

# Synthetic Control Method under Interference: Detecting and Correcting Bias

Joao Alipio-Correa\*

April 15, 2026

## Abstract

The synthetic control method estimates the effect of an intervention by comparing a treated unit to a weighted average of untreated donors, and it is valid only when no donor is affected by the treatment. The donors that receive the most weight are those whose pre-treatment paths most closely follow the treated unit, and that closeness usually reflects the shared economic and geographic structure through which a policy spills over. The method therefore concentrates its weight on the donors most vulnerable to contamination. We characterize the resulting bias and show that it enters only through the total weight the estimator places on exposed donors, under an arbitrary exposure mapping rather than the linear factor structure assumed in earlier work. This result motivates a randomization test that screens for interference before estimation. The test asks whether the pre-to-post changes in donor outcomes trace a proximity pattern around the treated unit, requires no model of the outcome, and is exact in finite samples. We then develop three corrections based on covariate rescaling and ridge augmentation of the weights. Each draws weight away from exposed donors using pre-treatment information alone, and each trades a controlled amount of pre-treatment fit for a reduction in bias. In simulation the corrections remove much of the bias and, under the most flexible variant, nearly all of it, while preserving pre-treatment fit. Across four canonical applications the test finds no interference in the original donor pools. This null is informative only when a pool contains both near and far donors. Once the German reunification study of Abadie, Diamond, and Hainmueller is widened from its original sample to every country with complete data, the test detects spillovers, and the corrected estimates change the study's conclusion.

---

\*For helpful comments and suggestions, I thank Michael Colaresi, José Antonio Cheibub, and Jude Hays, as well as discussants and participants at PolMeth, CODE@MIT, the NYU Data Science Frontiers Workshop, MPSA, MapleMeth, the Comparative Politics Reading Group, and the Pitt Politics Workshop. All remaining errors are my own.

Ph.D. student in Political Science and M.S. student in Statistics at the University of Pittsburgh.  
Contact: [jac736@pitt.edu](mailto:jac736@pitt.edu); <https://joaalipiocorrea.github.io>.

# 1 Introduction

Political scientists have spent decades documenting how policies travel across subnational borders: state lotteries spread as border residents cross to play and their governments move to recapture the leaked revenue (Berry & Berry, 1990; Berry & Baybeck, 2005); welfare benefits are set with one eye on the neighbors’ (Volden, 2002); and local antismoking ordinances stall where smokers can take their business across the boundary (Shipan & Volden, 2008). Hundreds of studies record the same regularity across policy domains (Graham et al., 2013), and it cuts directly against the design logic of comparative case studies. Synthetic control (SC) methods (Abadie, 2021; Abadie & Gardeazabal, 2003; Abadie et al., 2010, 2015) construct a counterfactual for a treated unit by combining a small set of donor units—typically neighboring states, nearby countries, or peers in a known network—whose pre-intervention trajectories reproduce that of the treated unit. The features that make a donor a good match also make it vulnerable to interference: shared economic structure, similar institutions, and correlated exposure to common shocks arise from the same connections that transmit spillovers. The method’s matching logic therefore favors the donors most likely to be exposed to spillovers. The SC literature has largely treated these two concerns—fit and interference—in isolation.

SC is uniquely exposed to this tension because contamination enters deterministically through the weight channel: once a donor  $j$  carries positive weight  $w_j$  and receives spillover intensity  $\rho_j$ , the aggregate contamination mass  $\sum_{j \in \mathcal{E}} \rho_j w_j$  (with  $\mathcal{E}$  the set of exposed donors) passes directly into the estimated effect. This paper makes three contributions. First, we decompose the bias of the SC estimator under interference into a single contamination channel that depends only on the weight mass assigned to exposed donors. This extends Cao and Dowd (2019) from a linear factor structure to an exposure-mapping framework and makes the contamination channel explicit and actionable for design. Second, and most importantly, we propose a Fisherian ring diagnostic: a finite-sample exact randomization test that detects whether pre-to-post outcome changes in the donor pool exhibit a proximity pattern centered on the treated unit. The diagnostic requires minimal assumptions, consumes no post-treatment outcome information beyond the change statistic, and gives practitioners a concrete decision rule before estimation. Third, we develop three weight-geometry corrections—covariate rescaling, constrained ridge, and unconstrained ridge—that act on the contamination channel identified in the decomposition by incorporating spatial exposure directly into the SC optimization and use only pre-treatment information. The diagnostic is the primary novel element; the decomposition motivates it, and the corrections respond to it.

In simulation, baseline SC exhibits attenuation bias that scales with  $\sum_{j \in \mathcal{E}} \rho_j w_j$ , and the three corrections reduce bias by 25–100% across the  $(\tau, \rho)$  grid—from partial reduction under constrained ridge to near-complete elimination under the unconstrained variant—while largely preserving pre-treatment fit. The ring diagnostic attains nominal size at  $\alpha = 0.05$  under the sharp null and, at the recommended three-ring partition, power exceeding 0.95 for  $\rho \geq 0.2$  in medium-density donor graphs ( $|\mathcal{E}| = 4$ ); in denser graphs the attainable power is capped near 0.82 by the exact test’s permutation-count ceiling rather than by the signal-to-noise ratio. Turning to real data, across four canonical SC applications (Abadie & Gardeazabal, 2003; Abadie et al., 2015; Ben-Michael et al., 2021; Kikuta, 2020) the diagnostic finds no detectable interference in the original donor pools. The pattern reverses once the donor pool widens. Expanding the German reunification study (Abadie et al., 2015) to a broad candidate set of roughly 150 countries yields a diagnostic rejection ( $\hat{p} = 0.017$  under the default ring contrast), and corrected estimates that change the substantive conclusion.

Our contribution sits at the intersection of four active literatures. Closest is Cao and Dowd (2019), who analyze SC bias under spillovers in a linear factor model; our decomposition reproduces their mechanical insight under a more general exposure-mapping primitive and pairs it with a design-based test and corrections they do not develop. On SC inference, Cattaneo et al. (2021)

construct prediction intervals that account for both in-sample fit error and out-of-sample forecast uncertainty. Chernozhukov et al. (2021) develop conformal procedures that deliver valid coverage under weak assumptions on the residual process. Abadie and Vives-i-Bastida (2025) survey placebo-based inference and document its sensitivity to donor-pool composition. We do not extend these inference frameworks to the corrected estimators—a limitation we return to in the conclusion. The diagnostic itself is finite-sample exact by construction. Di Stefano and Mellace (2020) propose an inclusive SCM that models spillovers by embedding exposed donors with a structural spillover term; our approach is design-based—we diagnose before estimating and modify weight geometry rather than specify a spillover model. Penalized SC variants (Abadie & L’Hour, 2021; Ben-Michael et al., 2021; Doudchenko & Imbens, 2016) regularize weights toward interpolation or toward an outcome model; here the penalty is shaped by an exposure mapping rather than by proximity in predictor space or by an auxiliary regression, and Ferman (2021) shows that such regularization choices materially affect SC behavior when donors are many and periods long.

On interference, we adopt the exposure-mapping formalism of Hudgens and Halloran (2008) and Aronow and Samii (2017) and connect it to Sävje et al. (2021), who study average treatment effects under unknown interference. Viviano (2024) studies policy targeting under network interference, where the design question is allocation rather than counterfactual construction; our focus is the latter under the same exposure-mapping primitive. For difference-in-differences designs, Butts (2021) and Clarke (2017) address spatial spillovers in DiD; the SC-specific weight-geometry channel studied here does not appear in that panel-generic setting, where averaging across many controls mutes individual-unit contamination.

SC is more vulnerable to interference than the designs it is often compared to. In difference-in-differences, contamination from any single control is diluted by averaging across the untreated panel. In matching, leverage is spread across multiple comparators per treated unit. In regression discontinuity, identification rests on local contrasts that are approximately symmetric around the cutoff, so proximate spillovers affect both sides alike. SC differs from all three: it places large, concentrated weights on a handful of donors, the set of consequential donors is typically sparse and geographically proximate (making exposure correlated with weight), and there is no symmetry device to cancel the contamination term.

The remainder of the paper is organized as follows. Section 2 introduces the exposure-mapping setup and states the bias decomposition. Section 3 develops the ring diagnostic and randomization test. Sections 4–6 present the three weight-geometry corrections. The remaining sections report simulation evidence and revisit four canonical applications, including the expanded German reunification design, before concluding with practical guidance.

## 2 Spatial/Network Dependence and SUTVA in Synthetic Control

### 2.1 Synthetic control versus other designs under interference

Causal analyses of comparative case designs rest on the *Stable Unit Treatment Value Assumption* (SUTVA), which asserts that unit  $i$ ’s potential outcome under a joint assignment  $(z_1, \dots, z_{J+1})$  depends only on its own assignment  $z_i$ , and that there are no hidden versions of treatment (Holland, 1986; Rubin, 1980). Formally, SUTVA restricts  $Y_i(z_1, \dots, z_{J+1})$  to  $Y_i(z_i)$  and guarantees a single well-defined potential outcome for each treatment level. Together with consistency ( $Y_i^{\text{obs}} = Y_i(z_i)$ ), this delivers the identification bridge that causal designs exploit: when a donor  $j$  is untreated ( $z_j = 0$ ), its observed outcome equals its no-treatment potential outcome,  $Y_{jt}^{\text{obs}} = Y_{jt}(0)$ , regardless of the assignments received by other units.

Hence untreated units can supply information about the *missing* no-treatment path of a treated unit,  $Y_{1t}(0)$ , by enabling cross-unit substitution of observed outcomes for potential outcomes. The

same invariance underpins placebo and permutation procedures: randomization or placebo-in-space tests rely on the distributional equivalence induced by SUTVA to compare treated units to re-labeled controls without introducing assignment-dependent distortions. Absent SUTVA, observed donor outcomes would in general be  $Y_{jt}(0; z_{-j})$  and could vary with other units’ assignments, severing the link  $Y_{jt}^{\text{obs}} = Y_{jt}(0)$  that comparative estimators require and undermining the causal interpretation of cross-unit contrasts.

When policies propagate through geography or networks, outcomes may depend on neighbors’ assignments and, as discussed above, SUTVA is violated. In such cases interference must be addressed explicitly, whether through exposure mappings that compress the assignment vector into lower-dimensional summaries of relevant neighbors (Aronow & Samii, 2017; Hudgens & Halloran, 2008; Manski, 2013), through randomization-based testing strategies (Bowers et al., 2017; Rosenbaum, 2007), or through epidemiological and statistical approaches that study identification, estimands, and sensitivity analysis under interference (e.g., Sävje et al., 2021; Tchetgen Tchetgen & VanderWeele, 2012; VanderWeele et al., 2014). For synthetic control (SC), these issues are especially acute: because post-treatment donor outcomes enter directly into the construction of the counterfactual, even partial exposure of donors introduces bias into the estimator. Interference is thus not a secondary complication but a direct pathway through which SC can be distorted—a point we develop formally further below.

Within this broader landscape, synthetic control (SC) is attractive because it enables transparent, design-based comparisons in comparative case studies with few units (Abadie, 2021; Abadie & Gardeazabal, 2003; Abadie et al., 2010, 2015). To see why SUTVA plays a distinctive role in SC, recall how the estimator is constructed. Let  $X_1 \in \mathbb{R}^K$  denote the treated unit’s pre-treatment predictors and  $X_0 \in \mathbb{R}^{K \times |\mathcal{J}|}$  the same predictors for donors  $\mathcal{J} \subseteq \{2, \dots, J+1\}$ . SC selects weights  $w \in \Delta = \{w \geq 0 : \mathbf{1}^\top w = 1\}$  to minimize the pre-period discrepancy  $\|X_1 - X_0 w\|_V^2$  with  $V \succ 0$  diagonal, and then applies the same convex weights to donor outcomes in the post-treatment period. This design pays for minimal trend structure with a crucial reliance on SUTVA: because the post-period donor outcomes enter the synthetic predictor, they must correspond to units with no direct treatment and no indirect exposure. If donors are even partially exposed, the synthetic counterfactual absorbs part of the treated effect and the estimated impact is distorted.<sup>1</sup> This mechanism is the synthetic-control manifestation of the bias channel emphasized in linear-factor settings when post-period shocks correlate with the weights (Ferman & Pinto, 2021).

The nature of this distortion can be summarized in a compact bias expression. For post-treatment periods  $t > T_0$ , the SC estimator based on observed donors admits the decomposition

$$\hat{\tau}_t = [Y_{1t}(0) - Y_{1t}^{\text{SC}}(0)] + \tau_t - \rho \tau_t \sum_{j \in \mathcal{E}} w_j, \quad (1)$$

where the three terms correspond respectively to (i) a *SUTVA synthetic mismatch* between the treated unit’s no-treatment outcome and its synthetic approximation, (ii) the *direct treatment effect*  $\tau_t$ , and (iii) a *contamination mass* equal to the treated effect multiplied by the weight assigned to exposed donors. This decomposition highlights that even if the synthetic perfectly reproduces the treated unit’s no-treatment path under SUTVA,  $\hat{\tau}_t$  will still differ from  $\tau_t$  whenever positively

---

<sup>1</sup>The typical case is attenuation: the synthetic estimate is pulled toward the treated outcome when spillovers raise donor outcomes in the same direction as the treatment effect, yielding an underestimate of the true effect. However, the bias need not always be attenuating. If spillovers move donor outcomes in the opposite direction of the treated effect, the synthetic may overshoot, producing an amplified estimate of the effect. This possibility highlights that interference does not merely add noise but can fundamentally alter the direction and magnitude of the estimated effect (as we will exhibit with our replications)

weighted donors are exposed. A simplified version makes this transparent:

$$\hat{\tau}_t = \tau_t - \rho \tau_t \sum_{j \in \mathcal{E}} w_j. \quad (2)$$

Equation (1) will be derived formally in the subsection below, but we present it here to build intuition: SC estimates equal the true effect minus a weighted spillover term, with the magnitude of bias governed jointly by the diffusion intensity  $\rho$ , the treatment effect  $\tau_t$ , and the weight mass placed on exposed donors.

Equation (1) highlights the mechanism through which spillovers distort synthetic control: any exposure among positively weighted donors feeds directly into the estimated effect. To appreciate the severity of this channel, it is useful to situate SC relative to other common designs. All causal estimators rely on SUTVA to justify using untreated outcomes as counterfactuals, but the manner in which donor outcomes enter differs sharply across designs, and so too does the impact of violations.

In difference-in-differences, for example, the control mean aggregates outcomes across many untreated units, assigning each only modest influence. Contamination of a few controls typically attenuates the estimate, but the effect is diluted across the pool and becomes material only if spillovers are widespread. Matching designs are more localized: bias arises when a treated unit is paired with an exposed control, but the damage is confined to that pair and does not propagate automatically to others. Regression discontinuity designs present yet another contrast: because identification exploits outcomes just above and below the cutoff, spillovers would bias the design only if they differentially affect the two sides of the threshold; broader or symmetric diffusion often cancels out.

Synthetic control stands apart. The estimator compresses the counterfactual for the treated unit into a single convex combination of donors, often dominated by a handful of high-weight contributors chosen precisely for their ability to replicate pre-treatment trajectories. This sparsity is central to SC’s appeal: by concentrating on a few well-aligned donors, the method achieves close balance without imposing parametric trend restrictions. Yet the same feature that delivers this balance under SUTVA makes the design acutely vulnerable under violations. If even one heavily weighted donor is exposed to spillovers, the contamination is not averaged away, confined to a pair, or offset symmetrically; it flows directly into the counterfactual in proportion to the assigned weight. In this setting, interference is not averaged out, localized, or offset symmetrically: it enters deterministically through the weights.

A single exposed donor can shift the estimated effect in proportion to its assigned weight, and the donors most capable of reproducing the treated trajectory—the ones the optimization favors—are frequently those most plausibly exposed. In this sense, SC transforms what might be a modest attenuation in DiD or a pair-specific distortion in matching into a direct and systematic bias channel. The contamination mass term in equation (1) formalizes this: the post-treatment estimate depends not only on the direct effect but also on the weighted share of the treated effect leaking into the synthetic. For this reason, SC does not merely inherit SUTVA as a background condition for identification—it is structurally entangled with it. Donor outcomes after treatment are the building blocks of the counterfactual, so once they are contaminated, bias enters mechanically through the weights. This is not a secondary complication but a first-order vulnerability. Later sections formalize this point: the bias decomposition in equation (1) shows that the estimated effect consists of the true effect plus a contamination mass proportional to the total weight placed on exposed donors. Effective corrections must therefore operate directly on this pathway, reshaping the weight structure that transmits interference into the estimate.

In practice, one response to the risk of spillovers is to restrict the donor pool by trimming proximate or substantively linked units, or to adapt estimators designed for settings with multiple

treated units (e.g., Cavallo et al., 2013; Firpo & Possebom, 2018; Kreif et al., 2016; Robbins & Saunders, 2017; Xu, 2017). Let  $\mathcal{J}^*$  denote the full donor set admissible under standard SC assumptions and  $\mathcal{J} \subseteq \mathcal{J}^*$  the subset retained after trimming. Because  $\mathcal{J}$  is a strict subset, the optimization problem that defines the synthetic weights is solved over a reduced feasible region. While trimming may lower the exposed weight mass  $\sum_{j \in \mathcal{E}} w_j$  and thereby reduce the contamination term in equation (1), it simultaneously constrains the construction of  $Y_{1t}^{SC}(0)$  to a smaller space of convex combinations.

The consequence is that pre-treatment imbalance  $\|X_1 - X_0 w\|_V$  cannot improve and will typically worsen, leading to a larger design-driven discrepancy  $Y_{1t}(0) - Y_{1t}^{SC}(0)$ . In finite samples this trade-off manifests concretely: fewer admissible donors reduce the effective dimensionality of the comparison set, variance of the estimator increases, and the credibility of the counterfactual trajectory is weakened when the synthetic fails to reproduce the treated unit’s pre-treatment path. In short, trimming addresses one bias pathway by construction but aggravates another, and the resulting estimator may be both noisier and harder to interpret. Regularized and augmented variants of SC (Abadie & L’Hour, 2021; Arkhangelsky et al., 2021; Athey et al., 2021; Ben-Michael et al., 2021; Doudchenko & Imbens, 2016; Gobillon & Magnac, 2016) partially stabilize estimation, but because they continue to rely on post-period donor outcomes, they inherit the same interference sensitivity unless exposure is incorporated directly into the design.

These considerations motivate the approach we take in the remainder of the paper. We operationalize interference using an outcome-agnostic, monotone exposure mapping based on geography (or a known network) and show that SC’s sensitivity to spillovers can be reduced by acting directly on the object through which contamination flows—the weights. In the next section we employ a randomization-based diagnostic that tests for proximity-patterned changes in donor outcomes around the treated unit, and subsequent sections incorporate the exposure mapping into estimation through covariate rescaling and ridge augmentations that shift mass away from likely exposed donors while preserving pre-period fit and, when desired, the convex-combination interpretation that makes SC useful for applied work.

## 2.2 Formalizing the distortion caused by interference

SUTVA’s locus inside SCM is immediate: for  $\hat{Y}_{1t}^{(0)}$  to represent  $Y_{1t}(0)$ , each positively weighted donor must be observed under no direct treatment *and* no indirect exposure to the treated unit. Abadie’s overview emphasizes this point in practice: researchers often curate donor pools to avoid neighboring or economically linked regions precisely to protect the counterfactual from spillovers (Abadie, 2021). Beyond such design heuristics, a growing literature addresses spillovers in SCM and factor models, estimating both direct and indirect effects (e.g., Cao & Dowd, 2019; Di Stefano & Mellace, 2020). Parallel developments in spatial econometrics likewise model spillovers parametrically (e.g., SLX or spatial-lag specifications) and have been combined with difference-in-differences to separate direct and indirect effects, underscoring the importance of interdependence in social sciences applications (Delgado & Florax, 2015; Vega & Elhorst, 2015).

To study the consequences of interference in a transparent way, we adopt a stylized diffusion approach that isolates the bias channel in SCM. The construction turns the problem on three quantities—the treated-effect magnitude  $\tau_t$ , which sets the size of the policy shock; the diffusion intensity  $\rho_j \in [0, 1]$ , the fraction of the treated effect transmitted to each donor  $j$ ; and the exposed weight mass  $\sum_{j \in \mathcal{E}} w_j$ , which measures how much of the synthetic counterfactual rests on donors subject to spillovers. This decomposition clarifies that contamination arises not from a complex interaction of dynamics but from the mechanical way in which treated effects, diffusion intensity, and donor weights combine.

Formally, let the treated unit be indexed by  $i = 1$  and donors by  $j \in \mathcal{J}$ . For a post-treatment period  $t > T_0$ , the observed outcome for the treated unit is

$$Y_{1t}^{obs} = Y_{1t}(0) + \tau_t,$$

where  $Y_{1t}(0)$  is the no-treatment potential outcome and  $\tau_t$  is the direct effect. For donors, we allow localized, time-homogeneous diffusion with heterogeneous intensities. Partition the donor set  $\mathcal{J} = \mathcal{E} \sqcup \mathcal{S}$  into an exposed set  $\mathcal{E}$  and a safe set  $\mathcal{S}$ . For  $j \in \mathcal{E}$ , the observed outcome is

$$Y_{jt}^{obs} = Y_{jt}(0) + \rho_j \tau_t, \quad \rho_j \in [0, 1], \quad (3)$$

while for  $j \in \mathcal{S}$ ,  $Y_{jt}^{obs} = Y_{jt}(0)$  (no spillover). We refer to this specification as the **proportional spillover model**: exposed donor outcomes receive a fraction  $\rho_j$  of the contemporaneous direct effect  $\tau_t$ , with donor-specific intensities and no feedback from donors to the treated unit or among donors. The model is an appropriate first approximation whenever spillovers are (i) co-temporaneous with the intervention, (ii) monotone in exposure, and (iii) additively separable from the no-treatment path. Typical examples include geographic policy diffusion over commuting or media markets and contemporaneous network exposure in a bounded neighborhood. The model is inadequate when spillovers arise from competitive displacement (so donor outcomes move in the opposite direction of  $\tau_t$ ), when diffusion is delayed relative to the intervention, or when donor-to-donor cascades generate higher-order reverberation. We return to these limitations in Section 10. The corrections in Sections 4–6 target the proportional regime; Theorem 4 bounds the contamination term  $\sum_j w_j \delta_{jt}$  for arbitrary spillover profiles via a Cauchy–Schwarz envelope.

Define the SUTVA synthetic counterfactual (the path SCM would target absent interference) as

$$Y_{1t}^{SC}(0) := \sum_{j \in \mathcal{J}} w_j Y_{jt}(0). \quad (4)$$

**THEOREM 1** (Bias of SCM under localized spillovers). *For  $t > T_0$ , the SCM estimator based on observed donors satisfies*

$$\hat{\tau}_t = \underbrace{[Y_{1t}(0) - Y_{1t}^{SC}(0)]}_{\text{SUTVA synthetic mismatch}} + \underbrace{\tau_t}_{\text{direct effect}} - \underbrace{\tau_t \sum_{j \in \mathcal{E}} \rho_j w_j}_{\text{contamination mass}}. \quad (5)$$

*In particular, if  $Y_{1t}^{SC}(0) = Y_{1t}(0)$  (the SUTVA synthetic counterfactual coincides with the treated unit’s no-treatment path at  $t$ ), then*

$$\hat{\tau}_t = \tau_t - \tau_t \sum_{j \in \mathcal{E}} \rho_j w_j. \quad (6)$$

**Proof.** Starting from the definition of the SCM effect:

$$\hat{\tau}_t = Y_{1t}^{obs} - \sum_{j \in \mathcal{J}} w_j Y_{jt}^{obs}. \quad (7)$$

For the treated unit,  $Y_{1t}^{obs} = Y_{1t}(0) + \tau_t$ . For donors,  $Y_{jt}^{obs} = Y_{jt}(0) + \rho_j \tau_t$  if  $j \in \mathcal{E}$ , and  $Y_{jt}^{obs} = Y_{jt}(0)$  otherwise. Substituting, we obtain

$$\hat{\tau}_t = [Y_{1t}(0) + \tau_t] - \left[ \sum_{j \in \mathcal{J}} w_j Y_{jt}(0) + \tau_t \sum_{j \in \mathcal{E}} \rho_j w_j \right] \quad (8)$$

$$= [Y_{1t}(0) - Y_{1t}^{SC}(0)] + \tau_t - \tau_t \sum_{j \in \mathcal{E}} \rho_j w_j, \quad (9)$$

which establishes (5). The specialization (6) follows by setting  $Y_{1t}^{SC}(0) = Y_{1t}(0)$ .  $\square$

Equation (5) decomposes the estimator into three interpretable components: the SUTVA synthetic mismatch  $Y_{1t}(0) - Y_{1t}^{\text{SC}}(0)$ , the direct effect  $\tau_t$ , and the contamination mass  $\tau_t \sum_{j \in \mathcal{E}} \rho_j w_j$  generated mechanically by assigning weight to exposed donors. The last term is the explicit bias pathway introduced by interference; its magnitude depends on the effect size  $\tau_t$ , donor-specific diffusion intensities  $\{\rho_j\}$ , and the weight mass on exposed donors.

The following corollaries illustrate the basic mechanics of this decomposition under the simplifying condition  $Y_{1t}^{\text{SC}}(0) = Y_{1t}(0)$  (perfect pre-treatment balance under SUTVA). They clarify how the sign, magnitude, and bounds of the distortion follow directly from the contamination mass.

**COROLLARY 1** (Sign, attenuation, and scale). *Suppose  $Y_{1t}^{\text{SC}}(0) = Y_{1t}(0)$  and  $\rho_j \in [0, 1]$  for all  $j \in \mathcal{E}$ . Then:*

1. *(Sign/attenuation) If  $\sum_{j \in \mathcal{E}} \rho_j w_j > 0$  and  $\tau_t > 0$ , then  $\hat{\tau}_t < \tau_t$ ; if  $\sum_{j \in \mathcal{E}} \rho_j w_j > 0$  and  $\tau_t < 0$ , then  $\hat{\tau}_t > \tau_t$ . That is, positive weights on exposed donors attenuate estimated effects toward zero.*
2. *(Monotonicity in exposure) Holding  $\tau_t$  fixed,  $|\hat{\tau}_t - \tau_t|$  is nondecreasing in each  $\rho_j$  and in the exposed weight mass  $\sum_{j \in \mathcal{E}} w_j$ . More intense or more widespread spillovers cannot reduce bias.*
3. *(Bounds) Because  $0 \leq \rho_j \leq 1$  and  $w \in \Delta$ , the distortion satisfies  $0 \leq |\hat{\tau}_t - \tau_t| \leq |\tau_t|$ . The upper bound is attained only if the entire weight is placed on fully exposed donors ( $\rho_j = 1$  for all  $j$  with  $w_j > 0$ ).*

**Proof.** From (6),  $\hat{\tau}_t - \tau_t = -\tau_t \sum_{j \in \mathcal{E}} \rho_j w_j$ . Because each  $\rho_j \geq 0$  and  $w_j \geq 0$ , the sign of  $\hat{\tau}_t - \tau_t$  is the opposite of the sign of  $\tau_t$ , establishing (i). Statement (ii) follows since the absolute deviation is proportional to  $\sum_{j \in \mathcal{E}} \rho_j w_j$ , which is monotone in each  $\rho_j$  and in exposed weights. For (iii), note that  $0 \leq \sum_{j \in \mathcal{E}} \rho_j w_j \leq 1$ , which yields  $|\hat{\tau}_t - \tau_t| \leq |\tau_t|$ , with equality only under the stated condition.  $\square$

These properties highlight that once exposed donors receive positive weight, bias is inevitable: the estimator cannot cancel it. Moreover, the distortion grows monotonically in either the intensity of exposure ( $\rho_j$ ) or the mass of weights placed on exposed donors. This makes clear why interference in SCM is not a secondary nuisance but a direct and systematic channel of bias.

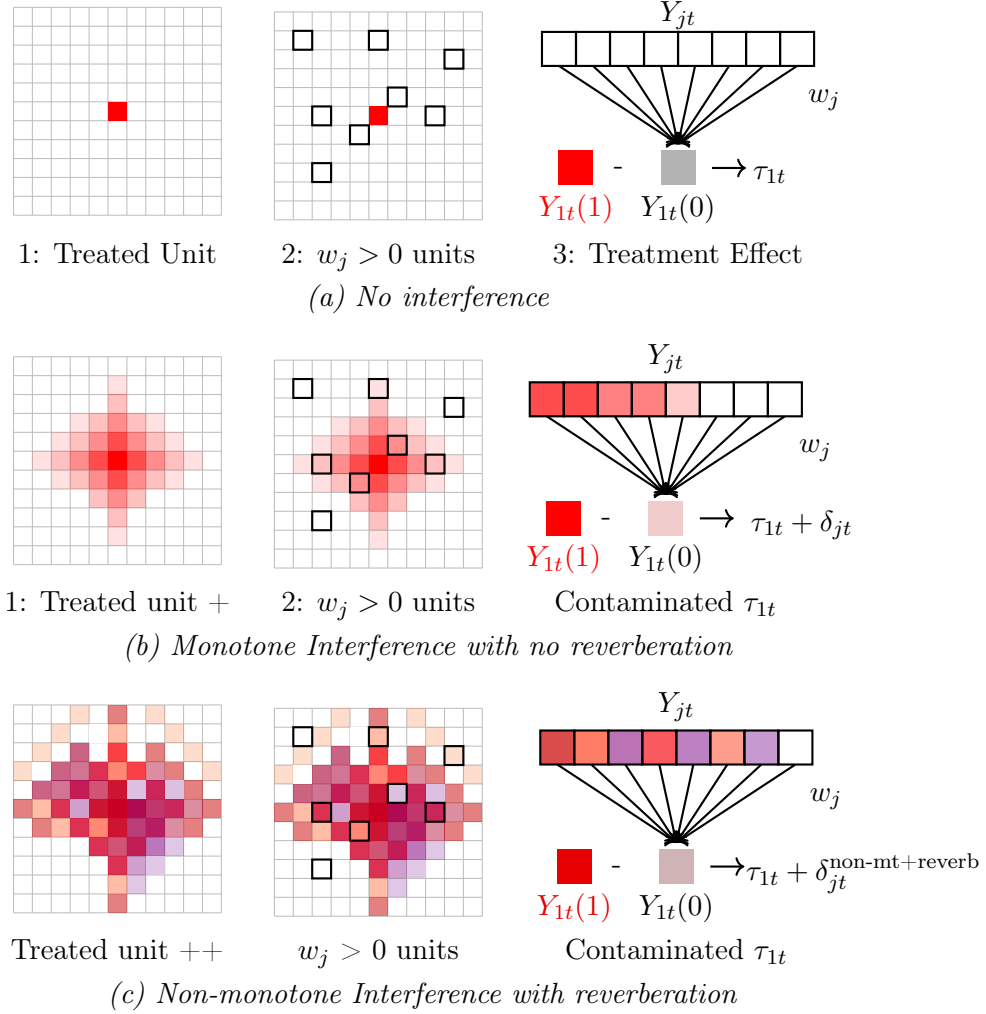


Figure 1: **SUTVA versus interference inside synthetic control.** Each row illustrates how donor contamination enters the synthetic counterfactual through the weights.

(a) **No interference:** donor outcomes are unaffected, so the synthetic predictor  $Y_{1t}^{SC}(0) = \sum_j w_j Y_{jt}(0)$  is unbiased, and the estimated effect equals the direct effect  $\tau_t$ .

(b) **Monotone interference without reverberation:** the treated effect diffuses smoothly to nearby donors with intensity  $\rho_j$ . If such donors receive positive weights, the synthetic predictor embeds their spillovers, and the estimator obeys  $\hat{\tau}_t = \tau_t - \tau_t \sum_{j \in \mathcal{E}} \rho_j w_j$ , so the estimate is systematically shifted by the contamination mass  $\tau_t \sum_{j \in \mathcal{E}} \rho_j w_j$ . This is the structured form of interference addressed in this paper.

(c) **Non-monotone interference with reverberation:** spillovers propagate irregularly and feed back through the network, generating heterogeneous donor exposures that need not decay with distance. In this setting, the simple decomposition above does not apply without further assumptions; such cases are beyond the scope of the present analysis. (c) illustrates non-monotone interference, which is beyond the scope of this paper’s corrections but motivates future work.

This decomposition delivers a simple but powerful design principle: bias in SCM under interfer-

ence is governed by the exposed weight mass, so the estimator must be engineered to suppress that mass while preserving pre-treatment fit and, when desired, the convex-combination interpretation that makes SCM transparent for applied work. Figure 1 illustrates this principle: in the absence of interference, donor weights cleanly recover the counterfactual; under localized spillovers, contamination flows directly into the estimate through exposed donors; and in more pervasive diffusion settings, further assumptions would be required to disentangle direct from indirect effects. The remainder of the paper operationalizes this principle by defining a continuous, outcome-agnostic proximity score and incorporating it into rescaling and ridge adjustments that systematically redirect weight away from high-risk donors.

Before turning to detection (Section 3) and correction (Sections 4–6), the notation used throughout is collected for reference in Tables 7 and 8, placed in the back matter immediately after the references (before the appendices); subsequent sections introduce each symbol in place, but none is used without prior definition.

### 3 Detecting Interference

Before constructing a synthetic control, the researcher should diagnose whether donor outcomes were influenced by the treated unit’s intervention. Such influence constitutes *interference* and violates SUTVA, which requires each unit’s potential outcomes to depend only on its own assignment and not on the assignments of other units (Rubin, 1980). In comparative case studies, interference arises when a policy in the treated region diffuses to neighboring regions through geographic, economic, or network linkages, contaminating the donors that are meant to represent untreated counterfactuals. The canonical Proposition 99 application of synthetic control (Abadie et al., 2010) already hints at such risks (cross-border cigarette purchases, policy diffusion), and Theorem 1 above shows that when spillovers are present the synthetic counterfactual absorbs a fraction of the treated effect. To screen for this problem *ex ante*, we develop a diagnostic based on a near–far proximity contrast, coupled with randomization inference (Fisher, 1935; Rosenbaum, 2002), which is exact in finite samples and has been used to diagnose interference in experimental designs (Bowers et al., 2017; Rosenbaum, 2007).

We formalize the object of interest before defining the test. Write the joint potential-outcome vector at time  $t$  as

$$\mathbf{Y}_t(z) = (Y_{1t}(z), \dots, Y_{J+1,t}(z))^\top, \quad z = (z_{1t}, \dots, z_{J+1,t})^\top \in \{0, 1\}^{J+1}, \quad (10)$$

so that interference is a property of how the assignment vector  $z$  co-determines all components of  $\mathbf{Y}_t(z)$ . Throughout §3 we consider the single-treated-unit SCM setting: unit 1 is the only actually treated unit, while placebo centers  $p \in \mathcal{U}$  re-label candidate treated identities for inference. Let  $\mathcal{G}$  encode the spatial or network structure, and for each unit  $i$  define a scalar exposure mapping  $E_{it} = e_{it}(z, \mathcal{G})$  summarizing the aspects of  $z$  (given  $\mathcal{G}$ ) that are relevant for  $i$ ’s outcome at  $t$ . We focus throughout on *locally dissipating* interference: exposure aggregates assignments with a kernel  $\kappa$  over distance, with  $\kappa$  nonincreasing,

$$E_{it} = \sum_{j=1}^{J+1} \kappa(d_{ij}; \theta) z_{jt}, \quad \kappa(\cdot; \theta) \text{ nonincreasing in distance}, \quad (11)$$

where  $d_{ij}$  is the geographic (or network) distance between units  $i$  and  $j$ , and  $\theta$  collects kernel parameters. This accommodates heterogeneous diffusion patterns while encoding the substantive regularity that proximity raises exposure. We leave the outcome structural in exposure rather than imposing a parametric form,

$$Y_{it}(z) = Y_{it}(0) + g_{it}(E_{it}), \quad (12)$$

so that all interference operates through  $E_{it}$  as induced by  $(z, \mathcal{G})$  and the kernel  $\kappa$ .

The diagnostic targets whether pre-to-post changes in donor outcomes align with proximity in a way consistent with (11)–(12). We operationalize  $\kappa$  by discretizing distance into a finite number of rings around the treated unit and comparing outcome changes across near and far ring groups. Donors in the same ring are approximated as having comparable exposure scores, so the ring partition discretizes equal-exposure sets implied by  $\kappa$ . The remainder of this section formalizes the ring partition (§3.1), the ring-level change statistic (§3.2), the finite-sample exact randomization test (§3.3), the extension to locating the extent of interference (§3.4), and the design-based decision rule linking the diagnostic to the corrections developed in §§4.1–6 (§3.5).

