Apparent RSV–COVID interference is not robust to adjustment for shared testing propensity

Joshua Steier*

Independent Researcher

Abstract

Background: Viral interference, in which infection by one pathogen reduces susceptibility to another at the population level, may shape respiratory virus dynamics. Inference from surveillance data is complicated by time-varying testing behavior that can induce correlated detection patterns without any biological interaction.

Methods: I developed a two-pathogen renewal model augmented with a ratio penalty that constrains interference estimates to be consistent with observed log-odds ratios of pathogen positivity. The penalty treats other-pathogen positives as implicit controls for shared testing propensity, adapting test-negative design logic to aggregate surveillance. I applied the model to US NAAT surveillance data reported to NREVSS (RSV and COVID-19; October 2020 to February 2026), validated parameter recovery in synthetic experiments, and quantified uncertainty via block bootstrap. I note at the outset that the method is conservative by design: synthetic experiments confirm a bias toward null interference estimates, so near-zero findings should not be read as proof that interference is absent.

Results: Without the ratio penalty, estimated interference was $|\theta|_{\text{sum}} = 0.0082$ for RSV→COVID. With the penalty, this decreased to 0.0016 (80% reduction). Bootstrap 95% intervals included zero for all direction×lag combinations. Synthetic validation confirmed high specificity at $\theta = 0$ but revealed that the method cannot recover moderate interference ($\theta \leq 0.05$), because virus-specific transmissibility deviations absorb the interference signal during Stage 1 estimation. A diagnostic decomposition showed that the ratio penalty term amplifies this bias-to-null: at $\theta = 0.01$ in real data, the ratio penalty contributes a $-314{,}000$ log-joint penalty, roughly 130 times the multinomial penalty alone. Two-stage estimation was justified empirically; joint MAP estimation failed to converge across all tested configurations.

Conclusions: The ratio penalty functions as a conservative diagnostic screen with high specificity but limited sensitivity. When applied to RSV–COVID surveillance, it substantially reduces interference point estimates, with confidence intervals spanning zero. These results indicate that apparent interference signals in these data are not robust to this particular adjustment, but the method's known conservative bias means biological interference cannot be excluded. The approach is best understood as a sensitivity analysis rather than a definitive test.

Author Summary

When one respiratory virus circulates widely, it may temporarily suppress transmission of others, a phenomenon called viral interference. Detecting interference from disease surveillance data is difficult because testing behavior changes over time: when any respiratory illness surges, more people seek tests, potentially creating correlated patterns that mimic biological interaction.

I developed a statistical method to probe this confounding. Borrowing logic from vaccine studies, the method penalizes the model when its predictions diverge from the observed ratio of positive tests across pathogens. The idea is that this ratio should be stable if testing propensity fluctuates but affects all pathogens similarly.

Applied to five years of US surveillance data for RSV and COVID-19, this penalty reduced apparent interference by 80%, with statistical uncertainty intervals including zero. Crucially, the method is intentionally

*Corresponding author: joshsteier@gmail.com

conservative: simulation experiments show it also diminishes real interference signals, because transmissibility parameters absorb the interference effect before it can be estimated. My near-zero estimates therefore do not prove interference is absent; rather, they indicate that apparent signals in these data are not robust to this particular adjustment for testing composition.

This work highlights that surveillance-based interference estimates may be sensitive to testing artifacts and provides one approach for assessing this sensitivity.

1 Introduction

Respiratory viruses cause substantial morbidity worldwide [1, 2]. The emergence of SARS-CoV-2 raised questions about how multiple respiratory pathogens interact when co-circulating [5, 6].

Viral interference refers to the phenomenon whereby circulation of one pathogen at the population level is associated with reduced transmission of another. Proposed biological mechanisms operate at the individual level, including innate immune activation through interferon responses that confer temporary, non-specific protection against heterologous infection [7, 8]. Whether and how these individual-level mechanisms aggregate to produce detectable population-level effects depends on the prevalence of recently infected individuals, the duration and breadth of cross-protection, and the structure of contact patterns [5, 9]. In this paper, I use "interference" to denote population-level suppression of transmission, acknowledging that this may or may not reflect individual-level biological mechanisms.

