ and $\pi^* + t\eta$ feasible for small $t > 0$). Differentiating as in Lemma S3,

$$\left. \frac{d}{dt} \Phi(\pi^* + t\eta) \right|_{t=0} = a'\Sigma^{-1}\mathbb{E}[xx'\eta]\Sigma^{-1}g + a'\Sigma^{-1}\mathbb{E}[xu\eta] = \mathbb{E}[\eta s(x, u)],$$

where $s = (r'_{a,-p}x)(u + x'b)$, evaluated at $(\Sigma, g) = (\Sigma(\pi^*), g(\pi^*))$. Optimality of π^* therefore makes it a maximiser of the *linear* functional $\pi \mapsto \mathbb{E}[\pi s]$ over $\{0 \leq \pi \leq 1, \mathbb{E}\pi = p\}$ —a fractional knapsack, whose solutions are $\pi = \mathbf{1}\{s > \xi\}$ together with arbitrary mass on $\{s = \xi\}$, where ξ is the $(1 - p)$ -quantile of s . By E2 the score s has no atom at ξ , so $\pi^* = \mathbf{1}\{s \geq \xi\}$ almost surely. Thus the relaxed maximiser is an honest region, the relaxed and integral suprema coincide, and π^* satisfies the three equations of Definition S2: it is an exact fixed point. The same argument applies to every relaxed maximizer and therefore to every maximizing region. $\square$

Assumption S4 (Exact-target regularity).

(E1) *Restricted eigenvalue.*

$$\kappa_p := \inf \{ \lambda_{\min}(\mathbb{E}[xx'\rho]) : 0 \leq \rho \leq 1, \mathbb{E}\rho = \tau \} > 0,$$

equivalently $\inf_{\|v\|=1} \int_0^\tau F_{(v'x)^2}^{-1}(s) ds > 0$, i.e. $\sup_{\|v\|=1} P(v'x = 0) < \tau$. This is the multivariate counterpart of Assumption 3, stated for retention intensities so it applies directly to the relaxation in Proposition S6. It bounds every retained region; under a nonatomic population, the intensity infimum also equals the infimum over regions of mass p , by splitting any tied leverage level.

(E2) *Score regularity at the cutoff.* For every multiplier (Σ, g) generated by a maximizer of the relaxed problem in the proof of Proposition S6, and at the directional pair (Σ_{-A^*}, g_{A^*}) , the cutoff $\xi(\Sigma, g)$ is a continuity point of the score distribution:

$P(s_{\Sigma,g} = \xi(\Sigma, g)) = 0$. (Absolute continuity of the joint law of (x, u) is sufficient, and it delivers the condition at every nonsingular (Σ, g) at once by Lemma S5.)

(E3) **Well-separated global target.** The value functional Φ of (S6) has a unique global maximiser region D_p among all measurable deletion regions of mass p , separated in value: for every $\epsilon > 0$,

$$\Phi(D_p) - \sup \{ \Phi(D') : P(D') = p, P(D' \Delta D_p) > \epsilon \} > 0.$$

By Proposition S6 the supremum is attained and every maximiser is an exact fixed point, so this is no stronger than it needs to be: restricting the separation to fixed points would not by itself connect the population target to the empirical maximum.

To verify the equivalences in E1, put $h(v) := \inf_{0 \leq \rho \leq 1, \mathbb{E}\rho = \tau} \mathbb{E}[(v'x)^2 \rho]$ for $\|v\| = 1$. The lower-tail allocation identity, allowing fractional allocation at a tied cutoff, gives $h(v) = \int_0^\tau F_{(v'x)^2}^{-1}(s) ds$, which is positive exactly when $P(v'x = 0) < \tau$. Under the maintained $\mathbb{E}\|x\|^2 < \infty$,

$$|h(v) - h(w)| \leq \mathbb{E} |(v'x)^2 - (w'x)^2| \leq 2\mathbb{E}\|x\|^2 \|v - w\|.$$

Thus h is continuous on the compact unit sphere and, if positive everywhere, has positive minimum κ_p . Also $v \mapsto P(v'x = 0)$ is upper semicontinuous by reverse Fatou, since $v_n \rightarrow v$ gives $\limsup_n \mathbf{1}\{v'_n x = 0\} \leq \mathbf{1}\{v'x = 0\}$; its maximum is attained, so pointwise strictness is equivalent to the displayed strict supremum bound. On an atomic population, the region infimum need not equal the intensity infimum; the intensity formulation avoids assuming that a tied atom can be split.

Condition E1 keeps retained designs nonsingular, as required by the first display of Proposition S8. Condition E2 supplies the no-atom premise used in the bang-bang step of Proposition S6, Proposition S7(iii), and the strictness clause of Proposition S5; the additional boundary regularity those results require is stated at their use. Condition E3 supplies the separated global target and, as a standing premise of Proposition S7, licenses the necessity clause read by the strictness clause of Proposition S5; index linearity alone gives fixed-point alignment, not global optimality.

S3.5 Alignment, stationarity and global agreement

Corresponding deviations from Equations (SN1) and (SN2) are measured by the Euclidean *contrast-rotation* and *numerator-rotation* angles

$$\varrho_p := \angle(\Sigma_{-A^*}^{-1}a, r_a), \quad \varsigma_p := \angle(\Sigma_{-A^*}^{-1}g_{A^*}, r_a), \quad (\text{S10})$$

so that Equation (SN1) holds if and only if $\varrho_p = 0$, and Equation (SN2) holds if and only if $\varsigma_p = 0$. Both conditions hold under Assumption S3, in particular under mean-zero ellipticity with $\mathbb{E}[x \mid r'_a x, u] = \mathbb{E}[x \mid r'_a x]$ (Lemma S6); ellipticity alone is not sufficient. These angles record departures from the two alignment conditions; the results below use the alignments themselves and make no quantitative claim bounding the target gap by either angle.

Recall that boundary-regularity requires positive ambient Lebesgue density on an open ball intersecting the directional cutoff surface. For any candidate direction $v \neq 0$, let $\Phi_a^{\text{sup}}(v)$ be the directional envelope in (S11) and $\Phi_a^* := \sup_{v \neq 0} \Phi_a^{\text{sup}}(v)$. Then $\Phi_a^{\text{sup}}(r_a) = \delta(p)$ and $\delta^{\text{ex}}(p) \leq \Phi_a^*$ by Proposition S8.

Proposition S7 (Stationarity, bounds, and target agreement). *Let Assumption 6 and Assumption S4(E1)–(E3) hold.*

- (i) (Stationarity.) *If Equations (SN1) and (SN2) hold, then A^* is an exact fixed point of (S7) with*

$$r_{a,-p} = B_a^{*-1} r_a, \quad b = \mu_Z(a)^{-1} \delta r_a, \quad \xi_p = B_a^{*-1} \xi_a^*, \quad \Phi(A^*) = \delta.$$

Conversely, if the design is boundary-regular, Equations (SN1) and (SN2) are necessary for A^* to be an exact fixed point.