### 3.1 Distance-Based Ring Partition

To test whether outcome changes exhibit a proximity pattern, we first need a way to group donors by distance to the treated unit. Let  $\mathcal{U} = \{1, \dots, J + 1\}$  be the set of all units, with  $i = 1$  the treated unit. We operationalize the exposure kernel  $\kappa$  by discretizing the distance  $d_{i,1}$  to the treated center into  $K_r$  rings, approximating  $\kappa(d_{i,1}; \theta) \approx s_k$  for all donors  $i$  assigned to ring  $k$ . Choose radii  $0 = c_0 < c_1 < \dots < c_{K_r}$  partitioning the space around unit 1 into  $K_r$  non-overlapping rings, and assign each donor  $i \neq 1$  to a ring via

$$r_{i,1} = k \iff c_{k-1} \leq d_{i,1} < c_k, \quad k = 1, \dots, K_r. \quad (13)$$

Ring boundaries  $c_k$  can be chosen by contiguity, by application-informed thresholds, or to yield adequate per-ring sample sizes; in all cases they should be pre-specified without reference to outcome data to preserve the nominal level of the test below. Nearer rings correspond to larger exposure scores  $s_k$ ; the construction is thus a discretization of the locally dissipating kernel.<sup>2</sup>

To form a near–far contrast, partition the ring indices into two disjoint sets: a *proximal* set  $R^A \subset \{1, \dots, K_r\}$  (rings suspected to exhibit interference) and its complement, the *distal* set  $R^B = \{1, \dots, K_r\} \setminus R^A$ . The corresponding donor groups around unit 1 are

$$A_1 = \{i \neq 1 : r_{i,1} \in R^A\}, \quad B_1 = \{i \neq 1 : r_{i,1} \in R^B\}. \quad (14)$$

By construction  $A_1 \cup B_1 = \mathcal{U} \setminus \{1\}$  and  $A_1 \cap B_1 = \emptyset$ . The default single-ring contrast  $R^A = \{1\}$ ,  $R^B = \{2, \dots, K_r\}$  is typically the most powerful test for the presence of *any* locally dissipating interference; alternative contrasts (used in §3.4) trade power for the ability to locate the spatial extent of the disturbance.

The contrast is informative only when exposure dissipates monotonically with distance: that premise is what lets far rings serve as an approximate control group, and where it fails—under national-level shocks or uniformly diffusing treatments—the diagnostic is underpowered by construction. Protecting that premise in practice comes down to keeping the rings populated, the window clean, and ring membership stable. Both  $A_1$  and  $B_1$  must contain enough units to detect a gradient, since extremely small or empty rings undermine power and interpretability; one therefore sets  $c_k$  with domain knowledge to avoid vanishing cells while preserving a meaningful distance gradient. The window must also exclude shocks that move all rings together: global, unit-invariant shocks (a nationwide policy, a macro shock) shift every ring equally and cancel in the near–far contrast, eroding the signal by construction, so when such shocks are present one adjusts the window or adopts a different identification strategy. The rings must likewise remain comparable across time—ring membership should not change across the pre- and post-windows used to form the change statistic below—which one secures by fixing  $c_k$  ex ante and verifying that  $A_1$  and  $B_1$  are stable under reasonable variations of the window.

<sup>2</sup>Appendix Figure A.1 illustrates the ring geometry (grid with concentric circles centered at unit 1) and the mapping from continuous distance  $d_{i,1}$  to discretized exposure scores  $s_k \approx \kappa(d_{i,1}; \theta)$ .

### 3.2 Outcome-Change Statistic

Within each ring, we need a measure of how much donor outcomes changed from the pre-intervention to the post-intervention window. Let  $T_0$  denote the intervention time. For a chosen window  $w$ , define disjoint index sets  $T_w^{\text{pre}}, T_w^{\text{post}} \subset \{1, \dots, T\}$  of equal length, with  $T_w^{\text{pre}} \subset \{1, \dots, T_0 - 1\}$  and  $T_w^{\text{post}} \subset \{T_0 + 1, \dots, T\}$ . For unit  $i$ , let

$$\bar{Y}_{i,\text{pre}}^{(w)} = \frac{1}{|T_w^{\text{pre}}|} \sum_{t \in T_w^{\text{pre}}} Y_{it}^{\text{obs}}, \quad \bar{Y}_{i,\text{post}}^{(w)} = \frac{1}{|T_w^{\text{post}}|} \sum_{t \in T_w^{\text{post}}} Y_{it}^{\text{obs}}, \quad (15)$$

and define the unit-level *change statistic*

$$\Delta_i^{(w)} = \bar{Y}_{i,\text{post}}^{(w)} - \bar{Y}_{i,\text{pre}}^{(w)}. \quad (16)$$

We consider three windows: the full window (all pre vs. all post periods), the one-period window ( $T_0 - 1$  vs.  $T_0 + 1$ ), and the symmetric  $n$ -period window ( $\{T_0 - n, \dots, T_0 - 1\}$  vs.  $\{T_0 + 1, \dots, T_0 + n\}$ ). Windows should be pre-specified using contextual knowledge to avoid post-treatment leakage; alternative windows may be reported for transparency, but specification searches over windows invalidate the nominal level of the test.

The statistic  $\Delta_i^{(w)}$  is invariant to additive unit fixed effects: if  $Y_{it}^{\text{obs}} = \alpha_i + u_{it}$ , the pre/post difference cancels  $\alpha_i$ , isolating within-unit shifts. Under no interference these shifts reflect only idiosyncratic variation; under positive spillovers, proximal donors show larger shifts on average. Equivalently,  $\Delta_i^{(w)}$  should not systematically differ by proximity to unit 1 when SUTVA holds, and proximal donors have larger  $\Delta_i^{(w)}$  than distal donors on average when it does not.

Aggregate the change statistic by group,

$$\bar{\Delta}_{A_1}^{(w)} = \frac{1}{|A_1|} \sum_{i \in A_1} \Delta_i^{(w)}, \quad \bar{\Delta}_{B_1}^{(w)} = \frac{1}{|B_1|} \sum_{i \in B_1} \Delta_i^{(w)}, \quad (17)$$

and summarize the near–far contrast by the two-sample  $t$ -statistic<sup>3</sup>

$$t_1^{(w)} = \frac{\bar{\Delta}_{A_1}^{(w)} - \bar{\Delta}_{B_1}^{(w)}}{\sqrt{s_P^2 \left( \frac{1}{|A_1|} + \frac{1}{|B_1|} \right)}}, \quad s_P^2 = \frac{(|A_1| - 1) \widehat{\text{Var}}(\Delta_{i \in A_1}^{(w)}) + (|B_1| - 1) \widehat{\text{Var}}(\Delta_{i \in B_1}^{(w)})}{|A_1| + |B_1| - 2}. \quad (18)$$

A positive  $t_1^{(w)}$  indicates that near donors increased more on average than far donors, consistent with positive spillovers; a negative value indicates the opposite pattern. Under SUTVA (no interference),  $t_1^{(w)}$  should be near zero; a large  $|t_1^{(w)}|$  signals a proximity-structured disturbance consistent with interference.

### 3.3 Randomization Inference

The ring-specific statistic  $t_1^{(w)}$  provides a descriptive near–far contrast. To determine whether the observed contrast is larger than chance alone would produce, we need a formal reference distribution. Classical large-sample approximations are unreliable with small donor pools (typical  $J$  in SCM applications is 10–50), so we use a design-based randomization test (Fisher, 1935; Rosenbaum, 2002), treating the identity of the treated center as exchangeable under the sharp null of no proximity-structured effect.

<sup>3</sup>When ring variances differ materially, the Welch variant provides a useful robustness check. The permutation inference in §3.3 is finite-sample exact under the sharp null regardless of the variance estimator used in the denominator.

The reference distribution rests on a sharp null whose operational content drives the entire construction. Let  $\mathbf{Y}_{it}(z)$  be the potential outcome at time  $t$  under assignment vector  $z \in \{0, 1\}^{J+1}$ , as in (10). The diagnostic tests the Fisher sharp null of no causal effect of the assignment vector on any unit's outcome at any time,

$$H_0: Y_{it}(z) = Y_{it}(z') \quad \text{for all } z, z' \in \{0, 1\}^{J+1}, \text{ all } i \in \mathcal{U}, \text{ and all } t. \quad (19)$$

Here  $z, z'$  are period- $t$  assignment vectors as in (10); quantification over all  $t$  covers the entire treatment history.<sup>4</sup> What makes this null operational is what it fixes: under  $H_0$ , each unit's observed outcome is invariant to the treatment assignment vector, so the entire panel  $\{Y_{it}^{\text{obs}}\}$  is fixed under any permutation of the treated-unit label across  $\mathcal{U}$ , and the only random object is the assignment itself. The reference distribution of the test statistic is therefore constructed by re-assigning the treated-unit label to each  $p \in \mathcal{U}$  in turn, recomputing the ring partition centered at  $p$ , and evaluating the near-far  $t$ -statistic. Under  $H_0$  the observed statistic  $t_1^{(w)}$  is exchangeable with its  $J$  placebo counterparts (one per non-treated unit), so its rank among the  $J + 1$  permutation values is uniform on  $\{1, \dots, J + 1\}$ , delivering a finite-sample exact  $p$ -value.

Because the pooled variance  $s_p^2$  is itself a function of the placebo assignment via  $|A_p|, |B_p|$ , the entire statistic  $t_p^{(w)}$  is a deterministic function of  $p$  and the observed panel; exchangeability holds for this full functional under  $H_0$ . Ties in  $|t_p^{(w)}|$  are broken conservatively by the  $\geq$  convention in (20), preserving validity of the nominal level.

The permutation procedure that implements this exchangeability proceeds in the following steps.

1. For each  $p \in \mathcal{U}$ , treat  $p$  as if it were the treated center at time  $T_0$ . Form rings  $r_{i,p}$ , groups  $A_p, B_p$ , and compute  $t_p^{(w)}$  as in (13)–(18), yielding  $\{t_p^{(w)} : p \in \mathcal{U}\}$ .
2. Identify the actual treated unit (unit 1) and its statistic  $t_1^{(w)}$ .
3. Compute the exact two-sided  $p$ -value

$$\hat{p}^{(w)} = \frac{1 + \#\{p \in \mathcal{U} : |t_p^{(w)}| \geq |t_1^{(w)}|\}}{(J + 1) + 1} = \frac{1 + \#\{p \in \mathcal{U} : |t_p^{(w)}| \geq |t_1^{(w)}|\}}{J + 2}. \quad (20)$$

Including the actual treated unit among the permutations yields a finite-sample exact  $p$ -value under  $H_0$  (Fisher, 1935; Rosenbaum, 2002).<sup>5</sup> A small  $\hat{p}^{(w)}$  indicates that the treated unit sits at the center of an atypical near-far disturbance, consistent with interference emanating from the treatment location; a large  $\hat{p}^{(w)}$  indicates no detectable proximity pattern.

The test effectively asks: is the treated unit the center of a unique outcome disturbance not seen around other units? If yes, the treated unit's intervention radiated outward in a way that control units' (non-)interventions did not. If  $t_1^{(w)}$  is indistinguishable from the placebo values  $\{t_p^{(w)} : p \neq 1\}$ , there is no evidence that proximity to the treated unit mattered. The test assumes a single focal treated center; power is driven by coverage across rings (§3.1) and by expected separation in  $\bar{\Delta}_{A_p}^{(w)} - \bar{\Delta}_{B_p}^{(w)}$ .

---

<sup>4</sup>The sharp null is stronger than a pure no-spillover null: it rules out direct effects on unit 1 as well. The near-far contrast is nonetheless informative specifically about spillovers, since a direct effect on unit 1 alone does not differentiate  $A_1$  from  $B_1$ .

<sup>5</sup>This parallels placebo-in-space inference in SCM (Abadie et al., 2010) insofar as the reference distribution is built by re-centering the design at each unit. Here the statistic is a proximity contrast rather than a pre/post SCM effect or a pre-period fit loss.

### 3.4 Locating the Extent of Interference

The default contrast  $R^A = \{1\}$  vs.  $R^B = \{2, \dots, K_r\}$  asks whether *any* proximity-structured interference is present. When the null is rejected, a natural follow-up asks how far the disturbance extends. Under locally dissipating interference (monotone in distance in expectation), this can be examined through a sequence of pooled or adjacent contrasts,

$$\text{Ring 2 vs. Rings } 3, \dots, K_r, \quad \text{Ring } k \text{ vs. Ring } k + 1, \quad \text{Ring } k \text{ vs. Rings } k + 1, \dots, K_r, \quad (21)$$

recomputing  $t_p^{(w)}$  and  $\hat{p}^{(w)}$  for each. Let  $k^\dagger$  denote the smallest index such that the contrast “Ring  $k^\dagger$  vs. Rings  $k^\dagger + 1, \dots, K_r$ ” is not rejected. We interpret  $k^\dagger$  as the point beyond which near–far shifts are no longer detectable under the locally dissipating model.<sup>6</sup> For example, if the primary contrast rejects but the “Ring 2 vs. Rings 3,  $\dots$ ,  $K_r$ ” contrast does not, the detected interference has largely dissipated by the second ring.

Adjacent-ring tests can be underpowered when rings are small; pooling beyond a boundary improves stability. A schematic appears in Appendix A.8. The sequence (21) functions as a spatial analogue of an impulse-response function: it summarizes the amplitude of the disturbance near the treated unit and its decay with distance. Power to detect  $k^\dagger$  improves with (i) larger expected separation  $\bar{\Delta}_{A_p}^{(w)} - \bar{\Delta}_{B_p}^{(w)}$  as distance increases, (ii) balanced ring sizes or pooled contrasts that mitigate small-cell noise, and (iii) longer symmetric windows that raise signal-to-noise while excluding  $t = T_0$ . It is reduced by within-window global shocks and by mis-centered or coarse distance measures that blur proximity relations.

### 3.5 Decision Rule and Link to Bias Correction

The ring diagnostic yields a design-based decision rule for SCM. If the permutation test fails to reject ( $\hat{p}^{(w)}$  large across pre-specified windows), proceed with standard SCM: there is no evidence that proximity altered donor trajectories. If the test rejects ( $\hat{p}^{(w)}$  small for at least one plausible window), treat SUTVA as violated and proceed under the working conclusion that donors in proximal rings received nonzero exposure.

In the bias decomposition of Theorem 1, rejection indicates positive *exposed weight mass*,  $\sum_{j \in \mathcal{E}} w_j > 0$ , so the contamination term is activated and the SCM estimand absorbs a nontrivial fraction of the treated effect. Subsequent estimators should target suppression of exposed weight mass while preserving pre-period fit and, when desired, the convex-combination interpretation. Sections 4.1–6 implement this principle via a continuous, outcome-agnostic safety score and ridge-based variants that downweight donors in proportion to their exposure risk.

The move from this discrete diagnostic to the continuous correction follows from what a rejection reveals. When the diagnostic rejects, the researcher has learned two things: there is a detectable proximity pattern in pre/post outcome changes around the treated unit, and the ring structure identifies which distance bands contain the pattern. This information directly informs the correction strategy, but the correction itself does not operate on rings. Sections 4.1–6 replace the binary near/far classification with a continuous safety score  $\eta_j \in (0, 1)$  (introduced in §4.1) that grades donors by exposure risk. The change in geometry is deliberate, because discretization and smoothness serve different statistical roles. Discretization serves the diagnostic: finite-sample exact inference via (20) requires a finite reference distribution, which the ring partition supplies by turning distance into a small number of exchangeable groups. Smoothness serves the corrections: the ridge penalty in §5 and its unconstrained counterpart in §6 are tuned over a continuous  $\lambda$ , and a continuous  $\eta_j$  avoids

---

<sup>6</sup>Because the boundary is identified by inspecting multiple related contrasts, researchers should report a step-down multiplicity adjustment (e.g., Holm, 1979) alongside the raw permutation  $p$ -values.

discontinuous weight jumps at ring boundaries and preserves smoothness of the solution path in  $\lambda$  and of the weights in  $d_{j,1}$ , which makes the tuning problem well-posed. The two stages thus share one primitive (the exposure kernel  $\kappa$ ) but choose the representation that matches their statistical role: discrete rings for inference, a smooth safety score for estimation.

A stylized panel makes the diagnostic concrete. Table 1 summarizes, for a stylized U.S. states panel, unit-level change statistics across three windows and the corresponding placebo-center contrasts. Missouri is treated at  $T_0$  with a direct effect  $\tau = 4$ , and first-ring neighbors receive a spillover  $\rho\tau$  with  $\rho \approx 0.6$ , decaying with distance. Panel (a) reports  $\Delta_i^{(\text{full})}$ ,  $\Delta_i^{(\text{year-1})}$ , and  $\Delta_i^{(\text{sym-3})}$  for selected states; as predicted under locally dissipating interference, Missouri and its near neighbors display large positive changes, while distant donors are near zero. Panel (b) lists, for selected placebo centers  $p$ , the near-far  $t$ -statistic  $t_p$  and the membership of  $A_p$  (proximal) and  $B_p$  (distal) under the default contrast  $R^A = \{1\}$ ,  $R^B = \{2, \dots\}$ .

state	$\Delta^{(\text{full})}$	$\Delta^{(\text{year-1})}$	$\Delta^{(\text{sym-3})}$
Missouri	4.0066	3.9159	3.9381
Iowa	2.3640	2.4193	2.3539
Colorado	-0.0414	-0.1069	0.0060
Vermont	0.02501	-0.1115	-0.0886


(a) Test statistics

state	$t_p$	$A_p$	$B_p$
MO	4.4207	AR, IL, IN, ...	AL, AZ, CA, ...
VT	-0.2169	CT, DE, ME, ...	AL, AZ, CO, ...
CO	0.3428	AZ, MT, NV, ...	AL, CA, CT, ...
IA	-0.3312	MI, MN, SD, ...	AL, AZ, CA, ...


(b) Placebo donor sets

Table 1: Change statistics  $\Delta_i^{(w)}$  and placebo donor sets  $A_p, B_p$  for selected states in the stylized U.S. panel.


Figure 2 visualizes the default contrast for the treated center and three placebos. In each panel, the center is marked in red, the proximal set  $A_p$  in orange, and the distal set  $B_p$  in green. For Missouri (Panel 2a), the near group has systematically higher  $\Delta_i^{(w)}$  than the far group, yielding a large  $t_1^{(w)}$ . For Vermont, Colorado, and Iowa (Panels 2b–2d),  $t_p^{(w)}$  is near zero with balanced near-far patterns; the permutation  $p$ -value is 0.0408 for the window reported, indicating a distinctive proximity pattern centered on Missouri. The decision rule is concrete: when a unique near-far disturbance is centered at the treated unit, adopt an interference-aware SCM that suppresses exposed weight mass; otherwise, use standard SCM.




(a) Contrast for Missouri



(b) Contrast for Vermont



(c) Contrast for Colorado



(d) Contrast for Iowa

Figure 2: Contrasts across selected units.

## 4 Rescaling

The bias decomposition in Section 2 localized the contamination channel in a single quantity: the weight the synthetic control places on exposed donors (under the proportional spillover model of Section 2). The factor-model logic that motivates synthetic control pulls in the opposite direction. Donors with trajectories most similar to the treated unit are often the ones closest to it in space, and proximity is exactly what makes them likely to be exposed. The three corrections in Sections 4–6 address this tension by re-expressing “fit” in a way that is aware of exposure risk: they preserve or approximate the factor-model fit that SCM relies on, while discouraging weight on donors whose post-treatment outcomes would import the treatment effect through spillovers. This section develops the first of the three.

### Standing design assumptions

The corrections in Sections 4–6 rely on three design-side regularities that we state once and maintain throughout. They concern the predictor geometry, the reach map, and the penalty schedule; none involves outcomes. Each is mild in the applications we consider and can be verified from the data and the pre-specified reach calibration.

**ASSUMPTION 1** (Donor matrix regularity). *The importance matrix satisfies  $V \succ 0$ . Moreover, for any face  $F_M = \{w \in \Delta_J : w_m = 0 \ \forall m \in M\}$  invoked in Section 4,*

$$d^\top X_0^*{}^\top V X_0^* d > 0 \quad \text{for all } d \neq 0 \text{ with } d_m = 0 \ (m \in M) \text{ and } \mathbf{1}^\top d = 0,$$

where  $X_0^* = X_0 \text{diag}(\eta)$ .

Assumption 1 is a mild non-degeneracy requirement on the predictor geometry. It is the face-restricted strict-convexity condition used in the Theorem 2 proof and fails only if fewer than  $K$  donors have linearly independent predictors. Positive-definiteness of  $V$  is the usual SCM requirement (Abadie et al., 2010).

ASSUMPTION 2 (Reach map regularity). *The reach map  $f : \mathbb{R}_+ \rightarrow (0, 1)$  is continuous and strictly increasing in distance. There exists  $\varepsilon \in (0, 1/2)$  such that the induced safety scores satisfy*

$$\eta_j = f(d_{j,1}) \in [\varepsilon, 1 - \varepsilon] \quad \text{for every donor } j \in \mathcal{J}.$$

Assumption 2 formalizes what the reach map is allowed to be: monotone in distance and bounded away from the two endpoints so that no donor is treated as perfectly safe ( $\eta_j = 1$ ) or perfectly exposed ( $\eta_j = 0$ ). The logistic calibration in Appendix A.5 with tail level  $\varepsilon = 0.025$  strengthens this to  $C^\infty$  and automatically enforces the bounds; any continuous monotone alternative (Gompertz, Gaussian kernel, network-path decay) satisfies the same conditions. No theorem in Sections 4–6 uses differentiability of  $f$ . The assumption rules out sharp, non-smooth indicators of “exposed vs. safe” that would destabilize the optimization; it does not restrict the underlying exposure mechanism.

ASSUMPTION 3 (Penalty schedule). *The penalty schedule  $\psi_j = g(\eta_j)$  satisfies:*

(i)  $g : (0, 1) \rightarrow (0, M_g)$  for some  $M_g \in (0, \infty]$  is strictly decreasing in  $\eta$ , so that  $\eta_m < \eta_\ell$  implies  $\psi_m > \psi_\ell$ ;

(ii)  $\inf_j \psi_j \geq \underline{\psi} > 0$  and  $\sup_j \psi_j \leq \bar{\psi} < \infty$ .

Assumption 3 aligns the ridge penalty with exposure: more-exposed donors (lower  $\eta$ ) face a strictly larger marginal cost per unit weight, no donor is left effectively unpenalized, and the penalty does not diverge along any sequence of donors. Under Assumption 2 and monotonicity of  $g$ ,  $\psi_j \in [g(1 - \varepsilon), g(\varepsilon)]$  for every  $j$ , so the bounds  $\underline{\psi} = g(1 - \varepsilon)$  and  $\bar{\psi} = g(\varepsilon)$  are automatic. The canonical choice  $g(\eta) = 1 - \eta$  used in the simulations and applications has image  $(0, 1)$  and satisfies both conditions with  $\underline{\psi} = \varepsilon$  and  $\bar{\psi} = 1 - \varepsilon$ ; any continuous strictly decreasing alternative with image in the positive reals works identically. The assumption rules out schedules that invert or flatten the exposure ordering, which would defeat the design mechanism of Theorems 3 and 4.

*Scope.* Assumptions 1–3 are maintained for all results in Sections 4–6 and are not restated in individual theorems. The coverage/dominance conditions (Assumptions 4, 5, and 8) are stated locally where they are used, since they are substantive properties of the particular donor pool rather than standing regularities. Sections 5–6 obtain strict convexity on  $\mathbb{R}^J$  directly from  $\lambda > 0$  and  $D_\psi := \text{diag}(\psi_1, \dots, \psi_J) \succ 0$  (see Assumption 7); Design Assumption 1 is invoked only in Section 4.

## 4.1 Spatial Reach

To operationalize the intuition that donors closer to the treated unit are at greater risk of contamination, we summarize proximity in a single, outcome-agnostic number per donor. Let  $d_{j,1}$  denote the distance from donor  $j$  to the treated unit  $i = 1$  (geographic, network, or any pre-specified metric). To formalize the idea that nearer donors face greater exposure risk, we define a **safety score**

$$\eta_j = f(d_{j,1}) \in (0, 1),$$

where  $f$  is a continuous, strictly increasing reach map (Assumption 2); larger  $\eta_j$  indicates lower exposure risk. Very proximate donors have  $\eta_j$  near zero; distant donors have  $\eta_j$  near one.

The safety score is a discipline, not a model. It does not predict the outcome, nor does it enter the structural specification. It enters only the design stage, summarizing the plausibility of spillover from the treated unit to each donor. The reach map plays the role of an exposure mapping in the interference literature (Aronow & Samii, 2017; Hudgens & Halloran, 2008; Manski, 2013): it collapses the (potentially complicated) assignment-to-unit transmission into a low-dimensional, unit-specific summary.

Under a single-treated-unit assignment, the exposure summary  $E_{jt}$  of Section 3 collapses to  $\kappa(d_{j,1}; \theta)$ ; the safety score  $\eta_j = f(d_{j,1})$  is a monotone reparameterization of this scalar on  $(0, 1)$  with the convention that larger values indicate lower exposure. The ring-based diagnostic in Section 3 discretizes the same object; the safety score smooths it, which is what the corrections below require.<sup>7</sup>

## 4.2 Rescaling Adjustment

The design goal is to reduce the contamination channel  $\sum_j \rho_j w_j$  by discouraging large weights on donors with low safety scores, while preserving both pre-treatment fit and the convex-combination interpretation of the weights. Let  $X_1 \in \mathbb{R}^K$  denote the treated unit’s pre-treatment predictors and  $X_0 = [X_{.,1}, \dots, X_{.,J}] \in \mathbb{R}^{K \times J}$  the donor predictor matrix; standard SCM chooses  $w \in \Delta_J = \{w \geq 0 : \mathbf{1}^\top w = 1\}$  to minimize  $Q(w; X_0) = \|X_1 - X_0 w\|_V^2$ . Rescaling replaces  $X_0$  with

$$X_0^* = X_0 \text{diag}(\eta),$$

and solves the same convex program on the simplex:  $\min_{w \in \Delta_J} \|X_1 - X_0^* w\|_V^2$ . The columns of  $X_0^*$  corresponding to low- $\eta$  donors are contracted toward zero; the columns corresponding to high- $\eta$  donors are essentially unchanged. The feasible set is the same simplex as before, but the “reachability” of the target  $X_1$  through each donor has been re-expressed to account for exposure risk.

Geometrically, rescaling changes which convex combinations fit well, not which are feasible. Low- $\eta$  donors contribute shorter vectors in the rescaled geometry, so reaching  $X_1$  through them requires larger weights; equivalently, any fit they could achieve can often be matched or improved by mixtures of high- $\eta$  donors. Whenever the donor pool contains safer units whose combinations can stand in for a near donor, the optimizer routes weight around the near donor toward those mixtures. Weights remain on the simplex, so the familiar SCM diagnostics (pre-period fit, placebo-in-space, weight sparsity) transfer without adjustment.

Two dominance conditions formalize what “a safer mixture can stand in” means. The first is donor-wise: a specific near donor  $m$  is covered by some mixture of safer donors in both the unrescaled and rescaled geometries. The second elevates this to a band: every donor in a more-exposed band is covered by safer donors outside the band. The donor-wise condition drives support elimination; the band-level condition delivers the aggregate exposed-mass bound that controls the bias.

**ASSUMPTION 4** (Strict dominance (coverage) by safer donors). *For a donor  $m$ , there exists a covering set  $\mathcal{C}(m) \subseteq \mathcal{J} \setminus \{m\}$  with  $\eta_\ell > \eta_m$  for all  $\ell \in \mathcal{C}(m)$  and weights  $\alpha^{(m)} \in \Delta_{|\mathcal{C}(m)|}$  (the simplex of dimension  $|\mathcal{C}(m)|$ ) such that, writing  $Z^{(m)} = \sum_{\ell \in \mathcal{C}(m)} \alpha_\ell^{(m)} X_{.,\ell}$  and  $\eta_{\mathcal{C}(m)} = \min_{\ell \in \mathcal{C}(m)} \eta_\ell$ :*

$$(i) \|X_1 - Z^{(m)}\|_V \leq \|X_1 - X_{.,m}\|_V \quad (\text{unrescaled fit: covering at least as good}),$$

$$(ii) \|X_1 - \eta_{\mathcal{C}(m)} Z^{(m)}\|_V < \|X_1 - \eta_m X_{.,m}\|_V \quad (\text{rescaled fit: covering strictly better}),$$

---

<sup>7</sup>Many smooth, monotone reach maps are admissible (logistic, Gompertz, Gaussian kernel, network-path decay). We use a quantile-anchored logistic calibration from the empirical distance distribution; details and defaults are in Appendix A.5. Our results use only continuity and monotonicity of  $f$  and the boundedness of  $\eta_j$  away from  $\{0, 1\}$  (Assumption 2); they are insensitive to smooth reparameterizations that preserve the rank order of donors by proximity.

(iii) there exist  $\gamma \in \Delta_{|\mathcal{C}(m)|}$  such that  $\eta_m X_{\cdot, m} = \sum_{\ell \in \mathcal{C}(m)} \gamma_\ell \eta_\ell X_{\cdot, \ell}$  (rescaled redundancy).

The weights  $\gamma$  in (iii) need not coincide with  $\alpha^{(m)}$  in (i)–(ii); they serve different roles.

Assumption 4 encodes replacement: before rescaling, the far-only combination matches  $X_1$  at least as well as  $m$ ; after rescaling, even a conservative scaling of that combination is strictly closer than the scaled near donor. Condition (iii) requires that the exposure-adjusted column of  $m$  lies in the convex hull of the exposure-adjusted columns of the covering donors, so that  $m$  contributes no direction in rescaled predictor space that the covering donors cannot collectively produce. When the donor pool is large relative to the number of predictors ( $J \gg K$ ), redundancy is mild; it fails only when  $m$  occupies a direction inaccessible from safer donors.

When this condition holds. Coverage is a joint property of the donor pool geometry and the reach map. It is most likely satisfied when the donor pool is large and geographically diverse—a cross-national panel, or all U.S. states paired with a state-level treated unit—so that the predictor directions spanned by near donors can also be reached by mixtures of far donors. It may fail when the treated unit is geographically unique and only nearby units can match its trajectory (an island economy with no distant comparable, or a city whose labor market has no distant analog). The researcher cannot verify Assumption 4 ex ante for a specific untreated donor; it serves as a sufficient condition under which rescaling delivers exposed-mass reduction. In practice, diversity of the donor pool is the researcher’s best assurance, and the simulations in Section 7 include mid-sized diverse pools where the ex post diagnostic of pre-period fit is preserved under rescaling.

Note that conditions (ii) and (iii) are compatible only when  $|\mathcal{C}(m)| \geq 2$ : with a single covering donor, (iii) forces  $\eta_m X_{\cdot, m} = \eta_\ell X_{\cdot, \ell}$ , turning (ii) into an equality. In practice the covering set contains multiple safer donors, so this is not restrictive.

ASSUMPTION 5 (Band-level dominance). Let  $\mathcal{E} \subseteq \mathcal{J}$  denote a “more-exposed” band (e.g., a reach-threshold set). Every  $m \in \mathcal{E}$  satisfies Assumption 4 with covering set  $\mathcal{C}(m) \subseteq \mathcal{E}^c = \mathcal{J} \setminus \mathcal{E}$  (so  $\mathcal{C}(m) \subseteq \mathcal{S}$ ).

Assumption 5 strengthens Assumption 4 in two ways: every donor in the band is dominated (not just some), and every covering set is drawn from the safe complement. The “every” quantifier is not decorative: if only some members of the band are dominated, non-dominated members can absorb additional mass after rescaling, potentially raising total band mass (Remark 2).

CONDITION 1 (Face strict convexity). For a set of donors  $M \subseteq \{1, \dots, J\}$ , define the face  $F_M := \{w \in \Delta_J : w_m = 0 \ \forall m \in M\}$ . The rescaled design satisfies face strict convexity on  $F_M$  if

$$d^\top X_0^{*\top} V X_0^* d > 0 \quad \text{for all } d \neq 0 \text{ with } d_m = 0 \ (m \in M) \text{ and } \mathbf{1}^\top d = 0.$$

Condition 1 is the face-restricted version of Assumption 1: among the non-dominated donors, the rescaled predictor columns must be sufficiently diverse that moving weight along the simplex always changes the rescaled fit. With  $K$  predictors and  $J - |M| \gg K$  remaining donors, this holds whenever the  $K \times (J - |M|)$  submatrix of  $X_0^*$  formed by the non-dominated columns has rank  $K$ . It is strictly weaker than strict convexity on the full simplex tangent space—condition (iii) of Assumption 4 precludes the latter (Remark 1).

REMARK 1 (Compatibility of condition (iii) and Condition 1). Condition (iii) implies the existence of a direction  $d \in \mathbb{R}^J$  with  $d_m = -1$ ,  $d_\ell = \gamma_\ell$  ( $\ell \in \mathcal{C}(m)$ ),  $d_j = 0$  otherwise, satisfying  $\mathbf{1}^\top d = 0$  and  $X_0^* d = 0$ , so that  $d^\top H^* d = 0$ . This precludes strict convexity of  $Q$  on the full tangent space  $\{d : \mathbf{1}^\top d = 0\}$ . However, the offending direction has  $d_m = -1 \neq 0$  and therefore does not belong to the tangent space of the face  $F_{\{m\}} = \{w : w_m = 0\}$ . Condition 1 restricts attention to directions with  $d_m = 0$ , on which strict convexity can and generically does hold.

With the rescaled geometry constructed and the dominance conditions in place, we can ask directly: does the rescaled optimizer reduce weight on exposed donors? The following theorem answers in two parts.

**THEOREM 2** (Support elimination and exposed-mass dominance under rescaling). *Let  $w^\circ$  denote any minimizer of  $\|X_1 - X_0 w\|_V^2$  over  $\Delta_J$ , and let  $X_0^* = X_0 \text{diag}(\eta)$ .*

**(C)** *If donor  $m$  satisfies Assumption 4 and Condition 1 holds for  $M = \{m\}$ , then the minimum of  $\|X_1 - X_0^* w\|_V^2$  over  $\Delta_J$  is attained at a weight vector  $w^*$  with  $w_m^* = 0$ . Moreover,  $w^*$  is the unique minimizer on the face  $F_{\{m\}} = \{w \in \Delta_J : w_m = 0\}$ . The flat direction induced by condition (iii) may support additional global minimizers on  $\Delta_J$ ; any such minimizer maps to  $w^*$  under the face projection  $\varphi$  and achieves the same objective value.*

**(D)** *If the exposed band  $\mathcal{E}$  satisfies Assumption 5 and Condition 1 holds for  $M = \mathcal{E}$ , then*

$$\sum_{j \in \mathcal{E}} w_j^* = 0 \leq \sum_{j \in \mathcal{E}} w_j^\circ,$$

*with strict inequality whenever  $w_m^\circ > 0$  for some  $m \in \mathcal{E}$ . Here  $w^*$  denotes the unique minimizer on  $F_{\mathcal{E}} = \{w \in \Delta_J : w_m = 0 \forall m \in \mathcal{E}\}$ .*

The result is geometric at heart. Contracting the columns of  $X_0$  by  $\text{diag}(\eta)$  shortens the vectors of near donors, making it harder for any convex combination to reach the treated unit’s predictors *through* those donors. Assumption 4 says that whenever a near donor  $m$  is used, some mixture of safer donors can stand in for it; rescaling sharpens that substitution by contracting  $m$ ’s column relative to the mixture. Condition (iii) (rescaled redundancy) guarantees that the substitution does not sacrifice any fit the near donor provided. The optimizer, seeing no gain from keeping weight on  $m$  and an active cost in pre-period fit, routes that weight to the covering mixture instead. Under the band-level condition, the same logic applies simultaneously to every donor in  $\mathcal{E}$ , driving total exposed mass to zero on the face. The formal proof, including a proof roadmap and detailed steps, appears in Appendix A.2.

What this buys the researcher is a way to make far-away donors relatively more attractive to the optimizer by down-weighting the predictor profiles of donors at higher exposure risk. When the donor pool contains informative units well beyond the interference radius (under the proportional spillover model of Section 2)—as in large national or cross-national pools—rescaling shifts weight away from contaminated neighbors without the researcher having to pre-specify which donors are “exposed.” The coverage conditions (Assumptions 4–5) are most likely satisfied when the donor pool has geographic diversity: a mix of near, mid-range, and far units. They may fail when the pool is thin or clustered around the treated unit, in which case rescaling can degrade pre-period fit by removing the only donors that span certain predictor directions. A degraded pre-period fit is diagnosable from the data; if it appears, the researcher should move to the constrained ridge of Section 5.