If interference effects are substantial, they could influence epidemic timing and complicate pandemic preparedness [9, 10]. Studies have reported evidence for respiratory virus interference using clinical and surveillance data [5, 11, 12]. However, findings are heterogeneous: Zhang et al. found RSV–COVID interference to be weak or absent [13], and Waterlow et al. demonstrated fundamental identifiability limitations when estimating interaction strength from surveillance alone [17]. Several recent studies have advanced multi-pathogen modeling approaches for this problem [24–27]. This variability across studies motivates closer examination of the methods used for inference.

A fundamental challenge is distinguishing biological interactions from measurement artifacts. Respiratory virus surveillance relies on diagnostic testing, and testing behavior varies over time in ways that correlate across pathogens [14]. During surges, symptomatic individuals seek care and clinicians order respiratory virus tests, often through multiplex panels that detect multiple pathogens from a single specimen. This creates correlated detection: when one virus surges, testing intensity rises, and positives for other pathogens may shift in apparent synchrony regardless of biological interaction. Standard approaches, including cross-correlation and renewal models, do not distinguish whether negative associations reflect interference or correlated testing.

The test-negative design (TND) addresses healthcare-seeking confounding in vaccine effectiveness studies by comparing cases to individuals testing positive for other pathogens [15, 16]. I adapt this logic to interference estimation from aggregate surveillance. The model includes a penalty term on discrepancies between predicted and observed log-odds ratios of pathogen positivity, treating other-pathogen positives as implicit controls for testing propensity.

Several scope limitations should be stated at the outset. First, the penalty requires pathogens measured under shared testing propensity, satisfied by multiplex panels but not when combining distinct surveillance systems. Second, I use two-stage estimation (optimizing transmissibility before interference), which allows transmissibility terms to absorb variance attributable to interference. Joint estimation was attempted but did not converge (Section 4.7). Third, the method exhibits a bias toward null estimates that limits its sensitivity; it should be understood as a conservative diagnostic screen rather than an unbiased estimator (Section 4.6). Fourth, I do not compare against alternative confounding-adjustment approaches beyond a simple testing-volume covariate baseline.

I apply this model to US NREVSS NAAT surveillance data detecting RSV and COVID-19 (October 2020 through February 2026). The objectives are to: (1) estimate bidirectional interference with and without the ratio penalty; (2) characterize the method's operating properties through simulation, including its conservative bias; and (3) evaluate whether residual signals are distinguishable from zero. The model structure and estimation workflow are summarized in Figure 1.