- (ii) (Two-sided envelope bound.) *Under Equation (SN1) alone, $\Phi(A^*) = \delta$, yielding the sandwich*

$$\delta = \Phi(A^*) \leq \delta^{\text{ex}} \leq \Phi_a^*.$$

- (iii) (Global agreement.) *Suppose Equation (SN1), boundary-regularity, and the separated-maximum condition E3 hold, and let D_p denote the unique global maximizer of $\Phi(D)$ in (S6). Then*

$$\delta^{\text{ex}} = \delta \quad \Longleftrightarrow \quad A^* = D_p \text{ up to a } P\text{-null set,}$$

and either statement implies Equation (SN2).

- (iv) (Sufficient globality via the envelope.) Under Equation (SN1), if r_a maximizes $\Phi_a^{\text{sup}}(v)$ over $\mathbb{R}^d \setminus \{0\}$, then $\Phi_a^* = \delta$, the sandwich in (ii) closes with $\delta^{\text{ex}} = \delta$, and Equation (SN2) follows automatically.

Corollary S3 (A shared target). Under Assumption S3, Assumption 6, and Assumption S4(E1), A^* is an exact fixed point with value δ . If, in addition, x is mean-zero elliptical with $u \perp x$, then r_a globally maximizes Φ_a^{sup} , the sandwich of Proposition S8 closes, A^* is the global maximizer in (S6), and $\delta^{\text{ex}} = \delta$.

S3.5.1 Proofs for the population exact-refitting bridge

Proof of Proposition S7. Write $t = r'_a x$, so $Y(a) = tu$, $Z(a) = t^2/\mu_Z(a)$ and $W_a^* = t\{u + \mu_Z(a)^{-1}\delta t\}$, and recall $A^* = \{W_a^* \geq \xi_a^*\}$.

(i) *Stationarity: sufficiency.* Assume condition SN1 and condition SN2. By condition SN1, $\Sigma_{-A^*}^{-1}a = c r_a$ for some $c > 0$; multiplying $a = c \Sigma_{-A^*} r_a$ by r'_a gives $\mu_Z(a) = a' r_a = c r'_a \Sigma_{-A^*} r_a = c \mathbb{E}[t^2 \mathbf{1}\{A^{*,c}\}] = c \mu_Z(a) B_a^*$, so $c = B_a^{*-1}$ and $r'_{a,-p} x = B_a^{*-1} t$. By condition SN2, $b = \Sigma_{-A^*}^{-1} g_{A^*} = \gamma r_a$ for a scalar γ , so $x'b = \gamma t$ and $s = (r'_{a,-p} x)(u + x'b) = B_a^{*-1} t(u + \gamma t)$. The value at A^* is

$$\Phi(A^*) = a' \Sigma_{-A^*}^{-1} g_{A^*} = B_a^{*-1} r'_a g_{A^*} = B_a^{*-1} \mathbb{E}[Y(a) \mathbf{1}\{A^*\}] = \delta,$$

the last step by the root identity $\mathcal{F}_p(W_a^*) = \delta$, which reads $\mathbb{E}[Y(a) \mathbf{1}\{A^*\}] + \delta \nu_{a,Z}^* = \delta$, i.e. $\mathbb{E}[Y(a) \mathbf{1}\{A^*\}] = \delta B_a^*$ by (6). Also $\Phi(A^*) = a'b = \gamma a' r_a = \gamma \mu_Z(a)$, so $\gamma = \delta/\mu_Z(a)$ and $s = B_a^{*-1} W_a^*$. Setting $\xi_p = B_a^{*-1} \xi_a^*$ therefore gives $\{s \geq \xi_p\} = A^*$ and $P(s \geq \xi_p) = p$, while $\Sigma_{-A^*} = \mathbb{E}[xx' \mathbf{1}\{A^{*,c}\}]$ and $g_{A^*} = \mathbb{E}[xu \mathbf{1}\{A^*\}]$ hold by definition: all three equations of Definition S2 are satisfied at A^* , with value δ .

Necessity. Suppose A^* is an exact fixed point, i.e. $A^* = \{s \geq \xi_p\}$ up to a null set, with $s = (r'_{a,-p} x)(u + x'b)$, $r_{a,-p} = \Sigma_{-A^*}^{-1} a$, $b = \Sigma_{-A^*}^{-1} g_{A^*}$, and let

$$\begin{aligned} P_1(x, u) &:= s - \xi_p = (r'_{a,-p} x)u + (r'_{a,-p} x)(x'b) - \xi_p, \\ P_2(x, u) &:= W_a^* - \xi_a^* = (r'_a x)u + \mu_Z(a)^{-1} \delta (r'_a x)^2 - \xi_a^*. \end{aligned}$$

The key structural fact is that *both are affine in u* , so it is enough to compare them section by section; no appeal to the geometry of quadric surfaces is needed.

Let $\mathcal{U}$ be the ball supplied by boundary-regularity. Since $\{P_1 \geq 0\}$ and $\{P_2 \geq 0\}$ agree up to a P -null set and the density is positive on $\mathcal{U}$, the two open sets $\{P_1 > 0 > P_2\} \cap \mathcal{U}$ and $\{P_2 > 0 > P_1\} \cap \mathcal{U}$ are null, hence empty: on $\mathcal{U}$ the two polynomials never take strictly opposite signs. Fix $(x_0, u_0) \in \mathcal{U}$ with $P_2(x_0, u_0) = 0$. Then $r'_a x_0 \neq 0$, for otherwise $W_a^*(x_0, u_0) = 0 = \xi_a^*$, contradicting $\xi_a^* > 0$ (Lemma 11); and $r_{a,-p} \neq 0$ because $\Sigma_{-A^*} > 0$ and $a \neq 0$. The root $u_2(x) = \{\xi_a^* - \mu_Z(a)^{-1} \delta(r'_a x)^2\} / (r'_a x)$ is continuous near x_0 , so its graph lies in $\mathcal{U}$ over an open neighborhood of x_0 . Because $r_{a,-p} \neq 0$, that neighborhood is not contained in the hyperplane $\{x : r'_{a,-p} x = 0\}$; choose a boundary graph point outside that hyperplane. Shrinking to a ball around this point, we may assume $r'_a x \neq 0$ and $r'_{a,-p} x \neq 0$ throughout, with $u_2(x)$ in the ball's u -section for every x in an open set $O \subset \mathbb{R}^d$.

