When Is It Safe to Borrow Across Treatments? Target-Aware Learning for Combinatorial Interventions

Anonymous authors
Paper under double-blind review

Abstract

Combinatorial interventions create an exponential treatment space from a modest number of components. The resulting sparsity motivates cross-treatment borrowing, but smoothing between neighboring bundles can attenuate interactions of interest. We introduce Treatment-Space Graph Selection (TSGS), which separates an immutable graph of feasible interventions from a target-indexed graph used only for statistical borrowing. The outcome surface is decomposed into prespecified structured variation and its orthogonal complement; only the latter is graph-regularized, and the graph and penalty are selected by an independent augmented inverse-propensity-weighted score for the desired contrast family. Our theory gives (i) an exact finite-sample contrast-risk decomposition, (ii) a sharp characterization of uniform dominance over all zero-sum causal contrasts by Loewner ordering on the contrast subspace, (iii) a Boolean spectral theorem explaining why edge effects and interaction effects weight the same smoothing error differently, and (iv) a finite-candidate oracle inequality. In projected graph-sequence experiments across 36 independently simulated support–roughness regimes, TSGS beats no borrowing in every cell, whereas indiscriminate full-graph smoothing is significantly harmful in five rough regimes; edge and synergy targets select different rules in 65.8% of repetitions. In the randomized ACTG 175 trial, target-selected borrowing reduces held-out contrast score by 0.7–1.7% relative to no borrowing across five clinical targets. The central conclusion is that treatment geometry specifies feasible interventions, but borrowing rules should be selected for the causal contrasts of interest.

1 Introduction

Many consequential interventions are bundles: a cancer regimen combines drugs, a public-health program combines services, and a digital policy activates several modules. With m binary components there are 2^m possible bundles, so direct support per treatment collapses long before the total sample size looks small. This is a central problem in causal learning with complex treatments: the treatment is an entire intervention bundle, but the scientifically reusable structure lies in components, interactions, and local modifications of that bundle (Zou et al., 2020; Qian et al., 2021; Tanimoto et al., 2021; Harada & Kashima, 2021).

Existing approaches attack this sparsity through treatment representations, shared outcome heads, augmentation, or fusion of empirically similar arms (Kaddour et al., 2021; Parbhoo et al., 2021; Ma et al., 2022; 2023; Zhu et al., 2025). These approaches use cross-treatment information sharing to address sparse support, but they do not determine whether the same sharing rule is valid for every causal question. Prediction loss averages errors over treatment nodes. By contrast, a scientific target is a signed contrast: adding one component, measuring a pairwise synergy, or comparing a combination with monotherapy. The same estimation error can cancel in one contrast and compound in another. Consequently, outcome-predictive similarity and target-relevant similarity are different statistical objects (Figure 1).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

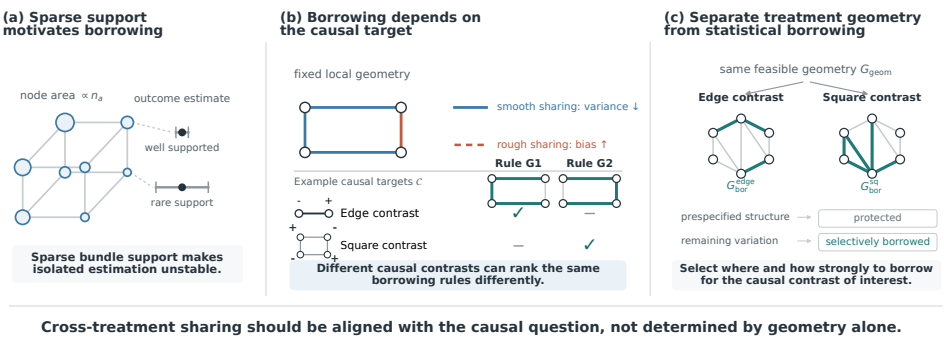

Figure 1: **Why borrowing must be target-aware.** **(a)** Sparse support motivates borrowing. **(b)** Borrowing trades variance against bias, so different contrasts can prefer different rules. **(c)** TSGS protects prespecified structure and selects target-aligned borrowing within fixed feasible geometry.

Before any borrowing decision, we distinguish domain-feasible bundles from those with marginal assignment support and from those satisfying positive propensity in the target population: not every combination is meaningful, observed, or overlap-safe.

ACTG 175 illustrates the distinction between treatment comparisons and statistical borrowing. The randomized trial compared zidovudine (AZT), didanosine (ddI), and the combinations AZT+ddI and AZT+ddC (Hammer et al., 1996). These four arms form a natural one-drug-addition graph, yet its edges encode clinically distinct questions: the benefit of adding ddI to AZT, adding ddC to AZT, or adding AZT to ddI. Pooling every adjacent arm can stabilize a rare contrast, but it can also shrink a regimen-specific interaction toward the wrong neighboring effect. Removing the edge would be equally problematic if that edge defines the scientific comparison. This exposes the research gap: *the graph that defines feasible treatment relations must not be conflated with the graph that determines statistical borrowing*.

We resolve this gap with *Treatment-Space Graph Selection* (TSGS). A fixed feasible graph G_{geom} records which interventions and local changes are scientifically meaningful. A second graph $G_{\mathrm{bor}} \subseteq G_{\mathrm{geom}}$ is a tunable statistical operator: it specifies where the learner may exchange information, never which effect is defined. TSGS represents the outcome surface as a prespecified structured component plus a projected complementary signal. Graph regularization acts only on this complementary signal. Finally, an independent augmented inverse-propensity-weighted (AIPW) validation score selects $(G_{\mathrm{bor}}, \lambda_G)$ for the prespecified contrast matrix C. This construction contains no-graph and full-graph learning as special cases, accommodates constrained treatment spaces, and remains meaningful when a flexible neural backbone replaces the restricted sequence model used for theory.

Adaptive and doubly robust treatment-fusion methods learn global arm partitions for policy optimization (Ma et al., 2022; 2023; Zhu et al., 2025); TSGS instead retains every node and allows the borrowing relation to change with the reported contrast. Its validation identity specializes established pseudo-outcome model selection (Alaa & Van Der Schaar, 2019; Saito & Yasui, 2020; Curth & Van Der Schaar, 2023; Lan & Syrgkanis, 2024). The contribution is their coupling with fixed scientific geometry and protected structured effects; Appendix B gives the full comparison.

Candidate risks need not order uniformly across causal contrasts. We show that uniform dominance over all zero-sum contrasts is equivalent to Loewner ordering of conditional error second moments after projection onto the zero-sum contrast subspace; an indefinite projected difference yields a rank-one reversal. On the Boolean cube, order-r face contrasts diagonalize in the Walsh basis: edges weight degree k by $2k$, whereas square synergies weight it by $4\binom{k}{2}$.

The evidence follows the theory. Across 1,440 support–roughness runs, target-selected borrowing improves on NoGraph in all 36 cells, while strong full-graph smoothing is significantly harmful in five rough cells; edge and square targets choose different rules in 65.8% of repetitions. Under fixed information per split, sharing gains grow from 64 to 1,024 bundles. In ACTG 175, target-selected borrowing improves held-out contrast score over NoGraph in 70–89% of splits across the local additions.

Contributions. Our contributions are fourfold.

(1) **Separating treatment geometry from statistical borrowing.** We separate feasible intervention geometry from target-dependent borrowing geometry and define a single framework that includes learning without graph regularization, fixed-graph smoothing, and target-selected borrowing.
(2) **A structured learner with contrast-aware graph selection.** We introduce a structured/complementary decomposition, graph-selective regularization, and direct AIPW model selection for arbitrary prespecified contrast families.
(3) **Exact contrast-risk analysis and target-dependent ordering.** We prove identification and non-identification boundaries, exact finite-sample contrast risk, projected-Loewner characterization of target-uniform dominance, the Boolean face-contrast spectrum, and a finite-candidate selection guarantee. Core algebraic claims are independently machine-checked in Lean 4.
(4) **Empirical tests of target-dependent borrowing.** We execute neural benchmarks, controlled graph-sequence experiments, support and roughness stress tests, target-switching experiments, constrained-space and propensity ablations, and a four-arm randomized clinical-trial analysis.

2 Problem Setup and Causal Targets

We first define the causal targets independently of any borrowing rule. Let $\mathcal{H} = \{0,1\}^m$ be the combinatorial treatment space and let $\mathcal{A}_f \subseteq \mathcal{H}$ denote the domain-feasible bundles. We observe i.i.d. data $O_i = (X_i, A_i, Y_i)$, $i = 1, \dots, n$, with $X_i \in \mathcal{X} \subseteq \mathbb{R}^d$, $A_i \in \mathcal{A}_f$, and $Y_i \in \mathbb{R}$. Let P_X be the marginal law of X_i, which is also the target-population covariate law throughout. For $a \in \mathcal{A}_f$, let $Y(a)$ denote the potential outcome and $\mu(x,a) = \mathbb{E}\{Y(a) \mid X = x\}$. Define the marginal assignment support

$$\mathcal{A}_s = \{a \in \mathcal{A}_f : \mathbb{P}(A = a) > 0\},$$

the bundle propensity $e_a(x) = \mathbb{P}(A = a \mid X = x)$, and, for a prespecified $\eta > 0$,

$$\mathcal{A}_{s,\eta} = \{a \in \mathcal{A}_s : e_a(X) \geq \eta \ \ P_X\text{-a.s.}\}.$$

Thus $\mathcal{A}_{s,\eta} \subseteq \mathcal{A}_s \subseteq \mathcal{A}_f \subseteq \mathcal{H}$: combinatorial possibility, domain feasibility, assignment support, and overlap are distinct objects.

Assumption 2.1 (Bundle-level causal identification)**.** *We assume: (i) consistency and no interference, so $A_i = a$ implies $Y_i = Y_i(a)$; (ii) bundle-level unconfoundedness, $\{Y(a) : a \in \mathcal{A}_f\} \perp A \mid X$; and (iii) identified validation and causal reporting use only contrasts supported on $\mathcal{A}_{s,\eta}$. We also assume $\mathbb{E}|Y(a)| < \infty$ for identification and invoke finite second moments only for squared-risk results.*

Fixed feasible geometry. The graph $G_{\text{geom}} = (\mathcal{A}_f, E_{\text{geom}})$ contains scientifically valid one-component transitions: $(a,b) \in E_{\text{geom}}$ only if $a, b \in \mathcal{A}_f$ and $\|a - b\|_0 = 1$, after removing domain-invalid toggles when needed. Geometry is fixed by intervention feasibility and semantics, not by outcomes or empirical support. It organizes local comparisons but does not define general endpoint effects: for any feasible a, b, $\mu(x,b) - \mu(x,a)$ is defined directly from the common potential-outcome surface, and path sums are only telescoping identities.

Contextual edge and face effects. For a valid activation edge $(a, a + \mathbf{e}_j) \in E_{\text{geom}}$, where $\mathbf{e}_j$ is the jth standard basis vector, define

$$\tau^{(1)}_{a,j}(x) = \Delta_j \mu(x,a) := \mu(x, a + \mathbf{e}_j) - \mu(x,a),$$

and, when the four required vertices are feasible, define the square-face synergy

$$\tau^{(2)}_{a,jk}(x) = \Delta_j \Delta_k \mu(x,a) := \mu(x, a + \mathbf{e}_j + \mathbf{e}_k) - \mu(x, a + \mathbf{e}_j) - \mu(x, a + \mathbf{e}_k) + \mu(x,a).$$

All such contrasts hold x fixed; heterogeneity arises because the within-x contrast varies across covariate values. For a fixed component j, the valid activation edges $\mathcal{E}_j = \{(a, a + \mathbf{e}_j) \in E_{\text{geom}} : a_j = 0\}$ therefore define a contextual effect family rather than a single context-free "effect of treatment j." Ordering the $K = |\mathcal{A}_f|$ feasible bundles, write $\mu(x) = (\mu(x,a) : a \in \mathcal{A}_f)$ and let $\mathbf{1}_K$ be the all-ones vector. A q_C-row causal contrast family is $C \in \mathbb{R}^{q_C \times K}$ with $C\mathbf{1}_K = 0$, targeting $\tau_C(x) = C\mu(x)$. General linear functionals need not be zero-sum, but the causal contrast theory in this paper is explicitly restricted to this contrast subspace. Let C_{edge} stack all valid Hamming-1 activation-edge contrasts and let C_{sq} stack all feasible square-face synergies; these are our two focal target families. Higher-order Boolean-face contrasts are defined analogously in Appendix G.

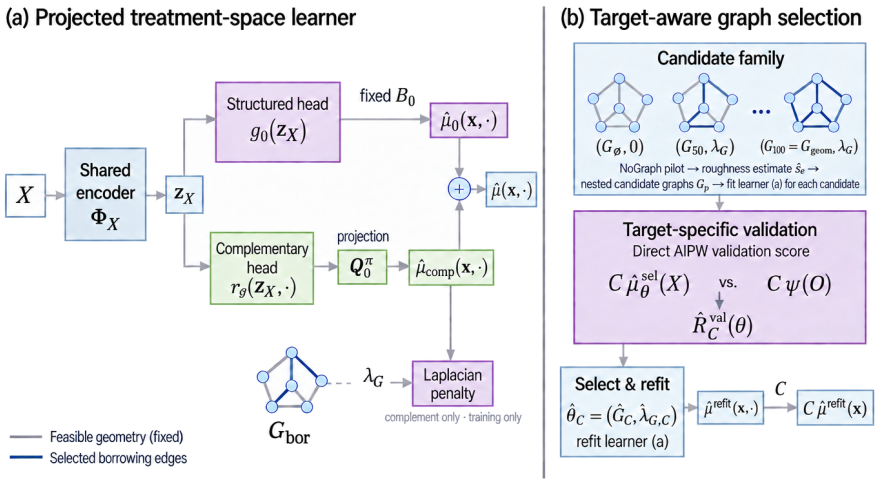

Figure 2: **Treatment-Space Graph Selection.** **(a)** Shared structured and projected-complement branches; Laplacian regularization acts only on the complement. **(b)** A training-only pilot builds nested graph candidates; independent target-specific AIPW validation selects the graph and penalty for final refit.

Theorem 2.1 (Identification of supported contrasts)**.** *Under Assumption 2.1, $\mu(x,a) = \mathbb{E}(Y \mid X = x, A = a)$ for every $a \in \mathcal{A}_{s,\eta}$, for P_X-almost every x in the target population. Hence a finite-difference contrast is identified whenever every treatment vertex entering it satisfies the required overlap condition.*

Theorem 2.2 (Non-identification outside overlap)**.** *If there exist $a_0 \in \mathcal{A}_f$ and a set B with $P_X(B) > 0$ such that $e_{a_0}(x) = 0$ for all $x \in B$, then $\mu(x, a_0)$ on B is not identified from the observed law under consistency and bundle-level unconfoundedness alone. Every contrast having a nonzero coefficient on that vertex is likewise unidentified there without additional structural assumptions.*

Proofs are in Appendix C. We validate and report only on $\mathcal{A}_{s,\eta}$. Strict positivity $e_a(X) > 0$ P_X-a.s. suffices for P_X-a.e. point identification; the stronger uniform floor stabilizes AIPW validation and finite-sample control. Bundles in $\mathcal{A}_s \setminus \mathcal{A}_{s,\eta}$ are supported but outside this region, whereas $\mathcal{A}_f \setminus \mathcal{A}_s$ are unsupported structural extrapolations; $\mathcal{A}_f$ alone defines the scientific geometry.

We next construct a shared learner that exploits recurring bundle structure while allowing local borrowing to remain selective.

3 Graph-Selective Treatment-Space Learner

TSGS separates treatment-space representation from borrowing: it decomposes the outcome surface into a prespecified structured span and its projected complement, regularizes only the latter over $G_{\text{bor}} \subseteq G_{\text{geom}}$, and selects the borrowing rule $(G_{\text{bor}}, \lambda_G)$ for a prespecified target family C by direct AIPW contrast validation (Figure 2).

3.1 Structured and Complementary Components

TSGS begins with a treatment dictionary fixed before fitting. Let $\phi_0(a) \in \mathbb{R}^{p_0}$ be a prespecified structured feature map and let $B_0 \in \mathbb{R}^{K \times p_0}$ stack its rows $\phi_0(a)^\top$ over the feasible nodes. On a Boolean bundle space our default choice exposes baseline, singleton, and pairwise structure:

$$\phi_{\le 2}(a) = \big(1, a_1, \ldots, a_m, \{a_j a_k\}_{j<k}\big)^\top, \qquad p_{\le 2} = 1 + m + \binom{m}{2},$$

with $\phi_0 = \phi_{\le 2}$ and $B_0 = B_{\le 2}$. Let $\Phi_X : \mathcal{X} \to \mathbb{R}^{d_z}$ be the shared covariate encoder and $g_0 : \mathbb{R}^{d_z} \to \mathbb{R}^{p_0}$ the structured head. Writing $z_x = \Phi_X(x)$ gives

$$\widehat{\mu}_0(x,a) = \phi_0(a)^\top g_0(z_x).$$

Degree two is the smallest explicit Boolean order that can represent pairwise synergy: every degree-one Boolean polynomial has zero square-face second difference, whereas $\Delta_j\Delta_k(a_ja_k) = 1$. This is a prespecified structured span, not a truncation assumption on the full outcome surface. The construction and theory below are expressed for B_0: the cube experiments use $B_0 = B_{\le 2}$, whereas ACTG 175 uses the shared treatment intercept $B_0 = \mathbf{1}$, leaving a rank-three contrast complement.

A flexible residual branch could otherwise relearn the same structured effects and make the decomposition functionally ambiguous. To keep its role distinct, enumerate $\mathcal{A}_f = \{a^{(1)}, \dots, a^{(K)}\}$ and, for a strictly positive reference probability mass function π on $\mathcal{A}_f$, let $W_\pi = \mathrm{diag}\{\pi(a^{(1)}), \dots, \pi(a^{(K)})\}$ and define

$$P_0^\pi = B_0(B_0^\mathsf{T} W_\pi B_0)^{-1} B_0^\mathsf{T} W_\pi, \qquad Q_0^\pi = I_K - P_0^\pi.$$

The reference measure π fixes the weighted notion of low-order versus complementary variation. We use the uniform measure $\pi(a) = 1/K$; other applications may prespecify a different reference distribution.

Proposition 3.1 (Unique low-order/complementary decomposition)**.** *Under* $\mathrm{rank}(B_0^\mathsf{T} W_\pi B_0) = p_0$, *every* $f \in \mathbb{R}^K$ *has a unique decomposition*

$$f = P_0^\pi f + Q_0^\pi f, \qquad B_0^\mathsf{T} W_\pi Q_0^\pi f = 0.$$

If B_0 is rank-deficient on a constrained feasible space, replacing it by any full-column-rank basis of $\mathrm{col}(B_0)$ yields the same functional projection. The weighted-projection proof and full-cube Walsh interpretation are in Appendix D. For the population surface vector $\mu(x)$, define $\mu_0(x) := P_0^\pi \mu(x)$ and $\mu_{\mathrm{comp}}(x) := Q_0^\pi \mu(x)$. The graph objective and sequence theory require only that the complement basis span $\mathrm{Null}(B_0^\mathsf{T} W_\pi)$; they do not require a Boolean cube or degree two.