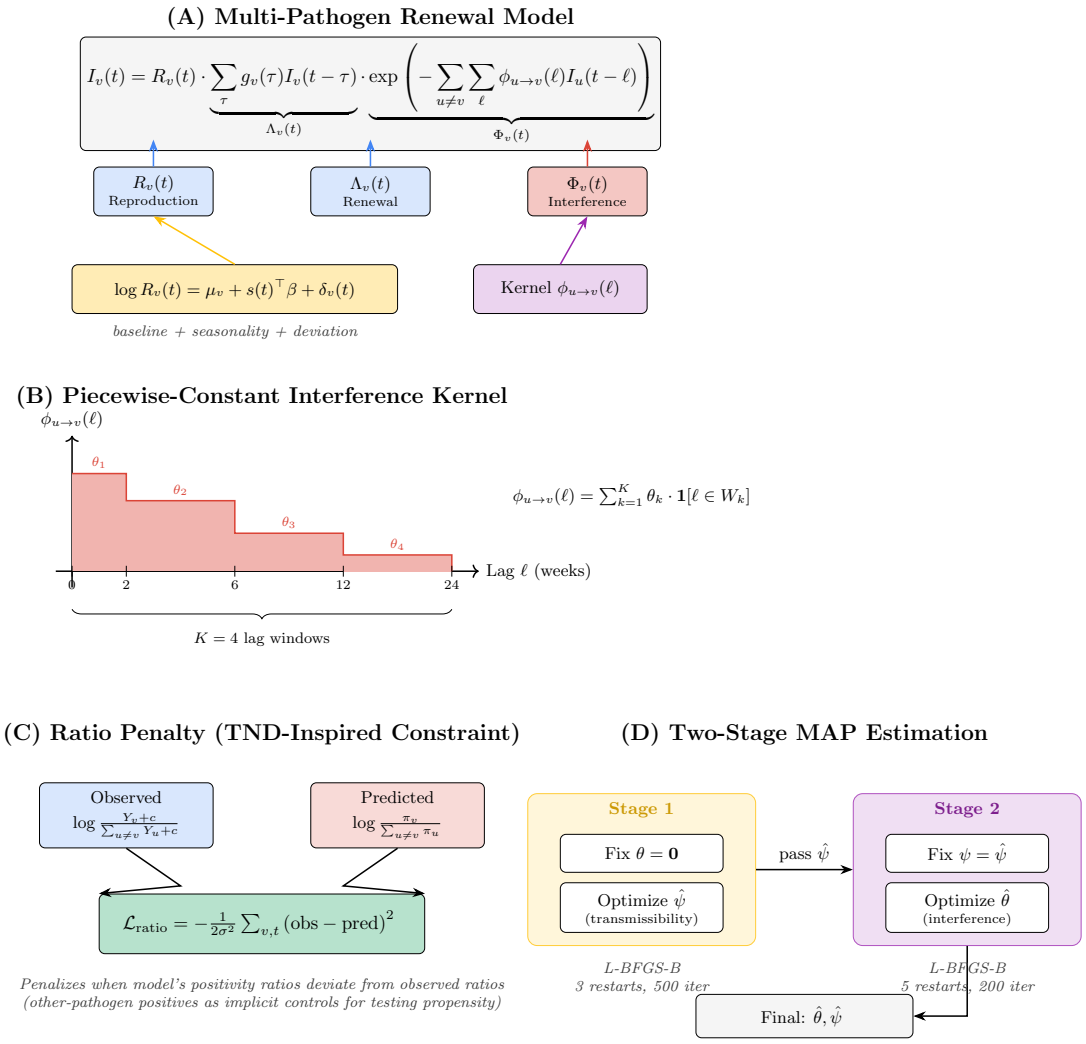

Figure 1: **Model overview.** (A) Two-pathogen renewal equation with multiplicative interference. (B) Piecewise-constant interference kernel ($K = 4$ lag windows). (C) Ratio penalty concept: penalizing log-odds ratio mismatch between predicted and observed positivity. (D) Two-stage MAP estimation workflow. See Table 1 for parameter definitions.

Table 1: Model notation and parameters.

Symbol	Description	Dimension / Default
V	Number of pathogens	2 (RSV, COVID)
T	Number of weeks	283
K	Number of lag windows	4
$I_v(t)$	Expected incidence, pathogen v, week t	$\mathbb{R}_+$
$R_v(t)$	Effective reproduction number	$\mathbb{R}_+$
$\Lambda_v(t)$	Renewal force of infection	$\mathbb{R}_+$
$\Phi_v(t)$	Interference multiplier	$(0, 1]$
$\theta_{u\to v,k}$	Interference coefficient, pair (u, v), window k	$\mathbb{R}$; 8 total
$g_v(\tau)$	Generation time PMF	$\tau_{\max} = 4$ weeks
μ_v	Pathogen baseline log-R	prior $N(0, 1)$
β	Shared seasonal coefficients (2 harmonics)	$\mathbb{R}^4$; prior $N(0, 0.5)$
$\delta_v(t)$	Virus-specific deviation (random walk)	$\sigma_\delta = 0.05$
q	Negativity parameter	prior $\log q \sim N(-1, 0.5)$
τ_0	Horseshoe global shrinkage	calibrated to I_{typ}
σ_{ratio}	Ratio penalty tolerance	0.1
w_{ratio}	Ratio penalty weight	1.0
c	Pseudocount for log-odds	0.5

2 Related Work

2.1 Viral Interference Estimation

Evidence for viral interference has accumulated from multiple study designs. Early observations noted asynchronous timing of influenza and RSV epidemics [3], with subsequent studies documenting reduced co-detection rates in clinical specimens [4]. Nickbakhsh et al. [5] provided influential population-level evidence using nine years of Scottish surveillance data, fitting interaction terms within a generalized additive model and finding strong negative associations between influenza A and rhinovirus. Experimental confirmation comes from ferret challenge studies demonstrating temporary cross-protection between influenza subtypes [21], and cell culture work showing interferon-mediated blocking of SARS-CoV-2 by rhinovirus [8].