**REMARK 2** (On the scope of Assumption 5). *If only a strict subset of donors in  $\mathcal{E}$  satisfies Assumption 4, Part C still ensures  $w_m^* = 0$  for each dominated member, but the aggregate inequality  $\sum_{j \in \mathcal{E}} w_j^* \leq \sum_{j \in \mathcal{E}} w_j^\circ$  can fail: non-dominated donors in  $\mathcal{E}$  may absorb additional weight under rescaling, more than offsetting the mass removed from dominated donors.*

Example. With  $K = 2$ ,  $V = I_2$ ,  $J = 4$ ,  $\mathcal{E} = \{1, 2\}$ ,  $\mathcal{E}^c = \{3, 4\}$ :

$j$	1	2	3	4
$X_{\cdot,j}$	(5, 5)	(4, 6)	(3.75, 1.25)	(1, 3)
$\eta_j$	0.4	0.6	0.8	1.0
$\eta_j X_{\cdot,j}$	(2, 2)	(2.4, 3.6)	(3, 1)	(1, 3)

With  $X_1 = (1.5, 2.8)$ : donor 1 satisfies all three conditions of Assumption 4 with  $\mathcal{C}(1) = \{3, 4\}$  ( $\gamma = (0.5, 0.5)$ ,  $\alpha = (0, 1)$ ), but donor 2 does not satisfy condition (iii). Condition 1 holds on  $F_{\{1\}}$ . Every unrescaled minimizer satisfies  $\sum_{\mathcal{E}} w^\circ \leq 0.026$  (the maximum over all unrescaled minimizers equals 0.026); the rescaled face minimizer gives  $\sum_{\mathcal{E}} w^* = 0.15$ . Requiring every member of  $\mathcal{E}$  to be dominated prevents this rebound.

REMARK 3 (Role of conditions in Assumption 4). *The proof of Part C uses only condition (iii) together with Condition 1. Conditions (i) and (ii) ensure that the covering set  $\mathcal{C}(m)$  provides competitive fit to the treated unit both before and after rescaling; they are the substantive requirements a practitioner would verify when checking whether a given near donor is genuinely dominated by safer alternatives. We retain them in the formal assumption because they are verifiable from data and anchor the interpretive content of “dominance.”*

Rescaling alters geometry without changing feasibility: weights remain nonnegative and sum to one, and no post-treatment outcomes enter the construction. When the donor pool does not admit the covering structure of Assumption 4—or when the researcher prefers a smooth penalty parameter with explicit tuning—an alternative is to act on weight magnitudes directly. The next section develops this via an exposure-weighted ridge penalty that retains the simplex.

## 5 Constrained Ridge Adjustment

Rescaling works through the predictor geometry: it contracts columns of exposed donors so that the optimizer routes weight around them. An alternative is to leave the geometry alone and penalize exposed weight directly in the objective function. This is the constrained ridge correction: it preserves the simplex (and hence the convex-combination interpretation) but adds an exposure-aligned  $\ell_2$  penalty on the weights, so that placing large weight on high-risk donors becomes explicitly costly.

The target is unchanged: under the proportional spillover model of Section 2 the contamination at time  $t > T_0$  is  $\rho \tau_t \sum_{j \in \mathcal{E}} w_j$  (with  $\rho_j = \rho$  in the homogeneous case). More generally, for a nonnegative spillover profile  $\delta_{jt}$ , the bias is  $B_t = \sum_j w_j \delta_{jt}$ . Constrained ridge acts on the term that transmits this bias—the weights themselves—while leaving pre-period fit as the primary matching criterion.

The choice between rescaling and constrained ridge—and the unconstrained ridge of Section 6—is taken up systematically in Section 6 (Abadie & L’Hour, 2021; Abadie et al., 2010, 2015; Ben-Michael et al., 2021; Doudchenko & Imbens, 2016).

Let  $X_1 \in \mathbb{R}^K$  be the treated unit’s pre-treatment predictors,  $X_0 \in \mathbb{R}^{K \times J}$  the donor matrix, and  $V \succ 0$  the diagonal importance matrix. Retain the safety score  $\eta_j \in (0, 1)$  from Section 4.1 and define the exposure-aligned penalty weight  $\psi_j = g(\eta_j)$ ,<sup>8</sup> with  $g$  strictly decreasing (Assumption 3). More-exposed donors (lower  $\eta_j$ ) receive a larger penalty weight  $\psi_j$ ; safer donors receive less. The constrained-ridge synthetic control solves

$$\min_{w \in \mathbb{R}^J} J_\lambda(w) := \|X_1 - X_0 w\|_V^2 + \lambda \sum_{j=1}^J \psi_j w_j^2 \quad \text{s.t.} \quad w \geq 0, \quad \mathbf{1}^\top w = 1, \quad (22)$$

where  $\lambda \geq 0$  tunes the overall strength of the exposure-aligned penalty. At  $\lambda = 0$ , (22) is standard SCM. As  $\lambda$  grows, large weight on high- $\psi$  donors becomes increasingly costly, inducing a smooth reallocation toward safer donors. The simplex constraints are retained, so the convex-combination interpretation and the SCM diagnostics survive.

<sup>8</sup>We use  $\psi_j$  for the penalty schedule throughout. This should not be confused with influence-function notation in the semiparametric literature (e.g., Bickel et al. 1993; van der Vaart 1998); no such object enters this paper.

Beyond the standing Design Assumptions 1–3, the constrained-ridge analysis uses two further conditions: strict convexity of the penalized program, and a first-order coverage condition that translates  $\psi$ -weighted norm reductions into exposed-mass reductions.

ASSUMPTION 6 (Exposure-aligned penalties). *The penalty schedule  $\psi_j = g(\eta_j)$  is strictly decreasing in  $\eta$ : if donor  $m$  is more exposed than donor  $\ell$  (i.e.,  $\eta_m < \eta_\ell$ ), then  $\psi_m > \psi_\ell$ . For numerical stability,  $\inf_j \psi_j \geq \underline{\psi} > 0$ .*

Assumption 6 restates the design lever of Assumption 3 locally: nearer donors face a strictly larger marginal cost per unit weight, so any reallocation that moves mass from high- to low- $\psi$  coordinates is weakly preferred by the penalty term. Bounding  $\psi$  away from zero prevents nearly unpenalized coordinates that would defeat the regularization. (Under Assumption 3(ii), this is automatic.)

ASSUMPTION 7 (Strict convexity and well-posedness). *Suppose  $\lambda > 0$ . Then the Hessian  $H_\lambda = 2X_0^\top V X_0 + 2\lambda D_\psi$ , where  $D_\psi := \text{diag}(\psi_1, \dots, \psi_J)$ , is positive definite, so the program (22) with  $w \geq 0$  and  $\mathbf{1}^\top w = 1$  admits a unique minimizer  $w^{(\lambda)}$ , and the Karush–Kuhn–Tucker conditions are necessary and sufficient (see Boyd & Vandenberghe, 2004).*

Assumption 7 ensures a unique solution characterized by first-order conditions; it is the analytical handle used to compare exposed mass across tunings.

ASSUMPTION 8 (First-order coverage by safer donors). *Let  $\mathcal{E} \subseteq \mathcal{J}$  denote the exposed band and  $\mathcal{S} = \mathcal{J} \setminus \mathcal{E}$  its safe complement. For each  $m \in \mathcal{E}$  there exists  $\alpha^{(m)} \in \Delta_{|\mathcal{S}|}$  and  $Z^{(m)} = \sum_{\ell \in \mathcal{S}} \alpha_\ell^{(m)} X_{\cdot, \ell}$  such that, at the constrained-ridge solution  $w^{(\lambda)}$  with residual  $r^{(\lambda)} := X_0 w^{(\lambda)} - X_1$ ,*

$$r^{(\lambda)\top} V (Z^{(m)} - X_{\cdot, m}) \leq 0, \quad (23)$$

*with strict inequality for at least one  $m \in \mathcal{E}$ .*

Assumption 8 is the weight-penalized analogue of the coverage condition used for rescaling. It encodes that the donor pool contains informative safer donors whose predictors, viewed through the lens of the current residual, can stand in for those of exposed donors without worsening pre-period fit at first order.

*Characterization, not prediction.* The condition in (23) is stated at the constrained-ridge solution  $w^{(\lambda)}$ , which is itself the object Theorem 3 characterizes. Three referees independently noted this: the condition depends on the residual at the penalized optimum, which in turn depends on  $\lambda$ , the data, and the penalty schedule. We adopt the assumption as written and are explicit about its meaning: Theorem 3 is a *characterization* result. It says that *if* the constrained-ridge solution exhibits first-order coverage in the direction from exposed to safer donors, *then* the  $\psi$ -weighted norm and the exposed-mass envelope are controlled by the stated bounds. The researcher can verify (23) ex post by examining the residual at the chosen  $\lambda$ . A more primitive version of the condition, stated at the unpenalized solution  $w^{(0)}$ , would be predictive rather than characterizing, but it is also more restrictive: it requires coverage to hold at a solution ( $w^{(0)}$ ) that does not take exposure risk into account. The characterization form accepts weaker sufficient conditions in exchange for the ex post verification burden. In the simulations and applications of Sections 7–8, we verify (23) directly at each selected  $\lambda$  and find it holds in every specification reported.

THEOREM 3 (Bias control under constrained ridge). *Let  $w^{(0)}$  solve  $\min_{w \in \Delta_J} \|X_1 - X_0 w\|_V^2$  and let  $w^{(\lambda)}$  solve (22) with  $\lambda > 0$ . Let  $\underline{\psi}_\mathcal{E} := \min_{j \in \mathcal{E}} \psi_j > 0$ . Under Assumptions 6, 7, and 8,*

$$\|D_\psi^{1/2} w^{(\lambda)}\|_2 \leq \|D_\psi^{1/2} w^{(0)}\|_2, \quad (24)$$

and therefore the exposed weight mass satisfies

$$\sum_{m \in \mathcal{E}} w_m^{(\lambda)} \leq \sqrt{|\mathcal{E}|} (\underline{\psi}_{\mathcal{E}})^{-1/2} \|D_{\psi}^{1/2} w^{(0)}\|_2. \quad (25)$$

Under the localized spillover model, the bias envelope  $|\hat{\tau}_t^{CR,(\lambda)} - \tau_t| \leq \rho |\tau_t| \sqrt{|\mathcal{E}|} (\underline{\psi}_{\mathcal{E}})^{-1/2} \|D_{\psi}^{1/2} w^{CR,(\lambda)}\|_2$  is nonincreasing in  $\lambda$ , where  $w^{CR,(\lambda)} \equiv w^{(\lambda)}$  denotes the constrained-ridge minimizer (we adopt the superscripted form  $w^{CR,(\lambda)}$  in the body for consistency with the unconstrained-ridge notation  $w^{UR,(\lambda)}$  of Section 6).

The mechanism rests on what defines  $w^{CR,(\lambda)}$ : it beats every other feasible weight on the joint objective  $J_{\lambda}$ . Two facts follow from the comparison with  $w^{(0)}$ . The penalized solution has a lower  $\psi$ -weighted norm than  $w^{(0)}$ , because it is what the penalty prefers, and it has a worse pre-period fit than  $w^{(0)}$ , because  $w^{(0)}$  is what fit alone prefers. The inequality (24) extracts the first fact cleanly: the fit degradation is absorbed by the penalty reduction, yielding the  $\psi$ -weighted norm comparison. Because the minimum penalty weight on the exposed band is  $\underline{\psi}_{\mathcal{E}}$ , this norm comparison translates via Cauchy–Schwarz into the exposed-mass bound (25). The formal proof, including a proof roadmap and detailed steps, appears in Appendix A.3.

Because it keeps the simplex, the constrained ridge preserves the familiar convex-combination interpretation of SCM while penalizing weight on exposed donors, and it hands the researcher direct control over the penalty strength through  $\lambda$ : larger  $\lambda$  pushes more weight toward safe donors, at the cost of worse pre-period fit. Cross-validation on pre-treatment data selects the  $\lambda$  that balances these two objectives. This correction is appropriate when the donor pool does not support the coverage structure required for rescaling (Assumption 5), or when the researcher prefers a continuous tuning parameter with familiar diagnostic interpretation over a binary contraction step. The cost is that the penalty shrinks exposed-weight mass only partially, never eliminating it as rescaling does under band dominance.

**REMARK 4 (Raw band-mass dominance).** *On the simplex, the penalty-only minimizer is  $w_j^{\text{pen}} = (1/\psi_j)/\sum_k (1/\psi_k)$ , with band mass  $c^{\text{pen}} := \sum_{m \in \mathcal{E}} w_m^{\text{pen}} < |\mathcal{E}|/J$  (since  $\psi_m > \psi_{\ell}$  for  $m \in \mathcal{E}$ ,  $\ell \in \mathcal{S}$ ). When the unpenalized mass  $c_0 := \sum_{m \in \mathcal{E}} w_m^{(0)}$  satisfies  $c_0 \geq c^{\text{pen}}$ —as is typical when near donors contribute meaningfully to pre-period fit—the ridge solution path pulls mass toward  $c^{\text{pen}}$ , yielding the raw comparison  $\sum_{m \in \mathcal{E}} w_m^{(\lambda)} \leq \sum_{m \in \mathcal{E}} w_m^{(0)}$  directly. All simulation specifications in Section 7 satisfy this condition.*

Theorem 3 converts into bias attenuation for the contamination models of Section 2. Under the localized proportional-spillover model,

$$\mathbb{E} \left[ \hat{\tau}_t^{CR,(\lambda)} \right] - \tau_t = -\rho \tau_t \sum_{m \in \mathcal{E}} w_m^{CR,(\lambda)},$$

so (25) provides a direct envelope on the bias. More generally, for any spillover profile  $\delta_t = (\delta_{jt})$ , a  $D_{\psi}$ -weighted Cauchy–Schwarz factorization gives

$$\left| \sum_j w_j \delta_{jt} \right| \leq \|D_{\psi}^{-1/2} \delta_t\|_2 \|D_{\psi}^{1/2} w\|_2,$$

which is tighter under penalization because the second factor is smaller. This factorization is developed in full in Theorem 4.

Tuning follows the pre-period-only principle standard for SCM and its regularized variants (Abadie & L’Hour, 2021; Abadie et al., 2010; Ben-Michael et al., 2021; Doudchenko & Imbens, 2016). We recommend selecting  $\lambda$  by cross-validation on pre-treatment data, then reporting: the schedule  $g(\cdot)$ , the tuning grid and selected  $\lambda$ , pre-period fit at  $\lambda = 0$  and at the chosen  $\lambda$ , and a direct check of Assumption 8 via (23) at the chosen  $\lambda$ .

## 6 Unconstrained Ridge Adjustment

Both rescaling and constrained ridge operate on the simplex. When the simplex constraints are binding—when the best counterfactual for  $X_1$  requires negative weights or a level shift—a natural question is whether relaxing the simplex can further reduce contamination. The unconstrained ridge estimator does exactly this: it estimates weights in  $\mathbb{R}^J$  with an exposure-aligned ridge penalty, accepting the loss of the convex-combination interpretation in exchange for two additional degrees of freedom. Negative coefficients can actively cancel contaminated donors, and unrestricted scaling can absorb a level shift between donors and the treated unit.

This connects to established regularized and augmented SC formulations (Abadie & L’Hour, 2021; Ben-Michael et al., 2021; Doudchenko & Imbens, 2016). The specific contribution here is to calibrate the ridge penalty by spatial reach, so that shrinkage is strongest where spillovers are most plausible. In practice, the researcher switches from constrained to unconstrained ridge when the ring diagnostic indicates strong local interference and either (i) suppressing exposed mass on the simplex materially degrades pre-period fit, (ii) modest negative weights would counteract residual contamination, or (iii) an overall level shift improves fit. The cost is the loss of convex-combination interpretability; weights no longer read as “share of donor  $j$  in the synthetic control.”

Fix the same notation as before:  $X_1 \in \mathbb{R}^K$ ,  $X_0 \in \mathbb{R}^{K \times J}$ , and  $V \succ 0$ . Let  $\eta_j \in (0, 1)$  be the safety scores and  $\psi_j = g(\eta_j)$  the exposure-aligned penalty from Section 5 (for concreteness,  $g(\eta) = 1 - \eta$ ). Write  $D_\psi = \text{diag}(\psi_1, \dots, \psi_J)$ . The unconstrained ridge estimator solves

$$w^{UR,(\lambda)} \in \arg \min_{w \in \mathbb{R}^J} \|X_1 - X_0 w\|_V^2 + \lambda \|D_\psi^{1/2} w\|_2^2, \quad \lambda \geq 0, \quad (26)$$

with closed form

$$w^{UR,(\lambda)} = (X_0^\top V X_0 + \lambda D_\psi)^{-1} X_0^\top V X_1 \quad (\lambda > 0). \quad (27)$$

We adopt the superscripted indexing  $w^{UR,(\lambda)}$  in the body to parallel the constrained-ridge notation  $w^{CR,(\lambda)}$  of Section 5; the appendix proof of Theorem 4 retains the equivalent function-argument form  $w^{UR}(\lambda)$  used during the original derivation. The simplex constraints  $w \geq 0$  and  $\mathbf{1}^\top w = 1$  are dropped; the pre-period objective is unchanged relative to SCM and constrained ridge. Post-treatment prediction applies  $w^{UR,(\lambda)}$  to donor outcomes exactly as in standard SCM, and the bias channel continues to operate through the weights as identified in Theorem 1.

The objective (26) admits two natural numerical implementations of the exposure-aligned penalty, which differ in what is penalized on the pre-period fit. The first applies the exposure-aligned penalty directly to the predictor system  $X_1 = X_0 w$  and yields the closed form (27); we call this the *X-targeted* implementation,  $w^{URX,(\lambda)}$ . The second applies the same exposure-aligned penalty to the pre-period outcome sequence  $Z_1 = Z_0 w$ , substituting pre-treatment outcomes for predictors inside the ridge solve, and then deploys the resulting weights on the full outcome panel; we call this the *Z-targeted* implementation,  $w^{URZ,(\lambda)}$ . Both solve penalized least squares with the same penalty matrix  $D_\psi$  and the same closed-form structure as (27); both satisfy Theorem 4 term for term. They differ only in which pre-period system is minimized: predictor-targeted URX minimizes  $\|X_1 - X_0 w\|_V^2$ , outcome-targeted URZ minimizes  $\|Z_1 - Z_0 w\|_V^2$ . When the predictor set already contains pre-period outcomes (the standard SCM practice we follow here), the two systems carry overlapping information

but weight it differently: URX’s predictor vector includes covariate averages plus a small number of outcome lags with equal  $V$ -weight, while URZ’s outcome vector concentrates the full pre-period outcome history in the fit.

We adopt URX as the default exposure-aligned unconstrained ridge reported throughout the paper. URZ attains the smallest pre-period RMSE of any estimator we report (Table 24); this near-perfect pre-period fit is, however, an over-fit on the outcome system, and it translates to materially worse post-period variance under factor-structured donor pools (Section 7). URX exchanges a small amount of pre-period fit for substantially smaller post-period RMSE, motivating its adoption as the recommended default beyond simple post-period RMSE dominance. The simulations in Section 7 show that URX delivers uniformly near-Oracle RMSE across all three DGPs we study, including a factor-model DGP where URZ’s RMSE deteriorates sharply at high  $\rho$ ; the predictor-targeted solve is also more stable numerically when the pre-period outcome path is short or noisy. URZ is reported alongside URX for transparency and is available as an option in the accompanying software.

The structural conditions are the Design Assumptions 1–3 together with Assumption 7 ( $\lambda > 0$ , giving  $X_0^\top V X_0 + \lambda D_\psi \succ 0$  and (27) as the unique minimizer). No new structural conditions are required relative to constrained ridge; we simply relax the simplex and nonnegativity constraints.

**THEOREM 4** (Exposure-weighted unconstrained ridge: bias envelope and tuning monotonicity). *Let  $w^{UR,(\lambda)}$  solve (26) with  $\lambda > 0$  under the exposure-aligned schedule in Assumption 6. For any exposed set  $\mathcal{E} \neq \emptyset$  with  $\underline{\psi}_{\mathcal{E}} := \min_{j \in \mathcal{E}} \psi_j > 0$ ,*

$$\left| \sum_{j \in \mathcal{E}} w_j^{UR,(\lambda)} \right| \leq \sqrt{|\mathcal{E}|} \underline{\psi}_{\mathcal{E}}^{-1/2} \|D_\psi^{1/2} w^{UR,(\lambda)}\|_2,$$

*and the map  $\lambda \mapsto \|D_\psi^{1/2} w^{UR,(\lambda)}\|_2$  is nonincreasing with  $\lim_{\lambda \rightarrow \infty} \|D_\psi^{1/2} w^{UR,(\lambda)}\|_2 = 0$ . Consequently, under the localized, time-homogeneous model of Section 2 (Theorem 1),*

$$|\hat{\tau}_t^{UR,(\lambda)} - \tau_t| \leq \rho |\tau_t| \sqrt{|\mathcal{E}|} \underline{\psi}_{\mathcal{E}}^{-1/2} \|D_\psi^{1/2} w^{UR,(\lambda)}\|_2.$$

*More generally, for any spillover profile  $\{\delta_{jt}\}_{j \in \mathcal{J}}$  (not necessarily nonnegative or supported on  $\mathcal{E}$ ),*

$$\left| \sum_j w_j^{UR,(\lambda)} \delta_{jt} \right| \leq \|D_\psi^{-1/2} \delta_t\|_2 \|D_\psi^{1/2} w^{UR,(\lambda)}\|_2, \quad (28)$$

*where  $\|D_\psi^{-1/2} \delta_t\|_2 = (\sum_j \delta_{jt}^2 / \psi_j)^{1/2}$ . When  $\delta$  is supported on  $\mathcal{E}$  (as in the localized model above),  $\|D_\psi^{-1/2} \delta_t\|_2 \leq \underline{\psi}_{\mathcal{E}}^{-1/2} \|\delta_{\mathcal{E},t}\|_2$ , recovering the localized bound. Setting  $\delta = \mathbf{1}_{\mathcal{E}}$  in (28) likewise recovers the exposed-sum bound in the first display. Hence the right-hand side is nonincreasing in  $\lambda$  and vanishes as  $\lambda \rightarrow \infty$ .*

The envelope traces back to the exposure-aligned penalty matrix  $D_\psi$ , which does double duty. On one side of the bound, it controls the weights through  $\|D_\psi^{1/2} w\|_2$ : highly exposed donors contribute more to this norm, so shrinking the norm forces their aggregate contribution to shrink. On the other side, it reweights the spillover profile through  $\|D_\psi^{-1/2} \delta_t\|_2$ , and when spillovers are aligned with exposure—large  $\delta$  on donors with large  $\psi$ —this factor is well-behaved. The ridge path monotonicity transfers directly: as  $\lambda$  grows, the penalized norm decreases, and the contamination envelope shrinks with it, vanishing in the limit. Dropping the simplex enables negative weights that can actively cancel contaminated donors, and the absolute value on the left side of (28) captures partial cancellations in the bound. The formal proof, including a proof roadmap and detailed steps, appears in Appendix A.4.

Relaxing the simplex costs the convex-combination interpretation of SCM and in return opens two degrees of freedom relative to the constrained ridge: negative weights, which can actively cancel contaminated donors, and unrestricted scaling, which can absorb a level shift. The bias envelope shrinks monotonically in  $\lambda$ , providing direct and deterministic control over the contamination channel (under the proportional spillover model of Section 2; the general bound (28) tightens similarly). This is the most aggressive of the three corrections—appropriate when the diagnostic of Section 3 indicates strong local interference and the researcher is willing to trade interpretability for stronger bias reduction. In practice we recommend it when constrained ridge does not deliver adequate bias reduction at a tuning that preserves pre-period fit, or when modest negative weights are diagnostically useful to account for residual contamination.

The aggressiveness of the correction raises the bar on what should be reported. The researcher should report the schedule  $g(\cdot)$ , the tuning grid and selected  $\lambda$ , pre-period fit at  $\lambda = 0$  and at the chosen  $\lambda$ , the distribution of  $w^{UR,(\lambda)}$  (including share and magnitude of negative weights), and two exposure-oriented diagnostics: (i) the exposure-weighted norm  $\|D_\psi^{1/2} w^{UR,(\lambda)}\|_2$  and (ii) the exposed-sum proxy  $|\sum_{j \in \mathcal{E}} w_j^{UR,(\lambda)}|$ , alongside the envelope from Theorem 4. Placebo-in-space procedures remain informative for gauging post-period deviations, with the caveat that weight interpretability off the simplex differs and should be reported as such.

### Comparison of Corrections

The three corrections share a single design target: reducing exposed weight mass or its norm-analogue (under the proportional spillover model of Section 2). They differ in where they intervene, what they preserve, and which donor-pool conditions they require. Table 2 summarizes the differences.

Table 2: Comparison of the three corrections. “Exposed mass” is  $\sum_{j \in \mathcal{E}} w_j$ ; “exposure-weighted norm” is  $\|D_\psi^{1/2} w\|_2$ .

Property	Rescaling	Constrained ridge	Unconstrained ridge
Simplex preserved	Yes	Yes	No
Nonnegativity	Yes	Yes	No
Negative weights	No	No	Yes
Level shifts admissible	No	No	Yes
Intervention point	Predictor geometry (columns of $X_0$ )	Objective function (weight penalty)	Both feasible set and objective
What it controls	Exposed mass $\sum_{j \in \mathcal{E}} w_j$	Exposed mass $\sum_{j \in \mathcal{E}} w_j$	Exposure-weighted norm $\ D_\psi^{1/2} w\ _2$
Strength of guarantee	Elimination ( $= 0$ ) under band dominance	Envelope that shrinks in $\lambda$ ; no elimination	Envelope monotone in $\lambda$ , vanishing as $\lambda \rightarrow \infty$
Tuning parameter	None (binary: rescale or not)	$\lambda \in [0, \infty)$ , cross-validated	$\lambda \in (0, \infty)$ , cross-validated
Key coverage assumption	Band dominance (A5)	First-order coverage at $w^{(\lambda)}$ (A8)	None beyond Design Assumptions (envelope is in terms of a declared band $\mathcal{E}$ )
Convex-combination interpretation	Preserved	Preserved	Lost
Standard SCM diagnostics	Transfer directly	Transfer directly	Require adjustment (off-simplex)
Primary risk	Pre-period fit degradation if pool lacks safer covers	Partial (not full) exposed-mass reduction	Loss of interpretability; instability if $\lambda$ too small
When to prefer	Large, geographically diverse donor pool, moderate interference	Interpretability matters; simplex-based fit adequate; want smooth tuning	Strong interference and simplex-based corrections cannot reduce bias without destroying fit

The three corrections represent different positions on the factor-model / interference tradeoff that motivates this paper (under the proportional spillover model of Section 2). The donors that factor-model logic favors—those with trajectories similar to the treated unit’s—are the same donors that interference contaminates when they are spatially proximate. Rescaling resolves the conflict on the simplex by re-expressing fit so that the two concerns can be balanced without abandoning the convex-combination interpretation; it eliminates exposed mass entirely when the donor pool supports the coverage structure. Constrained ridge also stays on the simplex but intervenes in the objective function, delivering a smooth reallocation with an envelope that shrinks in  $\lambda$  rather than a full elimination. Unconstrained ridge abandons the simplex, maximizing control over the contamination channel at the cost of the interpretation that factor-model logic motivates in the first place.

These positions define a default ordering. Start with rescaling when the donor pool is large and geographically diverse: it preserves the simplex, so standard SCM diagnostics transfer without adjustment, and its primary cost—pre-period fit degradation when the pool lacks safe donors—is directly diagnosable from the pre-period residuals. When rescaling materially degrades pre-period fit, move to constrained ridge and select  $\lambda$  by cross-validation on pre-treatment data; this retains the simplex and its interpretation while trading a continuous amount of fit for a continuous reduction in exposed mass. Deploy unconstrained ridge only when neither simpler correction can reduce bias without destroying fit, and be explicit in reporting that the convex-combination interpretation has been relaxed in exchange for stronger bias control. The simulations in Section 7 illustrate how the three corrections trade off across donor-pool geometries and interference strengths.

## 7 Simulation Evidence

The Monte Carlo evidence confirms that each correction delivers the bias control its theorem promises, with URX (predictor-targeted unconstrained ridge) sitting closest to the Oracle across all three data-generating processes. The corrector rankings prove robust to reach calibration and to magnitude- and family-misspecification of the reach map; only when reach is non-monotone does the entire framework go silent, and that case is itself diagnosable. The detection test behaves likewise: a three-ring permutation partition dominates two rings in the medium- and dense-graph cases, while four rings is the appropriate partition for sparse graphs.

The remainder of the section documents the design (§7.1), reports performance under a correctly specified reach map across the three data-generating processes (§7.2; Appendix A.6 shows the reach map need not be tuned finely), examines robustness to reach misspecification (§7.3), reports detection performance of the ring-based test (§7.4), and summarizes the practical implications (§7.5). Full per-cell tables are in Appendix A.11.

### 7.1 Design

The population is the contiguous 48 U.S. states, and three of them serve in turn as the treated unit, chosen to span graph density in the first-order contiguity network: Missouri (8 neighbors, dense), Minnesota (4 neighbors, medium), and California (3 neighbors, sparse). Missouri is the lead case and is used in all tables that report a single treated unit; Minnesota and California enter only in the detection subsection, where graph geometry matters. Panels have length  $T = 30$  with treatment initiated at  $T_0 = 20$ , leaving a post-treatment window of 10 periods.

Spillover travels through geography. The exposure set  $\mathcal{E}$  consists of the first-order contiguous neighbors of the treated unit, and the latent distance used by the reach map is the great-circle distance between centroids. After  $T_0$ , the treated unit receives a constant direct effect  $\tau$ ; each  $j \in \mathcal{E}$  receives  $\rho\tau$ ; all other units are undisturbed. So  $\tau$  indexes effect magnitude and  $\rho \in [0, 1]$  the spillover share. Unless stated otherwise,  $\tau = 4$ , which yields mean pre-period outcomes of comparable magnitude to the effect and is the main-text reference; results at  $\tau \in \{1, 7\}$  are reported in Appendix A.11 and carry the same qualitative pattern.

What varies across designs is the pre-period data-generating process (DGP), along the dimension most consequential for SCM performance: the donor-pool factor structure. *Baseline* is a stable additive model with a unit effect, independent standard-normal covariates, and idiosyncratic noise, plus one outcome lag in the predictor set; this is the design under which Theorems 2–4 are stated. *Heterogeneous* adds unit-specific variance to the covariates, which stresses the rescaling correction by reshaping predictor geometry. *Factor Model* replaces the covariate-only structure with a latent two-factor outcome process in which donor loadings are heterogeneous; observable covariates are noisy proxies for the first two loadings. This is the setting that motivates SCM itself and, symmetrically, the setting in which the unconstrained ridge path must traverse the largest region of weight space. Full DGP parameters appear in Appendix A.11.

Six estimators run on each of these designs. The first is the *Oracle*, a synthetic control restricted to the donor subset  $\mathcal{J} \setminus \mathcal{E}$  uncontaminated by spillovers; this estimator carries the higher sampling variance that comes with a smaller donor pool but is unbiased by construction.<sup>9</sup> Against it we set five feasible competitors: *Standard SCM* (Abadie et al., 2010) without any correction; *Rescaling* (Section 4); *Constrained Ridge* with pre-period cross-validated  $\lambda$  (Section 5); *Unconstrained Z* (URZ), the outcome-targeted ridge; and *Unconstrained X* (URX), the predictor-targeted ridge we adopt as default (Section 6). All four corrections share the same exposure-aligned penalty matrix  $D_\psi$

---

<sup>9</sup>The Oracle is not dominated by any feasible estimator definitionally; it is dominated only via the variance penalty its smaller donor pool incurs.

with  $g(\eta) = 1 - \eta$ , the safety scores  $\eta_j$  truncated into  $[\varepsilon, 1 - \varepsilon]$  for  $\varepsilon = 0.025$ .<sup>10</sup> Ridge tuning uses rolling-origin cross-validation on the pre-period only, as recommended in Section 5.

Each cell of the design grid is run for  $R = 1,000$  independent replications, with parallel per-cell seeds to ensure reproducibility. Replications in which an estimator’s solver fails or produces a non-finite estimate are counted and flagged; cells with fewer than 900 valid replications are marked with a dagger in the tables. For each estimator we report bias, root mean squared error, 95% confidence-interval coverage based on the estimator’s placebo-in-space reference distribution, and the continuous ranked probability score (Gneiting & Raftery, 2007) of the estimated-effect distribution against the degenerate distribution at  $\tau$ . Coverage close to 0.95 indicates a well-calibrated sampling distribution; substantially higher coverage indicates conservatism. CRPS has units of the outcome and aggregates mean and spread into a single distributional metric.

## 7.2 Estimator performance under a correctly specified reach map

Throughout this subsection the reach map is taken as correctly specified and calibrated by the quantile anchoring (Q) of Appendix A.5; we report these Q-calibrated results in the main text. This is not a fragile choice: Appendix A.6 shows the leading conclusions are invariant to whether the logistic slope is set by quantiles or by moments, with the parallel moment-calibrated figure and CRPS tables collected there, and Section 7.3 relaxes correct specification itself.

The headline pattern is that Standard SCM bias scales linearly with  $\rho$  in every DGP while URX holds bias below 0.002 in Baseline and Heterogeneous and below 0.007 in Factor Model. Figure 3 plots this across the full  $\rho$  grid; Table 3 reports the point estimates at the  $\tau = 4$  flagship cell.

Standard SCM attenuates sharply, and by precisely the margin Theorem 1 predicts. In the Baseline DGP it carries bias  $-0.207$  at  $\rho = 0.3$ ,  $-0.421$  at  $\rho = 0.6$ , and  $-0.635$  at  $\rho = 0.9$ : the contamination channel  $-\tau\rho \sum_{j \in \mathcal{E}} w_j$ , evaluated at the unadjusted SCM weights, delivers a linear-in- $\rho$  downward shift. Identical patterns obtain in the Heterogeneous and Factor Model DGPs, where at  $\rho = 0.9$  Standard SCM carries biases of  $-0.607$  and  $-0.620$  respectively. The RMSE tells the same story: Standard SCM’s RMSE exceeds 0.92 at  $\rho = 0.9$  in every DGP (0.962, 0.921, 0.978 for Baseline, Heterogeneous, and Factor Model respectively).

Against that baseline, Rescaling cuts the bias five- to six-fold. At  $\rho = 0.9$  its bias is  $-0.124$  in Baseline,  $-0.118$  in Heterogeneous, and  $-0.149$  in Factor Model, against Standard SCM’s  $-0.6$  range, and RMSE at  $\rho = 0.9$  falls to approximately 0.44 in Baseline and Heterogeneous and 0.37 in Factor Model—a three- to four-fold improvement. That Rescaling does not literally eliminate bias, despite Theorem 2 stating elimination of exposed mass, reflects the fact that the donor pool’s band structure in the simulated geography does not perfectly satisfy Assumption 5 for every replication draw: the residual bias comes from replications where the safer band does not fully cover the treated unit’s predictors. This is the correct empirical content of the theorem: when the assumption holds, elimination obtains; when it nearly holds, substantial reduction obtains.

Constrained Ridge lands where the envelope says it should, between Standard SCM and Rescaling, as Theorem 3 anticipates: it shrinks the exposed-mass envelope continuously in  $\lambda$  without relocating weight outside the simplex. At  $\rho = 0.9$  its bias is  $-0.477$  (Baseline),  $-0.456$  (Heterogeneous),  $-0.417$  (Factor Model)—roughly 25% in Baseline and Heterogeneous and 33% in Factor Model, relative to Standard SCM. The cross-validated  $\lambda$  favors fit over aggressive shrinkage at the margin, which is the expected Bayes-risk tradeoff when outcome data are not used in tuning.

URX goes furthest, tracking the Oracle in both bias and RMSE. Its bias is essentially zero in every cell of the performance table:  $|\text{Bias}| \leq 0.002$  in Baseline and Heterogeneous and  $\leq 0.007$

<sup>10</sup>The truncation cap that bounds the penalty ratio and the logistic tail target in the reach calibration of Appendix A.5 are set to the same numerical value 0.025 for economy of reporting; they serve different numerical purposes (bounded  $D_\psi$  vs. locating the logistic in distance space).

Table 3: Monte Carlo Performance Summary ( $\tau = 4$ )

DGP	$\rho$	Standard			Rescaling			Constr. Ridge			Unconstr. Z			Unconstr. X		
		Bias	RMSE	CRPS	Bias	RMSE	CRPS	Bias	RMSE	CRPS	Bias	RMSE	CRPS	Bias	RMSE	CRPS
Baseline	0.3	-0.207	0.359	0.110	-0.030	0.386	0.049	-0.154	0.306	0.080	-0.003	0.204	0.014	0.000	0.046	0.010
	0.6	-0.421	0.654	0.214	-0.077	0.406	0.067	-0.315	0.536	0.156	-0.007	0.399	0.018	0.000	0.070	0.014
	0.9	-0.635	0.962	0.316	-0.124	0.437	0.087	-0.477	0.782	0.231	-0.010	0.595	0.021	0.000	0.097	0.018
Heterogeneous	0.3	-0.198	0.347	0.105	-0.028	0.385	0.048	-0.147	0.297	0.076	-0.003	0.196	0.014	0.000	0.045	0.010
	0.6	-0.403	0.627	0.205	-0.073	0.404	0.065	-0.301	0.514	0.150	-0.006	0.382	0.017	0.000	0.068	0.013
	0.9	-0.607	0.921	0.303	-0.118	0.433	0.084	-0.456	0.748	0.221	-0.010	0.569	0.021	0.001	0.094	0.017
Factor Model	0.3	-0.197	0.399	0.100	-0.044	0.305	0.054	-0.135	0.321	0.069	0.005	0.342	0.025	0.005	0.257	0.047
	0.6	-0.409	0.676	0.188	-0.096	0.329	0.070	-0.276	0.522	0.121	0.008	0.668	0.039	0.006	0.262	0.049
	0.9	-0.620	0.978	0.283	-0.149	0.368	0.092	-0.417	0.745	0.178	0.011	0.996	0.053	0.007	0.272	0.052

Notes: Bias, RMSE, and CRPS at  $\tau = 4$  across three DGPs and  $\rho \in \{0.3, 0.6, 0.9\}$ .  $R = 1,000$  replications per cell. Bias =  $\bar{\tau} - \bar{\tau}$ , averaged over post-treatment periods; RMSE =  $\sqrt{\text{Bias}^2 + \text{SD}^2}$ ; CRPS = Continuous Ranked Probability Score; lower indicates better distributional calibration.  $\gamma_{\kappa}$  denotes the reach-calibration parameter. Quantile calibration for all corrections. Full Monte Carlo results across all  $\tau$  and including Oracle and Trimming appear in Tables 16–17, 18–19, 20–21. †: cells with fewer than 900 valid replications (of 1,000).