Fix $x \in O$. Both $P_1(x, \cdot)$ and $P_2(x, \cdot)$ are affine with nonzero slopes $r'_{a,-p} x$ and $r'_a x$, and $P_2(x, \cdot)$ changes sign at $u_2(x)$, an interior point of the section. Were $P_1(x, \cdot)$ to vanish elsewhere, or to have the opposite slope sign, the two would take strictly opposite signs on one side of $u_2(x)$. Hence the roots and the slope signs agree:

$$\frac{\xi_p - (r'_{a,-p} x)(x'b)}{r'_{a,-p} x} = \frac{\xi_a^* - \mu_Z(a)^{-1} \delta(r'_a x)^2}{r'_a x}, \quad \text{sign}(r'_{a,-p} x) = \text{sign}(r'_a x).$$

Cross-multiplying gives

$$(r'_a x)[\xi_p - (r'_{a,-p} x)(x'b)] = (r'_{a,-p} x)[\xi_a^* - \mu_Z(a)^{-1} \delta(r'_a x)^2]$$

for every $x \in O$; both sides are polynomials in x agreeing on an open set, hence identically on $\mathbb{R}^d$. Matching homogeneous parts, the degree-three part reads $(r'_a x)(r'_{a,-p} x)[x'b - \mu_Z(a)^{-1} \delta r'_a x] \equiv 0$, and since $\mathbb{R}[x]$ is an integral domain and $(r'_a x)(r'_{a,-p} x) \neq 0$,

$$b = \mu_Z(a)^{-1} \delta r_a,$$

which is condition [SN2](#). The degree-one part reads $\xi_p(r'_a x) = \xi_a^*(r'_{a,-p} x)$ for all x , i.e. $\xi_p r_a = \xi_a^* r_{a,-p}$; as $\xi_a^* > 0$ and $r_{a,-p} \neq 0$ this forces $\xi_p \neq 0$ and

$$r_{a,-p} = \kappa r_a, \quad \kappa = \xi_p / \xi_a^*,$$

which is condition [SN1](#), with $\kappa > 0$ by the slope-sign agreement.

(ii) *The one-sided bound.* The computation $c = B_a^{*-1}$ and $\Phi(A^*) = \delta$ in the sufficiency

part used condition [SN1](#) only. Since A^* is a feasible region of mass p , $\delta^{\text{ex}} \geq \Phi(A^*) = \delta$. The upper bound $\delta^{\text{ex}} \leq \Phi_a^*$ is the first display of Proposition [S8](#), which uses [E1](#) alone.

(iii) *Agreement.* Suppose $\delta^{\text{ex}} = \delta$. By (ii), $\Phi(A^*) = \delta = \delta^{\text{ex}} = \sup\{\Phi(D) : P(D) = p\}$, so A^* attains the supremum, and the uniqueness in [E3](#) identifies it with D_p up to a null set. Conversely, if $A^* = D_p$ up to a null set then $\delta^{\text{ex}} = \Phi(D_p) = \Phi(A^*) = \delta$ by (ii). For the implication, a maximiser is an exact fixed point by Proposition [S6](#), so under boundary-regularity the necessity clause of (i) applies to A^* and yields $b = \mu_Z(a)^{-1} \delta r_a$, which is condition [SN2](#). The two alignments do not settle the converse implication: conditions [SN1](#) and [SN2](#) place A^* at a stationary point of the exact problem and leave the global comparison open, as Remark [S3](#) states.

(iv) *A sufficient global condition.* If r_a maximises Φ_a^{sup} then $\Phi_a^* = \Phi_a^{\text{sup}}(r_a) = \Phi_a(r_a) = \delta$ by the identification above, and the sandwich of (ii) gives $\delta \leq \delta^{\text{ex}} \leq \delta$. Condition [SN2](#) is not needed as a hypothesis: $a'r_a = \mu_Z(a) > 0$, so on the neighbourhood $\{a'v > 0\}$ of r_a the truncation in ([S11](#)) can only raise the ratio, whence $\Phi_a(v) \leq \Phi_a^{\text{sup}}(v) \leq \Phi_a^{\text{sup}}(r_a) = \Phi_a(r_a)$ for every v there; r_a is therefore a local maximiser of Φ_a , which is differentiable at r_a by Lemma [5\(C5\)](#), so $g_a = \nabla_r \Phi_a(r_a) = 0$, and given condition [SN1](#) that is condition [SN2](#) by Lemma [S4](#). $\square$

Remark S3 (Stationarity is not optimality). *Conditions [SN1](#) and [SN2](#) place the directional set at a fixed point of the exact rule, but a fixed point need not be the global maximiser. This is why the agreement result states its globality condition separately.*

Proof of Proposition [S5](#). The computation $c = B_a^{*-1}$ and $\Phi(A^*) = \delta$ in the sufficiency part above used condition [SN1](#) only, not condition [SN2](#). (More generally, the same-set identity of Section [3.3](#) at $D = A^*$ reads $\Phi(A^*) = \delta + \Delta_{\text{rot}}$ with $\Delta_{\text{rot}} := g_{A^*}'\{\Sigma_{-A^*}^{-1}a - B_a^{*-1}r_a\}$, so the displayed ordering holds whenever this single scalar increment is nonnegative, condition [SN1](#) being the zero-increment special case.) Since A^* is a feasible region of mass p , $\delta^{\text{ex}} = \sup_{P(D)=p} \Phi(D) \geq \Phi(A^*) = \delta$. For strictness, suppose condition [SN2](#) fails and the design is boundary-regular. If A^* were equal to $\{s \geq \xi\}$ up to null sets, A^* would be an exact fixed point and the necessity clause of Proposition [S7](#), available under [E1](#)–[E3](#), would force $b = \mu_Z(a)^{-1} \delta r_a$, i.e. condition [SN2](#); so $A^* \neq \{s \geq \xi\}$ up to null sets. (Pointwise inequality of s and W_a^* is *not* enough on its own: two scores generate the same super-level sets whenever one is a strictly increasing transform of the other, and it is the polynomial argument in that necessity clause, and hence boundary-regularity, that rules this out.) Since s has no atom at its $(1-p)$ -quantile ξ by [E2](#), it follows that $P(A^* \cap \{s < \xi\}) > 0$ and $P(A^{*,c} \cap \{s > \xi\}) > 0$. Transfer

mass $\epsilon > 0$ from the first event to the second: by the first variation of Lemma S3, Φ increases by $\epsilon \mathbb{E}[s \mid A^{*,c} \cap \{s > \xi\}] - \epsilon \mathbb{E}[s \mid A^* \cap \{s < \xi\}] + O(\epsilon^2) > 0$ for small ϵ . This is a feasible intensity perturbation; by Proposition S6 under E1–E2, the relaxed and region suprema agree, so the improved relaxed value implies $\delta^{\text{ex}} > \Phi(A^*) = \delta$. Finally, in the strict-gap case just established, where condition SN2 fails and the design is boundary-regular, one has $\delta^{\text{ex}} > \delta$, so $\sqrt{N}(\hat{\delta} - \delta^{\text{ex}}) \rightarrow -\infty$ in probability and any interval of width $O_p(N^{-1/2})$ centered at $\hat{\delta}$ has coverage of δ^{ex} tending to zero. $\square$