3.2 Geometry, borrowing, and the shared learner

We regularize only the projected complementary component over the borrowing graph. For an oriented geometry edge $e = (a, b)$, define $\Delta_e f := f(b) - f(a)$ and write

$$\tau_e(x) = \Delta_e \mu(x) = \mu(x, b) - \mu(x, a), \qquad \delta_e(x) = \Delta_e \mu_{\mathrm{comp}}(x) = \mu_{\mathrm{comp}}(x, b) - \mu_{\mathrm{comp}}(x, a),$$

so $\tau_e = \Delta_e \mu_0 + \delta_e$, and define

$$\mathfrak{r}_e = \{\mathbb{E}_X[\delta_e(X)^2]\}^{1/2}.$$

For any fixed undirected weighted borrowing graph $G = (\mathcal{A}_f, E_G, w)$ with $w_e > 0$ for every $e \in E_G$, let L_G denote its weighted Laplacian; Appendix C.5 fixes the edge-counting convention. Then

$$\mathbb{E}_X[\mu_{\mathrm{comp}}(X)^\mathsf{T} L_G \mu_{\mathrm{comp}}(X)] = \sum_{e \in E_G} w_e \mathfrak{r}_e^2.$$

Thus the Laplacian penalizes complementary variation along borrowed relations, not the total causal effect along an edge; a large total effect alone does not imply unsafe borrowing. The identity is verified in Appendix C.5.

Candidate borrowing graphs. Rather than search arbitrary subgraphs, we construct a finite nested candidate family from a training-only NoGraph pilot. For each geometry edge $e = (a, b)$, estimate

$$\widehat{s}_e = \frac{1}{n} \sum_{i=1}^{n} \{\widehat{\mu}_{\mathrm{comp}}(X_i, b) - \widehat{\mu}_{\mathrm{comp}}(X_i, a)\}^2.$$

Here $\widehat{\mu}_{\mathrm{comp}}$ is the training-only NoGraph pilot output. Ranking edges by $\widehat{s}_e$ gives nested G_ρ, where $\rho \in \mathcal{R}$ is the retained-edge fraction. With $G_\emptyset$ and a finite penalty grid these form Θ; write $\theta = (G, \lambda_G) \in \Theta$. Pilot, graph construction, checkpoints, and the outcome nuisance use training data only; outer validation only selects, and the selected graph is frozen for final refit. The exact grid is in Appendix F.

For any candidate $G_{\mathrm{bor}} = (\mathcal{A}_f, E_{\mathrm{bor}}, w)$ with $E_{\mathrm{bor}} \subseteq E_{\mathrm{geom}}$, let $r_\vartheta(x) \in \mathbb{R}^K$ be a raw feasible-node signal parameterized by neural parameters ϑ, and define the learned complementary output $\widehat{\mu}_{\mathrm{comp}}(x) := Q_0^\pi r_\vartheta(x)$. The shared outcome surface is

$$\boxed{\widehat{\mu}(x, a) = \phi_0(a)^\mathsf{T} g_0(z_x) + [Q_0^\pi r_\vartheta(x)]_a.}$$

The branches may share the covariate trunk; output-space projection enforces functional separation. For nonempty E_{bor}, joint factual training uses the edge-normalized graph penalty

$$\mathcal{L} = \frac{1}{n}\sum_i \{Y_i - \widehat{\mu}(X_i, A_i)\}^2 + \lambda_G \frac{1}{n}\sum_i \frac{1}{|E_{\mathrm{bor}}|} \sum_{(a,b)\in E_{\mathrm{bor}}} w_{ab}\{\widehat{\mu}_{\mathrm{comp}}(X_i, a) - \widehat{\mu}_{\mathrm{comp}}(X_i, b)\}^2.$$

AdamW weight decay is reported separately in Appendix F and is not part of the graph-sequence equivalence below. The graph G determines *where* borrowing is imposed, while λ_G determines *how strongly* it is imposed. Its relation to the sequence-model penalty λ is given in Proposition 4.2. For $G_\emptyset$, we set $\lambda_G = 0$ and the graph penalty is identically zero by convention. The forward map and contrast family are graph-independent: removing an edge from E_{bor} changes statistical borrowing only, not the definition of the causal target. At inference, the same x is queried at different bundle nodes to form the target contrasts. Throughout the learner, the node set is the domain-feasible space $\mathcal{A}_f$; overlap restricts only which contrast coordinates can be used for identified validation and causal reporting, while feasible unsupported coordinates remain structural extrapolations.

Selection protocol. Training data determine the NoGraph pilot, nested candidate graphs, candidate fits, checkpoints, and outcome nuisance; an independent validation split scores the fixed candidates by $\widehat{R}_C^{\mathrm{val}}$. We retain the selected graph, λ_G, and epoch, then refit that rule once on training plus validation data. Full pseudocode is in Appendix E.9.

4 Contrast-Aware Graph Selection Theory

Selecting the borrowing graph matters because graph regularization can reduce variance by sharing information across bundle nodes while simultaneously shrinking differences along borrowed relations. Rough borrowed relations can therefore introduce bias in the causal contrasts of interest. We isolate this trade-off in a projected graph-sequence model, an exact restricted instance of the output-space borrowing objective used by the neural learner. Throughout this section, G abbreviates a candidate borrowing graph and λ denotes the sequence-model penalty.

4.1 Borrowing induces a target-specific bias–variance trade-off

Let

$$\mathcal{H}^{\pi}_{\mathrm{comp}} = \{h \in \mathbb{R}^K : B_0^{\mathsf{T}} W_\pi h = 0\},$$

write $d_{\mathrm{comp}} = \dim(\mathcal{H}^{\pi}_{\mathrm{comp}})$, and let $U \in \mathbb{R}^{K\times d_{\mathrm{comp}}}$ have Euclidean-orthonormal columns spanning this subspace. Let $n_a = \sum_{i=1}^n \mathbf{1}(A_i = a)$, $D_n = \mathrm{diag}(n_a)$, $h_0 \in \mathcal{H}^{\pi}_{\mathrm{comp}}$, let $Z \in \mathbb{R}^K$ be arbitrary at zero-count nodes, and define $\xi = D_n(Z - h_0)$. Conditional on D_n, assume

$$\mathbb{E}(\xi \mid D_n) = 0, \qquad \mathrm{Cov}(\xi \mid D_n) = \sigma^2 D_n.$$

At a positive-count node this is $Z_a = h_{0,a} + \varepsilon_a$ with variance σ^2/n_a, together with conditional uncorrelatedness across distinct node means. For an undirected positive-weight borrowing graph G and penalty $\lambda \geq 0$, set

$$A_{G,\lambda} = U^{\mathsf{T}}(D_n + \lambda L_G)U, \qquad S_{G,\lambda} = U A_{G,\lambda}^{-1} U^{\mathsf{T}} D_n.$$

Whenever $A_{G,\lambda} \succ 0$, the graph-sequence estimator is

$$\widehat{h}_{G,\lambda} = S_{G,\lambda} Z.$$

The corresponding admissibility condition and derivation are given in Appendix E.1.

For a fixed contrast matrix C, define

$$\mathcal{R}_C(G, \lambda; h_0 \mid D_n) = \mathbb{E}\left[\left\| C(\widehat{h}_{G,\lambda} - h_0)\right\|_2^2 \,\middle|\, D_n\right].$$

Theorem 4.1 (Exact conditional contrast risk). *For any admissible* (G, λ),

$$\mathcal{R}_C(G, \lambda; h_0 \mid D_n) = \underbrace{\|C(S_{G,\lambda} - I_K)h_0\|_2^2}_{\mathrm{Bias}_C^2} + \underbrace{\sigma^2 \operatorname{tr}\left[CUA_{G,\lambda}^{-1}U^\mathsf{T}D_nUA_{G,\lambda}^{-1}U^\mathsf{T}C^\mathsf{T}\right]}_{\mathrm{Variance}_C},$$

with

$$C(S_{G,\lambda} - I_K)h_0 = -\lambda CUA_{G,\lambda}^{-1}U^\mathsf{T}L_Gh_0.$$

Borrowing therefore changes two quantities at once: it reduces variance by pooling node means, while borrowed-edge roughness L_Gh_0 contributes to bias through the projected smoother and the target contrast operator. The contrast operator filters both terms. The next theorem turns this observation into a complete characterization.

Theorem 4.2 (Uniform causal-contrast ordering is projected Loewner ordering). *For a fitted candidate* θ *with complementary estimator* $\widehat{h}_\theta$, *let* $e_\theta = \widehat{h}_\theta - h_0$ *and* $M_\theta = \mathbb{E}(e_\theta e_\theta^\mathsf{T} \mid D_n)$, *with finite second moment, and define* $\mathcal{R}_C(\theta \mid D_n) = \mathbb{E}\{\|Ce_\theta\|_2^2 \mid D_n\}$. *Then*

$$\mathcal{R}_C(\theta \mid D_n) = \operatorname{tr}(CM_\theta C^\mathsf{T}).$$

Let $P_\perp = I_K - K^{-1}\mathbf{1}_K\mathbf{1}_K^\mathsf{T}$ *be the orthogonal projector onto the causal-contrast subspace* $\mathbf{1}_K^\perp$. *For any two candidates* θ_1, θ_2, *the following are equivalent:*

(i) $\mathcal{R}_C(\theta_1 \mid D_n) \le \mathcal{R}_C(\theta_2 \mid D_n)$ *for every finite contrast matrix* C *satisfying* $C\mathbf{1}_K = 0$;
(ii) the inequality holds for every rank-one contrast $C = v^\mathsf{T}$ *with* $v^\mathsf{T}\mathbf{1}_K = 0$;
(iii) $P_\perp(M_{\theta_2} - M_{\theta_1})P_\perp \succeq 0$.

If the projected difference in (iii) has both a positive and a negative direction in $\mathbf{1}_K^\perp$, *two zero-sum rank-one contrasts rank the candidates strictly in opposite orders.*

A candidate is uniformly optimal over all zero-sum contrasts if and only if its projected error second-moment matrix is no larger in Loewner order than that of every other candidate. An indefinite projected difference makes the ordering of that pair contrast-dependent. The proof is in Appendix E.3.

Theorem 4.3 (Spectrum of Boolean face contrasts). *For* $[m] := \{1, \dots, m\}$, $1 \le r \le m$, *and* $S \subseteq [m]$, *let* $u_S(a) = 2^{-m/2}\prod_{j\in S}(2a_j - 1)$ *and let* C_r *stack one canonical* $0 \to 1$ *orientation of every order-*r *Boolean-face difference. Then*

$$C_r^\mathsf{T}C_ru_S = 2^r\binom{|S|}{r}u_S, \qquad \|C_ru_S\|_2^2 = 2^r\binom{|S|}{r}.$$

In particular, edge effects weight Walsh degree $|S|$ *by* $2|S|$, *whereas square synergies weight it by* $4\binom{|S|}{2}$.

Corollary 4.1 (Exact target-dependent spectral risk). *With* $B_0 = B_{\le 2}$ *and uniform* π, *if* $D_n = I_K$, G *is the full Hamming graph,* $h_0 = \sum_{|S|>2}\alpha_Su_S$, *and the node noise is isotropic with covariance* $\sigma^2 I_K$, *then full-graph smoothing has degree-*k *shrinkage* $s_k(\lambda) = (1 + 2\lambda k)^{-1}$ *and*

$$\mathcal{R}_{C_r}(\lambda) = \sum_{|S|>2} 2^r\binom{|S|}{r}\left[\{1 - s_{|S|}(\lambda)\}^2\alpha_S^2 + \sigma^2 s_{|S|}(\lambda)^2\right].$$

The same smoother therefore exposes edge and synergy targets to different spectral bias–variance weights. A finite $m = 4$ reversal and the proofs of Theorem 4.3 and Corollary 4.1 are in Appendices E.7 and E.4.

Proposition 4.1 (Oracle shrinkage for one borrowed relation). *Let* $n_a, n_b > 0$, *let* d *be the true difference across one relation, and let* $v = \sigma^2(n_a^{-1} + n_b^{-1})$ *be the variance of its unsmoothed empirical difference. If borrowing multiplies that difference by* s, *then* $R(s) = d^2(1 - s)^2 + vs^2$. *If* $d^2 + v > 0$,

$$s^\star = \frac{d^2}{d^2 + v}, \qquad R(s) - R(s^\star) = (d^2 + v)(s - s^\star)^2.$$

If $d = v = 0$, *then* $R(s) \equiv 0$ *and every* s *is optimal.*

Thus rare support (v large) calls for stronger borrowing, while a genuine jump (d^2 large) calls for weaker borrowing. This exact local law motivates the support–roughness experiment in Section 5.1.

Connection to the neural objective.

Proposition 4.2 (Exact restricted connection). *Fix a structured predictor and define* $R_i = Y_i - \widetilde{\mu}_0(X_i, A_i)$ *and* $Z_a = n_a^{-1} \sum_{i:A_i=a} R_i$ *at positive-count nodes. Restricting to* x*-invariant* $h \in \mathcal{H}^{\pi}_{\mathrm{comp}}$*, for nonempty* G *the objective* $n^{-1} \sum_i (R_i - h_{A_i})^2 + \lambda_G |E_G|^{-1} h^{\mathsf{T}} L_G h$ *has the same minimizers as the graph-sequence objective with* $\lambda = n\lambda_G/|E_G|$*; for* $G_\emptyset$*, both penalties vanish.*

The proof is in Appendix E.6. This identity isolates the graph-shrinkage mechanism in the restricted model; λ_G is selected directly in the nonlinear learner and is not rescaled at deployment.

4.2 From target-specific validation to deployment

Target-specific validation without counterfactual supervision. The risk in Theorem 4.1 depends on C, but $C\mu(X)$ is not observed. TSGS therefore fixes the finite candidate family and outcome nuisance using training data, then uses an independent validation sample only to compare target contrasts. Let $\mathcal{F}_{\mathrm{tr}}$ collect these training-stage objects; every candidate and nuisance fit below is $\mathcal{F}_{\mathrm{tr}}$-measurable. Let $\widetilde{\mu}$ denote the training-only outcome nuisance, distinct from the fixed low-order predictor used only in Proposition 4.2. For validation, let $\mathcal{V}_C = \{a \in \mathcal{A}_f : C_{:,a} \neq 0\} \subseteq \mathcal{A}_{s,\eta}$. For $O = (X, A, Y)$ and $a \in \mathcal{V}_C$, define

$$\psi_a(O) = \widetilde{\mu}(X, a) + \frac{\mathbf{1}(A = a)}{e_a(X)} \{Y - \widetilde{\mu}(X, a)\}, \qquad \psi(O) = (\psi_a(O) : a \in \mathcal{A}_f).$$

Set $\psi_a(O) = 0$ for $a \notin \mathcal{V}_C$. With the true propensity, $\mathbb{E}\{C\psi(O) \mid X, \mathcal{F}_{\mathrm{tr}}\} = C\mu(X)$. This is the vector-valued, multi-arm specialization of the pseudo-outcome risk decomposition used in causal model selection (Saito & Yasui, 2020). With validation-independent estimated nuisances, exact unbiasedness is replaced by a product-of-nuisance-errors remainder; Appendix E.8 gives the expression.

Write $\widehat{\mu}^{\mathrm{sel}}_\theta$ for the training-fitted candidate used during selection, let $\mathcal{I}_{\mathrm{val}}$ denote the validation index set, and set $n_{\mathrm{val}} = |\mathcal{I}_{\mathrm{val}}|$. For $C \in \mathbb{R}^{q_C \times K}$, $q_C \geq 1$, set $\ell_C(\theta; O) := q_C^{-1} \left\| C\widehat{\mu}^{\mathrm{sel}}_\theta(X) - C\psi(O) \right\|_2^2$ and define the direct validation score

$$\widehat{R}^{\mathrm{val}}_C(\theta) = \frac{1}{n_{\mathrm{val}}} \sum_{i \in \mathcal{I}_{\mathrm{val}}} \ell_C(\theta; O_i), \qquad \widehat{\theta}_C \in \arg\min_{\theta \in \Theta} \widehat{R}^{\mathrm{val}}_C(\theta).$$

Define the corresponding conditional target risk

$$R^{\mathrm{sel}}_C(\theta \mid \mathcal{F}_{\mathrm{tr}}) = \mathbb{E}\left[q_C^{-1} \left\| C\widehat{\mu}^{\mathrm{sel}}_\theta(X) - C\mu(X) \right\|_2^2 \mid \mathcal{F}_{\mathrm{tr}} \right].$$

We suppress the explicit conditioning argument below. The same fitted candidate family can be scored for different prespecified contrast matrices without reconstructing the graph family.

Let $R^{\mathrm{refit}}_C(\theta)$ denote the corresponding conditional population risk after the prespecified final refit, and define

$$\Delta_{\mathrm{refit},C} = \sup_{\theta \in \Theta} \left| R^{\mathrm{refit}}_C(\theta) - R^{\mathrm{sel}}_C(\theta) \right|.$$

With independent validation observations and finite second moments, conditional unbiasedness gives

$$\mathbb{E}\left[\widehat{R}^{\mathrm{val}}_C(\theta) \mid \mathcal{F}_{\mathrm{tr}} \right] = R^{\mathrm{sel}}_C(\theta) + \kappa_C,$$

where κ_C does not depend on the candidate. The finite-candidate selector obeys the following deployment guarantee.

Theorem 4.4 (Contrast-aware finite-candidate oracle inequality). *For every* $\delta \in (0, 1)$*, conditionally on* $\mathcal{F}_{\mathrm{tr}}$*, let* $\zeta_{n_{\mathrm{val}},\delta}$ *be the Bernstein deviation term of Appendix E.8. Under its sub-exponential loss condition, any finite training-measurable* Θ *obeys with probability at least* $1 - \delta$*,*

$$\boxed{R^{\mathrm{refit}}_C(\widehat{\theta}_C) \leq \inf_{\theta \in \Theta} R^{\mathrm{refit}}_C(\theta) + 2\zeta_{n_{\mathrm{val}},\delta} + 2\Delta_{\mathrm{refit},C}.}$$

The deviation rate is $\zeta_{n_{\mathrm{val}},\delta} = O\{\sqrt{\log(|\Theta|/\delta)/n_{\mathrm{val}}}\}$ up to the usual Bernstein linear term. The theorem isolates selection error from the prespecified refit shift; the no-refit deployment has $\Delta_{\mathrm{refit},C} = 0$. The proof, exact constants, nuisance-product remainder, and pseudocode are in Appendices E.8 and E.9.

5 EXPERIMENTS

We test TSGS across Boolean, constrained, estimated-propensity, real-covariate, and randomized-clinical settings; Table 1 benchmarks neural estimators and Figure 3 isolates the borrowing mechanism. Full protocols are in Appendix F; NoGraph and Full denote $G_{\emptyset}$ and G_{geom}.

Table 1: **Neural edge/square PEHE** (mean ± SE, seeds 201–210). R1 is smooth; R2 has rough local transitions. Complete metrics and protocols: Table 4 and Appendix F.