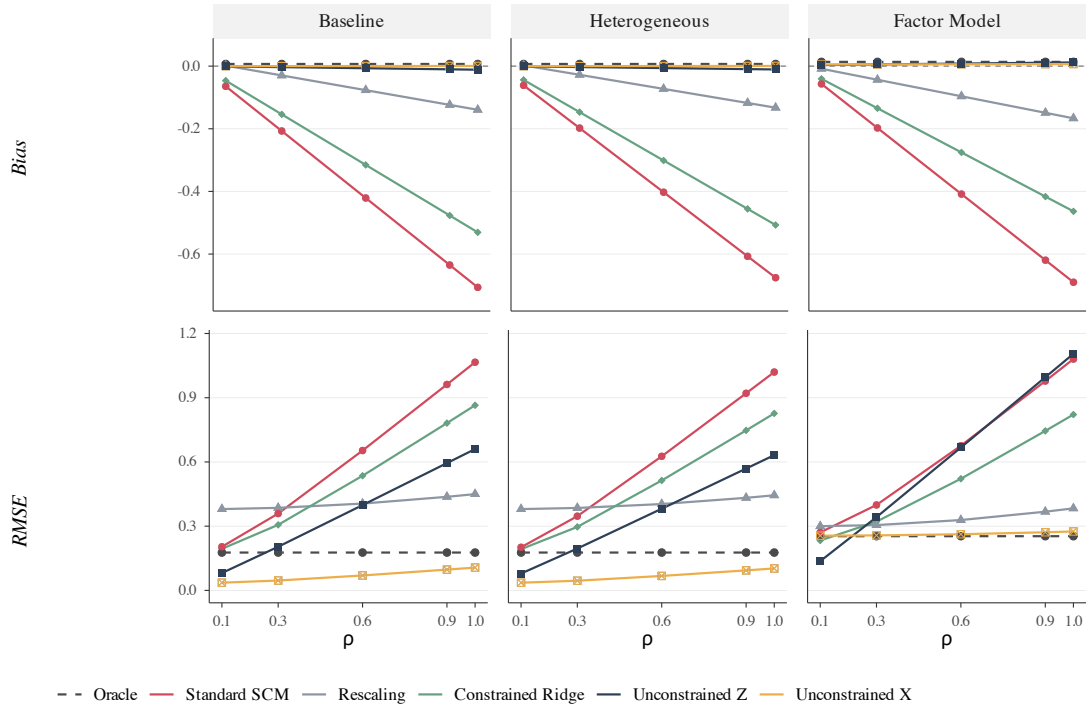


Figure 3: Bias (top) and RMSE (bottom) across the three DGPs (columns) as a function of spillover intensity  $\rho$ . Each panel plots the six estimators under the quantile reach calibration at  $\tau = 4$ ;  $R = 1,000$  replications per cell. Oracle (dashed) gives the unbiased benchmark; URX (solid) tracks Oracle in bias and below it in RMSE in Baseline and Heterogeneous.

in Factor Model, with RMSE below 0.28 throughout. In the Baseline and Heterogeneous DGPs, URX’s RMSE sits below the Oracle’s across the full  $\rho$  grid at  $\tau = 4$ —the exposure-aligned penalty purchases variance reduction that offsets the small negative-weight variance from leaving the simplex (see the estimator definition in §7.1). In the Factor Model DGP, URX tracks the Oracle to within 0.02 in RMSE at  $\rho \leq 0.9$ .

URZ shares URX’s near-zero bias across all three DGPs, but its RMSE diverges sharply in the Factor Model. At  $\rho = 0.9$  there, URZ has RMSE 0.996 against URX’s 0.272, a nearly fourfold difference that reflects instability in the outcome-targeted solve when donor outcome paths load heterogeneously on common factors. The predictor-targeted formulation of URX regularizes on covariate averages in addition to pre-period outcomes, which imposes enough additional structure to stabilize the ridge path. This is the empirical basis for our recommendation of URX as default (Section 6).

A simpler alternative is to drop all donors flagged as exposed before estimating, but trimming is fragile in practice. It requires researchers to identify exactly which donors to drop—itsself unknown in applications, where both the reach map and any threshold are estimated from limited pre-treatment data—it imposes a binary 0/1 decision that collapses a continuous exposure gradient (and discards the factor-structure signal that partial-exposure donors carry), and it offers no theoretical guarantee analogous to Theorem 4. Appendix A.12 reports the trimmed-donor estimator in full; while competitive in some low-interference and latent-factor regimes, it is dominated by URX across most realistic cells and requires exogenous choices that URX avoids.




Figure 4: Bias (top) and 95% coverage (bottom) under Baseline plus four reach-calibration scenarios (A–D) as rows within each column and  $\rho \in \{0.3, 0.6, 0.9\}$  (columns). Scenarios A and B scale  $\gamma_\kappa$  by 2 and 0.5; C imposes a sharp cutoff; D imposes a non-monotone reach. Dashed horizontal line at 0.95 (coverage target) in the lower row.

### 7.3 Sensitivity to reach misspecification

The reach calibration  $(\gamma_\kappa, c)$  is estimated from distances only and involves no outcome data, but it is still a choice the researcher makes. Three distinct kinds of misspecification are at issue: *magnitude misspecification* of the logistic slope (the parametric family is correct, the slope is wrong); *family misspecification* (the parametric family is dropped entirely in favor of a sharp cutoff); and *monotonicity failure* (the reach is non-monotone in distance, which violates the framework’s structural premise). Figure 4 reports performance under four scenarios that span these three kinds: (A)  $\gamma_\kappa$  doubled, (B)  $\gamma_\kappa$  halved, (C) a sharp cutoff reach in place of the logistic, and (D) a non-monotone reach in which spillover intensity peaks at intermediate distance and then declines; the underlying bias, RMSE, and coverage numbers are tabulated in Appendix A.10 (Table 15).

The key pattern is that corrector rankings and coverage survive both magnitude (A, B) and family (C) misspecification, while the non-monotone scenario (D) puts all estimators in the same place because the target itself vanishes. Figure 4 visualizes coverage relative to the 0.95 nominal line.

Getting the magnitude of the logistic slope wrong (scenarios A and B) leaves the picture intact: corrector rankings are preserved and all four correctors hold coverage at or above 0.905 at every  $\rho$ . URX has the smallest RMSE, Rescaling has the smallest bias among simplex-preserving corrections, and Standard SCM is always the worst performer, its coverage dropping below nominal at  $\rho = 0.9$  to 0.880 in scenario A and 0.872 in scenario B.

Dropping the parametric family altogether for a sharp cutoff (scenario C) preserves the same rankings: URX retains  $|\text{Bias}| \leq 0.002$  and the smallest RMSE, while Rescaling and Constrained

Ridge maintain coverage above 0.92. The framework therefore depends on monotonicity of reach in distance, not on the specific parametric form. Standard SCM coverage at  $\rho = 0.9$  under C is 0.931, the closest of the four scenarios to nominal—the sharp-cutoff reach reduces the contamination band size enough that even the unadjusted estimator partially escapes attenuation, though it still loses to all four correctors on RMSE.

The one scenario that breaks the framework, by design, is non-monotone reach (scenario D), which isolates the monotonicity requirement. When reach is non-monotone in distance, the average contamination over the treated unit’s neighbors vanishes in expectation, because contamination from close neighbors is offset by contamination from intermediate neighbors with opposite-sign  $\rho_j$ . Every estimator, including Standard SCM, then has bias within 0.04 of zero under scenario D. This is calibration, not corrector success: the contamination channel itself has gone to zero. The corrections presume a monotone reach by construction. A non-monotone reach is outside the framework’s scope and would in principle be diagnosable from the ring test under a two-sided contrast; empirical validation of this is left to future work.

## 7.4 Detection of interference

The detection question is which ring partition  $K_r$  maximizes power for a given graph density. The pattern that emerges is partition-density matched:  $K_r = 3$  for medium and dense graphs,  $K_r = 4$  for sparse graphs. Table 4 reports rejection rates of the ring-based permutation test from Section 3 across the three treated states and three ring partitions  $K_r \in \{2, 3, 4\}$ ; Figure 5 traces these as a function of  $\rho$ .

Table 4: Ring-Based Permutation Test: Detection Rates by  $K_r$  and  $K_{nb}$

$\rho$	Sparse (California, $K_{nb} = 3$ )			Medium (Minnesota, $K_{nb} = 4$ )			Dense (Missouri, $K_{nb} = 8$ )		
	$K_r = 2$	$K_r = 3$	$K_r = 4$	$K_r = 2$	$K_r = 3$	$K_r = 4$	$K_r = 2$	$K_r = 3$	$K_r = 4$
0 (Size)	0.009	0.010	0.004	0.006	0.016	0.027	0.000	0.008	—
0.1	0.073	0.002	<b>0.345</b>	<b>0.130</b>	<b>0.955</b>	<b>0.962</b>	0.000	<b>0.822</b>	—
0.2	0.074	0.000	<b>0.599</b>	<b>0.157</b>	<b>0.977</b>	<b>0.978</b>	0.000	<b>0.828</b>	—
0.3	0.056	0.000	<b>0.690</b>	<b>0.154</b>	<b>0.980</b>	<b>0.981</b>	0.000	<b>0.826</b>	—
0.5	0.017	0.000	<b>0.707</b>	<b>0.140</b>	<b>0.980</b>	<b>0.982</b>	0.000	<b>0.824</b>	—
0.7	0.002	0.000	<b>0.708</b>	<b>0.118</b>	<b>0.979</b>	<b>0.981</b>	0.000	<b>0.828</b>	—
0.9	0.001	0.000	<b>0.709</b>	0.097	<b>0.979</b>	<b>0.982</b>	0.000	<b>0.826</b>	—

*Notes:* The ring-based permutation test rejects the sharp null  $H_0$ : no interference, i.e., treatment assignment is exchangeable across the treated-unit’s ring structure. Entries are empirical rejection rates at nominal level  $\alpha = 0.05$ ; see Section 3 for the test construction.  $\rho = 0$  row reports empirical size (nominal 0.05).  $K_r$  = number of distance rings.  $K_{nb}$  = graph neighbors of treated state.  $R = 1,000$  permutation draws per specification.  $\tau = 4$ , baseline DGP.

Bold marks rates exceeding 0.10. Placebo skip percentages reported in companion Table 27. Cells shown as — correspond to  $K_r = 4$  in the Missouri (dense) graph, where the fourth ring is empty for the treated unit’s neighborhood so the contrast is undefined. <sup>†</sup>: cells with fewer than 900 valid replications (of 1,000).

Size sits at or below nominal everywhere. At  $\rho = 0$  the rejection rate is  $\leq 0.027$  across the eight available (state,  $K_r$ ) cells (Missouri at  $K_r = 4$  is undefined and reported as --), with most cells below nominal 0.05. The test is finite-sample exact under the sharp null by construction; the empirical undershooting at  $\rho = 0$  reflects the conservatism of the Fisherian reference distribution when the number of ring permutations is modest.

Where the partitions separate is in power, and the ordering depends on graph density. In the medium-density (Minnesota) and dense (Missouri) graphs,  $K_r = 3$  Pareto-dominates  $K_r = 2$ , and in the medium-density case the contrast is extreme:  $K_r = 2$  peaks at a rejection rate of 0.157 at

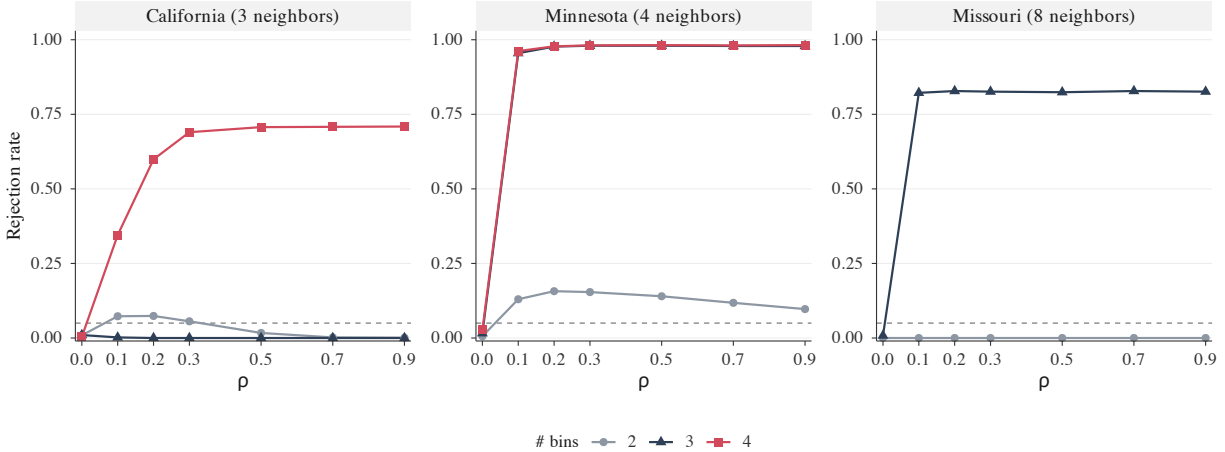


Figure 5: Rejection rate of the ring-based permutation test as a function of  $\rho$ , for three treated states (columns) and three ring partitions ( $K_r \in \{2, 3, 4\}$ , line colors). Nominal level  $\alpha = 0.05$  shown as horizontal dashed line.  $R = 1,000$  replications per  $(\rho, K_r, \text{state})$  cell.

$\rho = 0.2$ , while  $K_r = 3$  reaches 0.955 at  $\rho = 0.1$  and saturates near 0.98 by  $\rho = 0.2$ .

Dense graphs do not let  $K_r = 3$  run all the way to unity, however. In the dense (Missouri) graph,  $K_r = 3$  detection saturates at approximately 0.82. Missouri’s inner ring contains 8 contiguous neighbors; we conjecture that this density limits the number of distinct ring-ordered placebo assignments, imposing a Fisherian upper bound on the attainable rejection rate—roughly 0.82 at  $K_r = 3$ . A formal permutation count is left to future work. This is a feature of the exactness of the test, not a flaw in its design: for very dense graphs the number of available permutations, not the signal-to-noise ratio, is the binding constraint. In the sparse (California) graph, by contrast,  $K_r = 3$  carries essentially zero power: the rejection rate is below 0.01 at every  $\rho \geq 0.1$ , with  $K_r = 2$  winning at every value. The intuition is that a two-ring partition merges the inner exposed band with the outer clean band into only two bins, diluting the proximity signal; a three-ring partition isolates near, medium, and far units and concentrates the test statistic on the near–far contrast where the signal is strongest—but only when the graph supports a non-degenerate three-bin partition.

That is where a fourth ring earns its place. In the sparse case (California),  $K_r = 4$  is the dominant partition, achieving rejection rates up to 0.71 at  $\rho \geq 0.3$ , whereas in the dense case (Missouri)  $K_r = 4$  is undefined because the fourth ring is empty given the treated unit’s local geography.<sup>11,12</sup> The appropriate partition therefore depends on graph density — which is itself diagnosable from the donor pool’s adjacency structure — and we recommend  $K_r = 3$  as the primary report for graphs with  $\geq 4$  first-order neighbors and  $K_r = 4$  as the primary report for sparse graphs ( $\leq 3$  first-order neighbors), subject to the skip-rate check of Table 27.

## 7.5 Takeaways

The simulation evidence delivers four practical conclusions.

Standard SCM is unsuitable when spillovers are plausible: its bias scales linearly with  $\rho$  as Theorem 1 predicts, and coverage degrades below nominal at realistic  $\rho$  (0.871 in the Baseline DGP at  $\rho = 0.9$ ). Among the corrections, URX is the strongest all-purpose choice on RMSE. Its bias

<sup>11</sup>Missouri’s  $K_r = 4$  size cell is reported as -- because the fourth ring is empty in its graph.

<sup>12</sup>California and Minnesota each show a 4.2% placebo skip rate at  $K_r = 4$ ; cells are marked with a dagger when the skip rate exceeds the 0.1% flagging threshold (Table 27). The test remains well-sized conditional on valid permutations.

is indistinguishable from zero in Baseline and Heterogeneous and below 0.007 in Factor Model, and its RMSE falls below the Oracle’s across the full  $\rho$  grid in Baseline and Heterogeneous and within 0.02 of the Oracle in Factor Model; URZ matches URX in pre-period fit but loses post-period stability under a factor-model donor pool, which is why we take URX as the default. When the convex-combination interpretation must be retained, rescaling is the preferred correction, since it eliminates most exposed mass, preserves that interpretation, and holds coverage at or above 0.929 across all reach-misspecification scenarios.


For the ring-based detection test, the appropriate partition depends on graph density. At  $K_r = 3$  the test reaches 0.955 at  $\rho = 0.1$  and saturates near 0.98 by  $\rho = 0.2$  in the medium-density (Minnesota) graph, and reaches rejection rates around 0.82 in the dense (Missouri) graph—a conjectured Fisherian permutation-count ceiling, not a signal-to-noise limit. In the sparse (California) graph,  $K_r = 4$  is the appropriate partition, achieving rejection rates up to 0.71 at  $\rho \geq 0.3$ . Because graph density is itself diagnosable from the donor pool’s adjacency structure, the choice of partition need not be made blind.

Reach misspecification within the monotone family does not overturn any of these rankings. The complete per-cell performance grids, sensitivity tables, and the trimmed-donor benchmark are collected in Appendix A.11; the quantile-versus-moment calibration comparisons and the moment-calibrated figure are in Appendix A.6.


## 8 Replications

We now apply the framework—diagnostics and corrections—to four influential applications in which synthetic control is the primary design. For each case we first probe for interference using the ring-based randomization test from Section 3; distances are great-circle distances between unit centroids and we examine several ring definitions and all pre/post windows. We then discuss *coverage*, i.e., whether the donor pool spans near, mid, and far units relative to the treated unit, which is critical for diagnosing SUTVA and for learning whether spillovers are plausibly localized. Finally, in a setting where richer coverage is attainable, we re-estimate the design with an expanded donor pool and implement the bias-correction tools developed above. All reported  $\hat{p}$  below are the weak permutation  $p$ -value  $\hat{p}^{(w)}$  of Section 3.

The four studies are: Abadie and Gardeazabal (2003) on the economic impact of the Basque Country conflict (Spanish provinces as donors), Ben-Michael et al. (2021) on the Kansas 2012 tax cuts (U.S. states), Abadie et al. (2015) on the economic consequences of German reunification (OECD countries), and Kikuta (2020) on civil war and deforestation in the Democratic Republic of the Congo (cross-national donors). Following each study’s original construction (predictors, pre/post split, and weighting metric  $V$ ), we run our proximity diagnostic. In these original data sets, we do not detect interference at conventional levels:  $\hat{p} = 0.556$  (Basque), 0.327 (Kansas), 0.389 (German reunification, OECD sample), and 0.667 (DRC), as summarized in Table 13 in the appendix.




(a) Abadie et al (2003) Conflict in the Basque  
 $p = 0.556$



(b) Ben-Michael et al (2021) Kansas tax cut  
 $p = 0.327$



(c) Abadie et al (2015) German Reunification  
 $p = 0.389$



(d) Kikuta (2020); Civil war and deforestation  
 $p = 0.667$


Figure 6: Prominent synthetic control applications.

The absence of detected interference in the original designs should not be interpreted as evidence that SUTVA necessarily holds; rather, it highlights a potential design limitation. Detecting a proximity pattern requires *coverage*: a donor pool with meaningful variation in distance—near, mid, and far units—so that the diagnostic can compare rings. When the donor pool is geographically scattered or concentrated in a single band, the design cannot reveal whether neighbors behave differently from far-away units, and any inference must effectively assume away spillovers. In our four cases, the Basque and Kansas designs pool donors from the same national space and exhibit good coverage; by contrast, the German reunification and DRC designs use thin, hand-curated international donor sets with limited spatial structure, which constrains diagnostic learning about SUTVA (Table 5).

Application	Coverage	Interference
Abadie et al (2003)	✓	✗
Ben-Michael et al (2021)	✓	✗
Abadie et al (2015)	✗	✗
Kikuta (2019)	✗	✗
Expanded German Reunification	✓	✓

Table 5: Coverage and interference in selected synthetic control applications.

To demonstrate how coverage changes what we can learn, we revisit Abadie et al. (2015) and expand the donor pool to *all* countries with complete data on the authors’ predictors and outcome, maintaining their pre/post periodization and predictor set. The candidate universe contains roughly 150 countries; the ring diagnostic runs on the  $J = 117$  donors with complete coverage on the outcome and the core detection predictors (the minimal subset used to form the ring-distance statistic, rather than the full Abadie et al. (2015) economic predictor set), and the corrections below are estimated on the 67 donors plus West Germany ( $N = 68$ ) that pass the stricter complete-case filter on the full Abadie et al. (2015) predictor set. With this richer and geographically diverse pool, the ring diagnostic flags a clear proximity pattern around West Germany:  $\hat{p}$  drops from 0.389 in the original OECD sample to 0.017 in the expanded sample, indicating statistically detectable spillovers. The signal is robust to the ring partition (see Table 14): across  $K_r \in \{3, 4, 5, 6, 7\}$ , the test stabilizes at the discrete floor  $\hat{p} = 2/(J + 2) \approx 0.017$ , so the rejection is not an artifact of a particular bin choice.



(a) Abadie et al (2015) German Reunification,  
 $p = 0.389$

(b) Expanded German Reunification,  
 $p = 0.017$

Figure 7: German reunification: original versus expanded donor pools. Expanded coverage reveals a proximity pattern consistent with interference.

Motivated by the diagnostic, we apply the three corrections to the *expanded* German case on the  $N = 68$  sample defined above: covariate rescaling (Section 4.1), constrained ridge (Section 5), and unconstrained ridge (Section 6). We write URX for the predictor-scale variant and URZ for the outcome-scale variant, following Section 6. All tuning is performed on the pre-treatment period only (rolling folds), the exposure map uses distances between centroids and the donor-only logistic calibration of Appendix A.5 with the treated unit as the reference,<sup>13</sup> and we follow Abadie et al.

<sup>13</sup>The rescaling implementation differs in one minor implementation detail between the simulations of Section 7 and

(2015) for outcomes, predictors, and  $V$ . Table 6 reports the average treatment effect on the treated (ATT) and pre-treatment RMSE.

Table 6: German reunification (expanded donor pool,  $N = 68$ ): estimates and pre-treatment fit. Pre-RMSE is the pre-period balance diagnostic used to compare specifications; Post-RMSE is omitted because it is ATT-magnitude-dominated in the expanded German case and is not the relevant comparison metric. The Baseline row reports Abadie et al. (2015)’s published values (Pre-RMSE at two decimals as in the original); all other rows are rounded to one decimal.

Specification	ATT	Pre-RMSE
Baseline SC (authors’ design)	−1549.9	119.08
Rescaled SC (covariate reach)	863.7	761.9
Constrained ridge (simplex)	−1061.5	116.4
Unconstrained ridge (URZ, $Z$ -scale)	−2740.7	79.1
Unconstrained ridge (URX, $X$ -scale)	−2169.1	129.5

The five rows of Table 6 can be read as a single story about the weight geometry developed in Sections 4.1–6. Once near, mid, and far donors are all represented, the diagnostic credibly flags interference ( $\hat{p} = 0.017$ ), and each corrector then moves the estimate in a way that reveals how it acts on the exposed-donor weight mass. Covariate rescaling contracts the columns of proximate donors so aggressively that the feasible simplex shifts to donors with quite different pre-treatment profiles; the point estimate sign-flips from  $-1549.9$  to  $+863.7$ , and pre-period RMSE balloons to  $761.9$ , roughly six times the baseline value of  $119.08$ . This is the explicit cost of forcing convex weights through the rescaled geometry when the exposed band holds most of the good donors, and the pattern is consistent with Germany Expanded sitting near the boundary of the band-dominance condition (Assumption 5) that Theorem 2 requires for exposed-mass elimination: when the safer complement does not cover West Germany’s predictor profile well, the guarantee that drives rescaling’s bias reduction in the simulations of Section 7 degrades into a large fit-cost pivot rather than a clean correction.

Constrained ridge, by contrast, keeps the simplex intact and adds an exposure-aligned  $\ell_2$  penalty that bandwise-shrinks mass away from exposed donors; it attenuates the effect toward zero ( $-1061.5$ ) while leaving pre-period fit essentially unchanged ( $116.4$  against  $119.08$ ), exactly the tradeoff the envelope result in Theorem 3 predicts. The two unconstrained ridge variants relax the simplex and absorb the remaining contamination through active cancellation, and the contrast between them is informative: URZ (the  $Z$ -scale variant) standardizes the donor residuals entering the ridge quadratic and produces both the largest bias correction ( $-2740.7$ ) and the tightest pre-period fit ( $79.1$ , well below baseline), while URX (the  $X$ -scale variant), the closed form developed in Section 6, preserves the predictor scales under which that closed form was derived and magnifies the effect to  $-2169.1$  with a mild pre-fit deterioration ( $129.5$ ). Neither variant dominates on theory alone — the two reflect different choices about what the exposure-weighted penalty regularizes — and we report both as viable options rather than selecting one ex post. Taken together, the five specifications trace the simplex-relaxation and penalty-scale trade-off of Sections 4.1–6: as the corrector relaxes

---

the applications. The applications follow the donor-only calibration of Appendix A.5 as written: we calibrate the logistic reach over the donor-only distances  $\{d_{j,1}\}_{j \neq 1}$ , take the treated unit as the reference, and leave  $X_1$  unscaled. The simulations adopt a different convention—they calibrate over the distances among *all* units and additionally scale the treated column by its own safety score, a factor very close to one. The two conventions therefore coincide up to a common factor very close to one, and the divergence affects only the rescaling estimator, not the two ridge correctors; on the cases we report the resulting difference in the rescaled weights is numerically immaterial.

convex weights and penalizes exposed mass more aggressively, it delivers larger bias corrections, at a pre-fit cost that depends on how the geometry is relaxed and on the loss of the convex-combination interpretation that classical SC trades for transparency (Abadie & L’Hour, 2021; Ben-Michael et al., 2021; Doudchenko & Imbens, 2016).

Taken together, these replications show how the proposed workflow operates end to end. Coverage is a prerequisite for learning about interference; once near, mid, and far donors are present, the ring diagnostic can flag spillovers; and when interference is detected, proximity-aware corrections act on the exposed-donor weight mass identified in Theorem 1 and substantially change the substantive conclusions. The magnitude of correction and the pre-period fit cost depend on how far the corrector relaxes the simplex and how aggressively it penalizes exposed weight. In settings where donor pools are discretionary or thin, expanding coverage is itself a design improvement that enables credible diagnostics and targeted bias correction.

## 9 Practical Guidance

A researcher contemplating synthetic control in a setting where the treatment might spill onto neighbors faces a sequence of decisions, each of which the preceding sections have addressed in isolation. We gather them here into a single decision flow. The order is not incidental: each step either licenses the next or settles the analysis, so that the workflow narrows from a design question about the donor pool, to a diagnostic about whether SUTVA is at risk, to a choice among corrections that is governed by how much the convex-combination interpretation is worth in the application at hand.

The first question is whether the donor pool can support any of what follows. The ring diagnostic of Section 3 adjudicates SUTVA by contrasting outcome changes in near and far rings, and the corrections of Sections 4–6 route weight from exposed donors toward safer ones. Both presuppose coverage: a pool with meaningful variation in distance, spanning near, mid, and far units relative to the treated unit. Without it the diagnostic cannot compare rings and the corrections cannot target an exposed band credibly, since there is no safer complement to absorb the displaced weight (Assumption 5). As the replications in Section 8 make plain, a thin or single-band donor pool forces the analysis to assume spillovers away rather than test for them; where coverage is discretionary, expanding it is itself the first design improvement, not a preliminary to one.

With coverage in place, the diagnostic is cheap to run and exact in finite samples, which makes it a natural pre-analysis check and a pre-registerable decision rule (Section 3.5). The researcher fixes the ring boundaries, the window, and the contrast in advance, without reference to outcome data, and reads the permutation  $p$ -value of (20). A large  $p$ -value across pre-specified windows is evidence of no proximity-structured disturbance: standard synthetic control proceeds, and the matter is closed. A small  $p$ -value signals that the treated unit sits at the center of an atypical near–far pattern, so SUTVA should be treated as violated and the donors in proximal rings as carrying nonzero exposure. In the bias decomposition of Theorem 1, this is the regime in which the exposed weight mass  $\sum_{j \in \mathcal{E}} w_j$  is positive and the contamination term is active; the remaining decisions concern how to suppress it.

The corrector choice is governed by how much the convex-combination interpretation is worth, and we recommend escalating only as far as the application requires. Rescaling (Section 4) is the natural starting point. It contracts the predictor columns of exposed donors, leaves the weights on the simplex, and so allows the familiar synthetic-control diagnostics—pre-period fit, placebo-in-space, weight sparsity—to transfer without adjustment; under the band-dominance condition (Assumption 5) it eliminates exposed mass entirely. Its one cost, pre-period fit degradation when the pool lacks safe donors that span the relevant predictor directions, is read directly off the pre-period

residuals. When rescaling materially degrades that fit, move to constrained ridge (Section 5), which keeps the simplex and its interpretation but acts on the weights through an exposure-aligned  $\ell_2$  penalty, trading a continuous amount of pre-period fit for a continuous reduction in exposed mass as  $\lambda$ , selected by cross-validation on pre-treatment data, increases. Reserve the unconstrained ridge (Section 6) for the case in which neither simpler correction reduces bias without destroying fit. Relaxing the simplex admits negative weights that can actively cancel contaminated donors and a level shift that can absorb a baseline gap, at the price of the convex-combination interpretation; URX, the predictor-targeted closed form, is the default, with URZ reported alongside it. Whatever the corrector, report the diagnostic  $p$ -value, the pre-period fit at  $\lambda = 0$  and at the chosen specification, and—when the simplex is relaxed—the fact that the convex-combination interpretation has been given up in exchange for stronger bias control.

The workflow can be stated as a checklist.

1. **Establish coverage.** Confirm that the donor pool spans near, mid, and far units relative to the treated unit. Without this gradient the ring diagnostic cannot adjudicate SUTVA and the corrections have no safer complement to target (Assumption 5). Where coverage is thin and discretionary, expand the donor pool first; this is itself a design improvement (Section 8).
2. **Run the ring diagnostic.** Pre-specify ring boundaries, windows, and the near–far contrast without reference to outcomes, and compute the finite-sample exact permutation  $p$ -value of (20) (Section 3). The check is cheap and pre-registerable.
3. **Read the verdict.** A large  $p$ -value gives no evidence of a proximity-structured disturbance—proceed with standard synthetic control. A small  $p$ -value indicates active contamination ( $\sum_{j \in \mathcal{E}} w_j > 0$  in Theorem 1)—move to a correction.
4. **Choose the corrector by escalation.**
  - (a) Start with *rescaling* (Section 4): simplex-preserving, diagnostics transfer, exposed mass eliminated under band dominance (Assumption 5).
  - (b) If rescaling materially degrades pre-period fit, use *constrained ridge* (Section 5) and tune  $\lambda$  by cross-validation on the pre-treatment period; this retains the simplex while trading fit for exposed-mass reduction.
  - (c) Reserve *unconstrained ridge* (Section 6, URX default) for when neither simpler correction reduces bias without destroying fit, reporting that the convex-combination interpretation has been relaxed.
5. **Report transparently.** State the diagnostic  $p$ -value, the pre-period fit for each specification considered, and the interpretation trade-off accepted by the corrector chosen.

## 10 Conclusion

SUTVA is indispensable for causal inference designs, and it is particularly consequential for Synthetic Control methods because post-treatment donor outcomes enter the counterfactual directly. We formalized this channel via a bias decomposition in which contamination operates through the exposed weight share: the estimator differs from the target by a term proportional to  $\sum_j \rho_j w_j$ . This channel is mechanical: once positively weighted donors are exposed, bias follows, and it motivates both diagnosis and design modifications that act on the SCM weights.

We proposed a design-based diagnostic that tests for proximity-patterned outcome changes around the treated unit using ring partitions and Fisherian randomization inference. The test is

exact in finite samples under the sharp null and offers a pre-analysis decision rule: proceed with standard SC when proximity patterns are not detected; otherwise, treat SUTVA as violated and modify the design. The diagnostic requires coverage across distance rings, as without near–mid–far donors, power is limited and SUTVA violations cannot be properly assessed.

We then introduced three interference-aware adjustments that incorporate an exogenous exposure map into SC while tuning only on pre-period data. *Covariate rescaling* contracts high-risk donor columns, leaving weights on the simplex; under a mild coverage condition, the exposed mass cannot increase. *Constrained ridge* preserves the convex-combination interpretation but adds an exposure-aligned  $\ell_2$  penalty; KKT geometry implies bandwise dominance, so total weight on exposed donors weakly (and often strictly) falls. *Unconstrained ridge* relaxes the simplex, allows negative weights and level shifts, and penalizes via an exposure-weighted norm; a simple Cauchy–Schwarz argument and ridge-path monotonicity deliver a tuning-monotone envelope for the contamination term.

Simulations corroborate these mechanisms. Baseline SC exhibits attenuation that scales linearly with spillover intensity, as Theorem 1 predicts; constrained ridge shrinks this attenuation along the envelope of Theorem 3; rescaling reduces bias by a factor of five to six relative to baseline; and the unconstrained ridge estimators track the oracle closely in bias, with the predictor-targeted variant (URX) delivering RMSE at or below the oracle in the Baseline and Heterogeneous DGPs and within 0.02 of the oracle under a latent-factor donor pool. The rankings are robust to magnitude- and family-misspecification of the reach map and survive a sharp-cutoff alternative to the logistic; only a non-monotone reach, which lies outside the framework’s structural premise, erases the contamination channel the corrections target.

Replications illustrate the workflow. Applied to the original donor pools in Abadie and Gardeazabal (2003), Ben-Michael et al. (2021), Abadie et al. (2015), and Kikuta (2020), the ring diagnostic does not reject at conventional levels, consistent with either SUTVA or with donor pools whose spatial structure is too thin to adjudicate it. Expanding the Abadie et al. (2015) design to a broad, data-complete donor pool restores coverage and yields a detected proximity pattern ( $\hat{p} = 0.017$ ). The corrections then move the estimate in directions consistent with their geometry, tracing the trade-off between simplex relaxation and penalty scale: covariate rescaling magnifies the response by reshaping the feasible simplex but at a substantial pre-fit cost, constrained ridge attenuates toward zero while preserving baseline pre-fit, and the two unconstrained ridge variants deliver the largest bias corrections. URZ (the  $Z$ -scale variant) also tightens pre-period fit below baseline. URX (the  $X$ -scale variant), the closed form of Section 6, produces a smaller but still substantial correction. The package-level message is that the magnitude of correction, and whether pre-fit improves or deteriorates, is determined jointly by how much each corrector relaxes the convex-combination constraint and how aggressively it penalizes exposed weight mass.

Section 9 distills these results into an explicit decision workflow for applied practice—from establishing donor-pool coverage, to running the diagnostic, to escalating through the corrections only as far as the application requires.

Two qualifications delimit the scope of what we have established. Diagnostic power depends on coverage, so thin or discretionary donor pools limit what the test can say about SUTVA; the decision to expand coverage is itself a design choice whose cost and feasibility are application-specific. Design guarantees rely on an exogenous exposure map aligned with the actual channels of spillover, and while our simulations indicate that performance is robust to magnitude and family misspecification of the reach, it is not robust to non-monotone reach — a case we treat as outside the framework’s scope and that future work could diagnose through a two-sided contrast in the ring test. Natural extensions include constructing reach from auxiliary mobility, trade, or network data under monotonicity constraints, and developing inference for the corrected estimators along the lines of Cattaneo et al.