The envelope bound compares deletion regions, so the scalarisation it needs is a supremum over regions rather than the root of (27). For every $v \neq 0$ define the truncated directional supremum

$$\Phi_a^{\text{sup}}(v) := (a'v) \max \left\{ 0, \sup \left\{ \mathbb{E}[(v'x)u \mathbf{1}\{D\}] / \mathbb{E}[(v'x)^2 \mathbf{1}\{D^c\}] : P(D) = p \right\} \right\}. \quad (\text{S11})$$

By E1 every competitor has denominator at least $\kappa_p \|v\|^2 > 0$, while Cauchy–Schwarz and the standing model of Section 2, which gives $\mathbb{E} \|x\| |u| \leq (\mathbb{E} \|x\|^2 \mathbb{E} u^2)^{1/2} < \infty$, bound its numerator by $\|v\| \mathbb{E} \|x\| |u|$; hence $\Phi_a^{\text{sup}}(v)$ is well defined for every $v \neq 0$ under E1 alone, with no quantile or root condition charged on the direction-perturbed score, and, being of degree zero, obeys $\Phi_a^{\text{sup}}(v) \leq \|a\| \mathbb{E} \|x\| |u| / \kappa_p$ at every $v \neq 0$. It is homogeneous of degree zero under positive rescaling of v , non-negative on $\{a'v > 0\}$ and non-positive on $\{a'v < 0\}$. At any direction r for which the pair $((r'x)u, (r'x)^2)$ satisfies Assumption 2–Assumption 5, the Rockafellar–Uryasev dual (20) rewrites the criterion as $f(\lambda; r) = \sup_{P(D)=p} \{\mathbb{E}[(r'x)u \mathbf{1}\{D\}] - \lambda \mathbb{E}[(r'x)^2 \mathbf{1}\{D^c\}]\}$, and at a non-negative root the two-sided bound $\kappa_p \|r\|^2 \leq \mathbb{E}[(r'x)^2 \mathbf{1}\{D^c\}] \leq r'Qr$ forces that root to equal the inner supremum of (S11); a root is automatically non-negative, so the truncation does not bind there. At r_a these are the hypotheses of Proposition 4, so $\Phi_a^{\text{sup}}(r_a) = (a'r_a)\delta^{\text{sc}} = \delta = \Phi_a(r_a)$.

Proposition S8 (Directional envelope). *Take an arbitrary direction $v \neq 0$ and let Φ_a^{sup} be as in (S11), so that $\Phi_a^{\text{sup}}(r_a) = \Phi_a(r_a) = \delta$. Under Assumption S4(E1), for every region D with $P(D) = p$,*

$$\Phi(D) \leq \Phi_a^{\text{sup}}(\Sigma(D)^{-1}a), \quad \text{and therefore} \quad \delta^{\text{ex}} \leq \sup_{v \neq 0} \Phi_a^{\text{sup}}(v).$$

Under condition [SN1](#) the two bounds sandwich the exact target,

$$\delta = \Phi_a^{\text{sup}}(r_a) \leq \delta^{\text{ex}} \leq \sup_{v \neq 0} \Phi_a^{\text{sup}}(v) = \Phi_a^*.$$

Proof of Proposition [S8](#). Fix D with $P(D) = p$ and set $v := \Sigma(D)^{-1}a$, which is well defined and nonzero because $\Sigma(D) \succeq \kappa_p I$ by [E1](#). Then

$$\begin{aligned} \Phi(D) &= a' \Sigma(D)^{-1} g(D) = v' g(D) = \mathbb{E}[(v'x) u \mathbf{1}\{D\}], \\ a'v &= v' \Sigma(D) v = \mathbb{E}[(v'x)^2 \mathbf{1}\{D^c\}] > 0. \end{aligned}$$

Hence

$$\Phi(D) = (a'v) \cdot \frac{\mathbb{E}[(v'x)u \mathbf{1}\{D\}]}{\mathbb{E}[(v'x)^2 \mathbf{1}\{D^c\}]} \leq \Phi_a^{\text{sup}}(v),$$

because D is one competitor in the inner supremum of [\(S11\)](#) and $a'v > 0$, so the truncation $\max\{0, \cdot\}$ dominates the ratio. No root identity is invoked here, and therefore no quantile or tail condition is charged on the direction-perturbed score. Taking the supremum over D gives $\delta^{\text{ex}} \leq \sup_{v \neq 0} \Phi_a^{\text{sup}}(v) = \Phi_a^*$; the lower bound is Proposition [S5](#), together with $\Phi_a^{\text{sup}}(r_a) = \Phi_a(r_a) = \delta$ from [\(S11\)](#). $\square$

The same alignment governs the design layer of Theorem [1](#) and the fixed-point property of Proposition [S7](#)(i).

Lemma S4 (The design gradient is the numerator misalignment). Write $\delta^{\text{sc}} = \delta/\mu_Z(a)$ for the scalarised target as stated in [\(28\)](#). Without any assumption on the design,

$$m_a = a - \Sigma_{-A^*} r_a, \quad \text{hence} \quad g_a = \delta^{\text{sc}} a + B_a^{*-1} g_{A^*} - 2\delta^{\text{sc}} B_a^{*-1} \Sigma_{-A^*} r_a. \quad (\text{S12})$$

Under condition [SN1](#) these collapse to

$$m_a = \nu_{a,Z}^* a \quad \text{and} \quad g_a = B_a^{*-1} g_{A^*} - \delta^{\text{sc}} a,$$

so that, given condition [SN1](#), $g_a = 0$ if and only if condition [SN2](#) holds.