Method	R1 edge	R1 square	R2 edge	R2 square
MultiHotMLP	0.1853 ± 0.0009	0.1620 ± 0.0013	0.6961 ± 0.0093	0.9987 ± 0.0102
SharedBundleHead	0.4738 ± 0.0067	0.6442 ± 0.0111	0.4926 ± 0.0043	0.6742 ± 0.0075
SCP	0.2564 ± 0.0023	0.2246 ± 0.0035	0.7299 ± 0.0024	1.0295 ± 0.0026
SIN-vector	0.4105 ± 0.0068	0.2866 ± 0.0040	0.5966 ± 0.0170	0.6991 ± 0.0204
VSR	0.2550 ± 0.0022	0.1838 ± 0.0039	0.4539 ± 0.0297	0.5629 ± 0.0532
TSGS	0.2751 ± 0.0043	0.2389 ± 0.0034	0.3292 ± 0.0039	0.3670 ± 0.0054

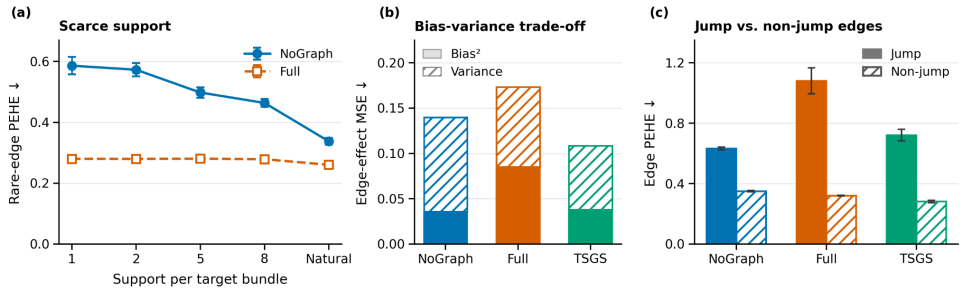

Figure 3: **When is cross-treatment borrowing safe? (a)** Scarce support rewards borrowing. **(b)** Full reduces variance but incurs rough-transition bias; TSGS has the lowest MSE. **(c)** NoGraph is best on jump edges; TSGS improves over NoGraph on non-jump edges. Whiskers in (a) and (c) show mean ± SE over five and ten seeds, respectively.

5.1 BORROWING SAFETY AND TARGET ALIGNMENT

Support and roughness. Figure 3 shows the mechanism: scarce support rewards borrowing, whereas rough transitions make indiscriminate borrowing biased. In the controlled 6×6 projected graph-sequence grid (40 repetitions/cell), TSGS stays below NoGraph throughout, while strong Full reaches 1.60–$1.93\times$ NoGraph in rough cells (Figure 6); scaling results are in Appendix F.13.

Comparators and robustness. MultiHotMLP is strongest in smooth R1; TSGS leads all five baselines in rough R2 (Table 1) and beats adaptive fusion in every rough benchmark cell. Against matched Full-Val, R2 edge/square MSE falls from $0.1562/0.2487$ to $0.1147/0.1454$ (Figure 4).

Target alignment and randomized trial. Edge and square validation choose different rules in 65.8% of 120 runs. On fresh R3 (amplitude 1), square-AIPW lowers square MSE to 0.1822 versus 0.2074 edge-selected and 0.2236 PredictiveCV, with similar edge risk (Figure 5). In ACTG 175, TSGS wins 70–89% of splits and improves held-out score by 0.7–1.7% (Figure 11).

6 DISCUSSION AND RESEARCH AGENDA

TSGS separates fixed treatment geometry from target-specific borrowing: the risk analysis explains target-dependent sharing, and neural, matched full-graph, cross-target, and ACTG 175 comparisons test graph selection and target-aligned validation. Future work includes scalable computation, partial identification for unsupported treatments, exposure mappings, and treatment histories.

AI USE STATEMENT

Generative AI tools assisted conceptual refinement, mathematical exposition, proof drafting, experimental feedback, implementation, literature analysis, and editing. The authors reviewed and validated all claims, proofs, code, experiments, citations, and text.

ETHICS STATEMENT

Our empirical study includes secondary analyses of publicly distributed benchmark data. The ACTG 175 experiment uses data from a randomized clinical trial; we did not recruit participants, intervene on participants, or collect new human-subject data. The original ACTG 175 trial was conducted with institutional review board approval and written informed consent from participants (Hammer et al., 1996). Our ACIC/LBIDD experiment uses publicly distributed benchmark covariates only for semi-synthetic evaluation; we remove the benchmark identifier, do not attempt re-identification, and do not release new individual-level data.

REPRODUCIBILITY STATEMENT

The anonymous code supplement provides executable Python code, frozen configurations, compact saved summaries and selected-rule artifacts for the neural synthetic, constrained-space, estimated-propensity, and ACIC/LBIDD semi-synthetic experiments, together with Figure 3 regeneration assets, numerical checks, and the kernel-checked Lean project. Appendices C–E.8 contain the analytical proofs, and Appendix F records the available specifications, protocols, and reported summaries for the remaining analyses. Execution artifacts for the projected graph-sequence, adaptive-fusion, and ACTG 175 analyses are not included in the supplement.

REFERENCES

Ahmed Alaa and Mihaela Van Der Schaar. Validating causal inference models via influence functions. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 191–201. PMLR, 2019. URL https://proceedings.mlr.press/v97/alaa19a.html.

Mikhail Belkin, Partha Niyogi, and Vikas Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7: 2399–2434, 2006. URL https://jmlr.org/papers/v7/belkin06a.html.

Vinod Kumar Chauhan, Jiandong Zhou, Ghadeer Ghosheh, Soheila Molaei, and David A. Clifton. Dynamic inter-treatment information sharing for individualized treatment effects estimation. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 3529–3537. PMLR, 2024. URL https://proceedings.mlr.press/v238/chauhan24a.html.

Vinod Kumar Chauhan, Lei Clifton, Gaurav Nigam, and David A. Clifton. Individualised treatment effects estimation with composite treatments and composite outcomes. In *2025 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 1–7. IEEE, 2025. doi: 10.1109/EMBC58623.2025.11252651.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. doi: 10.1111/ectj.12097.

Oriol Corcoll, Athanasios Vlontzos, Michael O'Riordan, and Ciarán M. Gilligan-Lee. Contrastive representations of structured treatments. *npj Artificial Intelligence*, 2:49, 2026. doi: 10.1038/s44387-026-00105-2.

Alicia Curth and Mihaela van der Schaar. On inductive biases for heterogeneous treatment effect estimation. In *Advances in Neural Information Processing Systems*, volume 34, pp. 15883–15894, 2021.

Alicia Curth and Mihaela Van Der Schaar. In search of insights, not magic bullets: Towards demystification of the model selection dilemma in heterogeneous treatment effect estimation. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 6623–6642. PMLR, 2023. URL https://proceedings.mlr.press/v202/curth23b.html.

Ahmed El Alaoui, Xiang Cheng, Aaditya Ramdas, Martin J. Wainwright, and Michael I. Jordan. Asymptotic behavior of ℓ_p-based laplacian regularization in semi-supervised learning. In *Proceedings of the 29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pp. 879–906. PMLR, 2016. URL https://proceedings.mlr.press/v49/elalaoui16.html.

Scott M. Hammer, David A. Katzenstein, Michael D. Hughes, Holly Gundacker, Robert T. Schooley, Richard H. Haubrich, W. Keith Henry, Michael M. Lederman, John P. Phair, Manette Niu, Martin S. Hirsch, and Thomas C. Merigan. A trial comparing nucleoside monotherapy with combination therapy in HIV-infected adults with CD4 cell counts from 200 to 500 per cubic millimeter. *New England Journal of Medicine*, 335(15):1081–1090, 1996. doi: 10.1056/NEJM199610103351501.

Shonosuke Harada and Hisashi Kashima. GraphITE: Estimating individual effects of graph-structured treatments. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*, pp. 659–668. ACM, 2021. doi: 10.1145/3459637.3482349.

Fredrik D. Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 3020–3029, 2016.

Fredrik D. Johansson, Uri Shalit, Nathan Kallus, and David Sontag. Generalization bounds and representation learning for estimation of potential outcomes and causal effects. *Journal of Machine Learning Research*, 23(166):1–50, 2022.

Jean Kaddour, Yuchen Zhu, Qi Liu, Matt J. Kusner, and Ricardo Silva. Causal effect inference for structured treatments. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Vassilis Kalofolias. How to learn a graph from smooth signals. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pp. 920–929. PMLR, 2016. URL https://proceedings.mlr.press/v51/kalofolias16.html.

Xiaopeng Ke, Yihan Yu, Ruyue Zhang, Zhishuo Zhou, Fangzhou Shi, Chang Men, and Zhengdan Zhu. Cross-treatment effect estimation for multi-category, multi-valued causal inference via dynamic neural masking. *arXiv preprint arXiv:2511.01641*, 2025. URL https://arxiv.org/abs/2511.01641.

Hui Lan and Vasilis Syrgkanis. Causal Q-aggregation for CATE model selection. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 4366–4374. PMLR, 2024. URL https://proceedings.mlr.press/v238/lan24a.html.

Haixu Ma, Donglin Zeng, and Yufeng Liu. Learning individualized treatment rules with many treatments: A supervised clustering approach using adaptive fusion. In *Advances in Neural Information Processing Systems*, volume 35, pp. 15956–15969. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/663865ea167425c6c562cb0b6bcf76c7-Abstract-Conference.html.

Haixu Ma, Donglin Zeng, and Yufeng Liu. Learning optimal group-structured individualized treatment rules with many treatments. *Journal of Machine Learning Research*, 24(102):1–48, 2023. URL https://www.jmlr.org/papers/v24/22-0520.html.

Abhirup Mondal, Anirban Majumder, and Vineet Chaoji. MEMENTO: Neural model for estimating individual treatment effects for multiple treatments. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pp. 3381–3390. ACM, 2022. doi: 10.1145/3511808.3557125.

Antonio Ortega, Pascal Frossard, Jelena Kovacevic, José M. F. Moura, and Pierre Vandergheynst. Graph signal processing: Overview, challenges, and applications. *Proceedings of the IEEE*, 106 (5):808–828, 2018. doi: 10.1109/JPROC.2018.2820126.

Sonali Parbhoo, Stefan Bauer, and Patrick Schwab. NCoRE: Neural counterfactual representation learning for combinations of treatments. *arXiv preprint arXiv:2103.11175*, 2021.

Zhaozhi Qian, Alicia Curth, and Mihaela van der Schaar. Estimating multi-cause treatment effects via single-cause perturbation. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Yuta Saito and Shota Yasui. Counterfactual cross-validation: Stable model selection procedure for causal inference models. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 8398–8407. PMLR, 2020. URL https://proceedings.mlr.press/v119/saito20a.html.

Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: Generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3076–3085, 2017.

Claudia Shi, David M. Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Xinyan Su, Jiacan Gao, Mingyuan Ma, Xiao Xu, Xinrui Wan, Tianqi Gu, Enyun Yu, Jiecheng Guo, and Zhiheng Zhang. Orthogonal uplift learning with permutation-invariant representations for combinatorial treatments. *arXiv preprint arXiv:2602.19851*, 2026. URL https://arxiv.org/abs/2602.19851.

Akira Tanimoto, Tomoya Sakai, Takashi Takenouchi, and Hisashi Kashima. Regret minimization for causal inference on large treatment space. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 946–954. PMLR, 2021. URL https://proceedings.mlr.press/v130/tanimoto21a.html.

Yingrong Wang, Haoxuan Li, Minqin Zhu, Anpeng Wu, Baohong Li, Keting Yin, Ruoxuan Xiong, Fei Wu, and Kun Kuang. Causal inference with complex treatments: A survey. *ACM Computing Surveys*, 58(9):230, 2026. doi: 10.1145/3789499.

Qingyan Xiang, Yubai Yuan, Dongyuan Song, Usman J. Wudil, Muktar H. Aliyu, C. William Wester, and Bryan E. Shepherd. Double machine learning to estimate the effects of multiple treatments and their interactions. *arXiv preprint arXiv:2505.12617*, 2025.

Ke Zhu, Jianing Chu, Ilya Lipkovich, Wenyu Ye, and Shu Yang. Doubly robust fusion of many treatments for policy learning. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 79772–79789. PMLR, 2025. URL https://proceedings.mlr.press/v267/zhu25e.html.

Hao Zou, Peng Cui, Bo Li, Zheyan Shen, Jianxin Ma, Hongxia Yang, and Yue He. Counterfactual prediction for bundle treatment. In *Advances in Neural Information Processing Systems*, volume 33, pp. 19705–19715, 2020.

Appendix Contents

A. Notation and Assumption Diagnostics 13
B. Extended Related Work 13
C. Identification and Decomposition Proofs 14
C.1 Identification 14
C.2 Non-identification 14
C.3 Minimal order for synergy 15
C.4 Projection proposition 15
C.5 Graph-roughness identity 15
D. Projection Details and Full-Cube Interpretation 15
E. Graph-Selection Theory Proofs 15
E.1 Projected linear smoother 15
E.2 Exact contrast risk 16
E.3 Projected Loewner ordering 16
E.4 Boolean contrast spectrum 17
E.5 Two-node oracle shrinkage 17
E.6 Restricted neural connection 17
E.7 Target-dependent counterexample 17
E.8 Validation and refit transfer 18
E.9 Practical selection algorithm 20
F. Experimental Details and Supplementary Results 20
F.1 Synthetic DGP and shared protocol 21
F.2 External baseline specifications 23
F.3 H1: rare supported bundles 24
F.4 H2: rough-transition borrowing 24
F.5 H3: projection ablations 25
F.6 H4: target-dependent borrowing 25
F.7 R3 selector confirmation 26
F.8 Selector choices and runtime 26
F.9 Constrained feasible spaces 27
F.10 Estimated-propensity robustness 28
F.11 Matched Full-Val and target transfer 29
F.12 ACIC semi-synthetic stress test 32
F.13 Projected graph-sequence experiments 33
F.14 Adaptive-fusion benchmark 35
F.15 ACTG 175 protocol 36
F.16 Formal verification in Lean 4 37
G. Boolean Finite-Difference Algebra 38

A Notation and Assumption Diagnostics

Table 2: **Core notation and its scientific meaning.**

Symbol	Meaning
$\mathcal{A}_f, \mathcal{A}_s, \mathcal{A}_{s,\eta}$	feasible, observed-support, and overlap-safe bundle sets
G_{geom}	fixed scientific graph of feasible one-component interventions
G_{bor}	selected subgraph used for statistical borrowing only
$\mu(x, a)$	conditional potential-outcome surface at covariates x and bundle a
$C, \tau_C(x)$	prespecified zero-sum contrast operator and causal target $C\mu(x)$
P_0^π, Q_0^π	prespecified structured projection and its weighted orthogonal complement
(G, λ_G)	neural borrowing rule: graph and penalty strength; λ denotes the graph-sequence penalty in Section 4

The causal conditions and reporting restriction above are operational. Consistency and no interference depend on intervention delivery and spillover mechanisms; if spillovers are scientifically relevant, the treatment node must be replaced by an exposure mapping. Randomization supports treatment exchangeability in ACTG 175, subject to the analyzed-sample and missingness assumptions; in observational settings, unconfoundedness motivates negative-control and sensitivity analyses, and failure calls for a different design rather than stronger smoothing. If overlap fails, the target population or reportable estimand/contrast set must be redefined. These conditions do not assert smoothness, structured-span truth, or correctness of any borrowing graph.

B Extended Related Work

Structured and multiple treatments. Recent work on complex treatments spans multi-valued, continuous, and bundle treatments (Wang et al., 2026). Beyond the closest bundle-treatment methods discussed in the Introduction, related approaches use treatment encoders, shared parameters, hypernetworks, or explicit cross-effect architectures to transfer information across treatment configurations (Kaddour et al., 2021; Curth & van der Schaar, 2021; Chauhan et al., 2024; 2025; Mondal et al., 2022; Ke et al., 2025). Recent contrastive treatment-representation work learns causally relevant factors from high-dimensional structured treatments (Corcoll et al., 2026). Orthogonal uplift learning for combinatorial treatments uses permutation-invariant treatment representations with Robinson-style nuisance orthogonalization (Su et al., 2026), while double-machine-learning methods target prespecified main and interaction effects of concurrent treatments through Neyman-orthogonal scores and

cross-fitting (Xiang et al., 2025; Chernozhukov et al., 2018). These approaches address representation, balancing, or nuisance estimation; TSGS focuses on selecting which prespecified local treatment relations should be used for borrowing.

Many treatments and fusion. Adaptive-fusion methods cluster treatments with similar response patterns and learn group-structured individualized rules (Ma et al., 2022; 2023); calibration-weighted fusion adds doubly robust group recovery and policy-regret guarantees (Zhu et al., 2025). These methods seek a global treatment partition. TSGS retains every feasible bundle as a distinct node, imposes continuous shrinkage only along scientifically admissible local relations, and lets the selected relation change with the reported contrast family. Appendix F.14 gives a matched specialization of the adaptive all-pairs fusion objective; under its balanced intercept-only design and the same projected parameter space, Zhu et al.'s calibration weights are uniform and its quadratic loss is proportional to the same node-mean least-squares loss, up to an additive constant.

Causal model selection. Influence-function validation, causal model selection, and aggregation provide related foundations (Alaa & Van Der Schaar, 2019; Curth & Van Der Schaar, 2023; Lan & Syrgkanis, 2024). The decomposition of Saito & Yasui (2020) states that expected squared error against a conditionally unbiased pseudo-outcome equals target risk plus a candidate-independent term. We turn this identity into a vector-valued, multi-arm selection principle for graph-valued borrowing rules under a prespecified contrast operator.

Graph regularization and causal representation learning. Laplacian and manifold regularization impose smoothness over a fixed geometry (Belkin et al., 2006; Ortega et al., 2018), graph-learning methods estimate graph structure from smooth signals (Kalofolias, 2016), and related statistical theory studies Laplacian-regularized regression (El Alaoui et al., 2016). TSGS uses a different role for the graph: the feasible intervention geometry is fixed, while a subgraph is selected as a statistical borrowing device for the causal contrasts of interest. This treatment-side question complements covariate-side causal representation learning, which primarily learns representations of X for counterfactual generalization (Johansson et al., 2016; Shalit et al., 2017; Johansson et al., 2022; Shi et al., 2019).

C Proofs for Identification and Decomposition Results

C.1 Proof of Theorem 2.1

For any $a \in \mathcal{A}_{s,\eta}$ and P_X-almost every target-population x at which $e_a(x) > 0$, consistency gives

$$\mathbb{E}(Y \mid X = x, A = a) = \mathbb{E}\{Y(a) \mid X = x, A = a\}.$$

Bundle-level unconfoundedness then yields $\mathbb{E}\{Y(a) \mid X = x, A = a\} = \mathbb{E}\{Y(a) \mid X = x\} = \mu(x, a)$. Positivity ensures that the observed conditional law for treatment a is defined on the stated target region. A finite-difference estimand is a finite signed sum of these identified conditional means, so it is identified whenever all of its vertices satisfy the required overlap condition. Integrating the identified conditional contrast over the target distribution of X identifies its population average.

C.2 Proof of Theorem 2.2

Let $B \subseteq \mathcal{X}$ satisfy $P_X(B) > 0$ and $e_{a_0}(x) = 0$ on B. Construct two full-data laws P_1 and P_2 with the same distribution of X, the same assignment law $P(A \mid X)$, and identical potential-outcome laws for every bundle except a_0 on B. On B, let

$$Y^{(r)}(a_0) = g_r(X) + \epsilon, \qquad r \in \{1, 2\},$$

where g_1 and g_2 are distinct bounded functions and ϵ has the same mean-zero distribution under both laws and is independent of assignment conditional on X. The two laws can therefore satisfy bundle-level unconfoundedness, and consistency is imposed in the usual way by setting the observed outcome to $Y(A)$. Because $A = a_0$ has zero probability on B, the altered potential-outcome law is never revealed there, so $P_1(O) = P_2(O)$ for the observed data $O = (X, A, Y)$. Yet $\mu_1(x, a_0) = g_1(x)$ and $\mu_2(x, a_0) = g_2(x)$ on B. Hence $\mu(x, a_0)$, and any contrast requiring it, is not identified on B from the observed law under the stated assumptions alone.