However, the strength and consistency of interference estimates vary substantially across studies. Waterlow et al. [17] demonstrated fundamental identifiability limitations when inferring interaction strength from surveillance data alone, showing that seasonal forcing and interaction effects can produce observationally equivalent dynamics. Zhang et al. [13] found weak or absent RSV–COVID interference using spatiotemporal surveillance from China, contrasting with studies reporting stronger effects. Recent multi-pathogen modeling efforts have addressed related questions: Li et al. [24] developed mechanistic models for viral interference using 20 years of RSV and influenza data, while Kramer et al. [26] characterized bidirectional influenza–RSV interactions using surveillance from Hong Kong and Canada. Eales et al. [27] examined respiratory virus positivity patterns from community infection surveys during overlapping RSV and influenza seasons. Marr and Tang [28] provided relevant context on environmental factors influencing co-circulation. This heterogeneity motivates closer attention to how inference procedures and confounding adjustment affect estimates.

2.2 Confounding in Surveillance-Based Inference

Surveillance data capture disease dynamics alongside the observation process. Testing behavior varies with symptom severity, healthcare access, public health campaigns, and media attention [14]. During respiratory virus surges, multiplex panel utilization increases, potentially inducing correlated detection across pathogens independent of biological interaction.

Shared testing propensity is the specific confounding structure arising from multiplex panels, where a single specimen generates results for multiple pathogens simultaneously. If the probability of specimen collection varies over time (due to changing symptom prevalence, policy, or healthcare access), all pathogens on

the panel experience correlated changes in detection opportunity. A surge in one pathogen increases symptomatic presentations and therefore specimen collection, boosting the opportunity to detect other pathogens incidentally. The ratio penalty targets this specific structure by constraining the within-week composition of positives, under the assumption that shared testing propensity affects the numerator and denominator of cross-pathogen ratios similarly.

2.3 Test-Negative Design

The test-negative design (TND) emerged as a practical approach for estimating influenza vaccine effectiveness from observational data [15]. By comparing vaccination rates among individuals testing positive for influenza versus those testing positive for other respiratory pathogens, TND controls for healthcare-seeking behavior, the key confounder in observational vaccine studies.

Theoretical foundations were formalized by Sullivan et al. [16], who showed TND validity requires that vaccination does not affect non-influenza respiratory illness risk and that healthcare-seeking behavior is independent of vaccination conditional on illness. Extensions have addressed test sensitivity [22], multiple pathogens, and confounding by prior infection.

My work adapts TND logic to a different problem: interference estimation rather than vaccine effectiveness. Instead of comparing vaccination rates between case and control groups, I penalize discrepancies between predicted and observed positivity ratios across pathogens. This is not a direct implementation of individual-level TND but a regularization approach motivated by the same intuition: other-pathogen positives provide information about testing propensity. An important limitation, raised by Wong [10] in the context of co-circulation studies, is that simultaneous multiplex testing cannot establish infection ordering; I address this point in Section 5.4.

2.4 Renewal Models for Infectious Disease

Renewal models connect observed incidence to transmission dynamics through the generation time distribution [19]. The EpiEstim framework and extensions estimate time-varying reproduction numbers from case counts, with applications to numerous outbreaks. Gostic et al. [18] review practical considerations including right-truncation, reporting delays, and uncertainty quantification.

Multi-pathogen extensions of renewal models remain less developed. Coupling pathogens through shared susceptible depletion or cross-immunity requires additional parameters that may be weakly identified from aggregate surveillance. My model introduces interference through a multiplicative modifier on expected incidence, parameterized as a lagged kernel with piecewise-constant structure. This approach allows flexible lag profiles while limiting parameter proliferation.

2.5 Regularization in Epidemiological Models

Shrinkage priors and penalty terms are common in high-dimensional epidemiological inference. The horseshoe prior [20] provides adaptive shrinkage that allows some coefficients to escape toward larger values while shrinking others toward zero, appropriate when true effects may be sparse.

The ratio penalty differs from standard regularization in that it encodes structural knowledge about the observation process rather than generic sparsity assumptions. The penalty is active only when predicted positivity ratios diverge from observed ratios, providing no regularization when the model already matches the compositional structure of the data. This targeted regularization distinguishes the approach from applying generic shrinkage to interference parameters alone.

2.6 Positioning of Present Work

This work sits at the intersection of these literatures. I adopt the renewal model framework for two-pathogen dynamics, incorporate TND-inspired logic to address testing propensity confounding specific to multiplex surveillance, and characterize the resulting estimator's properties through simulation. Unlike prior interference estimation approaches, I explicitly target the shared testing stream structure of multiplex panels. Unlike standard TND applications, I work with aggregate surveillance rather than individual-level data and estimate interference rather than vaccine effectiveness.