Proof of Lemma [S4](#). By definition $m_a = \mathbb{E}[x(r'_a x) \mathbf{1}\{A^*\}] = \mathbb{E}[xx' \mathbf{1}\{A^*\}] r_a = (Q - \Sigma_{-A^*}) r_a = a - \Sigma_{-A^*} r_a$, using $Q r_a = a$; substituting $m_a - a = -\Sigma_{-A^*} r_a$ into [\(8\)](#) gives the second identity in [\(S12\)](#). Under condition [SN1](#) the computation in the sufficiency part of the proof of Proposition [S7](#) gives $\Sigma_{-A^*} r_a = B_a^* a$, whence $m_a = (1 - B_a^*) a = \nu_{a,Z}^* a$

by (6) and

$$g_a = \delta^{\text{sc}} a + B_a^{*-1} g_{A^*} - 2\delta^{\text{sc}} a = B_a^{*-1} g_{A^*} - \delta^{\text{sc}} a.$$

Finally, the root identity $\mathcal{T}_p(W_a^*) = \delta$ gives $r'_a g_{A^*} = \mathbb{E}[Y(a)\mathbf{1}\{A^*\}] = \delta B_a^*$, so if $g_{A^*} = c a$ then $c \mu_Z(a) = \delta B_a^*$, i.e. $c = \delta^{\text{sc}} B_a^*$ and $g_a = 0$; conversely $g_a = 0$ forces $g_{A^*} = \delta^{\text{sc}} B_a^* a \propto a$. Under condition SN1, $g_{A^*} \propto a$ is exactly condition SN2. $\square$

Proof of Corollary S3. Under (S8) the active set A^* is a function of (t, u) alone, so $\Sigma_{-A^*} r_a = \mathbb{E}[x t \mathbf{1}\{A^{*,c}\}] = \mathbb{E}[\mathbb{E}[x | t, u] t \mathbf{1}\{A^{*,c}\}] = \mu_Z(a)^{-1} a \mathbb{E}[t^2 \mathbf{1}\{A^{*,c}\}] = B_a^* a$ and $g_{A^*} = \mathbb{E}[x u \mathbf{1}\{A^*\}] = \mu_Z(a)^{-1} a \mathbb{E}[t u \mathbf{1}\{A^*\}] = \mu_Z(a)^{-1} \delta B_a^* a$, which are condition SN1 and condition SN2 because $\Sigma_{-A^*} > 0$ by E1. The stationarity computation of Proposition S7(i) is direct here: $r_{a,-p} = B_a^{*-1} r_a$, $b = \mu_Z(a)^{-1} \delta r_a$ and $s = B_a^{*-1} W_a^*$, so the cutoff $B_a^{*-1} \xi_a^*$ recovers A^* and $\Phi(A^*) = \delta$, without invoking E2 or E3. For the elliptical statement write $x = Q^{1/2} \epsilon$ with ϵ spherical and $u \perp \epsilon$, and put $w = Q^{1/2} v$, so that $v'x = w'\epsilon$ and $v'Qv = \|w\|^2$. For every unit vector w the joint law of $(w'\epsilon, u)$ is the same, so the inner supremum in (S11) depends on w only through $\|w\|$; it is homogeneous of degree -1 under positive rescaling of v , and it is unchanged when v is replaced by $-v$ (the pair $((v'x)u, (v'x)^2)$ is then equal in law to the original, since $v'x$ and $-v'x$ are equidistributed and $u \perp x$), so it equals $\phi_p / \|w\|$ for a constant $\phi_p > 0$ — positive by the regular-root identification at r_a , where $\delta > 0$, and rotational invariance — and the truncation of (S11) is inactive. Therefore

$$\Phi_a^{\text{sup}}(v) = (a'v) \frac{\phi_p}{\|w\|} = \phi_p \frac{(Q^{-1/2}a)'w}{\|w\|} \leq \phi_p \|Q^{-1/2}a\| = \phi_p \sqrt{\mu_Z(a)},$$

with equality if and only if $w \propto Q^{-1/2}a$, i.e. $v \propto Q^{-1}a = r_a$. Hence $\Phi_a^* = \sup_{v \neq 0} \Phi_a^{\text{sup}}(v) = \Phi_a^{\text{sup}}(r_a) = \delta$, and the sandwich of Proposition S8 forces $\delta^{\text{ex}} = \delta$; since $\Phi(A^*) = \delta = \delta^{\text{ex}}$, A^* attains the supremum in (S6). $\square$

Lemma S5 (Atomlessness of the deletion score). *For a multiplier $\varpi = (r_{a,-p}, b)$, write $s_{\varpi}(x, u) = (r'_{a,-p}x)(u + x'b)$. If the law of (x, u) is absolutely continuous on $\mathbb{R}^{d+1}$, then $P(s_{\varpi} = w) = 0$ for every $w \in \mathbb{R}$ and every multiplier with $r_{a,-p} \neq 0$; in particular E2 holds at every score parameter the bridge evaluates.*

Proof. Fix ϖ with $r_{a,-p} \neq 0$ and $w \in \mathbb{R}$. The polynomial $(x, u) \mapsto (r'_{a,-p}x)(u + x'b) - w$ has $r'_{a,-p}x$ as the coefficient of u , and $r'_{a,-p}x$ is not the zero polynomial, so the polynomial is not identically zero. Its zero set is a proper algebraic hypersurface and

has Lebesgue measure zero, and absolute continuity of the law of (x, u) turns this into $P(s_{\overline{w}} = w) = 0$. $\square$

Lemma S6 (No rotation under elliptical designs). *Let Assumption 6 and Assumption S4(E1) hold, so that $B_a^* > 0$ and $\Sigma_{-A^*} > 0$. Suppose x is elliptically symmetric about the origin with $\mathbb{E}[xx'] = Q \succ 0$ and $\mathbb{E}[x \mid r'_a x, u] = \mathbb{E}[x \mid r'_a x]$ (in particular when $u \perp x$). Write $t := r'_a x$ and $\Sigma_{\perp} := Q - \mu_Z(a)^{-1}aa'$, so that $\Sigma_{\perp}r_a = 0$ and $Q = \Sigma_{\perp} + \mu_Z(a)^{-1}aa'$. Then:*

(i) $\mathbb{E}[x \mid t] = \mu_Z(a)^{-1}at$ and $\mathbb{E}[xx' \mid t] = w(t)\Sigma_{\perp} + \mu_Z(a)^{-2}t^2aa'$ for a scalar $w(\cdot) \geq 0$, with $w \equiv 1$ in the Gaussian case.

(ii) If additionally $u \perp x$, then $Q_{-A^*} = Q + c_{\perp}\Sigma_{\perp} + c_a aa'$ with $(Q_{-A^*} - Q)r_a = c_a\mu_Z(a)a$ and

$$c_a = \frac{1}{\mu_Z(a)} \left(\frac{B_a^*}{1-p} - 1 \right) \leq 0, \quad (\text{S13})$$

where $c_{\perp} = 0$ in the Gaussian case (so $\text{rank}(Q_{-A^*} - Q) \leq 1$). If moreover $|t|$ is non-degenerate and u has a strictly positive density on $\mathbb{R}$, then $B_a^* < 1 - p$ and $c_a < 0$, so $Q_{-A^*} \neq Q$ while Equation (SN1) ($Q_{-A^*}^{-1}a \propto Q^{-1}a$) holds identically.

Proof. (i) Elliptical symmetry gives $\mathbb{E}[x \mid t] = \text{Cov}(x, t)\text{Var}(t)^{-1}t = \mu_Z(a)^{-1}at$ and decomposes the conditional second moment into its r_a component and a scalar multiple $w(t)\Sigma_{\perp}$ on the orthogonal complement; Gaussianity gives $w \equiv 1$.