C.3 Minimal-order calculation for pairwise synergy

Let $f(a) = \beta_0 + \sum_{\ell=1}^{m} \beta_\ell a_\ell$. Applying Δ_k removes the intercept and every term not containing a_k, leaving the constant β_k; applying Δ_j with $j \neq k$ therefore gives zero. Thus $\Delta_j \Delta_k f(a) = 0$ for every degree-one Boolean polynomial. In contrast, $\Delta_j \Delta_k (a_j a_k) = 1$, so degree two is the minimal Boolean order capable of representing a nonzero square-face synergy.

C.4 Proof of Proposition 3.1

See Appendix D.

C.5 Complementary graph-roughness identity

For a weighted undirected graph with each edge counted once,

$$h^\top L_G h = \sum_{e=(a,b)\in E_G} w_e (h_b - h_a)^2.$$

Substituting $h = \mu_{\mathrm{comp}}(X)$ and taking expectation over X gives

$$\begin{aligned}\mathbb{E}_X\left[\mu_{\mathrm{comp}}(X)^\top L_G \mu_{\mathrm{comp}}(X)\right] &= \sum_{e=(a,b)\in E_G} w_e \, \mathbb{E}_X\left[\{\mu_{\mathrm{comp}}(X,b) - \mu_{\mathrm{comp}}(X,a)\}^2\right] \\ &= \sum_{e\in E_G} w_e \, \mathfrak{r}_e^2,\end{aligned}$$

which is the stated identity.

D Projection Details and Full-Cube Interpretation

For

$$\beta_f = (B_0^\top W_\pi B_0)^{-1} B_0^\top W_\pi f, \qquad h_f = f - B_0 \beta_f,$$

we have $B_0^\top W_\pi h_f = 0$, and invertibility of $B_0^\top W_\pi B_0$ gives uniqueness. This proves Proposition 3.1. If B_0 is rank-deficient, replacing it by any full-column-rank $\widetilde{B}_0$ with $\mathrm{col}(\widetilde{B}_0) = \mathrm{col}(B_0)$ gives the same unique W_π-orthogonal projection onto the structured functional span; only the coefficients in the redundant original dictionary lose uniqueness.

On the complete Boolean cube with uniform π and the default choice $B_0 = B_{\le 2}$, write $s_j = 2a_j - 1$ and $\chi_S(a) = \prod_{j\in S} s_j$. The span of $\{1, a_j, a_j a_k\}$ equals $\mathrm{span}\{\chi_S : |S| \le 2\}$, and Walsh characters are orthogonal under the uniform measure. Hence the complementary space is $\mathrm{span}\{\chi_S : |S| > 2\}$. On constrained feasible designs, Q_0^π is the reference-weighted complement of the structured span; raw higher-order $0/1$ Möbius monomials do not in general coincide with this subspace.

E Proofs for Graph-Selection Theory

E.1 Derivation of the projected graph linear smoother

Every $h \in \mathcal{H}^\pi_{\mathrm{comp}}$ has the unique representation $h = U\alpha$. Substitution gives

$$\begin{aligned}J(\alpha) &= (Z - U\alpha)^\top D_n (Z - U\alpha) + \lambda \alpha^\top U^\top L_G U \alpha \\ &= Z^\top D_n Z - 2\alpha^\top U^\top D_n Z + \alpha^\top A_{G,\lambda} \alpha.\end{aligned}$$

The Hessian is $2A_{G,\lambda}$. Admissibility therefore makes the objective strictly convex. The first-order condition is

$$A_{G,\lambda}\widehat{\alpha} = U^\top D_n Z,$$

which yields

$$\widehat{h}_{G,\lambda} = U\widehat{\alpha} = U A_{G,\lambda}^{-1} U^\top D_n Z.$$

The right-hand side depends on Z only through $D_n Z$, so zero-count coordinates of Z are irrelevant.

For the null-space statement, let $v \in \mathbb{R}^{d_{\text{comp}}}$ and $h = Uv$. Since D_n and L_G are positive semidefinite,

$$\begin{aligned} v^\top A_{G,\lambda} v &= h^\top D_n h + \lambda h^\top L_G h \\ &= \sum_a n_a h_a^2 + \lambda \sum_{(a,b)\in E_G} w_{ab}(h_a - h_b)^2. \end{aligned}$$

For $\lambda > 0$, this quantity is zero if and only if $h \in \mathrm{Null}(D_n) \cap \mathrm{Null}(L_G) \cap \mathcal{H}^{\pi}_{\text{comp}}$. Hence $A_{G,\lambda}$ is positive definite exactly when this intersection is $\{0\}$. If every connected component contains a positive-count node, $L_G h = 0$ makes h constant on each component and $D_n h = 0$ anchors each such constant at zero, proving the sufficient condition.

E.2 Proof of Theorem 4.1

Define

$$\xi = D_n(Z - h_0).$$

At positive-count nodes $\xi_a = n_a \varepsilon_a$, while at zero-count nodes $\xi_a = 0$. Conditional on D_n,

$$\mathbb{E}(\xi \mid D_n) = 0, \qquad \mathrm{Cov}(\xi \mid D_n) = \sigma^2 D_n.$$

Writing $M = U A_{G,\lambda}^{-1} U^\top$ and $S = M D_n$,

$$\begin{aligned} \widehat{h}_{G,\lambda} &= M D_n Z \\ &= M D_n h_0 + M\xi \\ &= S h_0 + M\xi. \end{aligned}$$

Therefore

$$C(\widehat{h}_{G,\lambda} - h_0) = C(S - I_K)h_0 + CM\xi.$$

The second term has conditional mean zero, so the expected cross term vanishes. Using

$$\mathbb{E}\,\|A\xi\|_2^2 = \mathrm{tr}\{A\,\mathrm{Cov}(\xi)A^\top\}$$

gives

$$\mathcal{R}_C(G, \lambda; h_0 \mid D_n) = \|C(S - I_K)h_0\|_2^2 + \sigma^2\,\mathrm{tr}(CMD_nM^\top C^\top).$$

Because $A_{G,\lambda}$ is symmetric, $M^\top = M$, which yields the displayed variance expression.

For the bias identity, write $h_0 = U\alpha_0$. Then

$$\begin{aligned} Sh_0 - h_0 &= U A_{G,\lambda}^{-1} U^\top D_n U \alpha_0 - U\alpha_0 \\ &= U A_{G,\lambda}^{-1}\{U^\top D_n U - A_{G,\lambda}\}\alpha_0 \\ &= -\lambda U A_{G,\lambda}^{-1} U^\top L_G U \alpha_0 \\ &= -\lambda U A_{G,\lambda}^{-1} U^\top L_G h_0. \end{aligned}$$

Left-multiplication by C proves the second claim.

E.3 Proof of Theorem 4.2

Let $D = M_{\theta_2} - M_{\theta_1}$ and $P_\perp = I_K - K^{-1}\mathbf{1}_K\mathbf{1}_K^\top$. If C has rows $c_1^\top, \ldots, c_{q_C}^\top$ and $C\mathbf{1}_K = 0$, then every $c_j = P_\perp c_j$ and

$$\mathcal{R}_C(\theta_2 \mid D_n) - \mathcal{R}_C(\theta_1 \mid D_n) = \mathrm{tr}(CDC^\top) = \sum_{j=1}^{q_C} c_j^\top P_\perp D P_\perp c_j.$$

If $P_\perp D P_\perp \succeq 0$, every summand is nonnegative, so (iii) implies (i); (i) immediately implies (ii). Conversely, (ii) states that $v^\top D v = v^\top P_\perp D P_\perp v \geq 0$ for every $v \in \mathbf{1}_K^\perp$. For arbitrary $x \in \mathbb{R}^K$, $P_\perp x \in \mathbf{1}_K^\perp$, and hence

$$x^\top P_\perp D P_\perp x = (P_\perp x)^\top D (P_\perp x) \geq 0.$$

Thus $P_\perp D P_\perp \succeq 0$, proving (ii)⇒(iii). Finally, if the restricted quadratic form has witnesses $v_+, v_- \in \mathbf{1}_K^\perp$ with $v_+^\top D v_+ > 0$ and $v_-^\top D v_- < 0$, the valid one-row contrasts $v_+^\top$ and $v_-^\top$ give the claimed strict reversal.

E.4 Proof of Theorem 4.3 and Corollary 4.1

Index an order-r face by a toggled coordinate set $R \subseteq [m]$, $|R| = r$, and a fixed assignment of the other $m - r$ coordinates. Its difference applied to u_S is zero unless $R \subseteq S$, because toggling a coordinate outside S cancels. When $R \subseteq S$, the absolute value is $2^{r-m/2}$. There are $\binom{|S|}{r}$ admissible choices of R and 2^{m-r} assignments of the untoggled coordinates, hence

$$\|C_r u_S\|_2^2 = \binom{|S|}{r} 2^{m-r} 2^{2r-m} = 2^r \binom{|S|}{r}.$$

Boolean translation operators commute and the Walsh characters are their joint eigenbasis. Because $C_r^{\mathsf{T}} C_r$ is a sum of the associated translation-invariant face-difference normal operators, it is diagonal in that basis, with the squared norm above as its eigenvalue. This proves Theorem 4.3.

For the corollary, take $B_0 = B_{\leq 2}$ and uniform π. The full-cube Hamming Laplacian obeys $L_H u_S = 2|S| u_S$, so $(I + \lambda L_H)^{-1} u_S = s_{|S|}(\lambda) u_S$. The Walsh basis is orthonormal, the node noise is isotropic, and Theorem 4.3 supplies the contrast weight. Adding squared bias and variance in each spectral coordinate gives

$$\sum_{|S|>2} 2^r \binom{|S|}{r} \left[\{1 - s_{|S|}(\lambda)\}^2 \alpha_S^2 + \sigma^2 s_{|S|}(\lambda)^2 \right],$$

as claimed.

E.5 Proof of Proposition 4.1

Completing the square yields

$$d^2(1-s)^2 + v s^2 = \frac{d^2 v}{d^2 + v} + (d^2 + v) \left(s - \frac{d^2}{d^2 + v} \right)^2.$$

When $d^2 + v > 0$, the second term has a unique zero at $s^\star = d^2/(d^2 + v)$, proving both optimality and the exact excess-risk identity. If $d = v = 0$, every s has zero risk.

E.6 Proof of Proposition 4.2

For every node a with $n_a > 0$,

$$R_i - h_a = (R_i - Z_a) + (Z_a - h_a).$$

Thus

$$\sum_{i:A_i=a} (R_i - h_a)^2 = \sum_{i:A_i=a} (R_i - Z_a)^2 + 2(Z_a - h_a) \sum_{i:A_i=a} (R_i - Z_a) + n_a (Z_a - h_a)^2.$$

The middle term is zero by the definition of Z_a. Summing over a gives

$$\sum_i (R_i - h_{A_i})^2 = \underbrace{\sum_a \sum_{i:A_i=a} (R_i - Z_a)^2}_{\text{independent of } h} + \sum_a n_a (Z_a - h_a)^2.$$

Multiplying the edge-normalized neural objective by n and setting $\lambda = n\lambda_G/|E_G|$ therefore leaves exactly the graph-sequence objective plus an h-independent constant, so the minimizers coincide for nonempty G. For $G_\emptyset$, both graph penalties vanish and the correspondence holds with $\lambda = 0$.

E.7 Spectral calculation for the target-dependent counterexample

On the four-dimensional Hamming cube, let G_H denote the full Hamming graph with Laplacian L_H, and take the complementary truth $h_0 = \alpha u_{[4]}$. The normalized Walsh characters satisfy

$$L_H u_S = 2|S| u_S.$$

The complementary subspace consists of four degree-three characters and one degree-four character. With unit counts, let R_I denote nodewise complementary-surface risk ($C = I_{16}$), while R_{edge} and R_{sq} use C_{edge} and C_{sq}. The empty-graph estimator has shrinkage factor one in all five directions. For G_H and $\lambda = 1/8$, the factors are

$$s_3 = \frac{1}{1 + 6/8} = \frac{4}{7}, \qquad s_4 = \frac{1}{1 + 8/8} = \frac{1}{2}.$$

Let C_{edge} be the incidence matrix containing every undirected Hamming edge once. Then

$$C_{\text{edge}}^{\mathsf{T}} C_{\text{edge}} = L_H, \qquad \|C_{\text{edge}} u_S\|_2^2 = 2|S|.$$

Let C_{sq} stack all square-face second differences. A square contrast is nonzero on u_S exactly when its two toggled coordinates lie in S; each such contrast has unit magnitude under the chosen normalization. Hence

$$\|C_{\text{sq}} u_S\|_2^2 = 4\binom{|S|}{2}.$$

No borrowing therefore has risks

$$R_I(G_\emptyset) = 5\sigma^2, \qquad R_{\text{edge}}(G_\emptyset) = 32\sigma^2, \qquad R_{\text{sq}}(G_\emptyset) = 72\sigma^2.$$

For full-Hamming borrowing, only the degree-four truth contributes bias:

$$R_I(G_H) = \frac{\alpha^2}{4} + \sigma^2 \left\{ 4\left(\frac{4}{7}\right)^2 + \left(\frac{1}{2}\right)^2 \right\} = \frac{\alpha^2}{4} + \frac{305}{196}\sigma^2,$$

$$R_{\text{edge}}(G_H) = 8\frac{\alpha^2}{4} + \sigma^2 \left\{ 24\left(\frac{4}{7}\right)^2 + 8\left(\frac{1}{2}\right)^2 \right\} = 2\alpha^2 + \frac{482}{49}\sigma^2,$$

$$R_{\text{sq}}(G_H) = 24\frac{\alpha^2}{4} + \sigma^2 \left\{ 48\left(\frac{4}{7}\right)^2 + 24\left(\frac{1}{2}\right)^2 \right\} = 6\alpha^2 + \frac{1062}{49}\sigma^2.$$

Comparison with the no-borrowing risks gives

$$\alpha^2 < \frac{675}{49}\sigma^2, \qquad \alpha^2 < \frac{543}{49}\sigma^2, \qquad \alpha^2 < \frac{411}{49}\sigma^2.$$

For example, $\alpha^2 = 10\sigma^2$ makes full-Hamming borrowing lower edge risk but raise square risk, giving the claimed target-dependent reversal; $\alpha^2 = 12\sigma^2$ raises both contrast risks while still lowering nodewise risk.

E.8 Proofs for target-specific validation and refit transfer

The following identity is the vector-contrast, multi-arm analogue of the conditionally unbiased pseudo-outcome decomposition of Saito & Yasui (2020).

Lemma E.1 (Direct AIPW validation differs by a candidate-independent constant)**.** *Let $q_C \geq 1$, $\mathcal{V}_C = \{a \in \mathcal{A}_f : C_{:,a} \neq 0\}$, and use the zero-extended ψ defined in Section 4. Conditional on $\mathcal{F}_{\text{tr}}$, suppose the finite candidate family and nuisance fits are measurable, the validation observations are independent of $\mathcal{F}_{\text{tr}}$, every treatment vertex used by C satisfies overlap, $\mathbb{E}\,\|C\psi(O)\|_2^2 < \infty$, and*

$$\mathbb{E}\{C\psi(O) \mid X, \mathcal{F}_{\text{tr}}\} = C\mu(X).$$

Then, for every $\theta \in \Theta$,

$$\mathbb{E}\left[\widehat{R}_C^{\text{val}}(\theta) \,\middle|\, \mathcal{F}_{\text{tr}}\right] = R_C^{\text{sel}}(\theta \mid \mathcal{F}_{\text{tr}}) + \kappa_C,$$

where κ_C does not depend on θ.

Proof. For $a \in \mathcal{V}_C$, the displayed AIPW score with the true propensity satisfies

$$\begin{aligned}
\mathbb{E}\{\psi_a(O) \mid X, \mathcal{F}_{\text{tr}}\} &= \widetilde{\mu}(X, a) + \frac{1}{e_a(X)}\mathbb{E}[\mathbf{1}(A = a)\{Y - \widetilde{\mu}(X, a)\} \mid X, \mathcal{F}_{\text{tr}}] \\
&= \widetilde{\mu}(X, a) + \mathbb{E}(Y \mid X, A = a) - \widetilde{\mu}(X, a) \\
&= \mu(X, a),
\end{aligned}$$

where the last equality follows from Assumption 2.1. Hence $\mathbb{E}\{C\psi(O) \mid X, \mathcal{F}_{\mathrm{tr}}\} = C\mu(X)$. Write

$$d_\theta(X) = C\widehat{\mu}_\theta^{\mathrm{sel}}(X) - C\mu(X), \qquad \eta_C(O) = C\psi(O) - C\mu(X).$$

Conditional on $\mathcal{F}_{\mathrm{tr}}$ and X, $\mathbb{E}\{\eta_C(O) \mid X, \mathcal{F}_{\mathrm{tr}}\} = 0$. Therefore

$$\begin{aligned}\mathbb{E}\{\ell_C(\theta; O) \mid \mathcal{F}_{\mathrm{tr}}\} &= \mathbb{E}\left[\frac{1}{q_C}\,\|d_\theta(X) - \eta_C(O)\|_2^2 \,\middle|\, \mathcal{F}_{\mathrm{tr}}\right] \\ &= R_C^{\mathrm{sel}}(\theta \mid \mathcal{F}_{\mathrm{tr}}) + \underbrace{\mathbb{E}\left[\frac{1}{q_C}\,\|\eta_C(O)\|_2^2 \,\middle|\, \mathcal{F}_{\mathrm{tr}}\right]}_{\kappa_C},\end{aligned}$$

and the final term is independent of θ. □

Remark E.1 (Estimated propensities). *Let $\widehat{e}_a$ and $\widetilde{\mu}$ be fitted without the validation observation and held fixed in the conditioning argument, with $\widehat{e}_a(X) > 0$ for $a \in \mathcal{V}_C$. On these coordinates, replacing e_a by $\widehat{e}_a$ gives*

$$\mathbb{E}\{\widehat{\psi}_a(O) \mid X, \mathcal{F}_{\mathrm{tr}}\} - \mu(X, a) = b_a(X) = \frac{\widehat{e}_a(X) - e_a(X)}{\widehat{e}_a(X)}\{\widetilde{\mu}(X, a) - \mu(X, a)\}.$$

Set $\widehat{\psi}_a(O) = b_a(X) = 0$ for $a \notin \mathcal{V}_C$. Let $b(X) = (b_a(X) : a \in \mathcal{A}_f)$, $d_\theta(X) = C\widehat{\mu}_\theta^{\mathrm{sel}}(X) - C\mu(X)$, and let $\widehat{\ell}_C$ denote the corresponding per-observation validation loss formed with $\widehat{\psi}$. Then

$$\mathbb{E}\{\widehat{\ell}_C(\theta; O) \mid \mathcal{F}_{\mathrm{tr}}\} = R_C^{\mathrm{sel}}(\theta) + \widehat{\kappa}_C + r_C(\theta), \qquad r_C(\theta) = -\frac{2}{q_C}\mathbb{E}\{d_\theta(X)^{\mathsf{T}} C b(X) \mid \mathcal{F}_{\mathrm{tr}}\},$$

where $\widehat{\kappa}_C$ is candidate-independent. Hence the same finite-family argument yields an additional $2\rho_C$ term when $\rho_C = \sup_{\theta\in\Theta} |r_C(\theta)|$. By Cauchy–Schwarz, ρ_C vanishes when candidate risks remain bounded and $\|Cb\|_{L_2} \to 0$; if $\widehat{e}_a$ is uniformly bounded away from zero, b_a is controlled by the product of propensity and outcome-nuisance errors. Cross-fitting supplies the independence used in this conditioning argument but does not make b_a exactly zero in finite samples.