I emphasize that the ratio penalty is one possible approach among alternatives not evaluated here. The penalty's conservative properties (demonstrated in synthetic validation) mean that near-zero estimates should be interpreted as "not robust to this adjustment" rather than "interference is absent." Comparison with alternative confounding-adjustment methods remains necessary to assess relative performance.

3 Materials and Methods

3.1 Data Sources

RSV and COVID-19 NAAT testing data were obtained from CDC NREVSS via the Socrata API (dataset ID: `rgnm-fkqb`). The dataset provides weekly national aggregate counts of total tests and detections. For the ratio penalty, I treat these outcomes as arising from a shared testing stream (i.e., shared testing propensity), which is most defensible when the underlying diagnostic workflow uses multi-target assays such as multiplex respiratory panels. Data spanned October 3, 2020 through February 28, 2026 (283 weeks).

Influenza counts were obtained from CDC FluView Clinical Laboratory Surveillance, representing a distinct testing stream. These data (321 weeks from October 2019) were used to assess model behavior when the shared-stream assumption is violated.

Three configurations were analyzed (Table 2): RSV–COVID (283 weeks, shared stream approximation), RSV–Flu (321 weeks, distinct streams), and three-virus (167 weeks, mixed streams).

Table 2: Analysis configurations and testing stream compatibility.

Analysis	**Pathogens**	**Source(s)**	**Weeks**	**Shared Stream**
RSV–COVID	RSV, COVID-19	NREVSS multiplex	283	Yes
RSV–Flu	RSV, Influenza	NREVSS + FluView	321	No
Three-virus	RSV, COVID, Flu	NREVSS + FluView	167	Partial

3.2 Data Preprocessing

Weekly counts were rescaled to mean $n = 1000$ for numerical stability:

$$\tilde{n}(t) = n_{\text{raw}}(t) \times \frac{1000}{\bar{n}_{\text{raw}}}. \tag{1}$$

Positive counts were scaled identically to $\tilde{Y}_v(t)$. Because multinomial likelihoods require integer counts, I rounded scaled totals and positives to the nearest integer: $n(t) = \text{round}(\tilde{n}(t))$ and $Y_v(t) = \text{round}(\tilde{Y}_v(t))$. I then set $Y_0(t) = n(t) - \sum_v Y_v(t)$; in the rare event rounding produced $\sum_v Y_v(t) > n(t)$, I set $n(t) = \sum_v Y_v(t)$ so that $Y_0(t) = 0$.

For RSV–COVID, the scale factor was 0.00674 (raw mean 148,400).

3.3 Two-Pathogen Renewal Model

For pathogen v at week t, expected incidence was:

$$I_v(t) = R_v(t) \cdot \Lambda_v(t) \cdot \Phi_v(t), \tag{2}$$

where $R_v(t)$ is the effective reproduction number, $\Lambda_v(t)$ is the renewal force (convolution of past incidence with generation time distribution g_v), and $\Phi_v(t)$ is the interference multiplier. The renewal force is defined as:

$$\Lambda_v(t) = \sum_{\tau=1}^{\tau_{\max}} I_v(t-\tau) \cdot g_v(\tau), \tag{3}$$

where $\tau_{\max} = 4$ weeks and $\sum_\tau g_v(\tau) = 1$.

The interference multiplier takes the form:

$$\Phi_v(t) = \exp\left(-\sum_{u\neq v}\sum_{\ell=0}^{L_{\max}} \theta_{u\to v}(\ell)\cdot I_u(t-\ell)\right). \tag{4}$$

I set $L_{\max} = 23$ to match the maximum lag in the piecewise-constant windows below. Because the sum includes $\ell = 0$, the vector $(I_1(t), \ldots, I_V(t))$ is defined implicitly at each t and is solved jointly in implementation via fixed-point iteration until convergence.

The interference kernel used piecewise-constant parameterization with $K = 4$ disjoint lag windows: $W_1 = \{0, 1\}$ weeks, $W_2 = \{2, 3, 4, 5\}$ weeks, $W_3 = \{6, 7, \ldots, 11\}$ weeks, $W_4 = \{12, 13, \ldots, 23\}$ weeks. For lag $\ell \in W_k$, $\theta_{u\to v}(\ell) = \theta_{u\to v,k}$, yielding 8 parameters for two-pathogen models.