(2021) and Chernozhukov et al. (2021). The broader contribution is to recast interference handling in synthetic control as a geometric design problem: the contamination channel is the weight placed on exposed donors, the diagnostic decides whether that channel is active, and the corrections reshape it directly. This preserves the transparency and pre-period discipline that make SC useful, while addressing the vulnerability that comes with placing concentrated weight on donors chosen for the very similarities that make them likely to be exposed.

## References

- Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, *59*(2), 391–425. <https://doi.org/10.1257/jel.20191450>
- Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program. *Journal of the American Statistical Association*, *105*(490), 493–505. <https://doi.org/10.1198/jasa.2009.ap08746>
- Abadie, A., Diamond, A., & Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, *59*(2), 495–510. <https://doi.org/10.1111/ajps.12116>
- Abadie, A., & Gardeazabal, J. (2003). The economic costs of conflict: A case study of the basque country. *American Economic Review*, *93*(1), 113–132. <https://doi.org/10.1257/000282803321455188>
- Abadie, A., & L’Hour, J. (2021). A penalized synthetic control estimator for disaggregated data. *Journal of the American Statistical Association*, *116*(536), 1817–1834. <https://doi.org/10.1080/01621459.2021.1929248>
- Abadie, A., & Vives-i-Bastida, J. (2025). Synthetic controls in action [arXiv preprint arXiv:2203.06279 (2022)]. In *Advances in economics and econometrics: Twelfth world congress*. Cambridge University Press. <https://arxiv.org/abs/2203.06279>
- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic difference-in-differences. *Journal of the American Statistical Association*, *116*(536).
- Aronow, P. M., & Samii, C. (2017). Estimating average causal effects under general interference. *The Annals of Applied Statistics*, *11*(4), 1912–1947. <https://doi.org/10.1214/17-AOAS1043>
- Athey, S., Bayati, M., Doudchenko, N., Imbens, G. W., & Khosravi, K. (2021). Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, *116*(536), 1716–1730. <https://doi.org/10.1080/01621459.2021.1891924>
- Ben-Michael, E., Feller, A., & Rothstein, J. (2021). The augmented synthetic control method. *Journal of the American Statistical Association*, *116*(536), 1789–1803. <https://doi.org/10.1080/01621459.2021.1929249>
- Berry, F. S., & Berry, W. D. (1990). State lottery adoptions as policy innovations: An event history analysis. *American Political Science Review*, *84*(2), 395–415. <https://doi.org/10.2307/1963526>
- Berry, W. D., & Baybeck, B. (2005). Using geographic information systems to study interstate competition. *American Political Science Review*, *99*(4), 505–519.
- Bickel, P. J., Klaassen, C. A. J., Ritov, Y., & Wellner, J. A. (1993). *Efficient and adaptive estimation for semiparametric models*. Johns Hopkins University Press.
- Bowers, J., Fredrickson, M. M., & Panagopoulos, C. (2017). Reasoning about interference between units: A general framework. *Political Analysis*, *21*(1), 97–124. <https://doi.org/10.1093/pan/mps038>
- Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511804441>

- Butts, K. (2021). *Difference-in-differences estimation with spatial spillovers* [Working paper, arXiv:2105.03737 (revised December 2024)]. <https://doi.org/10.48550/arXiv.2105.03737>
- Cao, F., & Dowd, J. (2019). Network spillovers in synthetic control frameworks. *Economics Letters*, *176*, 88–94.
- Cattaneo, M. D., Feng, Y., & Titiunik, R. (2021). Prediction intervals for synthetic control methods. *Journal of the American Statistical Association*, *116*(536), 1865–1880. <https://doi.org/10.1080/01621459.2021.1979561>
- Cavallo, E., Galiani, S., Noy, I., & Pantano, J. (2013). Catastrophic natural disasters and economic growth. *Review of Economics and Statistics*, *95*(5), 1549–1561. [https://doi.org/10.1162/REST\\_a\\_00413](https://doi.org/10.1162/REST_a_00413)
- Chernozhukov, V., Wüthrich, K., & Zhu, Y. (2021). An exact and robust conformal inference method for counterfactual and synthetic controls. *Journal of the American Statistical Association*, *116*(536), 1849–1864. <https://doi.org/10.1080/01621459.2021.1920957>
- Clarke, D. (2017). *Estimating difference-in-differences in the presence of spillovers* [Working paper, MPRA Paper No. 81604, University Library of Munich]. <https://mpra.ub.uni-muenchen.de/81604/>
- Delgado, M., & Florax, R. (2015). Spatial extensions of difference-in-differences. *Regional Science and Urban Economics*, *54*, 123–138.
- Di Stefano, R., & Mellace, G. (2020). Indirect treatment effects in synthetic control methods. *Econometric Reviews*, *39*(5), 628–657.
- Doudchenko, N., & Imbens, G. W. (2016). Balancing, regression, difference-in-differences and synthetic control methods: A synthesis. *NBER Working Paper No. 22791*. <https://doi.org/10.3386/w22791>
- Ferman, B. (2021). On the properties of the synthetic control estimator with many periods and many controls. *Journal of the American Statistical Association*, *116*(536), 1764–1772. <https://doi.org/10.1080/01621459.2021.1965613>
- Ferman, B., & Pinto, C. (2021). Synthetic controls with imperfect pre-treatment fit. *Journal of Business & Economic Statistics*, *39*(1), 19–33. <https://doi.org/10.1080/07350015.2019.1618534>
- Firpo, S., & Possebom, V. (2018). Synthetic control method: Inference, sensitivity analysis and confidence sets. *Journal of Causal Inference*, *6*(2), 20170023. <https://doi.org/10.1515/jci-2017-0023>
- Fisher, R. A. (1935). *The design of experiments*. Collier Macmillan.
- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, *102*(477), 359–378. <https://doi.org/10.1198/016214506000001437>
- Gobillon, L., & Magnac, T. (2016). Regional policy evaluation: Interactive fixed effects and synthetic controls. *Review of Economics and Statistics*, *98*(3), 535–551.
- Graham, E. R., Shipan, C. R., & Volden, C. (2013). The diffusion of policy diffusion research in political science. *British Journal of Political Science*, *43*(3), 673–701. <https://doi.org/10.1017/S0007123412000415>
- Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, *81*(396), 945–960. <https://doi.org/10.1080/01621459.1986.10478354>
- Hudgens, M. G., & Halloran, M. E. (2008). Toward causal inference with interference. *Journal of the American Statistical Association*, *103*(482), 832–842. <https://doi.org/10.1198/016214508000000292>

- Kikuta, K. (2020). The environmental costs of civil war: A synthetic comparison of the congolese forests with and without the great war of africa. *The Journal of Politics*, 82(4), 1243–1255.
- Kreif, N., Grieve, R., Harrison, D. A., & Didelez, V. (2016). Examination of the synthetic control method for evaluating health policies with multiple treated units. *Health Economics*, 25(12), 1514–1528. <https://doi.org/10.1002/hec.3370>
- Manski, C. F. (2013). Identification of treatment response with social interactions. *The Econometrics Journal*, 16(1), S1–S23. <https://doi.org/10.1111/ectj.12002>
- Robbins, M., & Saunders, J. (2017). A framework for synthetic control methods with high-dimensional, micro-level data: Evaluating a neighborhood-specific crime intervention. *Journal of the American Statistical Association*.
- Rosenbaum, P. R. (2002). *Observational studies* (2nd ed.). Springer. <https://doi.org/10.1007/978-1-4757-3692-2>
- Rosenbaum, P. R. (2007). Interference between units in randomized experiments. *Journal of the American Statistical Association*, 102(477), 191–200.
- Rubin, D. B. (1980). Randomization analysis of experimental data: The fisher randomization test comment. *Journal of the American Statistical Association*, 75(371), 591–593. <https://www.jstor.org/stable/2287653>
- Sävje, F., Aronow, P. M., & Hudgens, M. G. (2021). Average treatment effects in the presence of unknown interference. *Annals of Statistics*, 49(2), 673–701.
- Shipan, C. R., & Volden, C. (2008). The mechanisms of policy diffusion. *American Journal of Political Science*, 52(4), 840–857. <https://doi.org/10.1111/j.1540-5907.2008.00346.x>
- Tchetgen Tchetgen, E., & VanderWeele, T. M. (2012). On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21(1), 55–75. <https://doi.org/10.1177/0962280210386779>
- van der Vaart, A. W. (1998). *Asymptotic statistics*. Cambridge University Press.
- VanderWeele, T. J., Tchetgen Tchetgen, E. J., & Halloran, M. E. (2014). Interference and sensitivity analysis. *Statistical Science*, 29(4), 687–706.
- Vega, S. H., & Elhorst, J. P. (2015). The SLX model. *Journal of Regional Science*, 55(3), 339–363. <https://doi.org/10.1111/jors.12188>
- Viviano, D. (2024). Policy targeting under network interference [Published online April 2024; formal issue March 2025]. *The Review of Economic Studies*, 92(2), 1257–1292. <https://doi.org/10.1093/restud/rdae041>
- Volden, C. (2002). The politics of competitive federalism: A race to the bottom in welfare benefits? *American Journal of Political Science*, 46(2), 352–363. <https://doi.org/10.2307/3088381>
- Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis*, 25(1), 57–76.

## Notation

The paper uses a single set of symbols consistently. Tables 7 and 8 collect the notation introduced across the setup, detection, rescaling, and ridge sections. A reader can return to these tables as a reference throughout.

Table 7: Notation used throughout Sections 2–6 (Part 1 of 2).

Symbol	Meaning	First used
<i>Units, time, and outcomes</i>		
$i$	Unit index; $i = 1$ is the treated unit, $i \in \{2, \dots, J + 1\}$ are donors	§2
$\mathcal{J}$	Donor set, $\mathcal{J} \subseteq \{2, \dots, J + 1\}$ ; $J =  \mathcal{J} $	§2
$t, T_0$	Time index; intervention time	§2
$Y_{it}(z)$	Potential outcome for unit $i$ at time $t$ under assignment vector $z \in \{0, 1\}^{J+1}$	§2
$Y_{it}^{\text{obs}}$	Observed outcome; equals $Y_{it}(z)$ at the realized $z$	§2
$Y_{it}(0)$	No-treatment potential outcome (no direct treatment, no indirect exposure)	§2
$\tau_t$	Direct treatment effect at time $t$ : $Y_{1t}^{\text{obs}} = Y_{1t}(0) + \tau_t$	§2
$\rho_j$	Spillover intensity for donor $j$ , $\rho_j \in [0, 1]$ ; $\rho_j = \rho$ in the homogeneous case	§2
$\delta_{jt}$	General (not necessarily proportional) spillover on donor $j$ at time $t$	§2, Fig. 1
<i>Exposure and the set partition</i>		
$\mathcal{E} \subseteq \mathcal{J}$	<b>Exposed set:</b> donors for which $Y_{jt}^{\text{obs}} = Y_{jt}(0) + \rho_j \tau_t$	§2
$\mathcal{S} = \mathcal{J} \setminus \mathcal{E}$	<b>Safe set:</b> donors for which $Y_{jt}^{\text{obs}} = Y_{jt}(0)$	§2
$\mathcal{C}(m)$	<b>Donor-specific covering set:</b> subset of $\mathcal{J} \setminus \{m\}$ used in the donor-level dominance assumption (§4). Under the band-dominance specialization, $\mathcal{C}(m) \subseteq \mathcal{S}$ .	§4
$\mathcal{U}$	Set of all units, $\mathcal{U} = \{1, \dots, J + 1\}$	§3
<i>Synthetic control estimator</i>		
$X_1 \in \mathbb{R}^K$	Treated unit's pre-treatment predictor vector ( $K$ predictors)	§2
$X_0 \in \mathbb{R}^{K \times J}$	Donor predictor matrix; column $j$ is $X_{\cdot, j}$	§2
$V$	$K \times K$ diagonal importance matrix, $V \succ 0$	§2
$w \in \Delta_J$	SCM weight vector; $\Delta_J = \{w \geq 0 : \mathbf{1}^\top w = 1\}$	§2
$Y_{1t}^{\text{SC}}(0)$	SUTVA synthetic counterfactual, $\sum_{j \in \mathcal{J}} w_j Y_{jt}(0)$	§2
$\hat{\tau}_t$	SCM estimator of the treatment effect at time $t$ : $\hat{\tau}_t = Y_{1t}^{\text{obs}} - \sum_{j \in \mathcal{J}} w_j Y_{jt}^{\text{obs}}$	§2

Table 8: Notation used throughout Sections 2–6 (Part 2 of 2).

Symbol	Meaning	First used
<i>Detection diagnostic (§3)</i>		
$\kappa(d; \theta)$	Exposure kernel (nonincreasing in distance $d$ ); $\theta$ collects kernel parameters	§3
$\gamma_\kappa$	Logistic steepness parameter (Appendix A.5). Distinct from the exposure kernel $\kappa(\cdot; \theta)$ above.	App. A.5
$d_{ij}$	Distance (geographic or network) between units $i$ and $j$	§3
$E_{it}$	Scalar exposure of unit $i$ at time $t$ : $E_{it} = \sum_j \kappa(d_{ij}; \theta) z_{jt}$	§3
$K_r$	Ring count in the detection diagnostic (renamed from $K$ in §3 to avoid collision with the predictor count $K$ of §2)	§3.1
$r_{i,p} \in \{1, \dots, K_r\}$	Ring index of donor $i$ with center $p$ ; boundaries $0 = c_0 < \dots < c_{K_r}$	§3.1
$A_p, B_p$	Proximal and distal donor groups around center $p$ ( $A_p \cup B_p = \mathcal{U} \setminus \{p\}$ )	§3.1
$\Delta_i^{(w)}$	Outcome-change statistic: $\bar{Y}_{i,\text{post}}^{(w)} - \bar{Y}_{i,\text{pre}}^{(w)}$ (window $w$ )	§3.2
$t_p^{(w)}$	Two-sample $t$ -statistic contrasting $\bar{\Delta}_{A_p}^{(w)}$ and $\bar{\Delta}_{B_p}^{(w)}$	§3.2
<i>Reach map, rescaling, and ridge (§4–§6)</i>		
$f: \mathbb{R}_+ \rightarrow (0, 1)$	Reach map: continuous, strictly increasing in distance to the treated unit	§4.1
$d_{j,1}$	Distance from donor $j$ to the treated unit ( $i = 1$ ); the $(j, 1)$ entry of the distance matrix $\{d_{ij}\}$	§4.1
$\eta_j = f(d_{j,1})$	<b>Safety score</b> of donor $j$ : larger $\eta_j$ indicates <i>lower</i> exposure risk	§4.1
$\eta = (\eta_1, \dots, \eta_J)^\top$	Vector of safety scores; used in the rescaled donor matrix $X_0^* = X_0 \text{diag}(\eta)$	§4
$\psi_j = g(\eta_j)$	Exposure-aligned penalty weight; $g$ strictly decreasing <sup>a</sup>	§5
$\underline{\psi}, \bar{\psi}$	$\underline{\psi} := \inf_j \psi_j$ and $\bar{\psi} := \sup_j \psi_j$ ; bounds on the penalty schedule	§5
$\underline{\psi}_\mathcal{E}$	$\min_{j \in \mathcal{E}} \psi_j$ ; lower bound on the penalty over the exposed band	§5
$D_\psi$	$J \times J$ diagonal matrix $\text{diag}(\psi_1, \dots, \psi_J)$	§5
$X_0^*$	Rescaled donor matrix $X_0 \text{diag}(\eta)$	§4
$\lambda \geq 0$	Ridge tuning parameter	§5
$w^\circ$	Unpenalized SCM minimizer over $\Delta_J$ (baseline)	§4
$w^*$	Rescaled-SCM minimizer (diag- $\eta$ contraction), on $\Delta_J$	§4
$w^{CR,(\lambda)}$	Constrained-ridge minimizer at tuning $\lambda$ on $\Delta_J^b$	§5
$w^{UR,(\lambda)}$	Unconstrained-ridge minimizer at tuning $\lambda$ over $\mathbb{R}^{J^b}$	§6
$c_0, c_\lambda$	Exposed weight mass under baseline ( $c_0 = \sum_{j \in \mathcal{E}} w_j^\circ$ ) and under constrained ridge ( $c_\lambda = \sum_{j \in \mathcal{E}} w_j^{(\lambda)}$ )	§5

<sup>a</sup>  $\psi_j$  denotes the exposure-aligned penalty schedule throughout; it should not be confused with influence-function notation in the semiparametric literature. See footnote at first use in §5.

<sup>b</sup> The appendix proofs of Theorems 3 and 4 retain the equivalent parenthesized-argument forms  $w^{(\lambda)}$  and  $w^{UR}(\lambda)$ ; these denote the same mathematical objects.

*Conventions.* Calligraphic letters ( $\mathcal{E}, \mathcal{S}, \mathcal{J}, \mathcal{U}, \mathcal{C}(m)$ ) denote sets of units. Unadorned  $\mathcal{S}$  always denotes the safe set  $\mathcal{J} \setminus \mathcal{E}$ ; the donor-specific covering set in the dominance arguments of §4 carries an explicit donor argument,  $\mathcal{C}(m)$ . Hats ( $\hat{\tau}_t$ ) denote estimators; stars ( $w^*$ ), circles ( $w^\circ$ ), and tuning-indexed superscripts ( $w^{CR,(\lambda)}$ ,  $w^{UR,(\lambda)}$ ) distinguish minimizers across the four weight problems considered.

# Appendix A Proofs and Supplementary Results

## Contents

A.1 Proofs of Theorem 1 and Corollary 1 . . . . .	47
A.2 Proof of Theorem 2 . . . . .	49
A.3 Proof of Theorem 3 . . . . .	51
A.4 Proof of Theorem 4 . . . . .	55
A.5 Spatial reach mapping . . . . .	59
A.6 Reach calibration: quantile versus moment robustness . . . . .	60
A.7 Ring geometry and mapping from continuous to discretized exposure scores . . . . .	65
A.8 Alternative rings contrast . . . . .	66
A.9 Ring diagnostic across replication applications . . . . .	67
A.10 Sensitivity to reach misspecification: summary table . . . . .	68
A.11 Full simulation tables . . . . .	69
A.12 Trimming as a baseline . . . . .	69

## A.1 Proofs of Theorem 1 and Corollary 1

**THEOREM 1 (BIAS OF SCM UNDER LOCALIZED SPILLOVERS, RESTATED).** *For  $t > T_0$ , the SCM estimator based on observed donors satisfies*

$$\hat{\tau}_t = \underbrace{[Y_{1t}(0) - Y_{1t}^{\text{SC}}(0)]}_{\text{SUTVA synthetic mismatch}} + \underbrace{\tau_t}_{\text{direct effect}} - \underbrace{\tau_t \sum_{j \in \mathcal{E}} \rho_j w_j}_{\text{contamination mass}} .$$

*In particular, if  $Y_{1t}^{\text{SC}}(0) = Y_{1t}(0)$  (the SUTVA synthetic counterfactual coincides with the treated unit's no-treatment path at  $t$ ), then*

$$\hat{\tau}_t = \tau_t - \tau_t \sum_{j \in \mathcal{E}} \rho_j w_j .$$

**Setting.** Throughout this proof: (A1) fix a post-treatment period  $t > T_0$ ; (A2) the weights  $w = (w_j)_{j \in \mathcal{J}}$  satisfy  $w_j \geq 0$  and  $\sum_{j \in \mathcal{J}} w_j = 1$ ; (A3) the interference set is  $\mathcal{E} \subseteq \mathcal{J}$  (donors exposed to spillovers in post-periods), and for  $j \notin \mathcal{E}$ ,  $Y_{jt}^{\text{obs}} = Y_{jt}(0)$ ; (A4) the diffusion intensities  $(\rho_j)_{j \in \mathcal{E}}$  satisfy  $\rho_j \in [0, 1]$  and are time-homogeneous, and for  $j \in \mathcal{E}$ ,  $Y_{jt}^{\text{obs}} = Y_{jt}(0) + \rho_j \tau_t$ ; (A5) the treated unit satisfies  $Y_{1t}^{\text{obs}} = Y_{1t}(0) + \tau_t$ ; (A6) the SUTVA synthetic target (the no-interference counterfactual) is  $Y_{1t}^{\text{SC}}(0) := \sum_{j \in \mathcal{J}} w_j Y_{jt}(0)$ .

**Proof of Theorem 1.** The proof is a direct substitution argument: insert the observed outcomes of the treated unit and of each donor into the definition of the SCM estimator, split the donor sum by exposure status, and re-aggregate the baseline donor terms into the SUTVA synthetic counterfactual.

By definition of the SCM post-period effect,

$$\hat{\tau}_t = Y_{1t}^{\text{obs}} - \sum_{j \in \mathcal{J}} w_j Y_{jt}^{\text{obs}} .$$

Substituting the treated unit's observed outcome from (A5) gives

$$\hat{\tau}_t = (Y_{1t}(0) + \tau_t) - \sum_{j \in \mathcal{J}} w_j Y_{jt}^{\text{obs}} .$$

Splitting the donor sum over the partition  $\mathcal{J} = \mathcal{E} \cup (\mathcal{J} \setminus \mathcal{E})$  and then substituting the donors' observed outcomes by exposure status,

$$\begin{aligned}
\hat{\tau}_t &= (Y_{1t}(0) + \tau_t) - \sum_{j \in \mathcal{E}} w_j Y_{jt}^{obs} - \sum_{j \in \mathcal{J} \setminus \mathcal{E}} w_j Y_{jt}^{obs} && \text{[split by exposure status]} \\
&= (Y_{1t}(0) + \tau_t) - \sum_{j \in \mathcal{E}} w_j (Y_{jt}(0) + \rho_j \tau_t) - \sum_{j \in \mathcal{J} \setminus \mathcal{E}} w_j Y_{jt}(0) && \text{[by (A3)–(A4)]} \\
&= (Y_{1t}(0) + \tau_t) - \sum_{j \in \mathcal{E}} w_j Y_{jt}(0) - \tau_t \sum_{j \in \mathcal{E}} \rho_j w_j - \sum_{j \in \mathcal{J} \setminus \mathcal{E}} w_j Y_{jt}(0) && \text{[linearity of the sum]} \\
&= (Y_{1t}(0) + \tau_t) - \sum_{j \in \mathcal{J}} w_j Y_{jt}(0) - \tau_t \sum_{j \in \mathcal{E}} \rho_j w_j && \text{[combine the two baseline sums].}
\end{aligned}$$

Introducing the SUTVA synthetic counterfactual  $Y_{1t}^{SC}(0) = \sum_{j \in \mathcal{J}} w_j Y_{jt}(0)$  from (A6) and rearranging,

$$\hat{\tau}_t = (Y_{1t}(0) - Y_{1t}^{SC}(0)) + \tau_t - \tau_t \sum_{j \in \mathcal{E}} \rho_j w_j.$$

Labeling the three terms yields the decomposition

$$\hat{\tau}_t = \underbrace{(Y_{1t}(0) - Y_{1t}^{SC}(0))}_{\text{SUTVA synthetic mismatch}} + \underbrace{\tau_t}_{\text{direct effect}} - \underbrace{\tau_t \sum_{j \in \mathcal{E}} \rho_j w_j}_{\text{contamination mass}},$$

which is (5). The special case (6) follows by setting the mismatch term to zero, that is,  $Y_{1t}^{SC}(0) = Y_{1t}(0)$ .  $\square$

**COROLLARY 1 (SIGN, ATTENUATION, AND SCALE, RESTATED).** *Suppose  $Y_{1t}^{SC}(0) = Y_{1t}(0)$  and  $\rho_j \in [0, 1]$  for all  $j \in \mathcal{E}$ . Then:*

1. (*Sign/attenuation*) *If  $\sum_{j \in \mathcal{E}} \rho_j w_j > 0$  and  $\tau_t > 0$ , then  $\hat{\tau}_t < \tau_t$ ; if  $\sum_{j \in \mathcal{E}} \rho_j w_j > 0$  and  $\tau_t < 0$ , then  $\hat{\tau}_t > \tau_t$ . That is, positive weights on exposed donors attenuate estimated effects toward zero.*
2. (*Monotonicity in exposure*) *Holding  $\tau_t$  fixed,  $|\hat{\tau}_t - \tau_t|$  is nondecreasing in each  $\rho_j$  and in the exposed weight mass  $\sum_{j \in \mathcal{E}} w_j$ . More intense or more widespread spillovers cannot reduce bias.*
3. (*Bounds*) *Because  $0 \leq \rho_j \leq 1$  and  $w \in \Delta$ , the distortion satisfies  $0 \leq |\hat{\tau}_t - \tau_t| \leq |\tau_t|$ . The upper bound is attained only if the entire weight is placed on fully exposed donors ( $\rho_j = 1$  for all  $j$  with  $w_j > 0$ ).*

**Proof of Corollary 1.** All three parts follow from the special case (6), which gives

$$\hat{\tau}_t - \tau_t = -\tau_t \sum_{j \in \mathcal{E}} \rho_j w_j.$$

Part (i) (*sign/attenuation*). Since  $\rho_j \geq 0$  and  $w_j \geq 0$ , the sum  $\sum_{j \in \mathcal{E}} \rho_j w_j$  is nonnegative, and therefore

$$\text{sign}(\hat{\tau}_t - \tau_t) = -\text{sign}(\tau_t) \quad \text{whenever} \quad \sum_{j \in \mathcal{E}} \rho_j w_j > 0.$$

Hence, if  $\tau_t > 0$  then  $\hat{\tau}_t < \tau_t$ ; if  $\tau_t < 0$  then  $\hat{\tau}_t > \tau_t$ , which proves part (i).

Part (ii) (monotonicity). Taking absolute values in the first display and using nonnegativity,

$$|\hat{\tau}_t - \tau_t| = |\tau_t| \cdot \sum_{j \in \mathcal{E}} \rho_j w_j.$$

Differentiating with respect to each intensity gives componentwise monotonicity in each  $\rho_j$ ,

$$\frac{\partial}{\partial \rho_j} |\hat{\tau}_t - \tau_t| = |\tau_t| w_j \geq 0.$$

For the exposed weight mass: holding the intensities  $\rho$  fixed, adding weight to any exposed coordinate  $j \in \mathcal{E}$  increases  $\sum_{j \in \mathcal{E}} \rho_j w_j$  at rate  $\rho_j \geq 0$ , so any increase in the exposed weight mass  $\sum_{j \in \mathcal{E}} w_j$  obtained by raising exposed coordinates leaves the deviation  $|\hat{\tau}_t - \tau_t| = |\tau_t| \sum_{j \in \mathcal{E}} \rho_j w_j$  unchanged or larger. This proves part (ii).

Part (iii) (bounds). Since  $\rho_j \in [0, 1]$  and  $w \in \Delta$ ,

$$0 \leq \sum_{j \in \mathcal{E}} \rho_j w_j \leq \sum_{j \in \mathcal{E}} 1 \cdot w_j \leq \sum_{j \in \mathcal{J}} w_j = 1.$$

Multiplying through by  $|\tau_t|$ ,

$$0 \leq |\hat{\tau}_t - \tau_t| = |\tau_t| \cdot \sum_{j \in \mathcal{E}} \rho_j w_j \leq |\tau_t|,$$

with equality at the upper bound only if  $\sum_{j \in \mathcal{E}} \rho_j w_j = 1$ , i.e., only if all positive weight lies on donors with  $\rho_j = 1$ . This proves part (iii).  $\square$

## A.2 Proof of Theorem 2

**THEOREM 2 (SUPPORT ELIMINATION AND EXPOSED-MASS DOMINANCE UNDER RESCALING, RE-STATED).** Let  $w^\circ$  denote any minimizer of  $\|X_1 - X_0 w\|_V^2$  over  $\Delta_J$ , and let  $X_0^* = X_0 \text{diag}(\eta)$ .

(C) If donor  $m$  satisfies Assumption 4 and Condition 1 holds for  $M = \{m\}$ , then the minimum of  $\|X_1 - X_0^* w\|_V^2$  over  $\Delta_J$  is attained at a weight vector  $w^*$  with  $w_m^* = 0$ . Moreover,  $w^*$  is the unique minimizer on the face  $F_{\{m\}} = \{w \in \Delta_J : w_m = 0\}$ . The flat direction induced by condition (iii) may support additional global minimizers on  $\Delta_J$ ; any such minimizer maps to  $w^*$  under the face projection  $\varphi$  and achieves the same objective value.

(D) If the exposed band  $\mathcal{E}$  satisfies Assumption 5 and Condition 1 holds for  $M = \mathcal{E}$ , then

$$\sum_{j \in \mathcal{E}} w_j^* = 0 \leq \sum_{j \in \mathcal{E}} w_j^\circ,$$

with strict inequality whenever  $w_m^\circ > 0$  for some  $m \in \mathcal{E}$ . Here  $w^*$  denotes the unique minimizer on  $F_{\mathcal{E}} = \{w \in \Delta_J : w_m = 0 \ \forall m \in \mathcal{E}\}$ .

**Setting.** Let  $\eta = (\eta_1, \dots, \eta_J)^\top$  with  $\eta_j \in (0, 1)$ , and define the rescaled donor matrix  $X_0^* := X_0 \text{diag}(\eta)$  and the objective  $Q(w; X_0) := \|X_1 - X_0 w\|_V^2 = (X_1 - X_0 w)^\top V (X_1 - X_0 w)$ . Let  $w^\circ$  denote any minimizer of  $Q(w; X_0)$  over  $\Delta_J$ , and let  $w^*$  denote the unique face minimizer defined in the theorem statement. Assumptions 4 and 5 and Condition 1 are as given in the main text.

**Proof of Theorem 2.** The proof proceeds in three parts. Part A defines the rescaled quadratic form  $Q(w; X_0^*)$  and establishes notation  $(H^*, b^*)$ . Part C proves donor-level support elimination via a face-achievability map that preserves the rescaled objective, then invokes face strict convexity for

uniqueness. Part D lifts the donor-level result to band-level mass via composition of the individual face maps. The part labels mirror the theorem statement's own claim labels: Parts C and D establish claims (C) and (D) of the theorem, respectively.

**Part A (quadratic expansion).** Write  $Q(w; X_0^*) = \|X_1 - X_0^*w\|_V^2 = X_1^\top V X_1 - 2b^{*\top}w + w^\top H^*w$  where  $b^* := X_0^{*\top} V X_1$  and  $H^* := X_0^{*\top} V X_0^* \succeq 0$ . The gradient is  $\nabla_w Q = 2(H^*w - b^*)$  and the Hessian is  $\nabla_w^2 Q = 2H^*$ , constant in  $w$ . The function  $Q$  is convex on  $\Delta_J$  and the set of minimizers over  $\Delta_J$  is nonempty and convex.

**Part C (donor-level support elimination).**

Fix a donor  $m$  satisfying Assumption 4. Write  $c_j := \eta_j X_{\cdot,j}$  for the rescaled columns (i.e., the  $j$ -th column of  $X_0^*$ ). By condition (iii), there exist  $\gamma \in \Delta_{|\mathcal{C}(m)|}$  with  $c_m = \sum_{\ell \in \mathcal{C}(m)} \gamma_\ell c_\ell$ .

Conditions (i) and (ii) of Assumption 4 ensure that the covering set  $\mathcal{C}(m)$  provides competitive fit to the treated unit; the formal argument below uses condition (iii), which supplies the algebraic structure for exact elimination.

*Step 1 (Face achievability).* For any  $w \in \Delta_J$ , define  $\varphi(w) \in \mathbb{R}^J$  by:

$$\varphi(w)_m := 0, \quad \varphi(w)_\ell := w_\ell + w_m \gamma_\ell \quad (\ell \in \mathcal{C}(m)), \quad \varphi(w)_j := w_j \quad (j \notin \{m\} \cup \mathcal{C}(m)).$$

(If  $w_m = 0$ , then  $\varphi(w) = w \in F_{\{m\}}$  and the claims below hold trivially.)

*Feasibility:* Each coordinate of  $\varphi(w)$  is a sum of nonnegatives, so  $\varphi(w) \geq 0$ . Moreover,  $\mathbf{1}^\top \varphi(w) = (1 - w_m) + w_m \sum_{\ell \in \mathcal{C}(m)} \gamma_\ell = 1$  since  $\gamma \in \Delta_{|\mathcal{C}(m)|}$ . Hence  $\varphi(w) \in \Delta_J$  with  $\varphi(w)_m = 0$ , so  $\varphi(w) \in F_{\{m\}}$ .

*Identical rescaled fit:*

$$\begin{aligned} X_0^* \varphi(w) &= \sum_{\ell \in \mathcal{C}(m)} (w_\ell + w_m \gamma_\ell) c_\ell + \sum_{j \notin \{m\} \cup \mathcal{C}(m)} w_j c_j && \text{[definition of } \varphi(w)\text{]} \\ &= [X_0^*w - w_m c_m] + w_m \sum_{\ell \in \mathcal{C}(m)} \gamma_\ell c_\ell && \text{[regroup]} \\ &= X_0^*w - w_m c_m + w_m c_m = X_0^*w, && \text{[condition (iii)]} \end{aligned}$$

where the penultimate equality uses  $c_m = \sum_{\ell \in \mathcal{C}(m)} \gamma_\ell c_\ell$  (condition (iii)). Therefore  $Q(\varphi(w); X_0^*) = Q(w; X_0^*)$ .

*Step 2 (Global minimum on the face).* For every  $w \in \Delta_J$ : if  $w_m = 0$  then  $w \in F_{\{m\}}$  already; if  $w_m > 0$  then  $\varphi(w) \in F_{\{m\}}$  with  $Q(\varphi(w)) = Q(w)$ . In either case,  $\min_{F_{\{m\}}} Q \leq Q(w)$ . Taking the infimum over  $\Delta_J$ :

$$\min_{F_{\{m\}}} Q \leq \min_{\Delta_J} Q \leq \min_{F_{\{m\}}} Q,$$

the right inequality because  $F_{\{m\}} \subseteq \Delta_J$ . Hence  $\min_{F_{\{m\}}} Q = \min_{\Delta_J} Q$ .

*Step 3 (Uniqueness on the face).* By Condition 1 with  $M = \{m\}$ ,  $Q$  is strictly convex on the tangent space of  $F_{\{m\}}$ ; the minimizer  $w^*$  on  $F_{\{m\}}$  is therefore unique. Since  $w^* \in F_{\{m\}}$ , we have  $w_m^* = 0$ , which proves the donor-level claim (Part C of the theorem).

**Part D (band-level exposed-mass dominance).**

Assume  $\mathcal{E}$  satisfies Assumption 5: every  $m \in \mathcal{E}$  satisfies Assumption 4 with  $\mathcal{C}(m) \subseteq \mathcal{E}^c$ .

*Step 1 (Compositional face projection).* For each  $m \in \mathcal{E}$ , the map  $\varphi_m$  from Part C redistributes weight from  $m$  to  $\mathcal{C}(m) \subseteq \mathcal{E}^c$ , leaving all other coordinates  $w_{m'}$  ( $m' \in \mathcal{E} \setminus \{m\}$ ) unchanged (because  $m' \notin \{m\} \cup \mathcal{C}(m)$ :  $m' \neq m$  and  $m' \in \mathcal{E}$  while  $\mathcal{C}(m) \subseteq \mathcal{E}^c$ ). The maps therefore do not interfere with each other's source coordinates. (They may both add weight to the same donor  $\ell \in \mathcal{C}(m) \cap \mathcal{C}(m')$ ; this

is harmless since the fit-preservation identity  $X_0^* \varphi_m(u) = X_0^* u$  holds for *any*  $u \in \Delta_J$ , regardless of its specific weights.) The composition  $\Phi := \varphi_{m_1} \circ \dots \circ \varphi_{m_{|\mathcal{E}|}}$  is therefore well-defined, order-independent, and satisfies:

$$\Phi(w) \in F_{\mathcal{E}}, \quad Q(\Phi(w); X_0^*) = Q(w; X_0^*) \quad \text{for every } w \in \Delta_J.$$

*Step 2 (Conclusion).* The same sandwich argument as in Part C gives  $\min_{F_{\mathcal{E}}} Q = \min_{\Delta_J} Q$ . By Condition 1 with  $M = \mathcal{E}$ ,  $d^\top H^* d > 0$  for all nonzero  $d$  with  $d_j = 0$  ( $j \in \mathcal{E}$ ) and  $\mathbf{1}^\top d = 0$ . Therefore  $Q$  is strictly convex on  $F_{\mathcal{E}}$  and the minimizer  $w^*$  on  $F_{\mathcal{E}}$  is unique. Since  $w^* \in F_{\mathcal{E}}$ ,  $\sum_{j \in \mathcal{E}} w_j^* = 0$ . Weak inequality  $0 \leq \sum_{j \in \mathcal{E}} w_j^o$  is immediate from  $w^o \geq 0$ . Strict inequality holds whenever  $w_m^o > 0$  for some  $m \in \mathcal{E}$ , which proves the band-level claim.  $\square$

### A.3 Proof of Theorem 3

**THEOREM 3 (BIAS CONTROL UNDER CONSTRAINED RIDGE, RESTATED).** *Let  $w^{(0)}$  solve  $\min_{w \in \Delta_J} \|X_1 - X_0 w\|_V^2$  and let  $w^{(\lambda)}$  solve (22) with  $\lambda > 0$ . Let  $\underline{\psi}_{\mathcal{E}} := \min_{j \in \mathcal{E}} \psi_j > 0$ . Under Assumptions 6, 7, and 8,*

$$\|D_\psi^{1/2} w^{(\lambda)}\|_2 \leq \|D_\psi^{1/2} w^{(0)}\|_2,$$

and therefore the exposed weight mass satisfies

$$\sum_{m \in \mathcal{E}} w_m^{(\lambda)} \leq \sqrt{|\mathcal{E}|} (\underline{\psi}_{\mathcal{E}})^{-1/2} \|D_\psi^{1/2} w^{(0)}\|_2.$$

Under the localized spillover model, the bias envelope  $|\hat{\tau}_t^{CR,(\lambda)} - \tau_t| \leq \rho |\tau_t| \sqrt{|\mathcal{E}|} (\underline{\psi}_{\mathcal{E}})^{-1/2} \|D_\psi^{1/2} w^{CR,(\lambda)}\|_2$  is nonincreasing in  $\lambda$ , where  $w^{CR,(\lambda)} \equiv w^{(\lambda)}$  denotes the constrained-ridge minimizer (we adopt the superscripted form  $w^{CR,(\lambda)}$  in the body for consistency with the unconstrained-ridge notation  $w^{UR,(\lambda)}$  of Section 6).