Theorem E.1 (Finite-candidate target-specific validation). *Conditional on $\mathcal{F}_{\mathrm{tr}}$, suppose the validation observations are i.i.d., Θ is finite, and, writing $\|\cdot\|_{\psi_1}$ for the sub-exponential Orlicz norm,*

$$\sup_{\theta\in\Theta} \|\ell_C(\theta; O) - \mathbb{E}\{\ell_C(\theta; O) \mid \mathcal{F}_{\mathrm{tr}}\}\|_{\psi_1} \le v_C.$$

Let $\widehat{\theta}_C \in \arg\min_{\theta\in\Theta} \widehat{R}_C^{\mathrm{val}}(\theta)$. Then, with probability at least $1 - \delta$,

$$R_C^{\mathrm{sel}}(\widehat{\theta}_C) \le \inf_{\theta\in\Theta} R_C^{\mathrm{sel}}(\theta) + 2\zeta_{n_{\mathrm{val}},\delta},$$

where

$$\zeta_{n_{\mathrm{val}},\delta} = c_0 v_C \left\{\sqrt{\frac{\log(2|\Theta|/\delta)}{n_{\mathrm{val}}}} + \frac{\log(2|\Theta|/\delta)}{n_{\mathrm{val}}}\right\}$$

for a universal constant $c_0 > 0$.

Proof. For each fixed candidate, Bernstein's inequality gives

$$\mathbb{P}\left(\left|\widehat{R}_C^{\mathrm{val}}(\theta) - \mathbb{E}\{\widehat{R}_C^{\mathrm{val}}(\theta) \mid \mathcal{F}_{\mathrm{tr}}\}\right| > t \,\middle|\, \mathcal{F}_{\mathrm{tr}}\right) \le 2\exp\left[-cn_{\mathrm{val}} \min\left\{\frac{t^2}{v_C^2}, \frac{t}{v_C}\right\}\right]$$

for a universal constant $c > 0$. A union bound over Θ and inversion of the tail bound yield, with probability at least $1 - \delta$,

$$\sup_{\theta\in\Theta}\left|\widehat{R}_C^{\mathrm{val}}(\theta) - R_C^{\mathrm{sel}}(\theta) - \kappa_C\right| \le \zeta_{n_{\mathrm{val}},\delta},$$

where Lemma E.1 supplies the candidate-independent κ_C. On this event,

$$\begin{aligned}R_C^{\mathrm{sel}}(\widehat{\theta}_C) &\le \widehat{R}_C^{\mathrm{val}}(\widehat{\theta}_C) - \kappa_C + \zeta_{n_{\mathrm{val}},\delta} \\ &\le \inf_{\theta\in\Theta} \widehat{R}_C^{\mathrm{val}}(\theta) - \kappa_C + \zeta_{n_{\mathrm{val}},\delta} \\ &\le \inf_{\theta\in\Theta} R_C^{\mathrm{sel}}(\theta) + 2\zeta_{n_{\mathrm{val}},\delta}.\end{aligned}$$

□

Corollary E.1 (Refit-aware deployment guarantee)**.** *With*

$$\Delta_{\mathrm{refit},C} = \sup_{\theta\in\Theta} \left| R_C^{\mathrm{refit}}(\theta) - R_C^{\mathrm{sel}}(\theta) \right|,$$

the event in Theorem E.1 also implies

$$R_C^{\mathrm{refit}}(\widehat{\theta}_C) \le \inf_{\theta\in\Theta} R_C^{\mathrm{refit}}(\theta) + 2\zeta_{n_{\mathrm{val}},\delta} + 2\Delta_{\mathrm{refit},C}.$$

Proof. By definition of $\Delta_{\mathrm{refit},C}$,

$$\begin{aligned} R_C^{\mathrm{refit}}(\widehat{\theta}_C) &\le R_C^{\mathrm{sel}}(\widehat{\theta}_C) + \Delta_{\mathrm{refit},C} \\ &\le \inf_{\theta\in\Theta} R_C^{\mathrm{sel}}(\theta) + 2\zeta_{n_{\mathrm{val}},\delta} + \Delta_{\mathrm{refit},C} \\ &\le \inf_{\theta\in\Theta} R_C^{\mathrm{refit}}(\theta) + 2\zeta_{n_{\mathrm{val}},\delta} + 2\Delta_{\mathrm{refit},C}. \end{aligned}$$

□

Corollary E.1 does not impose a stability condition that makes $\Delta_{\mathrm{refit},C}$ small; it makes the refit shift explicit. Its deployment-stage implication is therefore strongest when the prespecified refit map is stable enough that candidate risks change little between selection and refitting.

E.9 Practical selection algorithm

Algorithm 1 Treatment-Space Graph Selection

Require: Training data D_{tr}, validation data D_{val}, geometry G_{geom}, target contrast family C, retention fractions $\mathcal{R}$, penalties Λ.

1: split D_{tr} into inner-fit and inner-checkpoint data;
2: fit a NoGraph pilot, checkpoint by factual MSE, and refit it on D_{tr};
3: rank geometry edges by complementary roughness and form G_ρ, $\rho \in \mathcal{R}$;
4: form $\Theta = \{(G_\emptyset, 0)\} \cup \{(G_\rho, \lambda_G) : \rho \in \mathcal{R}, \lambda_G \in \Lambda\}$ and fit a training-only outcome nuisance;
5: **for** $\theta = (G, \lambda_G) \in \Theta$ **do**
6: checkpoint the projected learner on inner-fit/inner-checkpoint data and refit θ on D_{tr} for its selected epoch;
7: score θ on D_{val} by $\widehat{R}_C^{\mathrm{val}}(\theta)$;
8: **end for**
9: retain the selected graph, λ_G, and epoch; refit that rule on $D_{\mathrm{tr}} \cup D_{\mathrm{val}}$ and evaluate once.

F Experimental Details and Supplementary Results

This appendix records the frozen data-generating processes, training protocols, baseline specifications, and supplementary results used in Section 5. Metrics are reported as means, means with standard errors, or risk ratios with confidence intervals as specified in the corresponding table captions.

Table 3: **Theoretical results and related empirical analyses.**

Theoretical statement	Experimental intervention	Primary diagnostic
Proposition 4.1	support × local roughness	risk-ratio phase boundary
Theorem 4.2	edge versus square target	selected-rule reversal
Theorem 4.3	mixed Walsh degrees	target-risk landscape
Theorem 4.1	fixed truth, repeated noise	squared bias versus variance
Theorem 4.4	independent AIPW validation	oracle rank/regret
Assumption 2.1	randomized ACTG 175	held-out target score

Table 4: **Main neural comparison under smooth and rough treatment-space structure.** Mean ± SE over independent seeds 201–210.

	R1: smooth complement			R2: rough local transitions		
Method	Edge PEHE ↓	Square PEHE ↓	Surface RMSE ↓	Edge PEHE ↓	Square PEHE ↓	Surface RMSE ↓
MultiHotMLP	**0.1853±0.0009**	**0.1620±0.0013**	**0.2733±0.0037**	0.6961±0.0093	0.9987±0.0102	0.5371±0.0057
SharedBundleHead	0.4738±0.0067	0.6442±0.0111	0.4260±0.0032	0.4926±0.0043	0.6742±0.0075	0.4345±0.0018
SCP	0.2564±0.0023	0.2246±0.0035	0.3624±0.0024	0.7299±0.0024	1.0295±0.0026	0.5799±0.0017
SIN-vector	0.4105±0.0068	0.2866±0.0040	0.6207±0.0110	0.5966±0.0170	0.6991±0.0204	0.7002±0.0140
VSR	0.2550±0.0022	0.1838±0.0039	0.3522±0.0033	0.4539±0.0297	0.5629±0.0532	0.4588±0.0160
TSGS	0.2751±0.0043	0.2389±0.0034	0.3507±0.0034	**0.3292±0.0039**	**0.3670±0.0054**	**0.3702±0.0028**

Data sizes and seed blocks. Unless otherwise stated, the canonical synthetic runs use 3000 training observations, an independent 3000-observation validation sample, and a common 5000-point counterfactual evaluation grid. H1, H3, and the oracle-landscape H4 use seeds 101–105; the main R1/R2 comparison, H2, and the original true-propensity R3 selector confirmation use seeds 201–210. The constrained treatment-space stress test uses seeds 301–310, estimated-propensity robustness uses seeds 401–410, and the ACIC semi-synthetic experiment uses seeds 501–510. VSR uses the same 201–210 and 501–510 blocks as the comparisons in which it appears.

F.1 Synthetic DGP and shared protocol

Covariates and low-order outcome surface. The synthetic experiments use the full $m = 6$ Boolean cube, so $K = 64$, with

$$X \sim \mathcal{N}(0, I_{20}).$$

Write the degree-≤ 2 outcome component as

$$\mu_{\leq 2}(x, a) = \beta_0(x) + \sum_{j=1}^{6} a_j \beta_j(x) + \sum_{1 \leq j < k \leq 6} a_j a_k \beta_{jk}(x),$$

where

$$\beta_0(x) = 0.7 \sin(x_1) + 0.3 x_2^2 - 0.2 x_3 x_4,$$
$$\beta_j(x) = c_j + 0.25 \tanh(x_j) + 0.15 x_{6+j}, \qquad c = (0.6, -0.4, 0.5, 0.3, -0.5, 0.4),$$

and

$$\beta_{jk}(x) = 0.25(-1)^{j+k} \left\{1 + 0.5 \tanh\left(x_{\ell(j,k)}\right)\right\}, \qquad \ell(j,k) = 13 + [(j + k - 2) \bmod 8].$$

Observational assignment and overlap. For $j = 1, \ldots, 6$, define

$$q_j(x) = 0.45 x_j - 0.30 x_{6+j} + 0.20 x_{12+j}.$$

For bundle a, the assignment logit is

$$\ell_a(x) = \sum_{j=1}^{6} a_j q_j(x) - 0.25|a| - 0.06|a|^2,$$

with softmax probability $\widetilde{e}_a(x)$. The actual propensity mixes this assignment with a uniform component,

$$e_a(x) = 0.8\, \widetilde{e}_a(x) + \frac{0.2}{64},$$

so every bundle has the global floor

$$e_a(x) \geq \eta = 0.003125.$$

No additional propensity clipping is applied.

Reference measure. For the Boolean experiments $B_0 = B_{\leq 2}$; in this subsection only, write $P^{\pi}_{\leq 2} = P^{\pi}_0$ and $Q^{\pi}_{\text{comp}} = Q^{\pi}_0$. All experiments use the uniform reference measure $\pi(a) = 1/64$ over the full feasible cube, so $W_\pi = 64^{-1} I_{64}$ and $P^{\pi}_{\leq 2}$ equals the ordinary Euclidean projector onto $\text{col}(B_{\leq 2})$. The resulting Q^{π}_{comp} is therefore fixed across regimes and seeds.

Complementary regimes. Let $s_j(a) = 2a_j - 1$. For R2, let $u \in \mathbb{R}^{64}$ have $u_{111000} = 1$, $u_{000111} = -1$, and zero elsewhere, and define the frozen projected, RMS-normalized rough direction

$$v = \frac{Q^{\pi}_{\mathrm{comp}} u}{\{64^{-1} \left\| Q^{\pi}_{\mathrm{comp}} u \right\|_2^2\}^{1/2}}.$$

For R3, let $\alpha_{\mathrm{R3}} \in \{0.5, 1, 2, 4\}$ denote the scalar amplitude used in H4; this is distinct from the treatment random variable A. The complementary surface is

$$\mu_{\mathrm{comp}}(x, a) = \begin{cases} 0, & \mathrm{R0}, \\ (0.6 + 0.25 \tanh x_1)\, 2^{-3} \prod_{j=1}^{3} s_j(a), & \mathrm{R1}, \\ (0.45 + 0.20 \tanh x_1)\, v_a, & \mathrm{R2}, \\ \alpha_{\mathrm{R3}}\, 2^{-3} \prod_{j=1}^{6} s_j(a), & \mathrm{R3}. \end{cases}$$

The observed outcome is

$$Y = \mu_{\leq 2}(X, A) + \mu_{\mathrm{comp}}(X, A) + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, 0.5^2).$$

The common counterfactual evaluation grid contains 5000 independent $\mathcal{N}(0, I_{20})$ covariate draws generated with seed 2027.

TSGS training and selection. For the final selector, the original training set has 3000 observations and is deterministically split into 2400 inner-fit and 600 inner-checkpoint observations. Neural fits use AdamW with learning rate 10^{-3}, weight decay 10^{-4}, batch size 256, and at most 200 epochs; factual checkpoint MSE is evaluated every five epochs and the minimizing checkpoint is retained. A training-only NoGraph pilot is selected on the inner checkpoint set and refit on all 3000 training observations before the five nested graph edge sets are constructed. On the full $m = 6$ cube, $(G_{50}, G_{75}, G_{90}, G_{95}, G_{100})$ retain respectively $(96, 144, 173, 182, 192)$ pilot-ranked edges; connectivity is not imposed and isolated nodes are allowed. Exact pilot-roughness ties are broken lexicographically by the integer-coded edge endpoints. Each of the 26 candidates is then checkpointed on the same training-only split and refit on the full 3000-observation training set for its selected epoch. A separate training-only MultiHotMLP outcome nuisance supplies $\widetilde{\mu}$ for AIPW scoring. The independent outer validation set has 3000 observations; in the synthetic experiments the score uses the exact DGP propensities. Candidate order is fixed as $(G_\emptyset, 0)$ first, followed by $G_{50}, G_{75}, G_{90}, G_{95}, G_{100}$ and, within each nonempty graph, $\lambda_G \in \{0.01, 0.03, 0.10, 0.30, 1.00\}$ in the displayed order; the first exact tie is retained. The minimum direct target-specific AIPW score is selected under this ordering. The selected actual graph edge set, λ_G, and training-selected epoch are retained, and the selected rule is refit once on the combined 6000 training-plus-validation observations for exactly that epoch, without a new checkpoint search or graph reconstruction.

Metrics and uncertainty. For an evaluation grid $X_1, \ldots, X_{n_{\mathrm{eval}}}$ and a contrast matrix $C \in \mathbb{R}^{q_C \times K}$,

$$\mathrm{PEHE}_C = \left[\frac{1}{n_{\mathrm{eval}} q_C} \sum_{i=1}^{n_{\mathrm{eval}}} \left\| C\{\widehat{\mu}(X_i) - \mu(X_i)\} \right\|_2^2 \right]^{1/2}.$$

Surface RMSE is

$$\mathrm{RMSE}_{\mathrm{surf}} = \left[\frac{1}{n_{\mathrm{eval}} K} \sum_{i=1}^{n_{\mathrm{eval}}} \sum_{a \in \mathcal{A}_f} \{\widehat{\mu}(X_i, a) - \mu(X_i, a)\}^2 \right]^{1/2}.$$

Deployment contrast MSE is the squared quantity inside the PEHE square root. The graph-sequence risk $\mathcal{R}_C(G, \lambda; h_0 \mid D_n)$ in Theorem 4.1 uses the unnormalized row sum; validation and empirical PEHE average over q_C, a fixed positive scaling that does not affect candidate ordering. Component RMSEs are the analogous grid-by-node RMSEs between the predicted and true structured or complementary components. If $\widehat{M}_{\mathrm{comp}} \in \mathbb{R}^{n_{\mathrm{eval}} \times K}$ stacks the predicted complementary surface over the evaluation grid, projection leakage is $\left\| B^{\mathsf{T}}_{\leq 2} W_\pi \widehat{M}^{\mathsf{T}}_{\mathrm{comp}} \right\|_F / (\left\| \widehat{M}_{\mathrm{comp}} \right\|_F + 10^{-12})$. Reported means and standard errors are computed across independent replicate seeds. Paired method differences are formed within seed before their standard error is computed. Figure whiskers marked "±SE" are one standard error, not confidence intervals.

PredictiveCV comparator. In the fresh R3 confirmation and estimated-propensity robustness blocks only, PredictiveCV uses the same 26 candidate rules and the training-selected epoch of each candidate. It selects the graph/penalty rule by unweighted factual MSE on the independent outer validation sample, then applies the same final train-plus-validation refit at the already selected epoch. The ACTG comparison is reported separately in Appendix F.15.

Final TSGS architecture and graph weights. The shared covariate encoder is $20 \to 128 \to 64$, with GELU after each affine layer. From the resulting 64-dimensional representation, the low-order head is a single affine map $64 \to 22$, and the complementary head is a single affine map $64 \to 64$ followed by the fixed right projection Q^{π}_{comp}; neither head has an additional hidden layer or post-head activation. The learner uses no normalization layer, dropout, residual connection, attention, or graph message passing. Candidate graphs therefore affect training only through the complementary-output penalty. Every retained undirected edge has unit weight, and for minibatch size b the implemented graph term is exactly

$$\frac{\lambda_G}{b|E_G|}\sum_{i=1}^{b}\sum_{(u,v)\in E_G}\{\widehat{\mu}_{\mathrm{comp}}(X_i,u)-\widehat{\mu}_{\mathrm{comp}}(X_i,v)\}^2.$$

For $G_\emptyset$, the graph term is zero. Pilot roughness determines graph membership only; it is not reused as an edge-specific weight.

F.2 External baseline specifications

The main R1/R2 comparison uses the same datasets, treatment ordering, counterfactual evaluation grid, and contrast operators for every method. All baseline hyperparameters are fixed across R1/R2 and fresh seeds; only prespecified validation-driven epoch or stopping rules depend on the fitted data.

MultiHotMLP. The factual input concatenates the 20-dimensional covariates and six-bit treatment vector. The scalar network is $26 \to 128 \to 64 \to 1$, with GELU after the two hidden affine layers and no dropout, normalization, residual connection, or output activation. Training minimizes factual MSE with AdamW, learning rate 10^{-3}, weight decay 10^{-4}, batch size 256, and a 200-epoch selection run. External-validation factual MSE is evaluated every five epochs and the minimum-MSE epoch is retained; there is no patience-based early stopping. The model is then reinitialized and refit on the combined 6000 training-plus-validation observations for exactly the selected epoch, with seed $s + 7100$ used for initialization and minibatch shuffling.

SharedBundleHead. A covariate-only encoder $20 \to 128 \to 64$ with GELU feeds a linear $64 \to 64$ bundle head; factual training gathers the coordinate corresponding to the observed bundle. It uses the same AdamW, learning-rate, weight-decay, batch-size, 200-epoch checkpoint, and 6000-observation selected-epoch refit protocol as MultiHotMLP, with seed $s + 7300$.

SCP. SCP follows the official two-stage single-cause perturbation procedure (Qian et al., 2021); its single development-frozen first-stage configuration uses representation widths $6/18$, batch size 50, learning rate 0.005, MMD bandwidth 1, and unit factual, propensity, and MMD weights, while the second-stage outcome regression uses batch size 100. This configuration is held fixed across confirmatory regimes and seeds.

SIN-vector. SIN-vector preserves the generalized Robinson construction of Kaddour et al. (2021) while using the six-bit treatment vector directly. Its four networks are $h(a) : 6 \to 100 \to 100 \to 100$, $e(x) : 20 \to 60 \to 60 \to 60 \to 100$, $g(x) : 20 \to 200 \to 200 \to 100$, and $m(x) : 20 \to 200 \to 200 \to 1$; ReLU follows each non-scalar output layer, while $m(x)$ has a linear scalar output. The prediction is

$$\widehat{\mu}(x,a) = \langle h(a) - e(x), g(x)\rangle + m(x).$$

All four networks use Adam without weight decay; the learning rate is 10^{-3} for h, e, and m, and 5×10^{-4} for g. Batch size is 500. Stage 1 fits $m(x)$ by outcome MSE. Stage 2 performs 15 outcome updates of h and g per minibatch with e and m fixed, followed by one propensity-representation update of e toward $h(a)$. Each stage allows at most 50,000 epochs, validates every epoch, and stops after 11 consecutive non-improvements. The current state at stopping is retained (the best state is not restored), and no train-plus-validation refit is performed. SIN-vector uses replicate seed s.