(ii) If $u \perp x$, put $\pi(\tau) = P(A^* \mid |t| = \tau)$. Integrating (i) gives the aa' coefficient of Q_{-A^*} as $B_a^*/\{\mu_Z(a)(1-p)\}$ and hence c_a in (S13); the $\Sigma_{\perp}$ coefficient is $(1-p)^{-1}\mathbb{E}[w(t)(1 - \mathbf{1}\{A^*\})]$. It equals one when $w \equiv 1$, so $c_{\perp} = 0$ in the Gaussian case.

For $t > 0$, selection is $u \geq \varphi(|t|)$ and for $t < 0$ it is $u \leq -\varphi(|t|)$, where $\varphi(\tau) = \{\xi_a^* - \mu_Z(a)^{-1}\delta\tau^2\}/\tau$ is strictly decreasing because $\xi_a^* > 0$ (Lemma 11) and $\delta \geq 0$. Independence implies that $\pi(\tau)$ is nondecreasing, strictly increasing when u has a positive density. Therefore $\mathbb{E}[t^2\mathbf{1}\{A^*\}]/\mu_Z(a) = \mathbb{E}[Z(a)\mathbf{1}\{A^*\}] \geq p$, strictly if $|t|$ is nondegenerate and the density is positive. Hence $B_a^* \leq 1 - p$ and $c_a \leq 0$, with strict inequality under the stated strictness conditions. $\square$

S4 Monte Carlo evidence

This section of the Online Supplement presents supporting exhibits for the four simulation designs in Section 5.

S4.1 Wald coverage across the budget

Figure S1 traces empirical Wald coverage across deletion shares $p \in [0.01, 0.20]$, complementing the pointwise calibration in Table 1.




FIGURE S1: Empirical Wald coverage across deletion budgets $k = \lfloor pN \rfloor$ for $p \in [0.01, 0.20]$ (Corollary 3). The $N = 125, p = 0.01$ ($k = 1$) cell illustrates small- k undercoverage where the fixed-fraction asymptotic approximation deteriorates. Shading marks nominal coverage plus or minus twice the maximum Monte Carlo standard error across cells, not a simultaneous uncertainty band.

S4.2 The fragility curve and the breakdown fraction

Table S2 reports simultaneous grid coverage and breakdown-fraction Wald intervals on the eleven-point grid $0.02, 0.03, \dots, 0.12$ at $N = 1,500$, over 20,000 replications. The

functional simulations evaluate the general feasible influence score at every design, so the fitted design correction is retained even where its population counterpart vanishes (Remark 4); the residual-corrected score is kept as an ablation. The paired coverage difference, general feasible minus residual-corrected, is $+0.0003$ in the Gaussian design and -0.0001 in the elliptical- t design (both inside one paired standard error), and $+0.0129$ with a paired standard error of 0.0008 in the skewed design. Band coverage is below nominal in every cell, and the shortfall is concentrated at the smallest deletion budgets: marginal coverage generally rises across the grid from about 0.93 – 0.95 at $p = 0.02$ to about 0.96 – 0.97 at $p = 0.12$. Repeating the design at $N = 3,000$ raises band coverage to 0.934 , 0.937 , 0.922 and 0.924 on the Gaussian, elliptical- t , skewed and heteroskedastic laws, so the deficit shrinks with sample size but remains visible at $N = 3,000$. Both conclusion thresholds are fixed oracle-reference levels, chosen at $p^\dagger = 0.06$ and 0.10 ; this experiment does not calibrate the paired estimated-margin band of Proposition 2 or continuum coverage on $\mathcal{P}$. Crossing estimates linearly interpolate between grid points rather than search all integer budgets; the reported coverage concerns this numerical procedure, with no asserted interpolation-error rate. The printed lengths are mean crossing-Wald interval widths, not mean curve-band widths.

The earlier single-sample illustration in Figure S2 is distinct from both this 20,000-replication table and the paired estimated-margin campaign below. Its saved coordinates come from the prior fixed-margin study; it illustrates the curve and band rather than supplying a new coverage result. Across the four fixed-margin designs, the envelope slope $\partial_p \delta(p) = \xi_a^*(p)/B_a^*(p)$ closely matches finite differences of the oracle curve.

S4.3 Estimated-margin paired-gap campaign

The separate campaign underlying Table 2 uses $x = (1, T, C)'$ with cell probabilities $(.30, .35, .10, .25)$, conditional means $(0, .2, .9, 2.3)$ and scales $(.5, 1.2, .8, 1.5)$ in the order $(T, C) = (0, 0), (0, 1), (1, 0), (1, 1)$. Errors are independent standard normal or t_3 draws; the contrast is $a_0 = (0, 1, 0)'$ and the conclusion threshold is $b = 1$. The sign and margin $|a_0' \hat{\beta} - b|$ are estimated per sample, with all 1,600 attempts per cell retained and scored against the fixed population orientation. Conditional-tail integrals supply population curves and crossings, rather than an empirical oracle draw; numerical root brackets are floating checks, not outward arithmetic certificates.

TABLE S2: Inference on the fragility curve: simultaneous-band coverage and grid-interpolated crossing Wald intervals.

Design	Curve band (simult. cov.)		Crossing	Breakdown fraction $\hat{p}^\dagger$		
	Multiplier	Pointwise		$p^\dagger$	Coverage	Length
Gaussian	0.929	0.882	interior	0.060	0.944	0.0155
			near-bound	0.100	0.948	0.0231
Elliptical-t	0.930	0.883	interior	0.060	0.949	0.0205
			near-bound	0.100	0.950	0.0292
Skewed	0.912	0.872	interior	0.060	0.935	0.0379
			near-bound	0.100	0.978	0.0450
Heteroskedastic	0.914	0.870	interior	0.060	0.938	0.0209
			near-bound	0.100	0.944	0.0297

Notes. 20,000 replications at $N = 1,500$, nominal $1 - \alpha = 0.95$, all designs evaluated with the general feasible influence score; the residual-corrected score is retained as an ablation and reported in the online supplement. *Curve band*: simultaneous coverage of the whole fragility curve over the grid $\mathcal{D}$ by the multiplier-bootstrap band vs. the naive pointwise band. *Breakdown fraction*: the simulation linearly interpolates the fitted grid curve to estimate the crossing, then uses the plug-in Wald standard error from Theorem 3, at an interior and a near-boundary reference level c . This numerical procedure does not search all integer budgets; its interpolation error is not bounded here. Coverage is of $p^\dagger$; Length is the mean Wald-interval width. The binomial Monte Carlo standard error of a reported coverage $\hat{r}$ is $\{\hat{r}(1 - \hat{r})/20,000\}^{1/2}$, about 0.002 at nominal coverage. Reference targets are computed on an oracle draw of 300,000 observations; their approximation error is not included in these Monte Carlo standard errors.