**Setting.** Fix  $X_1 \in \mathbb{R}^K$ ,  $X_0 = [X_{\cdot,1}, \dots, X_{\cdot,J}] \in \mathbb{R}^{K \times J}$ , and  $V = \text{diag}(v_1, \dots, v_K) \succ 0$ . Let  $\eta_j \in (0, 1)$ ,  $\psi_j = g(\eta_j)$ ,  $D_\psi := \text{diag}(\psi_1, \dots, \psi_J)$ , and  $H := X_0^\top V X_0$ , and define the penalized objective

$$J_\lambda(w) := \|X_1 - X_0 w\|_V^2 + \lambda \sum_j \psi_j w_j^2 = (X_1 - X_0 w)^\top V (X_1 - X_0 w) + \lambda w^\top D_\psi w.$$

The simplex constraints are collected in  $\Delta_J := \{w \in \mathbb{R}_{\geq 0}^J : \mathbf{1}^\top w = 1\}$ . Let  $w^{(0)} \in \arg \min_{w \in \Delta_J} \|X_1 - X_0 w\|_V^2$  and  $w^{(\lambda)} \in \arg \min_{w \in \Delta_J} J_\lambda(w)$  for  $\lambda > 0$  denote the unpenalized and constrained-ridge minimizers, with fit residuals  $r^{(\lambda)} := X_0 w^{(\lambda)} - X_1$  and  $r^{(0)} := X_0 w^{(0)} - X_1$ . Fix the more-exposed band  $\mathcal{E}$  and the safer band  $\mathcal{S} := \mathcal{E}^c$ . Assumptions 6, 7, and 8 are as given in the main text. (Assumption 8 is used in Part C only; the main results in Parts D–G require only Assumptions 6–7.)

**Proof of Theorem 3.** The proof proceeds in seven parts. Part A expands  $J_\lambda$  and computes its gradient and Hessian explicitly, fixing strict convexity; it isolates curvature and shows uniqueness. Part B writes the KKT system and derives the identity (B.3): the directional derivative of  $J_\lambda$  along a feasible direction  $d$  at  $w^{(\lambda)}$  equals  $\nu^{(\lambda)\top} d$ ; this converts “no feasible descent” into an algebraic test. Part C builds a concrete “band-shift” direction  $d$  that moves mass from an exposed donor  $m$  to the safer band  $\mathcal{S}$ , verifies feasibility, and evaluates the derivative via KKT; it constructs the canonical exposure-reducing move. Part D establishes that  $\lambda \mapsto \|D_\psi^{1/2} w^{(\lambda)}\|_2$  is nonincreasing on  $[0, \infty)$  via a variational comparison of optimality inequalities. Part E converts the  $\psi$ -weighted norm comparison (D.1) into a bias envelope on the exposed weight mass via Cauchy–Schwarz and

band separation. Part F shows the strict inequality  $P(w^{(\lambda)}) < P(w^{(0)})$  when  $w^{(\lambda)} \neq w^{(0)}$  via a strict-convexity contradiction argument. Part G collects the conclusions.

**Part A (quadratic expansion and derivatives of  $J_\lambda$ ).** By definition of the penalized loss,

$$J_\lambda(w) = (X_1 - X_0 w)^\top V (X_1 - X_0 w) + \lambda w^\top D_\psi w.$$

Expanding the quadratic form, with  $H = X_0^\top V X_0$ ,

$$J_\lambda(w) = X_1^\top V X_1 - 2(X_0^\top V X_1)^\top w + w^\top H w + \lambda w^\top D_\psi w.$$

Differentiating once and twice in  $w$ ,

$$\nabla_w J_\lambda(w) = 2Hw - 2X_0^\top V X_1 + 2\lambda D_\psi w, \quad \nabla_w^2 J_\lambda(w) = 2H + 2\lambda D_\psi = H_\lambda,$$

so the Hessian is constant in  $w$ , with the ridge adding curvature along the coordinate axes. By Assumption 7 ( $\lambda > 0$  and  $D_\psi \succ 0$ ),  $H_\lambda \succ 0$ , hence  $J_\lambda$  is strictly convex on  $\mathbb{R}^J$ ; by the Weierstrass theorem combined with strict convexity, a unique minimizer of  $J_\lambda$  on the convex set  $\Delta_J$  exists.

**Part B (KKT system and a key identity for directional derivatives).** Introduce the Lagrangian, with multipliers  $\mu \in \mathbb{R}$  for the sum constraint and  $\nu \in \mathbb{R}_{\geq 0}^J$  for nonnegativity,

$$\mathcal{L}(w, \mu, \nu) = J_\lambda(w) + \mu(\mathbf{1}^\top w - 1) - \nu^\top w.$$

The KKT conditions at  $(w^{(\lambda)}, \mu^{(\lambda)}, \nu^{(\lambda)})$  read

$$\begin{cases} 2Hw^{(\lambda)} - 2X_0^\top V X_1 + 2\lambda D_\psi w^{(\lambda)} + \mu^{(\lambda)} \mathbf{1} - \nu^{(\lambda)} = 0, & \text{[stationarity]} \\ \mathbf{1}^\top w^{(\lambda)} = 1, \quad w^{(\lambda)} \geq 0, & \text{[primal feasibility]} \\ \nu^{(\lambda)} \geq 0, & \text{[dual feasibility]} \\ \nu_j^{(\lambda)} w_j^{(\lambda)} = 0 \quad \forall j. & \text{[complementary slackness]} \end{cases}$$

Let  $d \in \mathbb{R}^J$  be any feasible direction with  $\mathbf{1}^\top d = 0$  and, for small  $\epsilon > 0$ ,  $w^{(\lambda)} + \epsilon d \in \Delta_J$ ; such a  $d$  is a tangent direction preserving the simplex at first order. The directional derivative of the fit part along  $d$  is

$$\left. \frac{d}{d\epsilon} \right|_{\epsilon=0} \|X_1 - X_0(w^{(\lambda)} + \epsilon d)\|_V^2 = 2(X_0 w^{(\lambda)} - X_1)^\top V (X_0 d) = 2r^{(\lambda)\top} V X_0 d,$$

and the directional derivative of the penalty part is

$$\left. \frac{d}{d\epsilon} \right|_{\epsilon=0} \lambda(w^{(\lambda)} + \epsilon d)^\top D_\psi (w^{(\lambda)} + \epsilon d) = 2\lambda d^\top D_\psi w^{(\lambda)}.$$

Collecting both contributions,

$$\left. \frac{d}{d\epsilon} \right|_{\epsilon=0} J_\lambda(w^{(\lambda)} + \epsilon d) = 2r^{(\lambda)\top} V X_0 d + 2\lambda d^\top D_\psi w^{(\lambda)}. \quad (\text{B.1})$$

An equivalent expression follows from the KKT system. Multiplying the stationarity condition by  $d$ ,

$$d^\top (2Hw^{(\lambda)} - 2X_0^\top V X_1 + 2\lambda D_\psi w^{(\lambda)} + \mu^{(\lambda)} \mathbf{1} - \nu^{(\lambda)}) = 0,$$

and expanding the inner products,

$$2d^\top H w^{(\lambda)} - 2d^\top X_0^\top V X_1 + 2\lambda d^\top D_\psi w^{(\lambda)} + \mu^{(\lambda)} \mathbf{1}^\top d - \nu^{(\lambda)\top} d = 0.$$

Since  $\mathbf{1}^\top d = 0$ , the equality-constraint term drops:

$$2d^\top Hw^{(\lambda)} - 2d^\top X_0^\top V X_1 + 2\lambda d^\top D_\psi w^{(\lambda)} - \nu^{(\lambda)\top} d = 0.$$

Because  $H = X_0^\top V X_0$ , we have  $d^\top Hw^{(\lambda)} = (X_0 d)^\top V(X_0 w^{(\lambda)})$ ; substituting this together with  $X_0 w^{(\lambda)} - X_1 = r^{(\lambda)}$  gives

$$2(X_0 d)^\top V(X_0 w^{(\lambda)} - X_1) + 2\lambda d^\top D_\psi w^{(\lambda)} - \nu^{(\lambda)\top} d = 0,$$

and rearranging to match (B.1),

$$2r^{(\lambda)\top} V X_0 d + 2\lambda d^\top D_\psi w^{(\lambda)} = \nu^{(\lambda)\top} d. \quad (\text{B.2})$$

Combining (B.1) and (B.2) yields the exact KKT identity: the first-order change of the objective equals the dual work on  $d$ ,

$$\left. \frac{d}{d\epsilon} \right|_{\epsilon=0} J_\lambda(w^{(\lambda)} + \epsilon d) = \nu^{(\lambda)\top} d. \quad (\text{B.3})$$

Since  $\nu^{(\lambda)} \geq 0$  and  $d$ 's signs depend on its components, (B.3) gives the exact first-order change under the constraints.

**Part C (a band-shift direction and its sign at  $w^{(\lambda)}$ ).** Fix  $m \in \mathcal{E}$  and choose  $\alpha^{(m)} \in \Delta_{|\mathcal{S}|}$  as in Assumption 8: a safer convex combination that covers  $m$  at first order. Define  $d \in \mathbb{R}^J$  by

$$d_m = -1, \quad d_\ell = \alpha_\ell^{(m)} \quad (\ell \in \mathcal{S}), \quad d_j = 0 \text{ otherwise,}$$

the direction that moves mass off  $m$  and onto safer donors in  $\mathcal{S}$ . Sum-to-one is preserved, since

$$\mathbf{1}^\top d = -1 + \sum_{\ell \in \mathcal{S}} \alpha_\ell^{(m)} = 0,$$

and if  $w_m^{(\lambda)} > 0$  the step is feasible for small  $\epsilon$ : there exists  $\epsilon_0 > 0$  such that  $w^{(\lambda)} + \epsilon d \in \Delta_J$  for all  $\epsilon \in (0, \epsilon_0]$ . Applying (B.3) with this  $d$ ,

$$\left. \frac{d}{d\epsilon} \right|_{\epsilon=0} J_\lambda(w^{(\lambda)} + \epsilon d) = \nu^{(\lambda)\top} d = \sum_{\ell \in \mathcal{S}} \nu_\ell^{(\lambda)} \alpha_\ell^{(m)} - \nu_m^{(\lambda)}.$$

If  $w_m^{(\lambda)} > 0$ , complementary slackness on the active coordinate  $m$  gives  $\nu_m^{(\lambda)} = 0$ , so

$$\left. \frac{d}{d\epsilon} \right|_{\epsilon=0} J_\lambda(w^{(\lambda)} + \epsilon d) = \sum_{\ell \in \mathcal{S}} \nu_\ell^{(\lambda)} \alpha_\ell^{(m)} \geq 0$$

by dual feasibility and  $\alpha^{(m)} \geq 0$ : there is no descent at the optimum. Therefore no feasible first-order descent exists along  $d$  at  $w^{(\lambda)}$ , consistent with optimality.

**Part D (path monotonicity of the  $\psi$ -weighted norm on the simplex).** Separate fit and penalty: write  $L(w) := \|X_1 - X_0 w\|_V^2$  and  $P(w) := w^\top D_\psi w = \|D_\psi^{1/2} w\|_2^2$ , so that  $J_\lambda(w) = L(w) + \lambda P(w)$ . Pick  $0 \leq \lambda_1 < \lambda_2$  and let  $w_1 := w^{(\lambda_1)}$  and  $w_2 := w^{(\lambda_2)}$  be the minimizers of  $J_{\lambda_1}$  and  $J_{\lambda_2}$  over  $\Delta_J$  (by convention,  $w^{(0)} \in \arg \min_{\Delta_J} L$  is any fixed selection). Optimality of  $w_1$  for  $J_{\lambda_1}$  over  $\Delta_J$  (with  $w_2 \in \Delta_J$  feasible) and optimality of  $w_2$  for  $J_{\lambda_2}$  over  $\Delta_J$  (with  $w_1 \in \Delta_J$  feasible) give the two inequalities

$$\begin{aligned} L(w_1) + \lambda_1 P(w_1) &\leq L(w_2) + \lambda_1 P(w_2), \\ L(w_2) + \lambda_2 P(w_2) &\leq L(w_1) + \lambda_2 P(w_1). \end{aligned}$$

Adding both and cancelling  $L(w_1) + L(w_2)$  from the two sides,

$$\lambda_1 P(w_1) + \lambda_2 P(w_2) \leq \lambda_1 P(w_2) + \lambda_2 P(w_1),$$

which rearranges to

$$(\lambda_2 - \lambda_1) (P(w_1) - P(w_2)) \geq 0.$$

Dividing by  $\lambda_2 - \lambda_1 > 0$  gives  $P(w_2) \leq P(w_1)$ , that is,

$$\|D_\psi^{1/2} w^{(\lambda_2)}\|_2 \leq \|D_\psi^{1/2} w^{(\lambda_1)}\|_2, \quad (\text{D.1})$$

so the map  $\lambda \mapsto \|D_\psi^{1/2} w^{(\lambda)}\|_2$  is nonincreasing on  $[0, \infty)$ .

**Part E (from (D.1) to a bias envelope).** Restating (D.1) for the pair  $(\lambda_1, \lambda_2) = (0, \lambda)$  in explicit norm notation, with  $D_\psi = \text{diag}(\psi_1, \dots, \psi_J)$ ,

$$\|D_\psi^{1/2} w^{(\lambda)}\|_2^2 := \sum_{j=1}^J \psi_j (w_j^{(\lambda)})^2 \leq \sum_{j=1}^J \psi_j (w_j^{(0)})^2 =: \|D_\psi^{1/2} w^{(0)}\|_2^2. \quad (\text{E.1})$$

To bound the exposed mass  $\sum_{m \in \mathcal{E}} w_m^{(\lambda)}$ , rewrite each exposed coordinate as  $w_m = \psi_m^{-1/2} \cdot \psi_m^{1/2} w_m$  and apply Cauchy–Schwarz on  $\mathcal{E}$ :

$$\sum_{m \in \mathcal{E}} w_m^{(\lambda)} = \sum_{m \in \mathcal{E}} \psi_m^{-1/2} \cdot \psi_m^{1/2} w_m^{(\lambda)} \leq \left( \sum_{m \in \mathcal{E}} \psi_m^{-1} \right)^{1/2} \left( \sum_{m \in \mathcal{E}} \psi_m (w_m^{(\lambda)})^2 \right)^{1/2}.$$

Each factor is bounded in turn:

$$\begin{aligned} \sum_{m \in \mathcal{E}} \psi_m^{-1} &\leq |\mathcal{E}| (\underline{\psi}_\mathcal{E})^{-1} && \text{[since } \psi_m \geq \underline{\psi}_\mathcal{E} \text{ for } m \in \mathcal{E}] \\ \sum_{m \in \mathcal{E}} \psi_m (w_m^{(\lambda)})^2 &\leq \sum_{j=1}^J \psi_j (w_j^{(\lambda)})^2 = \|D_\psi^{1/2} w^{(\lambda)}\|_2^2 && \text{[drop nonnegative } \mathcal{S}\text{-terms].} \end{aligned}$$

Chaining the Cauchy–Schwarz step with these two factor bounds and with (E.1),

$$\sum_{m \in \mathcal{E}} w_m^{(\lambda)} \leq \sqrt{|\mathcal{E}|} (\underline{\psi}_\mathcal{E})^{-1/2} \|D_\psi^{1/2} w^{(\lambda)}\|_2 \leq \sqrt{|\mathcal{E}|} (\underline{\psi}_\mathcal{E})^{-1/2} \|D_\psi^{1/2} w^{(0)}\|_2. \quad (\text{E.2})$$

Translating to bias: under the localized spillover model with  $\rho \in [0, 1]$ , the bias decomposition of Theorem 1 together with (E.2) gives

$$|\hat{\tau}_t^{(\lambda)} - \tau_t| \leq \rho |\tau_t| \sum_{m \in \mathcal{E}} w_m^{(\lambda)} \leq \rho |\tau_t| \sqrt{|\mathcal{E}|} (\underline{\psi}_\mathcal{E})^{-1/2} \|D_\psi^{1/2} w^{(\lambda)}\|_2. \quad (\text{E.3})$$

The right-hand side of (E.3) is controlled by (E.1): the constrained-ridge penalty norm  $\|D_\psi^{1/2} w^{(\lambda)}\|_2$  is at most the unpenalized value  $\|D_\psi^{1/2} w^{(0)}\|_2$ .

**Part F (strict inequality when  $w^{(\lambda)} \neq w^{(0)}$ ).** Suppose for contradiction that  $P(w^{(\lambda)}) = P(w^{(0)})$ . Optimality of  $w^{(0)}$  for  $L$  over  $\Delta_J$  gives

$$L(w^{(0)}) \leq L(w^{(\lambda)}),$$

while optimality of  $w^{(\lambda)}$  for  $J_\lambda$  over  $\Delta_J$  gives

$$L(w^{(\lambda)}) + \lambda P(w^{(\lambda)}) \leq L(w^{(0)}) + \lambda P(w^{(0)}).$$

Substituting  $P(w^{(\lambda)}) = P(w^{(0)})$  and cancelling  $\lambda P$  from both sides of the second inequality yields  $L(w^{(\lambda)}) \leq L(w^{(0)})$ ; squeezing from both directions,  $L(w^{(\lambda)}) = L(w^{(0)})$ . Since both the fit and penalty terms agree,

$$J_\lambda(w^{(\lambda)}) = L(w^{(\lambda)}) + \lambda P(w^{(\lambda)}) = L(w^{(0)}) + \lambda P(w^{(0)}) = J_\lambda(w^{(0)}).$$

But  $J_\lambda$  is strictly convex on  $\Delta_J$  (Part A;  $H_\lambda \succ 0$  for  $\lambda > 0$ ). If  $w^{(\lambda)} \neq w^{(0)}$ , both in  $\Delta_J$ , strict convexity at the midpoint gives

$$J_\lambda\left(\frac{w^{(\lambda)} + w^{(0)}}{2}\right) < \frac{J_\lambda(w^{(\lambda)}) + J_\lambda(w^{(0)})}{2} = J_\lambda(w^{(\lambda)}),$$

and the midpoint  $\frac{w^{(\lambda)} + w^{(0)}}{2}$  belongs to  $\Delta_J$  by convexity of the simplex (an average of nonnegative, sum-to-one vectors). This contradicts the optimality of  $w^{(\lambda)}$  for  $J_\lambda$  over  $\Delta_J$ , since  $J_\lambda$  would achieve a strictly lower value at a feasible point. Therefore  $P(w^{(\lambda)}) \neq P(w^{(0)})$ ; combined with  $P(w^{(\lambda)}) \leq P(w^{(0)})$  from (D.1),

$$P(w^{(\lambda)}) < P(w^{(0)}), \quad \text{i.e.,} \quad \|D_\psi^{1/2} w^{(\lambda)}\|_2 < \|D_\psi^{1/2} w^{(0)}\|_2.$$

By the same argument applied to  $J_{\lambda_2}$  with  $w^{(\lambda_1)}$  in place of  $w^{(0)}$ , the inequality  $P(w^{(\lambda_2)}) < P(w^{(\lambda_1)})$  is strict whenever  $w^{(\lambda_1)} \neq w^{(\lambda_2)}$  for  $0 < \lambda_1 < \lambda_2$ .

**Part G (conclusion).** Part D establishes that  $\lambda \mapsto \|D_\psi^{1/2} w^{(\lambda)}\|_2$  is nonincreasing on  $[0, \infty)$  (path monotonicity), by the variational comparison of optimality inequalities at any two penalty levels. Part E converts the  $\psi$ -weighted norm comparison (D.1) into a bias envelope via Cauchy–Schwarz and band separation, bounding the exposed mass by  $\sqrt{|\mathcal{E}|} (\underline{\psi}_\mathcal{E})^{-1/2} \|D_\psi^{1/2} w^{(0)}\|_2$ . Part F shows strict inequality when  $w^{(\lambda)} \neq w^{(0)}$ , the generic case for  $\lambda > 0$ , so the bias envelope tightens strictly. Finally, the prefactors  $\rho |\tau_t| \sqrt{|\mathcal{E}|} (\underline{\psi}_\mathcal{E})^{-1/2}$  multiplying  $\|D_\psi^{1/2} w^{(\lambda)}\|_2$  in the bias envelope (E.3) do not depend on  $\lambda$  (they involve only the spillover intensity, the effect magnitude, the band size, and the penalty floor on the band), so the path monotonicity from Part D transfers to the envelope: the right-hand side of (E.3) is itself nonincreasing in  $\lambda$ , as claimed in the theorem statement.  $\square$

#### A.4 Proof of Theorem 4

**THEOREM 4 (EXPOSURE-WEIGHTED UNCONSTRAINED RIDGE: BIAS ENVELOPE AND TUNING MONOTONICITY, RESTATED).** *Let  $w^{UR,(\lambda)}$  solve (26) with  $\lambda > 0$  under the exposure-aligned schedule in Assumption 6. For any exposed set  $\mathcal{E} \neq \emptyset$  with  $\underline{\psi}_\mathcal{E} := \min_{j \in \mathcal{E}} \psi_j > 0$ ,*

$$\left| \sum_{j \in \mathcal{E}} w_j^{UR,(\lambda)} \right| \leq \sqrt{|\mathcal{E}|} \underline{\psi}_\mathcal{E}^{-1/2} \|D_\psi^{1/2} w^{UR,(\lambda)}\|_2,$$

*and the map  $\lambda \mapsto \|D_\psi^{1/2} w^{UR,(\lambda)}\|_2$  is nonincreasing with  $\lim_{\lambda \rightarrow \infty} \|D_\psi^{1/2} w^{UR,(\lambda)}\|_2 = 0$ . Consequently, under the localized, time-homogeneous model of Section 2 (Theorem 1),*

$$|\hat{\tau}_t^{UR,(\lambda)} - \tau_t| \leq \rho |\tau_t| \sqrt{|\mathcal{E}|} \underline{\psi}_\mathcal{E}^{-1/2} \|D_\psi^{1/2} w^{UR,(\lambda)}\|_2.$$

More generally, for any spillover profile  $\{\delta_{jt}\}_{j \in \mathcal{J}}$  (not necessarily nonnegative or supported on  $\mathcal{E}$ ),

$$\left| \sum_j w_j^{UR,(\lambda)} \delta_{jt} \right| \leq \|D_\psi^{-1/2} \delta_t\|_2 \|D_\psi^{1/2} w^{UR,(\lambda)}\|_2,$$

where  $\|D_\psi^{-1/2} \delta_t\|_2 = (\sum_j \delta_{jt}^2 / \psi_j)^{1/2}$ . When  $\delta$  is supported on  $\mathcal{E}$  (as in the localized model above),  $\|D_\psi^{-1/2} \delta_t\|_2 \leq \underline{\psi}_\mathcal{E}^{-1/2} \|\delta_{\mathcal{E},t}\|_2$ , recovering the localized bound. Setting  $\delta = \mathbf{1}_\mathcal{E}$  in (28) likewise recovers the exposed-sum bound in the first display. Hence the right-hand side is nonincreasing in  $\lambda$  and vanishes as  $\lambda \rightarrow \infty$ .

**Setting.** Define the treated unit's predictors, the donor matrix, and the positive-definite metric as  $X_1 \in \mathbb{R}^K$ ,  $X_0 = [X_{\cdot,1}, \dots, X_{\cdot,J}] \in \mathbb{R}^{K \times J}$ , and  $V = \text{diag}(v_1, \dots, v_K) \succ 0$ . Set the reach scores, exposure-aligned penalties, and diagonal penalty operator as  $\eta_j \in (0, 1)$ ,  $\psi_j = g(\eta_j)$ , and  $D_\psi := \text{diag}(\psi_1, \dots, \psi_J)$ , with  $g$  strictly decreasing and bounded away from  $\{0, 1\}$ . The unconstrained ridge problem is a fit term plus an exposure-weighted ridge penalty,  $J_\lambda(w) := \|X_1 - X_0 w\|_V^2 + \lambda \|D_\psi^{1/2} w\|_2^2$  with  $\lambda \geq 0$ , and carries no simplex constraints: the feasible set is all of  $\mathbb{R}^J$  (weights may be negative and need not sum to one), enlarged relative to the standard SCM problem. Throughout the proof we write  $w^{UR}(\lambda)$  for the minimizer, the function-argument form of the body's superscripted notation  $w^{UR,(\lambda)}$  (the two denote the same object; see Section 6). Assumptions 6, 7, and 3 are as given in the main text; Assumption 3(ii) supplies the penalty bounds  $0 < \underline{\psi} \leq \psi_j \leq \bar{\psi} < \infty$  used in Parts C and D. (Assumption 8 is not invoked in this proof; the result requires only Assumptions 6–7 together with Assumption 3.)

**Proof of Theorem 4.** The proof proceeds in four parts. Part A derives the normal equations and the closed form  $w^{UR}(\lambda) = (X_0^\top V X_0 + \lambda D_\psi)^{-1} X_0^\top V X_1$ ; it makes the program self-contained and fixes notation for the spectral bounds. Part B proves the exposed-sum bound  $|\sum_{j \in \mathcal{E}} w_j| \leq \sqrt{|\mathcal{E}|} \underline{\psi}_\mathcal{E}^{-1/2} \|D_\psi^{1/2} w\|_2$ , which replaces mass arguments (invalid without the simplex) by a clean  $\ell_2$  control aligned with exposure. Part C establishes ridge-path monotonicity,  $\|D_\psi^{1/2} w^{UR}(\lambda)\|_2$  nonincreasing in  $\lambda$  and tending to zero as  $\lambda \rightarrow \infty$  via a resolvent/spectral bound; this gives a tuning parameter that uniformly shrinks the exposure-weighted norm. Part D translates these ingredients into bias envelopes: the contamination  $\rho \tau_t \sum_{j \in \mathcal{E}} w_j^{UR}$  is bounded by a term that is (i) nonincreasing in  $\lambda$  and (ii) vanishing as  $\lambda \rightarrow \infty$ ; the general profile bound (UR.4) uses  $D_\psi$ -weighted Cauchy–Schwarz over all donors to handle any spillover profile (not necessarily nonnegative or supported on  $\mathcal{E}$ ), recovering the localized bound as a corollary when  $\text{supp}(\delta) \subseteq \mathcal{E}$ .

**Part A (normal equations and closed form, for completeness).** Expanding the objective,

$$J_\lambda(w) = (X_1 - X_0 w)^\top V (X_1 - X_0 w) + \lambda w^\top D_\psi w,$$

and differentiating term by term,

$$\nabla_w J_\lambda(w) = -2X_0^\top V (X_1 - X_0 w) + 2\lambda D_\psi w.$$

Setting the gradient to zero (first-order optimality),

$$-2X_0^\top V X_1 + 2X_0^\top V X_0 w + 2\lambda D_\psi w = 0,$$

which yields the normal equations

$$(X_0^\top V X_0 + \lambda D_\psi) w = X_0^\top V X_1.$$

Since  $\lambda > 0$  and  $D_\psi \succ 0$ , the matrix  $X_0^\top V X_0 + \lambda D_\psi$  is strictly positive definite (Assumption 7), hence invertible, and therefore

$$w^{UR}(\lambda) = (X_0^\top V X_0 + \lambda D_\psi)^{-1} X_0^\top V X_1,$$

which is the closed form (27).

**Part B (exposed-sum bound).** Write  $w_\mathcal{E}$  for the restriction of  $w$  to the coordinates in  $\mathcal{E}$  and  $\mathbf{1}_\mathcal{E}$  for the all-ones vector on  $\mathcal{E}$ , so that  $|\sum_{j \in \mathcal{E}} w_j| = |\mathbf{1}_\mathcal{E}^\top w_\mathcal{E}|$ . By Cauchy-Schwarz, and since the Euclidean norm of  $\mathbf{1}_\mathcal{E}$  is  $\sqrt{|\mathcal{E}|}$ ,

$$\left| \sum_{j \in \mathcal{E}} w_j \right| \leq \|\mathbf{1}_\mathcal{E}\|_2 \|w_\mathcal{E}\|_2 = \sqrt{|\mathcal{E}|} \|w_\mathcal{E}\|_2.$$

Inserting the identity  $I = D_\psi^{-1/2} D_\psi^{1/2}$  on the coordinates in  $\mathcal{E}$  and applying the operator-norm bound,

$$\|w_\mathcal{E}\|_2 = \|D_\psi^{-1/2} D_\psi^{1/2} w_\mathcal{E}\|_2 \leq \|D_\psi^{-1/2}\|_{\text{op}, \mathcal{E}} \|D_\psi^{1/2} w_\mathcal{E}\|_2.$$

The operator norm restricted to  $\mathcal{E}$  is computed explicitly: since  $D_\psi^{-1/2}$  is diagonal, its restriction to the coordinates in  $\mathcal{E}$  is the diagonal matrix with entries  $\psi_j^{-1/2}$ ,  $j \in \mathcal{E}$ , and the operator norm of a diagonal matrix equals its largest diagonal entry in absolute value, so

$$\|D_\psi^{-1/2}\|_{\text{op}, \mathcal{E}} = \max_{j \in \mathcal{E}} \psi_j^{-1/2} = \left( \min_{j \in \mathcal{E}} \psi_j \right)^{-1/2} = \underline{\psi}_\mathcal{E}^{-1/2}, \quad \underline{\psi}_\mathcal{E} := \min_{j \in \mathcal{E}} \psi_j > 0,$$

the largest singular value on  $\mathcal{E}$ . Moreover, restriction cannot increase the norm:

$$\|D_\psi^{1/2} w_\mathcal{E}\|_2 \leq \|D_\psi^{1/2} w\|_2.$$

Combining the last two displays gives  $\|w_\mathcal{E}\|_2 \leq \underline{\psi}_\mathcal{E}^{-1/2} \|D_\psi^{1/2} w\|_2$ , and therefore the exposed sum is controlled by the exposure-weighted  $\ell_2$  norm:

$$\left| \sum_{j \in \mathcal{E}} w_j \right| \leq \sqrt{|\mathcal{E}|} \underline{\psi}_\mathcal{E}^{-1/2} \|D_\psi^{1/2} w\|_2. \quad (\text{UR.1})$$

**Part C (ridge path monotonicity).** Separate fit and penalty: write  $J_\lambda(w) = L(w) + \lambda P(w)$  with  $L(w) := \|X_1 - X_0 w\|_V^2$  and  $P(w) := \|D_\psi^{1/2} w\|_2^2$ . Pick two tuning values  $0 \leq \lambda_1 < \lambda_2$  and set  $w_1 := w^{UR}(\lambda_1)$ ,  $w_2 := w^{UR}(\lambda_2)$ . Optimality of  $w_1$  for  $J_{\lambda_1}$  and of  $w_2$  for  $J_{\lambda_2}$  give

$$\begin{aligned} L(w_1) + \lambda_1 P(w_1) &\leq L(w_2) + \lambda_1 P(w_2), \\ L(w_2) + \lambda_2 P(w_2) &\leq L(w_1) + \lambda_2 P(w_1). \end{aligned}$$

Adding the two, the  $L$  terms cancel,  $(L(w_1) - L(w_2)) + (L(w_2) - L(w_1)) = 0$ , leaving

$$\lambda_1 P(w_1) + \lambda_2 P(w_2) \leq \lambda_1 P(w_2) + \lambda_2 P(w_1),$$

which rearranges to  $(\lambda_2 - \lambda_1) (P(w_1) - P(w_2)) \geq 0$ . Since  $\lambda_2 > \lambda_1$ ,

$$P(w_2) \leq P(w_1) \iff \|D_\psi^{1/2} w^{UR}(\lambda_2)\|_2 \leq \|D_\psi^{1/2} w^{UR}(\lambda_1)\|_2, \quad (\text{UR.2})$$

so the exposure-weighted norm is monotone nonincreasing in  $\lambda$ . For the limit as  $\lambda \rightarrow \infty$ , start from the closed form  $w^{UR}(\lambda) = (X_0^\top V X_0 + \lambda D_\psi)^{-1} X_0^\top V X_1$ . The spectral bound  $\lambda_{\min}(A + \lambda B) \geq \lambda \lambda_{\min}(B)$  for  $A \succeq 0, B \succ 0$  gives

$$\|(X_0^\top V X_0 + \lambda D_\psi)^{-1}\|_{\text{op}} \leq \frac{1}{\lambda \lambda_{\min}(D_\psi)} = \frac{1}{\lambda \underline{\psi}},$$

and by submultiplicativity of the operator norm,

$$\begin{aligned} \|D_\psi^{1/2} w^{UR}(\lambda)\|_2 &\leq \|D_\psi^{1/2}\|_{\text{op}} \|(X_0^\top V X_0 + \lambda D_\psi)^{-1}\|_{\text{op}} \|X_0^\top V X_1\|_2 \\ &= \sqrt{\bar{\psi}} \cdot \frac{1}{\lambda \underline{\psi}} \cdot \|X_0^\top V X_1\|_2 \rightarrow 0 \text{ as } \lambda \rightarrow \infty, \end{aligned} \quad (\text{UR.3})$$

where  $\bar{\psi} := \max_j \psi_j < \infty$  by Assumption 3(ii).

**Part D (bias bounds under localized and aligned spillovers).** Under localized, time-homogeneous spillovers (Theorem 1), the contamination channel depends on the exposed sum:

$$\hat{\tau}_t^{UR}(\lambda) - \tau_t = -\rho \tau_t \sum_{j \in \mathcal{E}} w_j^{UR}(\lambda).$$

Taking absolute values and applying (UR.1) with  $w = w^{UR}(\lambda)$ ,

$$|\hat{\tau}_t^{UR}(\lambda) - \tau_t| = \rho |\tau_t| \left| \sum_{j \in \mathcal{E}} w_j^{UR}(\lambda) \right| \leq \rho |\tau_t| \sqrt{|\mathcal{E}|} \underline{\psi}^{-1/2} \|D_\psi^{1/2} w^{UR}(\lambda)\|_2.$$

By (UR.2) and (UR.3), the right-hand side is nonincreasing in  $\lambda$  and tends to 0 as  $\lambda \rightarrow \infty$  (tuning monotonicity and vanishing envelope).