VSR. We use a paper-based reimplementation of Zou et al. (2020). The three-stage procedure fits (i) a treatment VAE with latent dimension three, (ii) a latent density-ratio classifier that supplies observation weights, and (iii) a weighted factual outcome model. The VAE uses a $6 \to 6$ encoder followed by three-dimensional mean and log-scale heads and a sigmoid decoder; unspecified VAE training details follow the VSR implementation distributed with the official SCP repository rather than original VSR author code. The density classifier is $(d+3) \to 10 \to 10 \to 2$ with ELU activations. The outcome model uses separate $x \to 10$ and $a \to 5$ branches, concatenates them, and maps $15 \to 10 \to 1$ with ELU. The frozen protocol uses the paper's latent dimension and synthetic-regression outcome architecture, with no tuning against counterfactual truth. One fitted model is evaluated at all 64 bundles. Confirmatory weights remain close to one: effective sample size is 2995.98 ± 1.58 in the original synthetic comparison and 2999.95 ± 0.03 in ACIC, with no clipping or truncation.

AIPW outcome nuisance. The selector's training-only nuisance uses the same $26 \to 128 \to 64 \to 1$ MultiHotMLP architecture, factual-MSE objective, and AdamW settings as the external MultiHotMLP baseline, but a different data protocol: seed $s + 4000$ defines a deterministic $2400/600$ split of the original 3000 training observations for epoch selection, after which the nuisance is reinitialized and refit on all 3000 training rows for the selected epoch. The independent outer validation sample is not used to fit or checkpoint the nuisance; it is used only after nuisance fitting to form the AIPW validation targets.

All reported methods produce absolute potential-outcome surfaces and are evaluated by the same C_{edge}, C_{sq}, and surface metrics.

F.3 H1: RARE SUPPORTED BUNDLES

The eight fixed target bundles are

$$000000,\ 001001,\ 010010,\ 011011,\ 100100,\ 101101,\ 110110,\ 111111,$$

which induce 48 distinct incident geometry edges.

Controlled support intervention. For each seed s and target index t, H1 pools that target's positions across the original training and validation samples, permutes them once with seed $s + 8000 + t$, and retains the first $r \in \{1, 2, 5, 8\}$ positions. The retained sets are therefore nested across r. All other observations of the target are removed from whichever of the two fitting pools they occupy, while every non-target observation is unchanged. The factual test sample and the 5000-point counterfactual evaluation grid are never thinned. The Natural condition uses the untouched training and validation samples. Rare-node RMSE is surface RMSE restricted to the fixed eight target bundles. Rare-edge PEHE is evaluated on every geometry edge with at least one endpoint in that target set, giving 48 edges.

Table 5: Full H1 rare-support results.

Support	Method	Rare-node RMSE ↓	Rare-edge PEHE ↓
1	Projected-NoGraph	0.614 ± 0.028	0.586 ± 0.029
1	Projected-Full	0.385 ± 0.008	0.280 ± 0.005
2	Projected-NoGraph	0.608 ± 0.021	0.573 ± 0.022
2	Projected-Full	0.383 ± 0.004	0.280 ± 0.002
5	Projected-NoGraph	0.551 ± 0.018	0.498 ± 0.017
5	Projected-Full	0.381 ± 0.006	0.280 ± 0.004
8	Projected-NoGraph	0.510 ± 0.015	0.464 ± 0.013
8	Projected-Full	0.384 ± 0.009	0.279 ± 0.005
Natural	Projected-NoGraph	0.391 ± 0.008	0.338 ± 0.011
Natural	Projected-Full	0.353 ± 0.006	0.260 ± 0.006

F.4 H2: ROUGH-TRANSITION BORROWING

Jump partition and synthetic references. Let $\mathcal{J} = \{\texttt{000111}, \texttt{111000}\} = \{7, 56\}$. A geometry edge (u, v) is a jump edge exactly when one, but not both, endpoints lie in $\mathcal{J}$. This gives 12 jump

Table 6: H1 rare-edge bias–variance decomposition (MSE scale).

Support	Method	Bias^2	Variance	MSE	PEHE
1	Projected-NoGraph	0.249323	0.097372	0.346695	0.588808
1	Projected-Full	0.037466	0.040991	0.078458	0.280103
2	Projected-NoGraph	0.226798	0.103311	0.330108	0.574550
2	Projected-Full	0.037098	0.041150	0.078248	0.279728
5	Projected-NoGraph	0.157545	0.091586	0.249131	0.499130
5	Projected-Full	0.037004	0.041636	0.078641	0.280429
8	Projected-NoGraph	0.124775	0.090928	0.215703	0.464438
8	Projected-Full	0.037262	0.040472	0.077733	0.278807
Natural	Projected-NoGraph	0.036849	0.077611	0.114460	0.338319
Natural	Projected-Full	0.028626	0.039337	0.067963	0.260698

edges and 180 non-jump edges. Projected-Oracle removes precisely those 12 jump edges and retains all 180 non-jump edges. Projected-Wrong instead retains all 12 jump edges and 168 non-jump edges, removing the lexicographically first 12 non-jump edges:

$$(0,1),(0,2),(0,4),(0,8),(0,16),(0,32),$$
$$(1,3),(1,5),(1,9),(1,17),(1,33),(2,3).$$

Projected-Full, Projected-Oracle, and Projected-Wrong select λ_G independently from the same five-value grid

$$\lambda_G \in \{0.01, 0.03, 0.10, 0.30, 1.00\}$$

using the same 96-mask AIPW diagnostic reference protocol. This diagnostic protocol is not the practical TSGS selector. Projected-NoGraph fixes $\lambda_G = 0$, while TSGS jointly selects from its frozen 26-rule direct-AIPW family described above.

For each of the 96 frozen oriented masks, the target endpoint is removed from both masked inner-fit and inner-checkpoint subsets. The masked candidate is fit and checkpointed on those subsets, while the independent external validation sample is used only for the target-specific AIPW edge score. Scores are averaged across masks for each penalty, and the minimum selects the reference penalty before the usual unmasked selected-epoch final refit. This reference construction is used for Projected-Full, Projected-Oracle, and Projected-Wrong.

Table 7: Full H2 performance summary over independent seeds 201–210.

Method	Edge PEHE	Jump PEHE	Non-jump PEHE	Square PEHE	Surface RMSE	Comp. RMSE
Projected-NoGraph	0.3740 ± 0.0036	0.6327 ± 0.0096	0.3499 ± 0.0039	0.4657 ± 0.0066	0.3865 ± 0.0031	0.1889 ± 0.0022
Projected-Full	0.4141 ± 0.0143	1.0802 ± 0.0860	0.3203 ± 0.0027	0.5342 ± 0.0236	0.4049 ± 0.0074	0.2176 ± 0.0106
TSGS	0.3292 ± 0.0039	0.7219 ± 0.0384	0.2820 ± 0.0084	0.3670 ± 0.0054	0.3702 ± 0.0028	0.1478 ± 0.0024
Projected-Oracle	0.3147 ± 0.0035	0.7656 ± 0.0341	0.2564 ± 0.0051	0.3539 ± 0.0057	0.3566 ± 0.0026	0.1429 ± 0.0027
Projected-Wrong	0.4341 ± 0.0183	1.2049 ± 0.0995	0.3185 ± 0.0019	0.5667 ± 0.0294	0.4147 ± 0.0096	0.2318 ± 0.0131

F.5 H3: PROJECTION ABLATIONS

Tables 9–10 report component-recovery diagnostics and predictive/effect metrics for the low-order-only, projected, and unprojected variants. Low-order and complementary RMSE are the component RMSEs defined above; Leakage is the normalized weighted structured-span moment defined above.

F.6 H4: TARGET-DEPENDENT BORROWING

The original H4 amplitude landscape is an oracle analysis based on synthetic truth risk. The subsequent practical R3 selector confirmation uses external-validation AIPW or factual PredictiveCV scores instead.

The mean-risk minima are $G_{100}, \lambda_G = 1$ versus $G_{75}, \lambda_G = 1$ at $\alpha_{\mathrm{R3}} = 0.5$; $G_{50}, \lambda_G = 1$ versus $G_{50}, \lambda_G = 0.03$ at $\alpha_{\mathrm{R3}} = 1$; $G_{100}, \lambda_G = 0.01$ for both targets at $\alpha_{\mathrm{R3}} = 2$; and $G_{\emptyset}, \lambda_G = 0$ for both targets at $\alpha_{\mathrm{R3}} = 4$ (edge versus square, respectively).

Table 8: H2 edge-effect bias–variance decomposition on seeds 201–210 (MSE scale).

Method	Edge group	Bias^2	Variance	MSE
Projected-NoGraph	All	0.03581	0.10417	0.13998
Projected-NoGraph	Jump	0.22259	0.17849	0.40108
Projected-NoGraph	Non-jump	0.02336	0.09921	0.12257
Projected-Full	All	0.08499	0.08832	0.17331
Projected-Full	Jump	0.97799	0.25527	1.23326
Projected-Full	Non-jump	0.02545	0.07719	0.10264
TSGS	All	0.03790	0.07064	0.10854
TSGS	Jump	0.34448	0.18996	0.53444
TSGS	Non-jump	0.01747	0.06268	0.08015
Projected-Oracle	All	0.04364	0.05550	0.09914
Projected-Oracle	Jump	0.43820	0.15837	0.59657
Projected-Oracle	Non-jump	0.01733	0.04864	0.06597
Projected-Wrong	All	0.10411	0.08733	0.19144
Projected-Wrong	Jump	1.26723	0.27382	1.54105
Projected-Wrong	Non-jump	0.02657	0.07490	0.10147

Table 9: H3 component-recovery diagnostics.

Regime	Variant	Leakage	Low-order RMSE	Complementary RMSE
R0	LowOrderOnly	NA	0.318399 ± 0.003855	0.000000 ± 0.000000
R0	Unprojected-NoGraph	0.210857 ± 0.006806	0.381251 ± 0.006387	0.251607 ± 0.006637
R0	Projected-NoGraph	$(1.25 \pm 0.05) \times 10^{-8}$	0.334576 ± 0.003483	0.152472 ± 0.005087
R0	Unprojected-Full	0.082570 ± 0.003687	0.324516 ± 0.005438	0.100394 ± 0.003051
R0	Projected-Full	$(6.44 \pm 0.28) \times 10^{-9}$	0.322715 ± 0.001895	0.062380 ± 0.001097
R1	LowOrderOnly	NA	0.322093 ± 0.004359	0.077274 ± 0.000000
R1	Unprojected-NoGraph	0.210698 ± 0.006423	0.380493 ± 0.006300	0.251892 ± 0.006183
R1	Projected-NoGraph	$(1.37 \pm 0.06) \times 10^{-8}$	0.334472 ± 0.003544	0.153236 ± 0.005244
R1	Unprojected-Full	0.080905 ± 0.003420	0.324662 ± 0.005792	0.115629 ± 0.002751
R1	Projected-Full	$(6.54 \pm 0.27) \times 10^{-9}$	0.323730 ± 0.003003	0.086768 ± 0.000790
R2	LowOrderOnly	NA	0.399565 ± 0.005137	0.465593 ± 0.000000
R2	Unprojected-NoGraph	0.320884 ± 0.008602	0.456214 ± 0.004143	0.361559 ± 0.008343
R2	Projected-NoGraph	$(3.89 \pm 0.11) \times 10^{-8}$	0.343392 ± 0.004999	0.183865 ± 0.003439
R2	Unprojected-Full	0.287402 ± 0.007448	0.435541 ± 0.003372	0.331490 ± 0.007169
R2	Projected-Full	$(3.71 \pm 0.13) \times 10^{-8}$	0.342985 ± 0.004929	0.182106 ± 0.002334

F.7 True-propensity R3 selector confirmation at $\alpha_{\text{R3}} = 1$

The independent selector-confirmation block uses seeds 201–210 and the same 26-rule candidate family for both AIPW and PredictiveCV selectors. Within-family regret is the selected rule's deployment contrast MSE minus the minimum deployment contrast MSE over the same 26-rule family; it is computed from truth only after practical selection and used only as an evaluation diagnostic.

Edge-AIPW has lower regret in $5/10$ seeds, so the edge comparison is a near-tie. Square-AIPW has lower regret in $10/10$ seeds. Mean Spearman correlation with oracle validation risk is 0.897 ± 0.063 for Edge-AIPW versus 0.950 ± 0.028 for PredictiveCV, and 0.444 ± 0.131 for Square-AIPW versus 0.139 ± 0.112 for PredictiveCV. Edge- and square-AIPW select different graph/λ_G rules in $8/10$ seeds. Using the edge-selected rule for square effects yields deployment regret 0.03575 ± 0.00938, whereas using the square-selected rule for edge effects yields 0.00656 ± 0.00274.

F.8 Selector choices and runtime

In R2, TSGS selects G_{90} in nine seeds and G_{50} in one; $\lambda_G = 0.30$ and 1.00 are each selected in five seeds. Selected epochs range from 20 to 30 with mean 24. Across 18 complete R1/R2 timing records, the practical selector, including candidate selection and final refit, takes 54.763 ± 0.782 seconds per

Table 10: H3 supporting predictive/effect metrics.

Regime	Variant	Edge PEHE	Square PEHE	Surface RMSE
R0	LowOrderOnly	0.221132 ± 0.003730	0.140625 ± 0.003640	0.318399 ± 0.003855
R0	Unprojected-NoGraph	0.317711 ± 0.008755	0.363372 ± 0.011991	0.361341 ± 0.004035
R0	Projected-NoGraph	0.330870 ± 0.007560	0.393248 ± 0.011047	0.367780 ± 0.004439
R0	Unprojected-Full	0.242828 ± 0.006408	0.192435 ± 0.004600	0.327833 ± 0.004300
R0	Projected-Full	0.246566 ± 0.004030	0.202456 ± 0.001853	0.328697 ± 0.001885
R1	LowOrderOnly	0.250673 ± 0.002922	0.200208 ± 0.002386	0.331239 ± 0.004241
R1	Unprojected-NoGraph	0.318175 ± 0.008414	0.363990 ± 0.011551	0.361335 ± 0.003892
R1	Projected-NoGraph	0.331449 ± 0.007736	0.393944 ± 0.011328	0.368009 ± 0.004545
R1	Unprojected-Full	0.255917 ± 0.004916	0.219654 ± 0.003630	0.333061 ± 0.004694
R1	Projected-Full	0.260806 ± 0.005635	0.229336 ± 0.002837	0.335161 ± 0.002982
R2	LowOrderOnly	0.769375 ± 0.002679	1.055276 ± 0.001268	0.613587 ± 0.003332
R2	Unprojected-NoGraph	0.352890 ± 0.003552	0.424038 ± 0.005804	0.376850 ± 0.002166
R2	Projected-NoGraph	0.371742 ± 0.006544	0.457504 ± 0.008228	0.389548 ± 0.005568
R2	Unprojected-Full	0.350943 ± 0.003368	0.419813 ± 0.005049	0.376174 ± 0.002087
R2	Projected-Full	0.369447 ± 0.005553	0.453008 ± 0.006223	0.388344 ± 0.005228

Table 11: H4 target dependence over the frozen R3 amplitude grid. Regrets are MSE-scale mean $\pm$ SE.

α_{R3}	Different optima	Graph differs	λ_G differs	Edge–square Spearman	Edge regret using square opt.	Square regret using edge opt.
0.5	$4/5$	$4/5$	$0/5$	0.95884 ± 0.02620	0.000807 ± 0.000447	0.001931 ± 0.000751
1.0	$4/5$	$4/5$	$4/5$	0.01183 ± 0.21455	0.005298 ± 0.002635	0.015852 ± 0.004726
2.0	$1/5$	$1/5$	$1/5$	0.98332 ± 0.00658	0.000018 ± 0.000018	0.000490 ± 0.000490
4.0	$0/5$	$0/5$	$0/5$	0.99631 ± 0.00144	0	0

dataset. The synthetic H2 mechanism-reference computations are excluded from practical TSGS timing.

F.9 Constrained feasible spaces and structural extrapolation

To separate full-cube geometry from observational support, the stress test restricts the feasible set to

$$\mathcal{A}_f = \{a \in \{0,1\}^6 : |a| \le 4\}, \qquad |\mathcal{A}_f| = 57,$$

with the uniform reference measure on these 57 feasible bundles. The induced Hamming-1 geometry has 156 edges. In the fully supported setting, $\mathcal{A}_s = \mathcal{A}_f$. In the support-boundary setting,

$$\mathcal{A}_s = \{a : |a| \le 3\}, \qquad |\mathcal{A}_s| = 42,$$

so 15 feasible weight-four bundles are absent from training and validation; 60 geometry edges cross from supported to unsupported feasible nodes. Candidate construction and AIPW selection use supported observations only. Errors on the unsupported shell are computed from synthetic truth after selection and are interpreted as structural extrapolation.

Table 14: Support-boundary stress test. "Shell" denotes the 15 feasible but unsupported weight-four bundles; these columns evaluate structural extrapolation, not identified causal effects.

Regime	Method	Supported-edge PEHE	Shell RMSE	Boundary-edge PEHE	Surface RMSE
R1	Projected-NoGraph	0.3134	0.6763	0.6239	0.4531
R1	Projected-Full	0.2512	0.4414	0.3207	0.3514
R1	TSGS	0.2507	0.4467	0.3282	0.3529
R1	MultiHotMLP	**0.1731**	**0.3209**	**0.2169**	**0.2799**
R2	Projected-NoGraph	0.3399	0.7891	0.7561	0.5072
R2	Projected-Full	0.3402	0.8274	0.8044	0.5240
R2	TSGS	**0.2921**	**0.6253**	**0.5742**	**0.4317**
R2	MultiHotMLP	0.3948	0.9659	1.0192	0.6050

Table 12: R3 $\alpha_{\mathrm{R3}} = 1$ practical selector results. Risk and regret are on the MSE scale.

Target	Selector	PEHE	Risk	Within-family regret
Edge	Edge-AIPW	0.33315 ± 0.00412	0.11114 ± 0.00278	0.00384 ± 0.00110
Edge	PredictiveCV	0.33333 ± 0.00383	0.11124 ± 0.00258	0.00394 ± 0.00091
Square	Square-AIPW	0.42615 ± 0.00802	0.18218 ± 0.00660	0.01056 ± 0.00329
Square	PredictiveCV	0.47278 ± 0.00259	0.22358 ± 0.00247	0.05196 ± 0.00443

Table 13: Constrained feasible space with full support. Means over seeds 301–310.

Regime	Method	Edge PEHE	Surface RMSE
R1	Projected-NoGraph	0.3368	0.3600
R1	Projected-Full	0.2609	0.3272
R1	TSGS	0.2607	0.3272
R1	MultiHotMLP	**0.1817**	**0.2727**
R2	Projected-NoGraph	0.3711	0.3744
R2	Projected-Full	0.3758	0.3760
R2	TSGS	**0.3283**	**0.3522**
R2	MultiHotMLP	0.7247	0.5409

Table 15: **Constrained-space matched Full-Val checks.** PEHE, mean $\pm$ SE over seeds 301–310. "All feasible edges" evaluates the fully supported 57-node feasible space; "Supported edges only" evaluates identified edges in the support-boundary setting. Δ = TSGS $-$ Full-Val. These PEHE values are not pooled with full-cube deployment MSE.