FIGURE S2: Earlier fixed-margin directional-curve illustration, from the study whose coverage experiment used 300 replications, not the later Table S2 or paired-gap campaign. A single-sample estimate and multiplier band on the eleven-point grid are compared with the numerical population reference; dotted lines mark the fixed conclusion threshold and crossing. These saved coordinates are not continuum or estimated-margin calibration.

Every integer count intersecting $[0.02, 0.20]$ uses the full paired influence and the same 999 multipliers. Because the fitted band is constant within a count cell while the population curve is monotone, checking both ends, including one-sided limits, evaluates coverage over the continuous interval without interpolating the empirical crossing. Grid-state coverage exceeds this continuum coverage by 0.0013–0.0088; the latter is the main-text column. Table S3 retains censoring and available-only summaries separately from the unconditional coverage denominator in the main table. At $N \geq 800$, conditional crossing-Wald coverage ranges from 0.937 to 0.956; small-sample heavy-tail undercoverage is not tuned away.

S5 Additional Sample-Audit Detail

The following calculations extend the deterministic institutions–income audit in Section 6.1, using the original AJR Table 2, column 6 specification with latitude and continent controls. The authors’ replication command requests heteroskedasticity-robust standard errors; this exercise reproduces its OLS point estimates and supplies

TABLE S3: Paired-gap campaign: censoring, availability and conditional widths.

Law	N	Censored L/R	Wald available	Wald cond.	Inv. width	Wald width	Sign errors
Normal	200	408/1141	1499	0.960	0.141	0.163	0
Normal	800	0/162	1599	0.937	0.089	0.083	0
Normal	3200	0/0	1600	0.956	0.045	0.042	0
t_3	200	1175/638	1462	0.966	0.143	0.158	26
t_3	800	455/3	1589	0.948	0.086	0.081	0
t_3	3200	7/0	1600	0.951	0.045	0.041	0

Notes. Counts are out of 1,600 attempts; left and right inversion censoring can overlap. Wald conditional coverage and width use available intervals only; inversion width includes every fit. Sign errors are scored against the fixed population orientation, not a sample-redefined target. There were no fit refusals or unavailable orientations; no attempt was replaced.

no standard errors or significance claims. The directional statistic freezes the full-sample adjusted contrast direction, while each attained OLS refit re-estimates all six coefficients on its retained sample.

TABLE S4: Adjusted institutions–income audit: movements, refits and retained leverage.

k	$100k/N$	$\hat{\delta}$	AMIP	$\hat{\beta}_{1,-S_{\text{dir}}}$	$\hat{\beta}_{1,-S_{\text{AMIP}}}$	$\hat{B}_a^*$
3	4.69	0.0833	0.0683	0.3104	0.3104	0.819
6	9.38	0.1416	0.0954	0.2493	0.2547	0.624
12	18.75	0.2301	0.1246	0.1621	0.1866	0.524
24	37.50	0.3697	0.1424	0.0278	0.0951	0.325
26	40.62	0.3913	0.1432	0.0024	0.0799	0.310
27	42.19	0.4033	0.1434	−0.0023	0.0774	0.271
32	50.00	0.4829	0.1435	−0.0442	0.0447	0.177

Notes. AJR Table 2, column 6; $N = 64$, observed slope 0.4013. AMIP is $T_{N,p}(\hat{s})$; both refit columns re-estimate all six coefficients on their retained samples. Every selected retained design has rank six over $k = 0, \dots, 32$. The directional denominator is $\hat{B}_a^* = N^{-1} \sum_{i \in S^*} \hat{Z}_i$; the all-set diagnostic $h_Z = T_{N,p}(\hat{Z})$ is at most 0.9017 on this domain.

Table S5 separates the two ranking rules from their own objective values. At $k = 27$, the sets share 19 of 27 countries, yet the directional criterion evaluated on the AMIP set is 0.3167, versus its own optimum 0.4033; AMIP evaluated on the directional set is 0.1092, versus 0.1434 on its own set. Thus substantial overlap does not make the rules interchangeable near the directional crossing.

At $k = 3$, both rules select Singapore, Gabon and Bangladesh, giving the same attained refit slope of 0.3104, compared with the baseline 0.4013. At the directional crossing, deleting its 27 selected countries gives an actual OLS slope of −0.0023; the

TABLE S5: Cross-evaluation of directional and AMIP criteria with latitude and continent controls.

k	Shared rows	Directional on		AMIP on	
		S_{dir}	S_{AMIP}	S_{AMIP}	S_{dir}
3	3/3	0.0833	0.0833	0.0683	0.0683
12	10/12	0.2301	0.1974	0.1246	0.1206
27	19/27	0.4033	0.3167	0.1434	0.1092

Notes. The directional value on S is $N^{-1} \sum_{i \in S} \hat{\delta}_i / (1 - N^{-1} \sum_{i \in S} \hat{Z}_i)$; AMIP is $N^{-1} \sum_{i \in S} \hat{\delta}_i$. Each criterion is evaluated on both fixed selected sets at the same budget; shared rows count common deletions.

AMIP-selected deletion leaves it at 0.0774. The attained movement is approximately 0.4036, versus the directional 0.4033: a same-set rotation increment of about +0.0003 using the unrounded results as defined in (S4). The directional-selected refit is still positive at $k = 26$ (0.0024), and the AMIP-selected refits remain positive throughout the reported path. These are attained sample movements from two computable selection rules, not maximizations over unrestricted refits.

At the crossing, the directional-selected deletion leaves 27.07% of directional leverage. The rule ranks the fitted scores $\hat{\delta}_i + \hat{\delta}(p)\hat{Z}_i$, combining residual contributions with directional leverage, and recomputes the selected set at each budget. The cross-evaluation isolates criterion-specific ranking losses on these two sets, but does not decompose their actual-refit coefficient gap into selection and refitting components.

Finite-share adjustment changes the sample curve substantially, but the crossing requires deleting 42.2% of the countries. The 27-country set is a feasible actual reversal, not a certified smallest reversing set, and the small negative refit is not a significance claim. This deterministic OLS audit is not a robustness analysis of the IV estimand or population inference, which requires the conditions in Section 4.

S6 Conditional Illustration

S6.1 Cognitive-training completers

The source is the official supplementary participant CSV of Hardy et al. (2015). The observed frame comprises 5,045 post-test completers (2,667 Lumosity and 2,378 crosswords) of 9,919 randomized participants. No Grand-Index change value is missing within this release, but completion is post-randomization; observed-case

completeness does not identify an all-randomized ITT. Trial registration NCT02367898 was first posted on February 20, 2015, after primary completion in April 2014. The unadjusted assignment contrast was self-declared for this calculation, not justified by prospective registration of that analysis. The published ANCOVA is reconstructible using the released pre-test index on the crossover-removed frame (coefficient 2.7677, $t(4712) = 8.7324$); absence of that covariate is not the reason for the unadjusted illustration.