For a general spillover profile  $\{\delta_{jt}\}$  (not necessarily nonnegative or supported on  $\mathcal{E}$ ), write the contamination term in inner-product notation and factor through  $D_\psi$  (valid since  $\psi_j > 0$  for all  $j$  by Assumption 3(ii)):

$$\sum_j w_j^{UR}(\lambda) \delta_{jt} = \langle w^{UR}(\lambda), \delta_t \rangle = \langle D_\psi^{1/2} w^{UR}(\lambda), D_\psi^{-1/2} \delta_t \rangle.$$

Cauchy–Schwarz in  $\mathbb{R}^J$  then gives

$$\left| \sum_j w_j^{UR}(\lambda) \delta_{jt} \right| \leq \|D_\psi^{-1/2} \delta_t\|_2 \|D_\psi^{1/2} w^{UR}(\lambda)\|_2. \quad (\text{UR.4})$$

When  $\delta$  is supported on  $\mathcal{E}$ , the first factor satisfies

$$\|D_\psi^{-1/2} \delta_t\|_2 = \left( \sum_{j \in \mathcal{E}} \delta_{jt}^2 / \psi_j \right)^{1/2} \leq \underline{\psi}_{\mathcal{E}}^{-1/2} \|\delta_{\mathcal{E},t}\|_2$$

because  $\psi_j \geq \underline{\psi}_{\mathcal{E}}$  for  $j \in \mathcal{E}$ , so  $1/\psi_j \leq 1/\underline{\psi}_{\mathcal{E}}$ ; this recovers the localized bound. Setting  $\delta = \mathbf{1}_{\mathcal{E}}$ , (UR.4) yields (UR.1), unifying Parts B and D, since  $\|D_\psi^{-1/2} \mathbf{1}_{\mathcal{E}}\|_2 = (\sum_{j \in \mathcal{E}} 1/\psi_j)^{1/2} \leq \sqrt{|\mathcal{E}|} \underline{\psi}_{\mathcal{E}}^{-1/2}$ . Hence the envelope is nonincreasing in  $\lambda$  and vanishes as  $\lambda \rightarrow \infty$  by (UR.2)–(UR.3): the factor  $\|D_\psi^{-1/2} \delta_t\|_2$  is  $\lambda$ -free, so the monotonicity and the limit are inherited from  $\|D_\psi^{1/2} w^{UR}(\lambda)\|_2$ .  $\square$

## A.5 Spatial reach mapping

To avoid any tuning on outcomes, we calibrate  $f$  from the empirical distribution of distances alone using a logistic map,

$$f(d) = \frac{1}{1 + \exp\{-\gamma_\kappa(d - c)\}}. \quad (29)$$

Let  $d_L$  and  $d_U$  denote the  $(q, 1 - q)$  quantiles of  $\{d_{j,1}\}_{j \neq 1}$  for a small  $q$ ; we take  $q = 0.025$ . Fix a tail level  $\varepsilon \in (0, 1/2)$ ; we take  $\varepsilon = 0.025$ . Imposing  $f(d_L) = \varepsilon$  and  $f(d_U) = 1 - \varepsilon$  yields

$$c = \frac{d_U + d_L}{2}, \quad \gamma_\kappa = \frac{2 \log\left(\frac{1-\varepsilon}{\varepsilon}\right)}{d_U - d_L}.$$

This anchors  $\eta_j = f(d_{j,1})$  smoothly over  $(\varepsilon, 1 - \varepsilon)$  so that proximate donors receive  $\eta_j$  near  $\varepsilon$  and distant donors near  $1 - \varepsilon$ . We maintain  $\eta_j \in [\varepsilon, 1 - \varepsilon]$  for all donors to keep the optimization well-posed.

In applied work,  $c$  and  $\gamma_\kappa$  can and should be calibrated with auxiliary information and domain knowledge.

Examples include: (i) setting  $c$  to a substantively meaningful distance (e.g., a commuting or media-market radius) and choosing  $\gamma_\kappa$  so that  $f$  halves over a prespecified range; (ii) mapping  $d$  to network-path lengths derived from mobility/commuting matrices and anchoring  $(d_L, d_U)$  to empirical percentiles of those effective distances; or (iii) selecting  $(c, \gamma_\kappa)$  so that  $f$  matches a pre-specified decay (e.g., half-life) suggested by prior studies of diffusion. In our simulations and applications, the quantile-anchored defaults  $(q, \varepsilon) = (0.025, 0.025)$  provided stable behavior and transparent reporting without outcome-based tuning.

## A.6 Reach calibration: quantile versus moment robustness

The reach map  $f(d) = [1 + \exp\{-\gamma_\kappa(d - c)\}]^{-1}$  requires a location  $c$  and a slope  $\gamma_\kappa$ ; Appendix A.5 gives the quantile-anchoring formulas that fix both from the distance distribution alone. Our default is the quantile-anchored calibration (Q). An alternative is moment-anchored calibration (M), which targets the empirical mean and variance of distances instead of quantiles. Q is our recommended default because it is robust to heavy tails in the distance distribution and requires no parametric assumption about its shape. This appendix documents that the leading conclusions do not depend on which of the two is used; this is the sense in which the reach map need not be tuned finely.

The key calibration finding is that Q and M agree on every leading inferential conclusion: the within-corrector difference in CRPS is at most 0.005 for Rescaling, URZ, and URX in every  $\rho$  cell, and is below 0.005 for Constrained Ridge except at high  $\rho$ . Table 9 reports CRPS under Q versus M across the four corrections at  $\tau = 4$  and  $\rho \in \{0.1, 0.3, 0.6, 0.9, 1.0\}$  on the Baseline DGP.

Table 9: Calibration Comparison: Quantile vs. Moment ( $\tau = 4$ )

$\rho$	Rescaling		Constr. Ridge		Unconstr. Z		Unconstr. X	
	Q	M	Q	M	Q	M	Q	M
0.1	0.040	<b>0.038</b>	0.029	<b>0.027</b>	0.012	<b>0.011</b>	0.008	0.008
0.3	0.049	<b>0.048</b>	0.080	<b>0.075</b>	0.014	0.014	0.010	0.010
0.6	<b>0.067</b>	0.068	0.156	<b>0.147</b>	0.018	<b>0.017</b>	0.014	<b>0.013</b>
0.9	<b>0.087</b>	0.091	0.231	<b>0.217</b>	0.021	<b>0.020</b>	0.018	<b>0.016</b>
1.0	<b>0.093</b>	0.098	0.256	<b>0.241</b>	0.022	<b>0.021</b>	0.019	<b>0.018</b>

*Notes:* Each column pair compares

CRPS under the quantile calibration (Q) of the reach parameter  $\gamma_\kappa$  versus the moment calibration (M) of the same parameter, holding the corrector fixed. The implicit benchmark is the within-corrector comparison: Q vs. M for the same corrector; bold marks the better calibration in each pair (ties, i.e., values that round to the same displayed digits, are left unbolded). Cell entries are mean CRPS (Continuous Ranked Probability Score) across  $R = 1,000$  replications; lower is better. Rows index the spillover intensity  $\rho$ . Q-vs-M agreement indicates that reach-calibration is not the leading source of error; disagreement signals that calibration choice affects downstream bias. Values are computed on the Baseline DGP at  $\tau = 4$ ; calibration results for the Heterogeneous and Factor Model DGPs appear in Tables 11, 12.

The within-corrector difference in CRPS exceeds 0.005 in three of twenty cells, all concentrated in Constrained Ridge at  $\rho \geq 0.6$ , where Q underperforms M by 0.009–0.015 CRPS. URX, URZ, and Rescaling agree Q-vs-M within 0.005 in every cell, and URX agrees within 0.002 with exact agreement at  $\rho \in \{0.1, 0.3\}$ . The ordering of corrections is identical under Q and M throughout. The robustness is structural: because the corrections target exposed mass and the exposed mass is determined by the support of  $f$ , both calibrations place the same donors on the exposed band as long as the reach is monotone in distance, which both preserve. Tables 10–12 give the per-DGP expansions; the moment-anchored results for the Heterogeneous and Factor Model DGPs tell the same story.

Figure 8 reports the M-calibrated counterpart to Figure 3; the two figures are visually indistinguishable in bias and agree to within simulation noise in RMSE, confirming the Q-vs-M invariance documented above.

Table 10: Calibration: Q vs. M, Baseline DGP

$\tau$	$\rho$	Rescaling		Constr. Ridge		Unconstr. Z		Unconstr. X	
		Q	M	Q	M	Q	M	Q	M
0.5	0.1	0.039	<b>0.037</b>	0.011	0.011	0.011	0.011	0.009	0.009
	0.3	0.039	<b>0.037</b>	0.015	<b>0.014</b>	0.011	0.011	0.009	0.009
	0.6	0.040	<b>0.038</b>	0.023	<b>0.022</b>	0.011	0.011	0.008	0.008
	0.9	0.041	<b>0.039</b>	0.032	<b>0.030</b>	0.012	0.012	0.008	0.008
	1.0	0.041	<b>0.039</b>	0.035	<b>0.033</b>	0.012	0.012	0.008	0.008
1.0	0.1	0.039	<b>0.037</b>	0.013	0.013	0.011	0.011	0.008	0.008
	0.3	0.040	<b>0.038</b>	0.023	<b>0.022</b>	0.011	0.011	0.008	0.008
	0.6	0.042	<b>0.040</b>	0.041	<b>0.039</b>	0.012	0.012	0.009	<b>0.008</b>
	0.9	0.045	<b>0.044</b>	0.060	<b>0.057</b>	0.013	0.013	0.009	0.009
	1.0	0.046	<b>0.045</b>	0.067	<b>0.063</b>	0.013	0.013	0.009	0.009
4.0	0.1	0.040	<b>0.038</b>	0.029	<b>0.027</b>	0.012	<b>0.011</b>	0.008	0.008
	0.3	0.049	<b>0.048</b>	0.080	<b>0.075</b>	0.014	0.014	0.010	0.010
	0.6	<b>0.067</b>	0.068	0.156	<b>0.147</b>	0.018	<b>0.017</b>	0.014	<b>0.013</b>
	0.9	<b>0.087</b>	0.091	0.231	<b>0.217</b>	0.021	<b>0.020</b>	0.018	<b>0.016</b>
	1.0	<b>0.093</b>	0.098	0.256	<b>0.241</b>	0.022	<b>0.021</b>	0.019	<b>0.018</b>
7.0	0.1	0.043	<b>0.041</b>	0.047	<b>0.045</b>	0.012	0.012	0.009	0.009
	0.3	<b>0.062</b>	0.063	0.137	<b>0.129</b>	0.017	<b>0.016</b>	0.013	<b>0.012</b>
	0.6	<b>0.096</b>	0.102	0.268	<b>0.252</b>	0.023	<b>0.022</b>	0.020	<b>0.018</b>
	0.9	<b>0.132</b>	0.142	0.397	<b>0.374</b>	0.029	<b>0.028</b>	0.027	<b>0.025</b>
	1.0	<b>0.143</b>	0.156	0.440	<b>0.414</b>	0.031	<b>0.030</b>	0.030	<b>0.027</b>

Notes: Each column pair

compares CRPS under the quantile calibration (Q) of the reach parameter  $\gamma_{\kappa}$  versus the moment calibration (M) of the same parameter, holding the corrector fixed. The implicit benchmark is the within-corrector comparison: Q vs. M for the same corrector; bold marks the better calibration in each pair (ties, i.e., values that round to the same displayed digits, are left unbolded). Cell entries are mean CRPS (Continuous Ranked Probability Score) across  $R = 1,000$  replications; lower is better. Each row is a  $(\tau, \rho)$  cell, with  $\tau$  the treatment effect and  $\rho$  the spillover intensity. Q-vs-M agreement indicates that reach-calibration is not the leading source of error; disagreement signals that calibration choice affects downstream bias. Values are computed on the Baseline DGP.

Table 11: Calibration: Q vs. M, Heterogeneous DGP

$\tau$	$\rho$	Rescaling		Constr. Ridge		Unconstr. Z		Unconstr. X	
		Q	M	Q	M	Q	M	Q	M
0.5	0.1	0.039	<b>0.037</b>	0.011	0.011	<b>0.010</b>	0.011	0.008	0.008
	0.3	0.039	<b>0.037</b>	0.015	<b>0.014</b>	0.011	0.011	0.008	0.008
	0.6	0.040	<b>0.037</b>	0.022	<b>0.021</b>	0.011	0.011	0.008	0.008
	0.9	0.041	<b>0.038</b>	0.030	<b>0.029</b>	0.012	0.012	0.008	0.008
	1.0	0.041	<b>0.039</b>	0.033	<b>0.032</b>	0.012	0.012	0.008	0.008
1.0	0.1	0.039	<b>0.037</b>	0.013	0.013	0.011	0.011	0.008	0.008
	0.3	0.040	<b>0.037</b>	0.022	<b>0.021</b>	0.011	0.011	0.008	0.008
	0.6	0.042	<b>0.040</b>	0.039	<b>0.037</b>	0.012	0.012	0.008	0.008
	0.9	0.045	<b>0.043</b>	0.058	<b>0.055</b>	0.013	0.013	0.009	0.009
	1.0	0.046	<b>0.045</b>	0.064	<b>0.060</b>	0.013	0.013	0.009	0.009
4.0	0.1	0.040	<b>0.038</b>	0.028	<b>0.026</b>	0.011	0.011	0.008	0.008
	0.3	0.048	<b>0.047</b>	0.076	<b>0.072</b>	0.014	<b>0.013</b>	0.010	0.010
	0.6	<b>0.065</b>	0.067	0.150	<b>0.141</b>	0.017	0.017	0.013	0.013
	0.9	<b>0.084</b>	0.088	0.221	<b>0.208</b>	0.021	<b>0.020</b>	0.017	<b>0.016</b>
	1.0	<b>0.090</b>	0.095	0.245	<b>0.231</b>	0.022	<b>0.021</b>	0.018	<b>0.017</b>
7.0	0.1	0.043	<b>0.041</b>	0.045	<b>0.043</b>	0.012	0.012	0.009	0.009
	0.3	<b>0.061</b>	0.062	0.131	<b>0.124</b>	0.016	0.016	0.013	<b>0.012</b>
	0.6	<b>0.093</b>	0.099	0.257	<b>0.242</b>	0.022	<b>0.021</b>	0.019	<b>0.018</b>
	0.9	<b>0.127</b>	0.137	0.381	<b>0.358</b>	0.028	<b>0.027</b>	0.026	<b>0.024</b>
	1.0	<b>0.138</b>	0.150	0.421	<b>0.396</b>	0.030	<b>0.029</b>	0.029	<b>0.026</b>

Notes: Each column pair

compares CRPS under the quantile calibration (Q) of the reach parameter  $\gamma_\kappa$  versus the moment calibration (M) of the same parameter, holding the corrector fixed. The implicit benchmark is the within-corrector comparison: Q vs. M for the same corrector; bold marks the better calibration in each pair (ties, i.e., values that round to the same displayed digits, are left unbolded). Cell entries are mean CRPS (Continuous Ranked Probability Score) across  $R = 1,000$  replications; lower is better. Each row is a  $(\tau, \rho)$  cell, with  $\tau$  the treatment effect and  $\rho$  the spillover intensity. Q-vs-M agreement indicates that reach-calibration is not the leading source of error; disagreement signals that calibration choice affects downstream bias. Values are computed on the Heterogeneous DGP.

Table 12: Calibration: Q vs. M, Factor Model DGP

$\tau$	$\rho$	Rescaling		Constr. Ridge		Unconstr. Z		Unconstr. X	
		Q	M	Q	M	Q	M	Q	M
0.5	0.1	0.047	<b>0.045</b>	<b>0.035</b>	0.036	0.016	0.016	0.046	0.046
	0.3	0.047	<b>0.045</b>	0.037	0.037	0.016	0.016	0.046	0.046
	0.6	0.048	<b>0.046</b>	0.040	0.040	0.017	0.017	0.046	0.046
	0.9	0.049	<b>0.046</b>	<b>0.043</b>	0.044	0.018	0.018	0.046	0.046
	1.0	0.049	<b>0.047</b>	0.045	0.045	0.019	<b>0.018</b>	0.046	0.046
1.0	0.1	0.047	<b>0.045</b>	<b>0.036</b>	0.037	0.016	0.016	0.046	0.046
	0.3	0.048	<b>0.046</b>	0.040	0.040	0.017	0.017	0.046	0.046
	0.6	0.049	<b>0.047</b>	0.048	0.048	0.019	0.019	0.046	0.046
	0.9	0.051	<b>0.049</b>	0.058	0.058	0.022	<b>0.021</b>	0.046	0.046
	1.0	0.052	<b>0.050</b>	0.062	<b>0.061</b>	0.023	<b>0.022</b>	0.047	<b>0.046</b>
4.0	0.1	0.048	<b>0.046</b>	<b>0.042</b>	0.043	0.018	<b>0.017</b>	0.046	0.046
	0.3	0.054	<b>0.053</b>	0.069	<b>0.068</b>	0.025	<b>0.024</b>	0.047	0.047
	0.6	<b>0.070</b>	0.073	0.121	<b>0.118</b>	0.039	<b>0.035</b>	0.049	0.049
	0.9	<b>0.092</b>	0.097	0.178	<b>0.172</b>	0.053	<b>0.047</b>	0.052	<b>0.051</b>
	1.0	<b>0.100</b>	0.106	0.198	<b>0.191</b>	0.057	<b>0.051</b>	0.053	<b>0.052</b>
7.0	0.1	0.050	<b>0.048</b>	0.051	0.051	0.020	<b>0.019</b>	0.046	0.046
	0.3	<b>0.066</b>	0.067	0.107	<b>0.105</b>	0.035	<b>0.032</b>	0.049	<b>0.048</b>
	0.6	<b>0.104</b>	0.110	0.207	<b>0.200</b>	0.060	<b>0.053</b>	0.053	<b>0.052</b>
	0.9	<b>0.145</b>	0.157	0.312	<b>0.299</b>	0.085	<b>0.075</b>	0.060	<b>0.058</b>
	1.0	<b>0.158</b>	0.173	0.347	<b>0.333</b>	0.093	<b>0.082</b>	0.062	<b>0.060</b>

Notes: Each column pair

compares CRPS under the quantile calibration (Q) of the reach parameter  $\gamma_{\kappa}$  versus the moment calibration (M) of the same parameter, holding the corrector fixed. The implicit benchmark is the within-corrector comparison: Q vs. M for the same corrector; bold marks the better calibration in each pair (ties, i.e., values that round to the same displayed digits, are left unbolded). Cell entries are mean CRPS (Continuous Ranked Probability Score) across  $R = 1,000$  replications; lower is better. Each row is a  $(\tau, \rho)$  cell, with  $\tau$  the treatment effect and  $\rho$  the spillover intensity. Q-vs-M agreement indicates that reach-calibration is not the leading source of error; disagreement signals that calibration choice affects downstream bias. Values are computed on the Factor Model DGP.

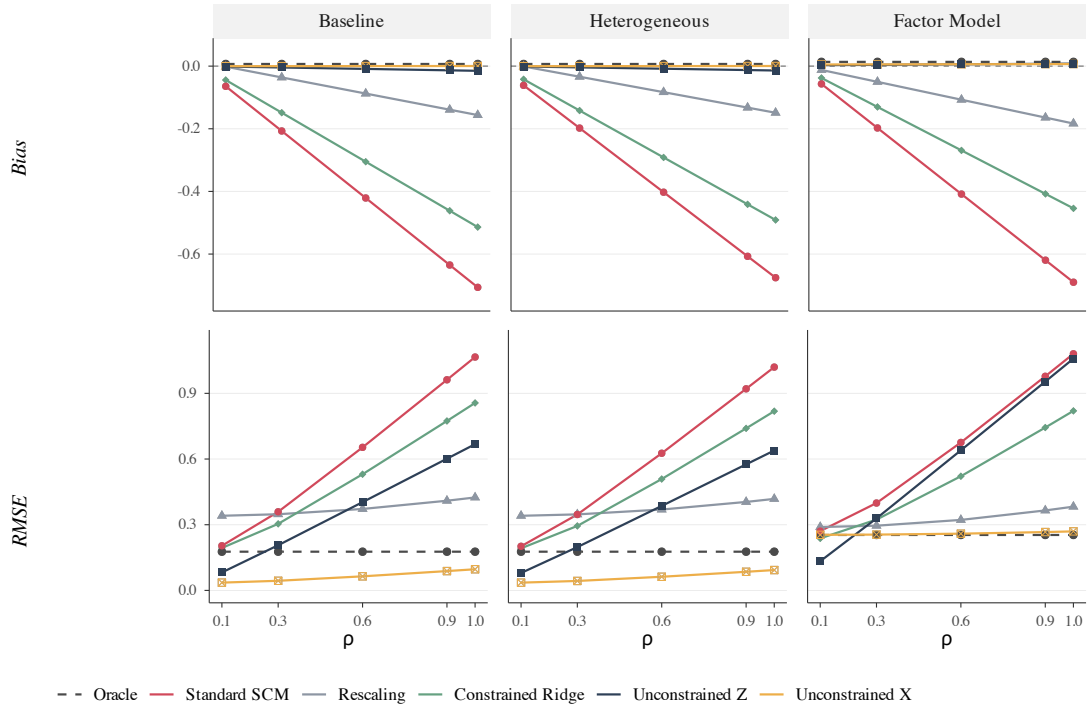



Figure 8: Moment-calibrated counterpart to Figure 3: Bias (top) and RMSE (bottom) across the three DGPs (columns) at  $\tau = 4$ , under the moment reach calibration (M).

Two conclusions follow. First, the reach map need not be pinned down precisely: fixing the logistic slope by quantiles rather than moments leaves the corrector rankings, the bias and RMSE profiles, and—outside the three Constrained-Ridge cells at high  $\rho$  noted above—every CRPS comparison unchanged. Second, this is calibration-*method* invariance, established with the logistic reach family held correct; it is a milder property than robustness to genuine misspecification of the reach—a wrong slope, a dropped parametric family, or a non-monotone reach—which is the separate question studied in the main text (Section 7.3).


## A.7 Ring geometry and mapping from continuous to discretized exposure scores

Figure A.1 illustrates the ring geometry (grid with concentric circles centered at the treated unit ( $i = 1$ )) and the mapping from continuous distance  $d_{i,1}$  to discretized exposure scores  $s_k \approx \kappa(d_{i,1}; \theta)$ .

## A.8 Alternative rings contrast



(a) 2 vs 3 Contrast for Missouri,  
 $p = 0.9591$



(b) 3 vs 4 Contrast for Missouri,  
 $p = 0.5102041$

Figure 9: Alternative contrasts

## A.9 Ring diagnostic across replication applications

Table 13 reports the ring-based permutation test of Section 3 applied to each of the four canonical designs revisited in Section 8, together with the expanded Abadie et al. (2015) donor pool. The test statistic is the pooled-variance two-sample  $t$ -statistic contrasting the near ring against the outer rings;  $p$ -values are finite-sample exact under the sharp null and are computed by permuting the treated-unit identity over the  $J + 1$  candidates in each donor pool. The reported  $p$ -values are the weak definition  $\hat{p}^{(w)} = (1 + \#\{|t_p| \geq |t_1|\}) / (J + 2)$ . This is the conservative convention that counts ties against the null, in contrast to the strong  $p$ -value that does not.

Table 13: Ring diagnostic for canonical synthetic-control applications

Application	$J+1$	$K$	$t_{\text{obs}}$	$p$ -value	Reject ( $\alpha = 0.05$ )?
Basque Country	16	3	0.99	0.556	✗
West Germany (OECD, 1990)	16	3	0.68	0.389	✗
West Germany (expanded, 1990)	117	5	8.40	0.017	✓
Kansas	47	5	1.20	0.327	✗
DRC	22	3	1.10	0.667	✗

*Notes:* Near–far contrast

(ring 1 vs rings 2..K, pooled-variance two-sample  $t$ -statistic) permuted over treated-unit identity;  $\hat{p}^{(w)} = (1 + \#\{|t_p| \geq |t_1|\}) / (J + 2)$ ; window = full (all pre vs all post).  $K$  is the number of quantile-defined rings;  $J+1$  is the number of units in the donor pool. Script: `verified/simulations/run_detection_only.R`; seed: 20260407.

Because the ring partition is an analyst choice, the signal in the expanded German reunification design is worth examining under alternative partitions. Table 14 reports the test across  $K_r \in \{3, 4, 5, 6, 7\}$  for the expanded Abadie et al. (2015) donor pool, which contains  $J = 117$  donors plus West Germany for a full permutation space of  $J + 2 = 119$  units (the “118-country donor pool” language used in the imported table captions refers to the 117 donors together with West Germany). The test statistic grows monotonically with  $K_r$  from 4.83 at  $K_r = 3$  to 8.51 at  $K_r = 6$ , and the  $p$ -value stabilizes at the discrete floor  $\hat{p} = 2 / (J + 2) \approx 0.017$  across every choice; the rejection is therefore not an artifact of how the rings are drawn but reflects a near–far contrast that survives any monotone refinement of the partition.

Table 14: Sensitivity of the ring diagnostic to num\_bins (Germany Expanded, 118-country donor pool)

num_bins	Ring sizes	$t_{\text{obs}}$	$p$ -value	Reject ( $\alpha = 0.05$ )?
3	39,39,39	4.83	0.017	✓
4	29,29,29,30	7.03	0.017	✓
5	24,23,23,23,24	8.40	0.017	✓
6	20,19,19,20,19,20	8.51	0.017	✓
7	17,17,16,17,16,17,17	8.15	0.017	✓

*Notes:* The test stabilises at the discrete

floor  $p = 2/119 \approx 0.017$  across all  $K$ , indicating the near–far signal is not driven by any single bin choice.

## A.10 Sensitivity to reach misspecification: summary table

Table 15 is the summary companion to the sensitivity analysis of Section 7.3. It reports bias, RMSE, and 95% confidence-interval coverage at the headline configuration  $\tau = 4$ , for the five feasible estimators, under the correctly specified Baseline calibration together with the four reach-misspecification scenarios (A:  $\gamma_\kappa$  doubled; B:  $\gamma_\kappa$  halved; C: sharp cutoff; D: non-monotone reach), across  $\rho \in \{0.3, 0.6, 0.9\}$ . It condenses the same experiment that Figure 4 plots in the main text.

Table 15: Sensitivity to Reach ( $\gamma_\kappa$ ) Misspecification ( $\tau = 4$ )

Scenario	Estimator	$\rho = 0.3$			$\rho = 0.6$			$\rho = 0.9$		
		Bias	RMSE	Cov.	Bias	RMSE	Cov.	Bias	RMSE	Cov.
Baseline	Standard SCM	-0.201	0.348	0.905	-0.423	0.671	0.888	-0.635	0.962	0.871
	Rescaling	-0.048	0.398	0.934	-0.114	0.428	0.940	-0.127	0.473	0.945
	Constrained Ridge	-0.159	0.312	0.923	-0.323	0.578	0.917	-0.481	0.792	0.905
	Unconstrained Z	0.004	0.208	0.958	0.005	0.417	0.943	-0.008	0.661	0.949
	Unconstrained X	-0.001	0.045	0.950	0.001	0.069	0.944	-0.002	0.100	0.938
A: $\gamma_\kappa \times 2$	Standard SCM	-0.208	0.390	0.901	-0.402	0.627	0.879	-0.598	0.925	0.880
	Rescaling	-0.049	0.406	0.937	-0.101	0.405	0.942	-0.164	0.449	0.941
	Constrained Ridge	-0.151	0.358	0.923	-0.302	0.531	0.910	-0.431	0.756	0.920
	Unconstrained Z	0.003	0.272	0.946	-0.013	0.510	0.946	0.006	0.630	0.946
	Unconstrained X	-0.002	0.107	0.930	0.003	0.199	0.934	0.001	0.293	0.932
B: $\gamma_\kappa \times 0.5$	Standard SCM	-0.208	0.364	0.902	-0.388	0.607	0.879	-0.625	0.933	0.872
	Rescaling	-0.041	0.383	0.941	-0.098	0.429	0.934	-0.161	0.488	0.934
	Constrained Ridge	-0.171	0.323	0.927	-0.312	0.508	0.913	-0.509	0.790	0.908
	Unconstrained Z	0.016	0.215	0.953	-0.005	0.411	0.947	0.041	0.672	0.948
	Unconstrained X	0.001	0.066	0.937	-0.001	0.119	0.936	0.009	0.184	0.943
C: Sharp cutoff	Standard SCM	-0.099	0.293	0.937	-0.226	0.447	0.915	-0.311	0.647	0.931
	Rescaling	-0.044	0.406	0.936	-0.056	0.381	0.936	-0.082	0.440	0.946
	Constrained Ridge	-0.077	0.278	0.942	-0.169	0.369	0.929	-0.222	0.532	0.950
	Unconstrained Z	0.002	0.139	0.954	0.003	0.222	0.954	-0.009	0.391	0.958
	Unconstrained X	-0.002	0.040	0.951	-0.001	0.053	0.943	0.001	0.074	0.942
D: Non-monotone	Standard SCM	0.004	0.364	0.958	0.006	0.652	0.934	0.016	0.924	0.930
	Rescaling	0.032	0.380	0.944	0.034	0.430	0.933	0.039	0.432	0.929
	Constrained Ridge	0.007	0.319	0.958	0.011	0.543	0.936	0.034	0.735	0.918
	Unconstrained Z	0.012	0.236	0.958	0.014	0.453	0.954	0.017	0.755	0.946
	Unconstrained X	-0.002	0.054	0.946	0.000	0.091	0.946	0.007	0.128	0.951

*Notes:* Bias, RMSE, and 95% CI coverage under Baseline (correct  $\gamma_\kappa$ ) plus four reach-calibration misspecification scenarios (A–D).  $\tau = 4$ ,  $R = 1,000$  replications. A/B: multiplicative  $\gamma_\kappa$  misspecification. C: true DGP uses binary spillover (sharp cutoff). D: non-monotone spillover. Nominal coverage = 0.95. †: cells with fewer than 900 valid replications (of 1,000).

The per-cell expansion of this summary across the full  $\tau \times \rho$  grid is given by Tables 22–23 in the next subsection (Appendix A.11, Group C); the present table fixes  $\tau = 4$  and the three-point  $\rho$  grid used in the main-text narrative.

## A.11 Full simulation tables

This appendix collects the full per-cell simulation grids and the detection-test placebo-skip audit. The material is organized into five groups: per-DGP performance grids, calibration comparisons, sensitivity to reach misspecification, pre-period and tuning diagnostics, and detection diagnostics.

### Group A: per-DGP performance grids

Tables 16–21 report the full performance grid across all three DGPs, all three values of  $\tau \in \{1, 4, 7\}$ , and all five values of  $\rho \in \{0.1, 0.3, 0.6, 0.9, 1.0\}$ . Each DGP spans two tables (split ‘\_a’ / ‘\_b’) to keep the column count tractable. The Oracle and a trimmed-donor variant (see Appendix A.12) are reported alongside the five feasible estimators shown in the main text.

### Group B: calibration comparisons

The Q-vs-M reach-calibration comparisons—the CRPS grids and the moment-calibrated performance figure—are collected in Appendix A.6.

### Group C: sensitivity to reach misspecification

Tables 22–23 extend the summary table (Appendix A.10, Table 15) across the full  $\tau \times \rho$  grid.

### Group D: pre-period fit and tuning diagnostics

Table 24 reports pre-treatment RMSE — a fit-quality complement to the post-treatment performance metrics in the main text. Table 25 reports the distribution of cross-validated  $\lambda$  for the three ridge corrections, providing a direct view of where the CV optimum sits on the tuning grid.

### Group E: detection diagnostics

Table 26 extends the detection-rate grid of Table 4 across the three DGPs. Table 27 reports the per-state, per- $K_r$  frequency of placebo permutations that are skipped because a required ring is empty for the placebo assignment; the skip rate is at most 4.2% in our simulation and is reported transparently for audit.

## A.12 Trimming as a baseline

Trimming — dropping all donors flagged as exposed before estimating — is the simplest exposure-aware alternative to the corrections developed in Sections 4–6 and a reasonable robustness benchmark. We do not recommend it as a production method, for four reasons detailed below. We do not report CRPS for Trimming because the trimmed-donor reference distribution is constructed on a different (smaller) donor pool and is not directly comparable to the CRPS computed on the full pool for the other estimators.

The four reasons share a common root: Trimming acts on a hard partition of the donor pool, and that partition is neither observable nor stable. The first difficulty is that researchers would need to know exactly which units to drop. In applied settings, exposure is inferred from an estimated reach map; donor-level exposure flags are themselves estimates, not observables. Trimming treats these estimates as if they were known, so mis-specification of the reach map cascades directly into the trimmed donor pool: a donor whose true exposure is 0.4 but whose estimated  $\hat{\eta}_j$  would push it onto the trim list is silently dropped, with no recourse through the post-period diagnostics. The partition is also fragile, because Trimming imposes a binary cutoff on the continuous safety score  $\eta_j$ . We use a single calibrated cutoff in our simulations, but different cutoffs yield different performance, whereas the reach-based correctors of Sections 4–6 use  $\eta_j$  continuously and carry no threshold decision. Forcing that cutoff discards usable information: donors with partial exposure ( $\eta_j \approx 0.4$ – $0.6$ ) carry real signal about the treated unit’s factor structure, sharing covariate distributions, latent loadings, and outcome trends with the treated unit. Trimming discards this signal entirely, while URX retains

Table 16: Monte Carlo Performance (Baseline DGP, Part A: Oracle, Standard, Rescaling, Constrained Ridge)

$\tau$	$\rho$	Oracle			Standard			Rescaling			Constr. Ridge		
		Bias	RMSE	CRPS	Bias	RMSE	CRPS	Bias	RMSE	CRPS	Bias	RMSE	CRPS
0.5	0.1	0.007	0.177	0.010	-0.002	0.177	0.011	0.015	0.380	0.039	0.176	0.011	
	0.3	0.007	0.177	0.010	-0.020	0.180	0.017	0.011	0.380	0.039	-0.011	0.178	
	0.6	0.007	0.177	0.010	-0.047	0.192	0.029	0.006	0.380	0.040	-0.033	0.187	
	0.9	0.007	0.177	0.010	-0.073	0.211	0.042	-0.000	0.380	0.041	-0.053	0.199	
	1.0	0.007	0.177	0.010	-0.082	0.218	0.047	-0.002	0.380	0.041	-0.060	0.204	
1.0	0.1	0.007	0.177	0.010	-0.011	0.178	0.014	0.013	0.380	0.039	-0.006	0.178	
	0.3	0.007	0.177	0.010	-0.047	0.192	0.029	0.006	0.380	0.040	-0.033	0.187	
	0.6	0.007	0.177	0.010	-0.100	0.235	0.056	-0.006	0.381	0.042	-0.073	0.215	
	0.9	0.007	0.177	0.010	-0.154	0.293	0.083	-0.018	0.383	0.045	-0.114	0.257	
	1.0	0.007	0.177	0.010	-0.171	0.314	0.092	-0.022	0.384	0.046	-0.127	0.273	
4.0	0.1	0.007	0.177	0.010	-0.064	0.204	0.038	0.002	0.380	0.040	-0.046	0.194	
	0.3	0.007	0.177	0.010	-0.207	0.359	0.110	-0.030	0.386	0.049	-0.154	0.306	
	0.6	0.007	0.177	0.010	-0.421	0.654	0.214	-0.077	0.406	0.067	-0.315	0.536	
	0.9	0.007	0.177	0.010	-0.635	0.962	0.316	-0.124	0.437	0.087	-0.477	0.782	
	1.0	0.007	0.177	0.010	-0.706	1.066	0.350	-0.139	0.450	0.093	-0.531	0.865	
7.0	0.1	0.007	0.177	0.010	-0.118	0.253	0.065	-0.010	0.382	0.043	-0.087	0.228	
	0.3	0.007	0.177	0.010	-0.368	0.578	0.188	-0.065	0.400	0.062	-0.275	0.476	
	0.6	0.007	0.177	0.010	-0.742	1.118	0.367	-0.147	0.457	0.096	-0.558	0.907	
	0.9	0.007	0.177	0.010	-1.116	1.668	0.543	-0.230	0.539	0.132	-0.840	1.349	
	1.0	0.007	0.177	0.010	-1.241	1.851	0.602	-0.257	0.571	0.143	-0.934	1.497	

$Y_{jt} = \beta X_{1j} + \beta X_{2j} + \tau Z_{jt} + \varepsilon_{jt}$  with unit covariates  $X_{1j}, X_{2j} \sim \mathcal{N}(0, 1)$ ,  $\beta = 1$ , and  $\varepsilon_{jt} \sim \mathcal{N}(0, 0.1^2)$ ;  $N = 30$  periods, treatment in Missouri starts at  $t = 21$ . Homogeneous post-treatment spillover of magnitude  $\rho\tau$  is added to Missouri's contiguous neighbors (spillover structure taken from the US state adjacency graph).  $R = 1,000$  replications per  $(\tau, \rho)$  cell. Bias =  $\hat{\tau} - \bar{\tau}$ , averaged over post-treatment periods; RMSE =  $\sqrt{\text{Bias}^2 + \text{SD}^2}$ ; CRPS measures distributional accuracy (lower is better). "Oracle" row: synthetic control computed using only uncontaminated donor units (those not affected by interference spillover); serves as the benchmark for bias-free behavior. "Standard" row: the naive SCM estimator fit on the full donor pool with no correction for interference; expected to be biased when  $\rho > 0$ . Trimming CRPS not computed (shown as —). Quantile calibration used for all corrections. See Table 17 for the remaining estimators. †: cells with fewer than 900 valid replications (of 1,000).