Evaluation	Regime	Full-Val	TSGS	Δ
All feasible edges	R1	0.26086 ± 0.00375	0.26073 ± 0.00381	$-1.344e-04 \pm 2.255e-04$
All feasible edges	R2	0.37715 ± 0.00797	0.32830 ± 0.00420	-0.048855 ± 0.008074
Supported edges only	R1	0.25124 ± 0.00227	0.25070 ± 0.00234	$-5.451e-04 \pm 2.607e-04$
Supported edges only	R2	0.34134 ± 0.00396	0.29213 ± 0.00381	-0.049205 ± 0.003888

F.10 Estimated-propensity robustness

The practical robustness check replaces the exact synthetic propensity in validation by a training-only 64-class neural propensity model. The network maps X through $20 \to 128 \to 64 \to 64$ logits with GELU, is selected by training-only checkpoint loss, and is refit on all 3000 training observations before candidate scoring. No clipping, calibration, truncation, or stabilization is applied. Across seeds 401–410, coordinate propensity RMSE averages 0.00888 and $\mathrm{KL}(e \,\|\, \widehat{e})$ averages 0.1102; the mean p99 inverse weight is 392.7. All weights are finite.

Table 16: Deployment performance with true and estimated propensities. Entries are means over seeds 401–410; regret is relative to the best candidate in the same frozen family.

Regime	Target	Selector	PEHE	Regret
R2	Edge	TrueProp-AIPW	0.33843	0.00514
R2	Edge	EstimatedProp-AIPW	0.33865	0.00535
R2	Square	TrueProp-AIPW	0.38087	0.01115
R2	Square	EstimatedProp-AIPW	0.37898	0.00997
R3, $\alpha_{\mathrm{R3}} = 1$	Edge	TrueProp-AIPW	0.33040	0.00335
R3, $\alpha_{\mathrm{R3}} = 1$	Edge	EstimatedProp-AIPW	0.32989	0.00300
R3, $\alpha_{\mathrm{R3}} = 1$	Edge	PredictiveCV	0.33469	0.00610
R3, $\alpha_{\mathrm{R3}} = 1$	Square	TrueProp-AIPW	0.42209	0.00791
R3, $\alpha_{\mathrm{R3}} = 1$	Square	EstimatedProp-AIPW	0.42166	0.00761
R3, $\alpha_{\mathrm{R3}} = 1$	Square	PredictiveCV	0.47418	0.05382

Estimated- and true-propensity selectors need not choose the identical candidate: near-equivalent rules can exchange rank while deployment risk remains stable. The result is therefore evidence for the tested nuisance-estimation regime, not for arbitrary propensity misspecification or poor overlap.

F.11 Matched full-graph tuning and target-transfer reanalyses

The matched Full-Val and cross-target analyses reuse existing validation scores and final-refit evaluations without additional model fitting. Full-Val restricts the frozen family to $(G_\emptyset, 0)$ and the five G_{100} penalties, then uses the same validation target as TSGS; TSGS retains access to all 26 rules. Full-cube tables report deployment contrast MSE, while the constrained-space check reports PEHE separately and is not pooled with those risks.

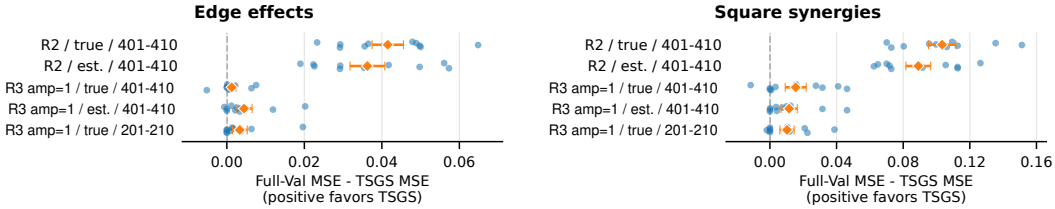

Figure 4: **Matched full-graph tuning.** Seed-wise Full-Val minus TSGS deployment contrast MSE; diamonds are means ± one paired SE, so positive favors TSGS. Using the same target-specific validation score, Full-Val searches NoGraph and five G_{100} penalties; TSGS searches all 26 rules. R3 uses $\alpha_{\mathrm{R3}} = 1$; robustness strata share fits. Tables 17–18 give values.

Table 17: **Matched Full-Val comparison for edge contrasts.** Deployment contrast MSE, mean ± SE over each ten-seed block; all R3 rows use $\alpha_{\mathrm{R3}} = 1$. NoGraph is the exact zero-penalty rule; Full-Val selects among NoGraph and the five G_{100} penalties; TSGS selects among all 26 rules. Δ = TSGS − Full-Val, so negative values favor TSGS.

Block	NoGraph	Full-Val	TSGS	Δ
R2 / true / 401–410	0.1530±0.0043	0.1562±0.0053	0.1147±0.0024	−0.04152±0.00404
R2 / estimated / 401–410	0.1530±0.0043	0.1511±0.0040	0.1149±0.0030	−0.03627±0.00450
R3 / true / 401–410	0.1272±0.0060	0.1106±0.0027	0.1094±0.0030	−0.00126±0.00113
R3 / estimated / 401–410	0.1272±0.0060	0.1134±0.0018	0.1090±0.0030	−0.00443±0.00209
R3 / true / 201–210	0.1306±0.0056	0.1145±0.0030	0.1111±0.0028	−0.00335±0.00192

Table 18: **Matched Full-Val comparison for square contrasts.** Deployment contrast MSE, mean ± SE over each ten-seed block; all R3 rows use $\alpha_{\mathrm{R3}} = 1$. NoGraph is the exact zero-penalty rule; Full-Val selects among NoGraph and the five G_{100} penalties; TSGS selects among all 26 rules. Δ = TSGS − Full-Val, so negative values favor TSGS.

Block	NoGraph	Full-Val	TSGS	Δ
R2 / true / 401–410	0.2379±0.0081	0.2487±0.0111	0.1454±0.0044	−0.10336±0.00823
R2 / estimated / 401–410	0.2379±0.0081	0.2333±0.0077	0.1442±0.0063	−0.08907±0.00745
R3 / true / 401–410	0.1900±0.0114	0.1945±0.0131	0.1790±0.0080	−0.01547±0.00630
R3 / estimated / 401–410	0.1900±0.0114	0.1903±0.0110	0.1787±0.0083	−0.01157±0.00487
R3 / true / 201–210	0.1937±0.0099	0.1924±0.0100	0.1822±0.0066	−0.01026±0.00428

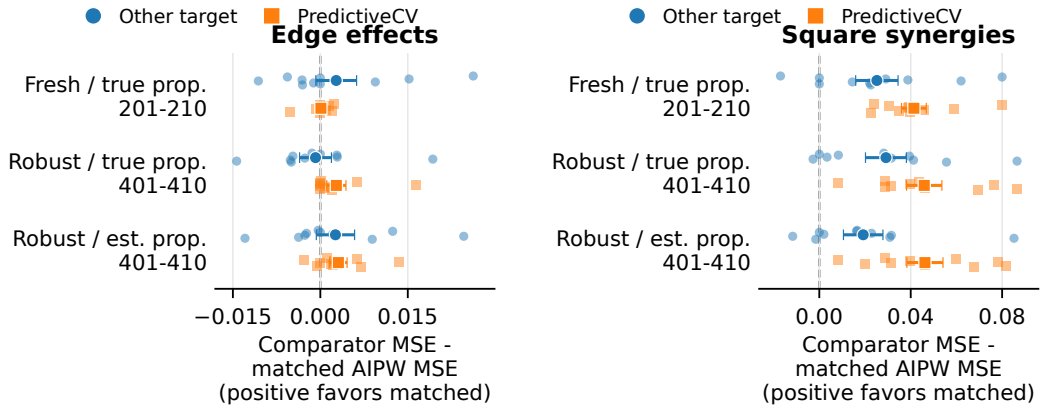

Figure 5: **Target-aligned selection at R3** $\alpha_{R3} = 1$**.** Each point is a seed-wise deployment contrast-MSE gap, comparator minus target-matched AIPW selection; diamonds show the mean and whiskers one paired SE, not a 95% confidence interval. Positive values favor the matched selector. "Other target" is Square-AIPW for edge evaluation and Edge-AIPW for square evaluation. The fresh block uses seeds 201–210; the true- and estimated-propensity robustness blocks use seeds 401–410, share candidate fits and PredictiveCV choices, and are not independent samples.

Table 19: **Fresh R3** $\alpha_{R3} = 1$ **cross-target risk matrix.** Rows identify the validation selector and columns the evaluation target. Entries are final-refit deployment contrast MSE, mean ± SE over seeds 201–210. The candidate is fixed by the row selector before evaluation in either column.

Selector	Edge MSE	Square MSE
Edge-AIPW	0.1111±0.0028	0.2074±0.0059
Square-AIPW	0.1139±0.0044	0.1822±0.0066
PredictiveCV	0.1112±0.0026	0.2236±0.0025

Table 20: **R3** $\alpha_{R3} = 1$ **edge deployment contrast MSE.** Mean ± SE over the stated ten-seed block. Robustness true- and estimated-propensity rows share candidate fits and PredictiveCV choices and are not independent samples.

Block	Edge-AIPW	Square-AIPW	PredictiveCV
Fresh / true / 201–210	0.1111±0.0028	0.1139±0.0044	0.1112±0.0026
Robust / true / 401–410	0.1094±0.0030	0.1085±0.0032	0.1121±0.0021
Robust / estimated / 401–410	0.1090±0.0030	0.1116±0.0028	0.1121±0.0021

Table 21: **R3** $\alpha_{R3} = 1$ **square deployment contrast MSE.** Mean ± SE over the stated ten-seed block. Robustness true- and estimated-propensity rows share candidate fits and PredictiveCV choices and are not independent samples.

Block	Edge-AIPW	Square-AIPW	PredictiveCV
Fresh / true / 201–210	0.2074±0.0059	0.1822±0.0066	0.2236±0.0025
Robust / true / 401–410	0.2082±0.0091	0.1790±0.0080	0.2249±0.0026
Robust / estimated / 401–410	0.1979±0.0095	0.1787±0.0083	0.2249±0.0026

Table 22: **R3** $\alpha_{\mathrm{R3}} = 1$ **target-transfer gaps.** Comparator minus target-matched AIPW deployment contrast MSE, mean $\pm$ paired SE. Positive values favor matched selection. For edge evaluation, other target is Square-AIPW; for square evaluation, it is Edge-AIPW. Bars in Figure 5 show one paired SE, not 95% confidence intervals.

Block	Target	Other target	PredictiveCV
Fresh / true / 201–210	Edge	0.00272 ± 0.00350	0.00010 ± 0.00065
Fresh / true / 201–210	Square	0.02519 ± 0.00927	0.04140 ± 0.00542
Robust / true / 401–410	Edge	-0.00083 ± 0.00273	0.00275 ± 0.00164
Robust / true / 401–410	Square	0.02917 ± 0.00895	0.04592 ± 0.00772
Robust / estimated / 401–410	Edge	0.00259 ± 0.00328	0.00310 ± 0.00149
Robust / estimated / 401–410	Square	0.01921 ± 0.00864	0.04622 ± 0.00791

Table 23: **ACIC semi-synthetic stress test.** Mean ± SE over seeds 501–510; TSGS uses the same frozen estimated-propensity selector in both regimes.

	Smooth			Rough		
Method	Edge PEHE ↓	Square PEHE ↓	Surface RMSE ↓	Edge PEHE ↓	Square PEHE ↓	Surface RMSE ↓
TSGS	0.1069±0.0021	0.1030±0.0019	0.1309±0.0026	0.1518±0.0071	0.1656±0.0074	0.1583±0.0067
Projected-NoGraph	0.1270±0.0035	0.1460±0.0056	0.1392±0.0029	**0.1485±0.0082**	0.1662±0.0094	0.1562±0.0069
Projected-Full	0.1070±0.0019	0.1035±0.0019	0.1309±0.0024	0.1508±0.0071	**0.1640±0.0072**	0.1578±0.0068
MultiHotMLP	**0.0942±0.0005**	0.0971±0.0003	0.1357±0.0069	0.3305±0.0010	0.3374±0.0008	0.2992±0.0021
SharedBundleHead	0.1458±0.0032	0.2004±0.0055	0.1293±0.0018	0.1548±0.0021	0.2068±0.0035	**0.1389±0.0013**
SCP	0.0979±0.0008	0.09753±0.00003	0.1276±0.0020	0.3281±0.0023	0.3346±0.0016	0.2943±0.0022
SIN-vector	0.1933±0.0087	0.1370±0.0065	0.3779±0.0142	0.2769±0.0093	0.2398±0.0077	0.4074±0.0082
VSR	0.0945±0.0014	**0.0956±0.0005**	**0.1202±0.0027**	0.2709±0.0112	0.2843±0.0098	0.2670±0.0044

F.12 ACIC SEMI-SYNTHETIC REAL-COVARIATE STRESS TEST

Covariates and split. We draw 11,000 distinct rows from the ACIC 2018/LBIDD covariates and fix a $3000/3000/5000$ train/validation/test split across replicates. Only the identifier is removed. Continuous missing values are imputed by training medians, categorical missingness receives an explicit level, categorical variables are one-hot encoded, and continuous variables are standardized using training statistics. The resulting covariate dimension is $d = 452$.

Assignment. A fixed random map forms $z(x) = \tanh(W_0 x/\sqrt{d}) \in \mathbb{R}^{32}$. Component and pair scores are

$$s_j(x) = \tanh(w_j^\mathsf{T} z/\sqrt{32}), \qquad r_{jk}(x) = \tanh(v_{jk}^\mathsf{T} z/\sqrt{32}),$$

and the six-bit bundle logit is

$$\ell_a(x) = 0.45 \sum_j a_j s_j(x) + 0.10 \sum_{j<k} a_j a_k r_{jk}(x) - 0.15|a| - 0.03|a|^2,$$

followed by a 64-way softmax. There is no propensity floor or clipping.

Outcome teacher. The outcome surface is generated independently of the TSGS projection and borrowing graph. Let $t(a) = 2a - 1$, and define a two-layer treatment network

$$r_1(a) = \tanh(W_1 t(a)/\sqrt{6} + b_1) \in \mathbb{R}^{48}, \qquad r_2(a) = \tanh(W_2 r_1(a)/\sqrt{48} + b_2) \in \mathbb{R}^{32},$$

with $g_x = \tanh(W_g z/\sqrt{32} + b_g)$. For fixed teacher coefficient vectors $c_0, c_1, c_2, g_R \in \mathbb{R}^{32}$, let $b_0(x) = 0.40 c_0^\mathsf{T} z/\sqrt{32}$. The Smooth surface is

$$\mu_S(x, a) = b_0(x) + \frac{0.80}{\sqrt{32}} c_1^\mathsf{T}\{g_x \odot r_2(a)\} + \frac{0.40}{\sqrt{32}} c_2^\mathsf{T} r_2(a).$$

Rough adds

$$\gamma(x) J(a), \qquad J(a) = \mathbf{1}(a_1 = a_2 = a_3 = 1) - \mathbf{1}(a_4 = a_5 = a_6 = 1),$$

where $\gamma(x) = 0.70 + 0.25 \tanh(g_R^\mathsf{T} z/\sqrt{32})$. All teacher matrices and coefficient vectors are drawn once from DGP seed 20260911 and then fixed; observation noise is $\mathcal{N}(0, 0.5^2)$. The fraction of centered outcome energy outside the degree-≤ 2 span is 5.52% in Smooth and 13.45% in Rough.

Selection behavior and nuisance diagnostics. The TSGS encoder is mechanically adapted to the 452-dimensional input, and the primary selector uses a training-estimated 64-class propensity model. Smooth selects G_{100} for the edge target in $8/10$ seeds and for the square target in $8/10$, with $\lambda_G = 1$ in all edge runs and $9/10$ square runs. In Rough, edge selection is G_{100} in six seeds, G_{50} in two, and $G_\emptyset$ in two; square selection is G_{100} in seven, G_{50} in one, and $G_\emptyset$ in two, with all nonzero penalties at most 0.30. Exact edge–square rule agreement is $8/10$ in both regimes. Propensity coordinate RMSE is 0.003478 ± 0.000078 and $\mathrm{KL}(e \,\|\, \widehat{e}) = 0.023939 \pm 0.001115$; the inverse-weight maximum is 279.1 ± 15.8 with no stabilization.

On the 48 prespecified Rough edges where $J(a)$ changes, NoGraph has PEHE 0.1843 ± 0.0098, Full 0.1935 ± 0.0102, and TSGS 0.1951 ± 0.0101. On the remaining 144 edges, the corresponding values are 0.1342 ± 0.0084, 0.1330 ± 0.0072, and 0.1336 ± 0.0072. These labels are used only after selection for mechanism diagnosis.

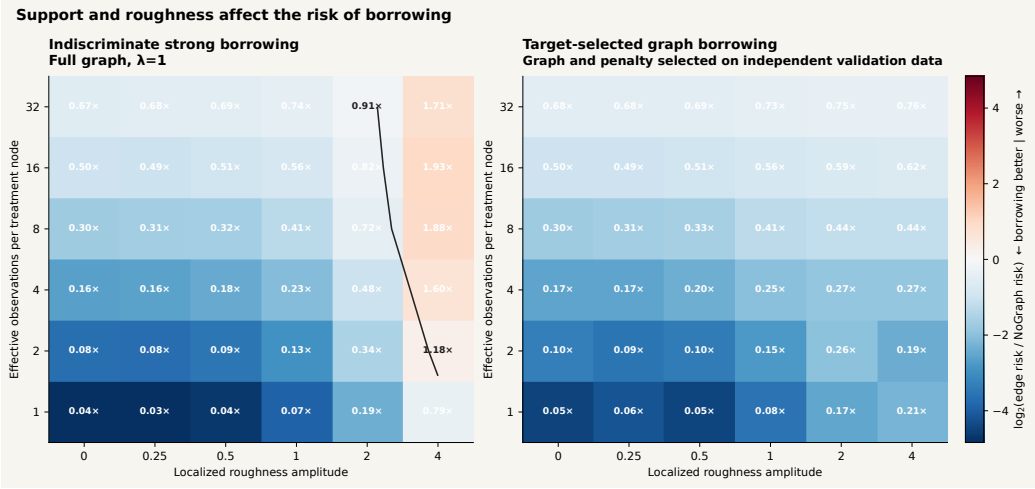

Figure 6: **Support and roughness affect the risk of borrowing.** Each cell reports mean edge risk relative to NoGraph over 40 independent repetitions. Left: fixing Full at $\lambda = 1$ crosses the parity contour in the high-information, high-roughness region. Right: target-selected TSGS reduces smoothing strength or removes edges and has lower mean risk than NoGraph throughout the tested grid.

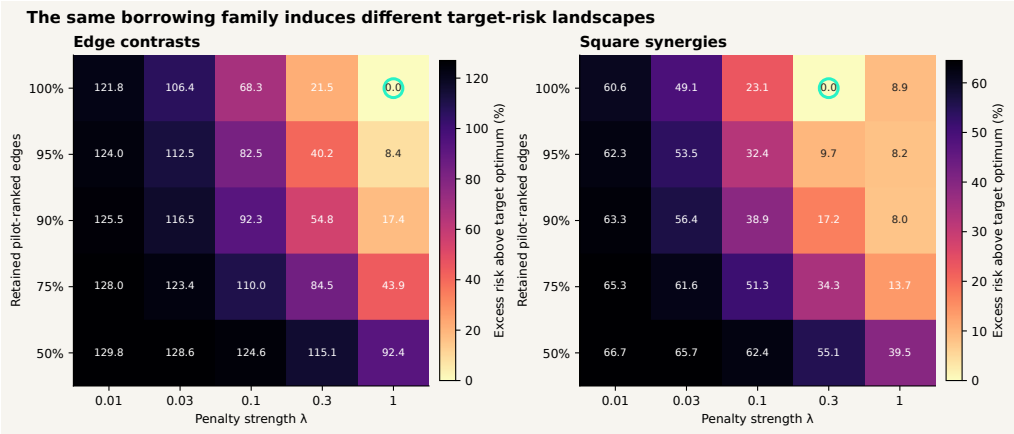

Figure 7: **The same borrowing family induces different target-risk landscapes.** Entries are percentage excess risk over the best mean candidate. Mint rings mark the lowest mean-risk candidate for each target.