The working model maintains iid sampling from a completer superpopulation and the LP cutoff conditions. Released scores are rounded; all observed change values being distinct does not verify positive population cutoff density. Exact reconstruction of percentile-rank norm tables is unavailable. The directional curve is not the unrestricted mean-contrast refit curve: actual selected-set refits are separate checks, and the directional and unrestricted-optimal arm allocations differ at the six checked budgets. The calculation supplies neither randomization-based coverage nor a missingness correction.

The displayed bands use integer counts 101, ..., 1009 and 99,999 common Gaussian multipliers with seed 20261019; the 95% and 99% critical values are 2.293468954 and 2.899102585, respectively. The zero-budget anchor is excluded from band inference. At the observed frozen contrast 3.09716777, the directional movement is 3.09701265 at $k = 326$ and 3.10480894 at $k = 327$, so the first integer sign crossing is 327/5045, not 326/5045. Monotonicity-based grid inversion of the 95% curve band reports [311/5045, 342/5045] without endpoint censoring. This holds the reported threshold fixed: it is not the paired estimated-margin procedure of Proposition 2, and no calibrated interval for the fixed population sign-margin radius is claimed. The crossover-removed ablation has $n = 4,715$, contrast 3.14867537 and bracket [289/4715, 322/4715]. The earlier 999-draw primary had 95% critical value 2.3604202182 and bracket [310/5045, 343/5045]; its separate multiplier-precision ablation used 9,999 draws and critical value 2.2657193861. Neither ablation validates empirical sampling coverage. The supplementary-file redistribution license is not established by the article's CC BY license; the computed figure does not redistribute the participant CSV.

S7 Population Geometry at Matched Signal-to-Noise

Standardized signal-to-noise does not determine the directional population curve or its zero-margin crossing, even within the finite-covariate class verified in the Online Supplement. Table S6 compares four deliberately calibrated specifications, not samples or a coverage experiment.

There are four strata s and a binary treatment T , with $x = (1, T, \mathbf{1}\{s = 2\}, \mathbf{1}\{s = 3\}, \mathbf{1}\{s = 4\})'$ and contrast $a = (0, 1, 0, 0, 0)'$. Write $w_s = P(s)$ and $q_s = P(T = 1 \mid s)$, so the eight cell probabilities are $w_s(1 - q_s)$ and $w_s q_s$. Outcomes obey $y = b_s + \kappa g_s T + \sigma_s e$, where e is independent of (s, T) , has mean zero and variance one, and $b = (.40, -.25, .70, -.35)$. For D0, D2 and D3, set $w = (.30, .30, .20, .20)$, $q = (.5, .3, .7, .4)$ and $g = (.6, -.1, .9, -.4)$. D0 uses standard Gaussian e and $\sigma = (1, 1, 1, 1)$; D2 changes only e to a logistic law with scale $\sqrt{3}/\pi$; D3 uses Gaussian e and $\sigma = (.55, .8, 1.25, 1.85)$. D4 uses Gaussian e , $\sigma = (1, 1, 1, 1)$, $w = (.43, .43, .07, .07)$, $q = (.52, .48, .97, .03)$ and $g = (.35, -.05, 4.5, -4.5)$. In each specification, β is the population linear projection, $u = y - x'\beta$, $\vartheta = a'\beta$ and $V = \mathbb{E}[(a'Q^{-1}xu)^2]$. Choose the positive effect scale κ to solve $\vartheta/\sqrt{V/1000} = 3$: approximately .810110894316 for D0 and D2, .918218422370 for D3 and 2.410023043738 for D4. The reference size 1000 expresses a population calibration, not a realized t -statistic.

All cell probabilities are positive and the support spans $\mathbb{R}^5$. Writing $H = \sum_s w_s q_s (1 - q_s) > 0$, the cell index is $x'Q^{-1}a = (T - q_s)/H \neq 0$. Gaussian and logistic conditional densities are continuous and strictly positive on $\mathbb{R}$, with finite second moments; hence all four laws satisfy that finite-covariate verification, including its positive retained-denominator guarantee. The entries solve the population modified-score cutoff and ratio equations of Section A, up to the reported numerical optimization tolerance.

D0 and D2 share the covariate law, conditional means and conditional variances, matching ϑ and V separately within 10^{-12} , yet D2 has a larger curve at 0.02 and 0.10, a smaller curve at 0.20 and an earlier zero-margin crossing. D3 and D4 match only standardized signal-to-noise; D4 has the largest movement at .10 but the latest crossing among these specifications, whose margins differ. D4 jointly changes stratum weights, treatment shares and effect patterns and recalibrates the scale; it illustrates joint geometry rather than isolating a leverage mechanism. These calibrated directional-population examples establish non-sufficiency of standardized signal-to-noise, not unrestricted-refit ordering, sampling coverage or empirical prevalence.

TABLE S6: Directional population geometry at matched standardized signal-to-noise.

Design	ϑ	$\sqrt{V}$	$\delta(0.02)$	$\delta(0.10)$	$\delta(0.20)$	$p_{\text{dir}}^\dagger$
D0 Gaussian baseline	0.2036	2.1461	0.1211	0.4523	0.7987	0.0375
D2 logistic twin	0.2036	2.1461	0.1293	0.4540	0.7839	0.0356
D3 heteroskedastic	0.2308	2.4325	0.1534	0.4985	0.8282	0.0342
D4 leverage-concentrated	0.3548	3.7396	0.2023	0.5523	0.9293	0.0523

Notes: $\vartheta = a\beta$ and $V = \mathbb{E}[(aQ^{-1}xu)^2]$. The zero-margin crossing is $p_{\text{dir}}^\dagger = \inf\{p : \delta(p) \geq \vartheta\}$. All rows are deliberately calibrated to $\vartheta/\sqrt{V/1000} = 3$ (absolute tolerance 10^{-12}); 1000 is a reference size, not a realized sample or t -statistic. D2 also matches D0 in ϑ and V separately within 10^{-12} ; D3 and D4 match only the ratio. Values are numerical population optimizations, not arithmetic enclosures or sampling uncertainty.

References

Hardy, J. L. et al. (2015). “Enhancing Cognitive Abilities with Comprehensive Training: A Large, Online, Randomized, Active-Controlled Trial”. In: *PLOS ONE* 10.9, e0134467. doi: [10.1371/journal.pone.0134467](https://doi.org/10.1371/journal.pone.0134467).