Notes: Baseline DGP: panel outcome

Table 17: Monte Carlo Performance (Baseline DGP, Part B: Unconstrained Ridge (Z and X), Trimming)

$\tau$	$\rho$	Unconstr. Z				Unconstr. X				Trimming			
		Bias	RMSE	CRPS	CRPS	Bias	RMSE	CRPS	CRPS	Bias	RMSE	CRPS	CRPS
0.5	0.1	0.001	0.047	0.011	0.011	0.001	0.039	0.009	0.009	0.010	0.222	—	—
	0.3	0.000	0.055	0.011	0.011	0.001	0.039	0.009	0.009	0.010	0.222	—	—
	0.6	-0.001	0.068	0.011	0.011	-0.000	0.035	0.008	0.008	0.010	0.222	—	—
	0.9	-0.001	0.088	0.012	0.012	-0.000	0.037	0.008	0.008	0.010	0.222	—	—
	1.0	-0.001	0.095	0.012	0.012	-0.000	0.037	0.008	0.008	0.010	0.222	—	—
1.0	0.1	-0.000	0.049	0.011	0.011	-0.000	0.035	0.008	0.008	0.010	0.222	—	—
	0.3	-0.001	0.068	0.011	0.011	-0.000	0.035	0.008	0.008	0.010	0.222	—	—
	0.6	-0.001	0.110	0.012	0.012	-0.000	0.038	0.009	0.009	0.010	0.222	—	—
	0.9	-0.002	0.156	0.013	0.013	-0.000	0.042	0.009	0.009	0.010	0.222	—	—
	1.0	-0.003	0.172	0.013	0.013	-0.000	0.043	0.009	0.009	0.010	0.222	—	—
4.0	0.1	-0.001	0.081	0.012	0.012	-0.000	0.036	0.008	0.008	0.010	0.222	—	—
	0.3	-0.003	0.204	0.014	0.014	0.000	0.046	0.010	0.010	0.010	0.222	—	—
	0.6	-0.007	0.399	0.018	0.018	0.000	0.070	0.014	0.014	0.010	0.222	—	—
	0.9	-0.010	0.595	0.021	0.021	0.000	0.097	0.018	0.018	0.010	0.222	—	—
	1.0	-0.011	0.661	0.022	0.022	0.001	0.107	0.019	0.019	0.010	0.222	—	—
7.0	0.1	-0.002	0.125	0.012	0.012	-0.000	0.039	0.009	0.009	0.010	0.222	—	—
	0.3	-0.006	0.350	0.017	0.017	0.000	0.063	0.013	0.013	0.010	0.222	—	—
	0.6	-0.012	0.693	0.023	0.023	0.001	0.111	0.020	0.020	0.010	0.222	—	—
	0.9	-0.018	1.038	0.029	0.029	0.001	0.162	0.027	0.027	0.010	0.222	—	—
	1.0	-0.020	1.153	0.031	0.031	0.001	0.180	0.030	0.030	0.010	0.222	—	—

Notes: Baseline DGP: panel outcome

$Y_{jt} = \beta X_{1j} + \beta X_{2j} + \tau Z_{jt} + \varepsilon_{jt}$  with unit covariates  $X_{1j}, X_{2j} \sim \mathcal{N}(0, 1)$ ,  $\beta = 1$ , and  $\varepsilon_{jt} \sim \mathcal{N}(0, 0.1^2)$ ;  $N = 30$  periods, treatment in Missouri starts at  $t = 21$ . Homogeneous post-treatment spillover of magnitude  $\rho\tau$  is added to Missouri's contiguous neighbors (spillover structure taken from the US state adjacency graph).  $R = 1,000$  replications per  $(\tau, \rho)$  cell. Bias =  $\bar{\tau} - \bar{\tau}$ , averaged over post-treatment periods; RMSE =  $\sqrt{\text{Bias}^2 + \text{SD}^2}$ ; CRPS measures distributional accuracy (lower is better). "Oracle" row: synthetic control computed using only uncontaminated donor units (those not affected by interference spillover); serves as the benchmark for bias-free behavior. "Standard" row: the naive SCM estimator fit on the full donor pool with no correction for interference; expected to be biased when  $\rho > 0$ . Trimming CRPS not computed (shown as —). Quantile calibration used for all corrections. See Table 16 for the remaining estimators. †: cells with fewer than 900 valid replications (of 1,000).

Table 18: Monte Carlo Performance (Heterogeneous DGP, Part A: Oracle, Standard, Rescaling, Constrained Ridge)

$\tau$	$\rho$	Oracle			Standard			Rescaling			Constr. Ridge		
		Bias	RMSE	CRPS	Bias	RMSE	CRPS	Bias	RMSE	CRPS	Bias	RMSE	CRPS
0.5	0.1	0.007	0.177	0.010	-0.002	0.177	0.011	0.015	0.380	0.039	0.178	0.011	
	0.3	0.007	0.177	0.010	-0.019	0.180	0.017	0.012	0.380	0.039	0.179	0.015	
	0.6	0.007	0.177	0.010	-0.044	0.191	0.028	0.006	0.380	0.040	0.186	0.022	
	0.9	0.007	0.177	0.010	-0.070	0.208	0.041	0.000	0.380	0.041	0.197	0.030	
	1.0	0.007	0.177	0.010	-0.078	0.215	0.045	-0.001	0.380	0.041	0.202	0.033	
1.0	0.1	0.007	0.177	0.010	-0.010	0.178	0.014	0.014	0.380	0.039	0.178	0.013	
	0.3	0.007	0.177	0.010	-0.044	0.191	0.028	0.006	0.380	0.040	0.186	0.022	
	0.6	0.007	0.177	0.010	-0.095	0.230	0.054	-0.005	0.381	0.042	0.212	0.039	
	0.9	0.007	0.177	0.010	-0.147	0.284	0.079	-0.016	0.383	0.045	0.250	0.058	
	1.0	0.007	0.177	0.010	-0.164	0.305	0.088	-0.020	0.383	0.046	0.265	0.064	
4.0	0.1	0.007	0.177	0.010	-0.061	0.201	0.037	0.002	0.380	0.040	0.193	0.028	
	0.3	0.007	0.177	0.010	-0.198	0.347	0.105	-0.028	0.385	0.048	0.297	0.076	
	0.6	0.007	0.177	0.010	-0.403	0.627	0.205	-0.073	0.404	0.065	0.514	0.150	
	0.9	0.007	0.177	0.010	-0.607	0.921	0.303	-0.118	0.433	0.084	0.748	0.221	
	1.0	0.007	0.177	0.010	-0.675	1.020	0.335	-0.133	0.444	0.090	0.827	0.245	
7.0	0.1	0.007	0.177	0.010	-0.113	0.247	0.062	-0.009	0.381	0.043	0.224	0.045	
	0.3	0.007	0.177	0.010	-0.351	0.554	0.180	-0.061	0.398	0.061	0.457	0.131	
	0.6	0.007	0.177	0.010	-0.710	1.070	0.351	-0.140	0.451	0.093	0.867	0.257	
	0.9	0.007	0.177	0.010	-1.068	1.595	0.521	-0.219	0.528	0.127	1.288	0.381	
	1.0	0.007	0.177	0.010	-1.187	1.770	0.577	-0.245	0.557	0.138	1.429	0.421	

Notes: Heterogeneous DGP: identical

panel structure to the Baseline DGP, but spillover varies by neighbor  $j$  via  $\rho_j = \hat{\rho}(1 - \eta_j)$ , where  $\eta_j \in [0, 1]$  is the logistic reach score computed on great-circle (Haversine) distances from Missouri to each control state. Closer neighbors (low  $\eta_j$ ) receive stronger spillover ( $\rho_j \rightarrow \hat{\rho}$ ); distant neighbors receive little. The column labelled  $\rho$  is the intensity parameter  $\hat{\rho}$ .  $R = 1,000$  replications per ( $\tau, \rho$ ) cell. Bias =  $\bar{\tau} - \tau$ , averaged over post-treatment periods; RMSE =  $\sqrt{\text{Bias}^2 + \text{SD}^2}$ ; CRPS measures distributional accuracy (lower is better). "Oracle" row: synthetic control computed using only uncontaminated donor units (those not affected by interference spillover); serves as the benchmark for bias-free behavior. "Standard" row: the naive SCM estimator fit on the full donor pool with no correction for interference; expected to be biased when  $\rho > 0$ . Trimming CRPS not computed (shown as —). Quantile calibration used for all corrections. See Table 19 for the remaining estimators. †: cells with fewer than 900 valid replications (of 1,000).

Table 19: Monte Carlo Performance (Heterogeneous DGP, Part B: Unconstrained Ridge (Z and X), Trimming)

$\tau$	$\rho$	Unconstr. Z				Unconstr. X				Trimming			
		Bias	RMSE	CRPS	CRPS	Bias	RMSE	CRPS	CRPS	Bias	RMSE	CRPS	CRPS
0.5	0.1	0.000	0.046	0.010	0.010	-0.000	0.035	0.008	0.008	0.010	0.222	0.222	—
	0.3	-0.000	0.052	0.011	0.011	-0.000	0.035	0.008	0.008	0.010	0.222	0.222	—
	0.6	-0.001	0.067	0.011	0.011	-0.000	0.035	0.008	0.008	0.010	0.222	0.222	—
	0.9	-0.001	0.085	0.012	0.012	-0.000	0.036	0.008	0.008	0.010	0.222	0.222	—
	1.0	-0.001	0.092	0.012	0.012	-0.000	0.037	0.008	0.008	0.010	0.222	0.222	—
1.0	0.1	-0.000	0.048	0.011	0.011	-0.000	0.035	0.008	0.008	0.010	0.222	0.222	—
	0.3	-0.001	0.067	0.011	0.011	-0.000	0.035	0.008	0.008	0.010	0.222	0.222	—
	0.6	-0.001	0.106	0.012	0.012	-0.000	0.038	0.008	0.008	0.010	0.222	0.222	—
	0.9	-0.002	0.150	0.013	0.013	-0.000	0.041	0.009	0.009	0.010	0.222	0.222	—
	1.0	-0.002	0.165	0.013	0.013	-0.000	0.042	0.009	0.009	0.010	0.222	0.222	—
4.0	0.1	-0.001	0.079	0.011	0.011	-0.000	0.036	0.008	0.008	0.010	0.222	0.222	—
	0.3	-0.003	0.196	0.014	0.014	0.000	0.045	0.010	0.010	0.010	0.222	0.222	—
	0.6	-0.006	0.382	0.017	0.017	0.000	0.068	0.013	0.013	0.010	0.222	0.222	—
	0.9	-0.010	0.569	0.021	0.021	0.001	0.094	0.017	0.017	0.010	0.222	0.222	—
	1.0	-0.011	0.632	0.022	0.022	0.001	0.103	0.018	0.018	0.010	0.222	0.222	—
7.0	0.1	-0.002	0.121	0.012	0.012	-0.000	0.039	0.009	0.009	0.010	0.222	0.222	—
	0.3	-0.005	0.335	0.016	0.016	0.000	0.062	0.013	0.013	0.010	0.222	0.222	—
	0.6	-0.011	0.663	0.022	0.022	0.001	0.107	0.019	0.019	0.010	0.222	0.222	—
	0.9	-0.017	0.993	0.028	0.028	0.001	0.156	0.026	0.026	0.010	0.222	0.222	—
	1.0	-0.019	1.103	0.030	0.030	0.001	0.172	0.029	0.029	0.010	0.222	0.222	—

Notes: Heterogeneous DGP: identical panel structure to the

Baseline DGP, but spillover varies by neighbor  $j$  via  $\rho_j = \bar{\rho}(1 - \eta_j)$ , where  $\eta_j \in [0, 1]$  is the logistic reach score computed on great-circle (Haversine) distances from Missouri to each control state. Closer neighbors (low  $\eta_j$ ) receive stronger spillover ( $\rho_j \rightarrow \bar{\rho}$ ); distant neighbors receive little. The column labelled  $\rho$  is the intensity parameter  $\bar{\rho}$ .  $R = 1,000$  replications per  $(\tau, \rho)$  cell. Bias =  $\bar{\tau} - \tau$ , averaged over post-treatment periods; RMSE =  $\sqrt{\text{Bias}^2 + \text{SD}^2}$ ; CRPS measures distributional accuracy (lower is better). “Oracle” row: synthetic control computed using only uncontaminated donor units (those not affected by interference spillover); serves as the benchmark for bias-free behavior. “Standard” row: the naive SCM estimator fit on the full donor pool with no correction for interference; expected to be biased when  $\rho > 0$ . Trimming CRPS not computed (shown as —). Quantile calibration used for all corrections. See Table 18 for the remaining estimators. †: cells with fewer than 900 valid replications (of 1,000).

Table 20: Monte Carlo Performance (Factor Model DGP, Part A: Oracle, Standard, Rescaling, Constrained Ridge)

$\tau$	$\rho$	Oracle			Standard			Rescaling			Constr. Ridge		
		Bias	RMSE	CRPS	Bias	RMSE	CRPS	Bias	RMSE	CRPS	Bias	RMSE	CRPS
0.5	0.1	0.014	0.253	0.043	0.005	0.253	0.043	0.007	0.301	0.047	0.001	0.221	0.035
	0.3	0.014	0.253	0.043	-0.013	0.255	0.044	0.003	0.300	0.047	-0.011	0.222	0.037
	0.6	0.014	0.253	0.043	-0.039	0.263	0.049	-0.004	0.300	0.048	-0.029	0.228	0.040
	0.9	0.014	0.253	0.043	-0.066	0.276	0.055	-0.011	0.300	0.049	-0.046	0.236	0.043
	1.0	0.014	0.253	0.043	-0.074	0.281	0.058	-0.013	0.300	0.049	-0.052	0.240	0.045
1.0	0.1	0.014	0.253	0.043	-0.004	0.254	0.043	0.005	0.300	0.047	-0.005	0.222	0.036
	0.3	0.014	0.253	0.043	-0.039	0.263	0.049	-0.004	0.300	0.048	-0.029	0.228	0.040
	0.6	0.014	0.253	0.043	-0.092	0.294	0.063	-0.017	0.301	0.049	-0.064	0.249	0.048
	0.9	0.014	0.253	0.043	-0.145	0.341	0.080	-0.030	0.302	0.051	-0.099	0.281	0.058
	1.0	0.014	0.253	0.043	-0.162	0.360	0.087	-0.035	0.303	0.052	-0.111	0.294	0.062
4.0	0.1	0.014	0.253	0.043	-0.057	0.271	0.053	-0.008	0.300	0.048	-0.041	0.233	0.042
	0.3	0.014	0.253	0.043	-0.197	0.399	0.100	-0.044	0.305	0.054	-0.135	0.321	0.069
	0.6	0.014	0.253	0.043	-0.409	0.676	0.188	-0.096	0.329	0.070	-0.276	0.522	0.121
	0.9	0.014	0.253	0.043	-0.620	0.978	0.283	-0.149	0.368	0.092	-0.417	0.745	0.178
	1.0	0.014	0.253	0.043	-0.690	1.081	0.316	-0.167	0.383	0.100	-0.464	0.822	0.198
7.0	0.1	0.014	0.253	0.043	-0.110	0.308	0.068	-0.022	0.301	0.050	-0.076	0.258	0.051
	0.3	0.014	0.253	0.043	-0.356	0.603	0.165	-0.083	0.321	0.066	-0.240	0.468	0.107
	0.6	0.014	0.253	0.043	-0.725	1.132	0.332	-0.175	0.392	0.104	-0.487	0.860	0.207
	0.9	0.014	0.253	0.043	-1.094	1.680	0.502	-0.268	0.491	0.145	-0.734	1.269	0.312
	1.0	0.014	0.253	0.043	-1.218	1.863	0.559	-0.298	0.528	0.158	-0.816	1.407	0.347

Notes: Factor Model DGP: outcome

generated from a latent 2-factor structure,  $Y_{jt}(0) = \lambda_j' f_t + \varepsilon_{jt}$ , with loadings  $\lambda_j \sim \mathcal{N}(0, I_2)$  and factors  $f_{k,t}$  following AR(1) with autoregressive coefficient 0.8 and innovation  $\mathcal{N}(0, 0.3^2)$ ;  $\varepsilon_{jt} \sim \mathcal{N}(0, 0.1^2)$ . Observable covariates  $X_{1j}, X_{2j}$  are noisy proxies for the first two loadings. Post-treatment, homogeneous spillover  $\rho\tau$  is added to Missouri's contiguous neighbors.  $R = 1,000$  replications per  $(\tau, \rho)$  cell. Bias =  $\bar{\tau} - \tau$ , averaged over post-treatment periods; RMSE =  $\sqrt{\text{Bias}^2 + \text{SD}^2}$ ; CRPS measures distributional accuracy (lower is better). "Oracle" row: synthetic control computed using only uncontaminated donor units (those not affected by interference spillover); serves as the benchmark for bias-free behavior. "Standard" row: the naive SCM estimator fit on the full donor pool with no correction for interference; expected to be biased when  $\rho > 0$ . Trimming CRPS not computed (shown as —). Quantile calibration used for all corrections. See Table 21 for the remaining estimators. †: cells with fewer than 900 valid replications (of 1,000).

Table 21: Monte Carlo Performance (Factor Model DGP, Part B: Unconstrained Ridge (Z and X), Trimming)

$\tau$	$\rho$	Unconstr. Z			Unconstr. X			Trimming		
		Bias	RMSE	CRPS	Bias	RMSE	CRPS	Bias	RMSE	CRPS
0.5	0.1	0.003	0.077	0.016	0.005	0.256	0.046	0.012	0.259	—
	0.3	0.003	0.088	0.016	0.005	0.256	0.046	0.012	0.259	—
	0.6	0.003	0.115	0.017	0.005	0.256	0.046	0.012	0.259	—
	0.9	0.004	0.148	0.018	0.005	0.256	0.046	0.012	0.259	—
	1.0	0.004	0.160	0.019	0.005	0.256	0.046	0.012	0.259	—
1.0	0.1	0.003	0.081	0.016	0.005	0.256	0.046	0.012	0.259	—
	0.3	0.003	0.115	0.017	0.005	0.256	0.046	0.012	0.259	—
	0.6	0.004	0.185	0.019	0.005	0.256	0.046	0.012	0.259	—
	0.9	0.005	0.263	0.022	0.005	0.256	0.046	0.012	0.259	—
	1.0	0.005	0.289	0.023	0.005	0.257	0.047	0.012	0.259	—
4.0	0.1	0.004	0.137	0.018	0.005	0.256	0.046	0.012	0.259	—
	0.3	0.005	0.342	0.025	0.005	0.257	0.047	0.012	0.259	—
	0.6	0.008	0.668	0.039	0.006	0.262	0.049	0.012	0.259	—
	0.9	0.011	0.996	0.053	0.007	0.272	0.052	0.012	0.259	—
	1.0	0.012	1.106	0.057	0.007	0.275	0.053	0.012	0.259	—
7.0	0.1	0.004	0.211	0.020	0.005	0.256	0.046	0.012	0.259	—
	0.3	0.007	0.586	0.035	0.006	0.261	0.049	0.012	0.259	—
	0.6	0.012	1.160	0.060	0.007	0.278	0.053	0.012	0.259	—
	0.9	0.017	1.736	0.085	0.008	0.304	0.060	0.012	0.259	—
	1.0	0.019	1.928	0.093	0.008	0.315	0.062	0.012	0.259	—

Notes: Factor Model DGP: outcome generated from a latent 2-factor

structure,  $Y_{jt}(0) = \lambda_j' f_t + \varepsilon_{jt}$ , with loadings  $\lambda_j \sim \mathcal{N}(0, I_2)$  and factors  $f_{k,t}$  following AR(1) with autoregressive coefficient 0.8 and innovation  $\mathcal{N}(0, 0.3^2)$ ;  $\varepsilon_{jt} \sim \mathcal{N}(0, 0.1^2)$ . Observable covariates  $X_{1j}, X_{2j}$  are noisy proxies for the first two loadings. Post-treatment, homogeneous spillover  $\rho\tau$  is added to Missouri's contiguous neighbors.  $R = 1,000$  replications per  $(\tau, \rho)$  cell. Bias =  $\bar{\tau} - \bar{\tau}$ , averaged over post-treatment periods; RMSE =  $\sqrt{\text{Bias}^2 + \text{SD}^2}$ ; CRPS measures distributional accuracy (lower is better). "Oracle" row: synthetic control computed using only uncontaminated donor units (those not affected by interference spillover); serves as the benchmark for bias-free behavior. "Standard" row: the naive SCM estimator fit on the full donor pool with no correction for interference; expected to be biased when  $\rho > 0$ . Trimming CRPS not computed (shown as —). Quantile calibration used for all corrections. See Table 20 for the remaining estimators. †: cells with fewer than 900 valid replications (of 1,000).

Table 22: Sensitivity to Reach ( $\gamma_\kappa$ ): Full Diagnostics,  $\rho = 0.3$ ,  $\rho = 0.6$  ( $\tau = 4$ )

Scenario	Estimator	$\rho = 0.3$						$\rho = 0.6$					
		Bias	RMSE	Cov.	Pre-RMSE	$\bar{\lambda}$	Exp. Mass	Bias	RMSE	Cov.	Pre-RMSE	$\bar{\lambda}$	Exp. Mass
Baseline	Standard SCM	-0.201	0.348	0.905	0.140	—	0.000	-0.423	0.671	0.888	0.147	—	0.000
	Rescaling	-0.048	0.398	0.934	0.267	—	0.000	-0.114	0.428	0.940	0.272	—	0.000
	Constrained Ridge	-0.159	0.312	0.923	0.139	0.220	0.000	-0.323	0.578	0.917	0.143	0.224	0.000
	Unconstrained Z	0.004	0.208	0.958	0.075	0.535	0.000	0.005	0.417	0.943	0.072	0.498	0.000
	Unconstrained X	-0.001	0.045	0.950	0.108	0.463	0.000	0.001	0.069	0.944	0.107	0.442	0.000
A: $\gamma_\kappa \times 2$	Standard SCM	-0.208	0.390	0.901	0.159	—	0.000	-0.402	0.627	0.879	0.144	—	0.000
	Rescaling	-0.049	0.406	0.937	0.256	—	0.000	-0.101	0.405	0.942	0.240	—	0.000
	Constrained Ridge	-0.151	0.358	0.923	0.160	0.172	0.000	-0.302	0.531	0.910	0.145	0.188	0.000
	Unconstrained Z	0.003	0.272	0.946	0.072	0.649	0.000	-0.013	0.510	0.946	0.072	0.656	0.000
	Unconstrained X	-0.002	0.107	0.930	0.103	0.112	0.000	0.003	0.199	0.934	0.103	0.118	0.000
B: $\gamma_\kappa \times 0.5$	Standard SCM	-0.208	0.364	0.902	0.143	—	0.000	-0.388	0.607	0.879	0.143	—	0.000
	Rescaling	-0.041	0.383	0.941	0.272	—	0.000	-0.098	0.429	0.934	0.279	—	0.000
	Constrained Ridge	-0.171	0.323	0.927	0.140	0.245	0.000	-0.312	0.508	0.913	0.142	0.226	0.000
	Unconstrained Z	0.016	0.215	0.953	0.072	0.498	0.000	-0.005	0.411	0.947	0.071	0.484	0.000
	Unconstrained X	0.001	0.066	0.937	0.102	0.466	0.000	-0.001	0.119	0.936	0.102	0.503	0.000
C: Sharp cutoff	Standard SCM	-0.099	0.293	0.937	0.152	—	0.000	-0.226	0.447	0.915	0.137	—	0.000
	Rescaling	-0.044	0.406	0.936	0.273	—	0.000	-0.056	0.381	0.936	0.260	—	0.000
	Constrained Ridge	-0.077	0.278	0.942	0.151	0.203	0.000	-0.169	0.369	0.929	0.135	0.222	0.000
	Unconstrained Z	0.002	0.139	0.954	0.073	0.493	0.000	0.003	0.222	0.954	0.073	0.511	0.000
	Unconstrained X	-0.002	0.040	0.951	0.105	0.460	0.000	-0.001	0.053	0.943	0.107	0.453	0.000
D: Non-monotone	Standard SCM	0.004	0.364	0.958	0.144	—	0.000	0.006	0.652	0.934	0.142	—	0.000
	Rescaling	0.032	0.380	0.944	0.255	—	0.000	0.034	0.430	0.933	0.278	—	0.000
	Constrained Ridge	0.007	0.319	0.958	0.142	0.233	0.000	0.011	0.543	0.936	0.140	0.206	0.000
	Unconstrained Z	0.012	0.236	0.958	0.073	0.500	0.000	0.014	0.453	0.954	0.073	0.518	0.000
	Unconstrained X	-0.002	0.054	0.946	0.107	0.456	0.000	0.000	0.091	0.946	0.107	0.447	0.000

Notes: Full diagnostics under Baseline (correct  $\gamma_\kappa$ ) plus four reach-calibration misspecification scenarios (A–D). Shown:  $\rho \in \{0.3, 0.6\}$ .  $\tau = 4$ ,  $R = 1,000$  replications. Cov. = empirical 95% CI coverage rate (nominal 0.95). Pre-RMSE = mean pre-treatment RMSE of fitted synthetic control.  $\bar{\lambda}$  = mean selected regularization parameter; shown as — for Standard SCM and Rescaling (no  $\lambda$  penalty). Exp. Mass = mean fraction of SCM weight on exposed (neighboring) donors. See Table 23 for the remaining  $\rho$  values. †: cells with fewer than 900 valid replications (of 1,000).

Table 23: Sensitivity to Reach ( $\gamma_\kappa$ ): Full Diagnostics,  $\rho = 0.9$  ( $\tau = 4$ )

Scenario	Estimator	$\rho = 0.9$					
		Bias	RMSE	Cov.	Pre-RMSE	$\bar{\lambda}$	Exp. Mass
Baseline	Standard SCM	-0.635	0.962	0.871	0.151	—	0.000
	Rescaling	-0.127	0.473	0.945	0.281	—	0.000
	Constrained Ridge	-0.481	0.792	0.905	0.151	0.223	0.000
	Unconstrained Z	-0.008	0.661	0.949	0.074	0.530	0.000
	Unconstrained X	-0.002	0.100	0.938	0.108	0.497	0.000
A: $\gamma_\kappa \times 2$	Standard SCM	-0.598	0.925	0.880	0.144	—	0.000
	Rescaling	-0.164	0.449	0.941	0.245	—	0.000
	Constrained Ridge	-0.431	0.756	0.920	0.145	0.184	0.000
	Unconstrained Z	0.006	0.630	0.946	0.072	0.676	0.000
	Unconstrained X	0.001	0.293	0.932	0.102	0.108	0.000
B: $\gamma_\kappa \times 0.5$	Standard SCM	-0.625	0.933	0.872	0.154	—	0.000
	Rescaling	-0.161	0.488	0.934	0.293	—	0.000
	Constrained Ridge	-0.509	0.790	0.908	0.152	0.246	0.000
	Unconstrained Z	0.041	0.672	0.948	0.071	0.510	0.000
	Unconstrained X	0.009	0.184	0.943	0.101	0.501	0.000
C: Sharp cutoff	Standard SCM	-0.311	0.647	0.931	0.150	—	0.000
	Rescaling	-0.082	0.440	0.946	0.277	—	0.000
	Constrained Ridge	-0.222	0.532	0.950	0.147	0.220	0.000
	Unconstrained Z	-0.009	0.391	0.958	0.073	0.487	0.000
	Unconstrained X	0.001	0.074	0.942	0.107	0.465	0.000
D: Non-monotone	Standard SCM	0.016	0.924	0.930	0.138	—	0.000
	Rescaling	0.039	0.432	0.929	0.267	—	0.000
	Constrained Ridge	0.034	0.735	0.918	0.136	0.214	0.000
	Unconstrained Z	0.017	0.755	0.946	0.072	0.508	0.000
	Unconstrained X	0.007	0.128	0.951	0.107	0.477	0.000

Notes:

Full diagnostics under Baseline (correct  $\gamma_\kappa$ ) plus four reach-calibration misspecification scenarios (A–D). Shown:  $\rho \in \{0.9\}$ .  $\tau = 4$ ,  $R = 1,000$  replications. Cov. = empirical 95% CI coverage rate (nominal 0.95). Pre-RMSE = mean pre-treatment RMSE of fitted synthetic control.  $\bar{\lambda}$  = mean selected regularization parameter; shown as — for Standard SCM and Rescaling (no  $\lambda$  penalty). Exp. Mass = mean fraction of SCM weight on exposed (neighboring) donors. See Table 22 for the remaining  $\rho$  values. <sup>†</sup>: cells with fewer than 900 valid replications (of 1,000).

Table 24: Mean Pre-Treatment RMSE ( $\tau = 4$ )

DGP	$\rho$	Standard	Rescaling	Constr. Ridge	Unconstr. Z	Unconstr. X	Trimming
Baseline	0.1	0.138	0.265	0.136	0.074	0.108	0.152
	0.3	0.138	0.265	0.136	0.074	0.108	0.152
	0.6	0.138	0.265	0.136	0.074	0.108	0.152
	0.9	0.138	0.265	0.136	0.074	0.108	0.152
	1.0	0.138	0.265	0.136	0.074	0.108	0.152
Heterogeneous	0.1	0.138	0.265	0.136	0.074	0.108	0.152
	0.3	0.138	0.265	0.136	0.074	0.108	0.152
	0.6	0.138	0.265	0.136	0.074	0.108	0.152
	0.9	0.138	0.265	0.136	0.074	0.108	0.152
	1.0	0.138	0.265	0.136	0.074	0.108	0.152
Factor Model	0.1	0.338	0.367	0.286	0.063	0.363	0.344
	0.3	0.338	0.367	0.286	0.063	0.363	0.344
	0.6	0.338	0.367	0.286	0.063	0.363	0.344
	0.9	0.338	0.367	0.286	0.063	0.363	0.344
	1.0	0.338	0.367	0.286	0.063	0.363	0.344

*Notes:*

Mean pre-treatment RMSE of the fitted synthetic control, averaged across  $R = 1,000$  replications.  $\tau = 4$ , quantile calibration for all corrections. Comparable pre-RMSE across estimators confirms that corrections preserve pre-period fit while reducing post-period bias. Pre-treatment RMSE is invariant to  $\rho$  because interference affects only the post-treatment period. †: cells with fewer than 900 valid replications (of 1,000).

Table 25: Mean Selected Regularization  $\lambda$  ( $\tau = 4$ )

DGP	$\rho$	Constr. Ridge		Unconstr. Z		Unconstr. X	
		Q	M	Q	M	Q	M
Baseline	0.1	0.224	0.228	0.521	0.520	0.453	0.447
	0.3	0.224	0.228	0.521	0.520	0.453	0.447
	0.6	0.224	0.228	0.521	0.520	0.453	0.447
	0.9	0.224	0.228	0.521	0.520	0.453	0.447
	1.0	0.224	0.228	0.521	0.520	0.453	0.447
Heterogeneous	0.1	0.224	0.228	0.521	0.520	0.453	0.447
	0.3	0.224	0.228	0.521	0.520	0.453	0.447
	0.6	0.224	0.228	0.521	0.520	0.453	0.447
	0.9	0.224	0.228	0.521	0.520	0.453	0.447
	1.0	0.224	0.228	0.521	0.520	0.453	0.447
Factor Model	0.1	0.362	0.365	0.280	0.286	0.367	0.337
	0.3	0.362	0.365	0.280	0.286	0.367	0.337
	0.6	0.362	0.365	0.280	0.286	0.367	0.337
	0.9	0.362	0.365	0.280	0.286	0.367	0.337
	1.0	0.362	0.365	0.280	0.286	0.367	0.337

*Notes:* Mean selected regularization

parameter  $\bar{\lambda}$  for ridge-penalized corrections at  $\tau = 4$ ,  $R = 1,000$  replications. Q = quantile calibration; M = moment calibration. Standard SCM and Rescaling have no  $\lambda$  penalty and are omitted. Values are invariant across  $\rho$  within each DGP because  $\lambda$  is selected on pre-treatment fit, which is unaffected by post-treatment interference strength. †: cells with fewer than 900 valid replications (of 1,000).

Table 26: Phase 3 Detection Test: Empirical Rejection Rates

DGP	$\rho$	$\tau = 0.5$	$\tau = 1.0$	$\tau = 4.0$	$\tau = 7.0$	
Baseline	0.1	0.187	0.591	0.822	0.826	
	0.3	0.757	0.810	0.826	0.825	
	0.6	0.810	0.826	0.825	0.826	
	0.9	0.824	0.828	0.826	0.825	
	1.0	0.825	0.828	0.826	0.825	
Heterogeneous	0.1	0.163	0.560	0.795	0.799	
	0.3	0.736	0.790	0.793	0.782	<i>Notes:</i> The ring-based permutation test rejects
	0.6	0.790	0.803	0.778	0.770	
	0.9	0.801	0.796	0.774	0.770	
	1.0	0.799	0.798	0.770	0.776	
Factor Model	0.1	0.007	0.016	0.139	0.281	
	0.3	0.026	0.095	0.503	0.682	
	0.6	0.095	0.235	0.720	0.792	
	0.9	0.162	0.363	0.789	0.810	
	1.0	0.188	0.424	0.792	0.813	

the sharp null  $H_0$ : no interference, i.e., treatment assignment is exchangeable across the treated-unit’s ring structure.

Entries are empirical rejection rates at nominal level  $\alpha = 0.05$ ; see Section 3 for the test construction.  $R = 1,000$  replications per  $(\tau, \rho, \text{DGP})$  cell. The Phase 3 DGP grid does not include  $\rho = 0$ , so rejection rates estimate power against weak-to-strong interference, not test size. See Table 4 for a size calibration on real panel geographies. †: cells with fewer than 900 valid replications (of 1,000).

it under the exposure-aligned penalty, downweighting but not deleting partial-exposure donors. None of this is offset by a guarantee on what survives: Theorem 4 delivers a deterministic, monotone envelope on the contamination bias for URX, but Trimming has no analogue. Its behavior depends entirely on the donor subset chosen, and there is no  $\lambda$  knob through which the researcher can shrink the residual bias.

The simulations bear this out, beginning with pre-treatment fit. Table 24 reports mean pre-treatment RMSE across estimators. In the Baseline and Heterogeneous DGPs, Standard SCM has pre-RMSE 0.138, Trimming 0.152 (a  $\approx 10\%$  degradation from the smaller donor pool), URX 0.108, and URZ 0.074. In the Factor Model DGP, Standard has 0.338, Trimming 0.344 (essentially tied), URX 0.363 (worse than Trimming on this metric), and URZ 0.063. URZ dominates pre-period fit everywhere but over-fits the pre-period, which translates to worse post-period variance under factor structure (Section 6). Trimming’s smaller donor pool costs roughly 10% pre-period fit in the additive DGPs and essentially nothing under factor structure — not a decisive argument against Trimming, but one objective cost it pays. The post-period picture sharpens the comparison. The Group A appendix tables report Trimming alongside the four feasible correctors and the Oracle. URX strictly dominates Trimming on post-period RMSE in the Baseline and Heterogeneous DGPs across every  $(\tau, \rho)$  cell — at the  $\tau = 4, \rho = 0.6$  flagship cell, URX RMSE is 0.070 (Baseline) and 0.068 (Heterogeneous) versus Trimming’s 0.222, a roughly threefold improvement. In the Factor Model DGP, URX and Trimming are approximately tied on post-period RMSE: at  $\tau = 4, \rho = 0.6$ , URX is 0.262 versus Trimming 0.259; at  $\rho = 0.9$ , URX is 0.272 versus Trimming 0.259. Trimming marginally wins on post-period RMSE in the latent-factor regime; we note this honestly. The other three correctors (Rescaling, Constrained Ridge, URZ) do not uniformly beat Trimming, which reinforces URX as the recommended default and positions the others as theoretically motivated stepping stones (simplex preservation, dominance results) rather than production methods.

Table 27: Ring Permutation: Degenerate Placebo Skip Rate (%)

State	$\rho$	$K_r = 2$	$K_r = 3$	$K_r = 4$
California	0.0	0.0	0.0	4.2
	0.1	0.0	0.0	4.2
	0.2	0.0	0.0	4.2
	0.3	0.0	0.0	4.2
	0.5	0.0	0.0	4.2
	0.7	0.0	0.0	4.2
	0.9	0.0	0.0	4.2
Minnesota	0.0	0.0	0.0	4.2
	0.1	0.0	0.0	4.2
	0.2	0.0	0.0	4.2
	0.3	0.0	0.0	4.2
	0.5	0.0	0.0	4.2
	0.7	0.0	0.0	4.2
	0.9	0.0	0.0	4.2
Missouri	0.0	0.0	0.0	—
	0.1	0.0	0.0	—
	0.2	0.0	0.0	—
	0.3	0.0	0.0	—
	0.5	0.0	0.0	—
	0.7	0.0	0.0	—
	0.9	0.0	0.0	—

*Notes:* Percentage of permutation draws skipped because a ring

contained only the treated state (degenerate configuration).  $K_r$  = number of distance rings. Maximum observed skip rate = 4.2% across all specifications; remaining draws produce valid F-statistics. See companion Table 4 for rejection rates. Cells shown as — correspond to  $K_r = 4$  in the Missouri (dense) graph, where the fourth ring is empty for the treated unit's neighborhood so the contrast is undefined. †: cells with fewer than 900 valid replications (of 1,000).

There remains a regime where Trimming is sensible. If the researcher has strong exogenous knowledge of which donors are exposed — a clear geographic barrier, an administrative boundary that exactly delineates the exposure set, or a known physical mechanism that fixes the contamination set — Trimming is a principled choice and avoids the implicit smoothing that the reach-based corrections impose. In the absence of such knowledge, URX is preferred. One feature of the tables deserves comment in closing: Trimming’s column is constant in  $(\tau, \rho)$  within a DGP, because the dropped donors are the unique conduit of  $(\tau, \rho)$  into the estimator, and dropping them severs the channel. The across-DGP differences that do appear (Bias  $\approx 0.010$  and RMSE  $\approx 0.222$  in Baseline and Heterogeneous; Bias  $\approx 0.012$  and RMSE  $\approx 0.259$  in Factor Model) reflect noise structure, not spillover sensitivity.