F.13 Projected graph-sequence experiments

This suite instantiates exactly the complementary graph-sequence model of Section 4, independently of the neural benchmark tables. A fixed root seed of 20260916 generated all observations. The run comprised 1,440 phase-diagram repetitions, 120 target-reversal repetitions, 1,920 scaling repetitions, and 200 bias–variance repetitions. Six numerical checks were used to verify Walsh orthonormality, annihilation of all degree-≤ 2 modes, $L_H u_S = 2|S|u_S$, the fast Walsh transform, complement preservation, and output integrity.

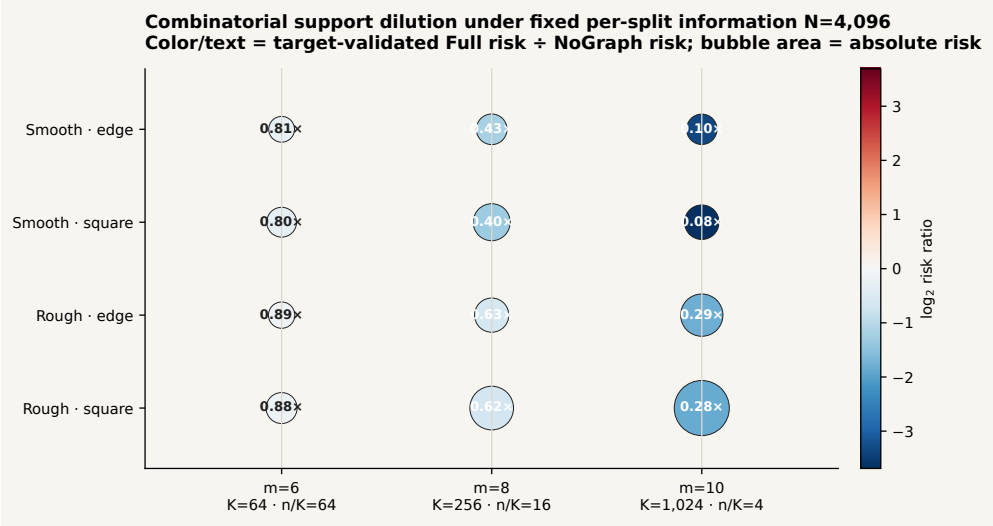

Figure 8: **Combinatorial support dilution at fixed information per split.** Color and printed values give selected-Full risk divided by NoGraph risk; bubble area gives absolute risk. As 2^m grows under fixed $N = 4096$, information per node falls and the variance advantage of borrowing increases in both smooth and localized-rough regimes.

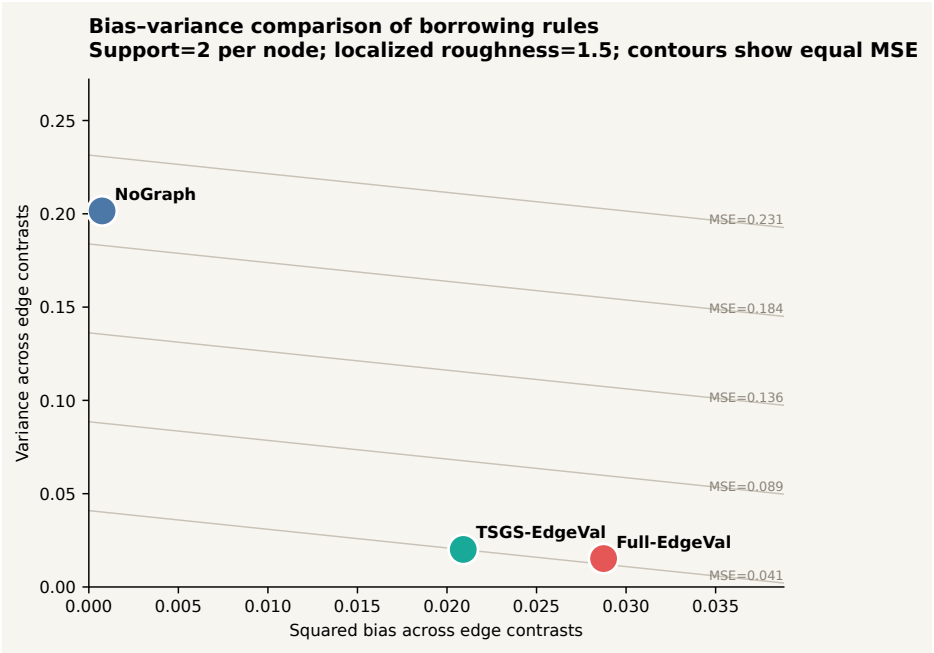

Figure 9: **Bias–variance comparison of borrowing rules.** At support two and roughness 1.5, NoGraph has negligible squared bias but large variance; Full sharply reduces variance at the cost of bias; TSGS has lower MSE in this setting. Contours show equal total MSE.

F.14 CLOSEST-WORK ADAPTIVE-FUSION BENCHMARK

We compare TSGS with the adaptive all-pairs fusion objective underlying the treatment-clustering approach of Ma et al. (2022). This experiment uses the exact projected Gaussian sequence model, so every method estimates the same complementary surface $h = U\beta$ on the $m = 6$ cube. For projected training coefficients z and common effective support n_{eff}, the comparator solves

$$\widehat{\beta} = \operatorname*{arg\,min}_{\beta \in \mathbb{R}^{d_{\text{comp}}}} \left\{ \frac{n_{\text{eff}}}{2} \|\beta - z\|_2^2 + \lambda_f \sum_{1 \le u < v \le K} \omega_{uv} \left| (U\beta)_u - (U\beta)_v \right| \right\},$$

$$\omega_{uv} \leftarrow \min\left\{ 20, \frac{1}{\max\{|h_u - h_v|, 10^{-8}\}} \right\}.$$

Each frozen-weight convex problem is solved by ADMM, followed by at most 40 adaptive-weight updates. The fusion grid is $\lambda_f \in \{0, 10^{-5}, 3{\times}10^{-5}, 10^{-4}, 3{\times}10^{-4}, 10^{-3}, 3{\times}10^{-3}, 10^{-2}, 3{\times}10^{-2}\}$. TSGS uses graph densities $\{0.25, 0.50, 0.75, 0.90, 1\}$ and $\lambda \in \{0.01, 0.03, 0.10, 0.30, 1, 3\}$, together with NoGraph. Both families are selected by the same independent target-specific validation loss and evaluated on the same truth, comparing borrowing strategies under a shared data and evaluation protocol.

The design crosses smooth versus localized-rough truth with $n_{\text{eff}} \in \{2, 16\}$ and edge versus square targets. Every cell uses 120 paired train/validation/test draws generated from seed root 43109. Under equal treatment counts and an intercept-only calibration vector, the calibration weights in Zhu et al. (2025) have a uniform minimum-dispersion solution. On the same projected parameter space, their weighted quadratic loss is proportional to $n_{\text{eff}} \|\beta - z\|_2^2$, up to an additive constant. The benchmark here uses the adaptive all-pairs fusion penalty specified above.

Across the eight cells, TSGS beats NoGraph eight times and AdaptiveFusion six times. Its largest separation occurs under rough scarce support: edge and square risks are 0.206 and 0.201 of NoGraph, versus 0.599 and 0.617 for AdaptiveFusion. All 3,840 positive-penalty frozen-weight convex fits converged; 97.7% of adaptive paths also met the outer tolerance before the prespecified 40-update cap. The maximum normalized KKT residual was 6.90×10^{-5}, the 99th percentile was 2.35×10^{-5}, and the maximum relative primal–dual gap was 8.38×10^{-5}. AdaptiveFusion attains the lowest mean risk among the compared methods in the smooth scarce-support cells, whereas TSGS has lower mean risk than AdaptiveFusion in all rough cells and the smooth ample-support cells.

Table 24: **Matched closest-work benchmark over 120 paired repetitions per cell.** Entries are mean target risk divided by mean NoGraph risk, with paired 95% bootstrap intervals; lower is better.

Truth	Support / node	Target	FullGraph	AdaptiveFusion	TSGS
Smooth	2	Edge	0.058 [0.056, 0.061]	**0.056** [0.054, 0.058]	0.065 [0.060, 0.070]
Smooth	2	Square	0.051 [0.044, 0.059]	**0.044** [0.038, 0.054]	0.053 [0.047, 0.058]
Smooth	16	Edge	**0.299** [0.289, 0.311]	0.462 [0.439, 0.486]	0.300 [0.290, 0.311]
Smooth	16	Square	**0.246** [0.238, 0.255]	0.330 [0.309, 0.354]	0.246 [0.238, 0.255]
Rough	2	Edge	0.659 [0.641, 0.678]	0.599 [0.571, 0.631]	**0.206** [0.192, 0.221]
Rough	2	Square	0.668 [0.646, 0.690]	0.617 [0.578, 0.661]	**0.201** [0.185, 0.217]
Rough	16	Edge	0.932 [0.916, 0.947]	0.847 [0.826, 0.869]	**0.493** [0.472, 0.515]
Rough	16	Square	0.932 [0.912, 0.953]	0.837 [0.817, 0.859]	**0.464** [0.442, 0.486]

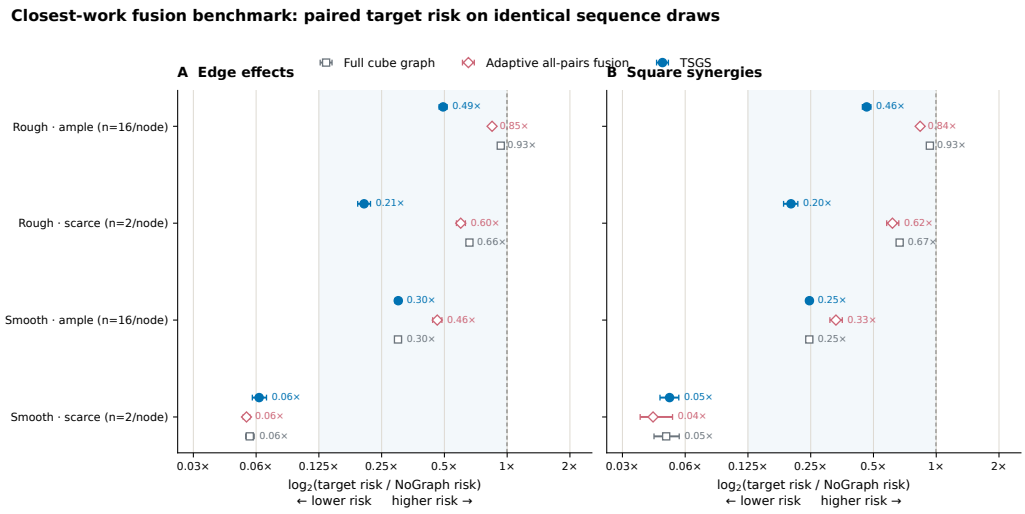

Figure 10: **Target risk against the closest fusion baseline.** Points show ratios to NoGraph; whiskers are paired 95% bootstrap intervals. Adaptive fusion has the lowest mean risk among these methods in the smooth scarce-support cells. TSGS has lower mean risk than adaptive fusion in all four rough cells and both smooth ample-support cells.

F.15 ACTG 175 PROTOCOL AND COMPLETE SELECTOR COMPARISON

For the four nonrectangular arms, let $b_0 = \mathbf{1}_4/2$, let $U \in \mathbb{R}^{4\times 3}$ have orthonormal columns spanning $\mathbf{1}_4^{\perp}$, and let d_x denote the outcome-feature dimension. Every arm-by-feature coefficient matrix $H \in \mathbb{R}^{4\times d_x}$ has the unique form $H = b_0\gamma_0^{\mathsf{T}} + U\Gamma^{\mathsf{T}}$, with $\gamma_0 \in \mathbb{R}^{d_x}$ and $\Gamma \in \mathbb{R}^{d_x\times 3}$. Since $L_G b_0 = 0$, the graph quadratic is

$$\operatorname{tr}(H^{\mathsf{T}}L_G H) = \operatorname{tr}\{\Gamma(U^{\mathsf{T}}L_G U)\Gamma^{\mathsf{T}}\} = \operatorname{tr}\{(U^{\mathsf{T}}L_G U)\Gamma^{\mathsf{T}}\Gamma\},$$

and contains no γ_0 term. Thus the implemented $[b_0, U]$ parameterization is exactly TSGS with structured dictionary $B_0 = \mathbf{1}_4$: the shared treatment intercept is unpenalized and only its rank-three contrast complement is smoothed.

The public data are taken from `BART` version 2.9.10; the source archive has SHA256 digest `a70b64ef465e9fa6f267e1269ca9be14d39642f4c55361d57141a7d37f371e84`. Continuous baseline variables are standardized from the training fold and expanded with squares and hyperbolic tangents; binary variables and treatment-history strata enter directly. Every rule uses the same feature matrix, the explicit $[\mathbf{1}_4/2, U]$ treatment basis, and the same ridge floor. Only the nested edge set and complementary graph penalty vary. The true randomized propensity $1/4$ and one training-only NoGraph outcome nuisance are shared by every candidate in both validation and untouched test scoring.

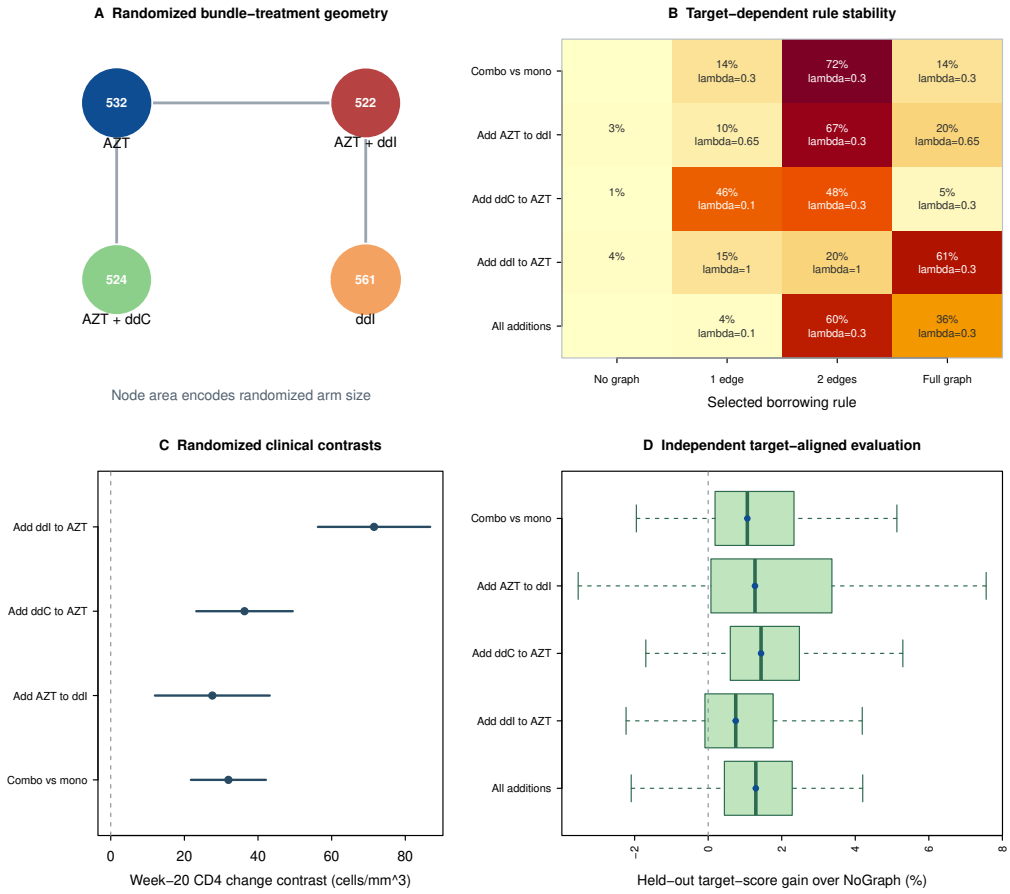

Figure 11: **Target-dependent borrowing in ACTG 175.** **A**, the randomized four-arm bundle geometry; node area encodes arm size. **B**, selection frequency and median penalty across 100 stratified splits. **C**, randomization-based week-20 CD4 contrasts with 95% confidence intervals. **D**, distribution of held-out target-score gain over NoGraph; blue points mark medians.

Table 25: **ACTG 175 held-out target scores over 100 repeated stratified random splits.** Lower is better. The difference compares TSGS with NoGraph within split; SE is the across-split standard error of that paired difference, and the last column is the percentage of splits won by TSGS.

Target	TSGS	PredictiveCV	NoGraph	TSGS–NoGraph	SE	Win %
All local additions	108,987	109,171	110,487	$-1,500$	137	86
Add ddI to AZT	111,961	112,354	112,764	-803	192	70
Add ddC to AZT	88,479	88,535	89,899	$-1,420$	141	89
Add AZT to ddI	126,627	126,626	128,798	$-2,172$	293	76
Combination vs monotherapy	53,922	53,973	54,674	-753	99	79

PredictiveCV nearly matches TSGS on the individual clinical targets, and both selectors have lower mean held-out contrast scores than NoGraph across the reported targets.

F.16 Formal verification in Lean 4

The supplementary `LeanProofs` project formalizes uniform risk ordering over every zero-sum contrast, its equivalence to nonnegativity of the projected error-moment difference, the zero-sum rank-one reversal, Proposition 4.1, and the affine-versus-pairwise Boolean finite-difference identities. It was

checked with Lean 4.33 and mathlib using both `lake build` and the strict `check-proof.sh` harness. The project contains no `sorry`, `admit`, custom `axiom`, or `unsafe`; `#print axioms` reports only the standard mathlib foundations `propext`, `Classical.choice`, and `Quot.sound`. The probabilistic results are proved analytically in Appendix E; the Lean project checks the algebraic statements listed above.

G Boolean Finite-Difference Algebra

For $S \subseteq [m]$, suppose $a_j = 0$ for $j \in S$ and every vertex of the corresponding Boolean face is feasible. The higher-order finite difference is

$$\Delta_S \mu(x, a) = \sum_{B \subseteq S} (-1)^{|S|-|B|} \mu(x, a + \mathbf{1}_B), \qquad \mathbf{1}_B = \sum_{j \in B} \mathbf{e}_j.$$

On the complete Boolean cube every real-valued function admits a unique Möbius expansion

$$f(a) = \sum_{S \subseteq [m]} \beta_S \prod_{j \in S} a_j, \qquad \beta_S = \sum_{B \subseteq S} (-1)^{|S|-|B|} f(\mathbf{1}_B).$$

The coefficient for $S = \{j, k\}$ equals the baseline square-face finite difference. Möbius inversion motivates the Boolean interaction hierarchy; on constrained feasible designs, the reference-weighted projection—not raw Möbius degree alone—defines the unique low-order/complementary split used by the learner.