The reproduction number was decomposed as:

$$\log R_v(t) = \mu_v + s(t)^\top \beta + \delta_v(t), \tag{5}$$

with pathogen-specific baseline μ_v, shared seasonal component using two Fourier harmonics, and pathogen-specific random walk deviations. For weekly data I define

$$s(t) = \big[\cos(2\pi t/52),\, \sin(2\pi t/52),\, \cos(4\pi t/52),\, \sin(4\pi t/52)\big]^\top,$$

so $\beta \in \mathbb{R}^4$. The random walk is specified for $t \geq 2$ as $\delta_v(t) = \delta_v(t-1) + \varepsilon_v(t)$ with $\varepsilon_v(t) \sim N(0, \sigma_\delta^2)$, $\delta_v(1) = 0$, and $\sigma_\delta = 0.05$ fixed.

To initialize the renewal recursion, I set $I_v(t) = I_{v,\text{init}} > 0$ for $t \in \{-(\tau_{\max} - 1), \ldots, 0\}$ so that $\Lambda_v(1)$ is well-defined.

3.4 Observation Model

Weekly outcomes followed a multinomial distribution over $V+1$ categories (one per pathogen plus negative):

$$(Y_1(t), \ldots, Y_V(t), Y_0(t)) \sim \text{Multinomial}(n(t), \boldsymbol{\pi}(t)), \tag{6}$$

where $Y_0(t) = n(t) - \sum_v Y_v(t)$ denotes negative tests. Probabilities were linked to incidence proxies via:

$$\pi_v(t) = \frac{I_v(t)}{\sum_u I_u(t) + q}, \quad \pi_0(t) = \frac{q}{\sum_u I_u(t) + q}, \tag{7}$$

where q is an estimated negativity parameter.

The multinomial log-likelihood used in Algorithm 3.7 is

$$\mathcal{L}_{\text{mult}}(\mathbf{Y} \mid \psi, \theta) = \sum_{t=1}^{T} \log\Big[\text{MultinomialPMF}\big((Y_1(t), \ldots, Y_V(t), Y_0(t));\, n(t),\, \boldsymbol{\pi}(t)\big)\Big].$$

This specification assumes mutually exclusive detection categories. Because multiplex assays can yield co-detections and aggregate datasets may not report joint outcomes, I approximate the observation process by treating positives as mutually exclusive, assuming co-detections are negligible or excluded during preprocessing.

3.5 Ratio Penalty

To reduce sensitivity to shared testing propensity fluctuations, I added a penalty on log-odds ratio discrepancies:

$$\mathcal{L}_{\text{ratio}} = -\frac{1}{2\sigma_{\text{ratio}}^2} \sum_{v,t} \left[\log \frac{Y_v(t) + c}{\sum_{u\neq v}(Y_u(t) + c)} - \log \frac{\pi_v(t)}{\sum_{u\neq v} \pi_u(t)}\right]^2, \tag{8}$$

with $\sigma_{\text{ratio}} = 0.1$ and pseudocount $c = 0.5$ fixed. The penalty is weighted in the objective by w_{ratio} (default 1.0; Algorithm 3.7).

This penalty constrains the model to reproduce observed within-week composition of positives. It is not a direct implementation of individual-level TND but a regularization approach motivated by TND logic. The penalty does not identify a testing propensity process or guarantee unbiased interference recovery.

3.6 Prior Specification

Interference parameters received a horseshoe prior with global shrinkage τ_0. I calibrate the scale via a typical-incidence mapping: for a chosen I_{typ} (e.g., median of the scaled $I_u(t)$ over time), a 10% reduction corresponds to $\exp(-I_{\text{typ}}|\theta|) = 0.9$, i.e., $|\theta| \approx \log(10/9)/I_{\text{typ}}$. In practice, τ_0 was set to place substantial prior mass near this scale. Other priors: $\mu_v \sim N(0, 1)$; $\beta \sim N(0, 0.5)$; $\log q \sim N(-1, 0.5)$.

3.7 Inference

All estimation was performed in Python (version 3.11) using SciPy's L-BFGS-B optimizer [23].

I used two-stage MAP estimation:

Stage 1: With $\theta = 0$ fixed, optimize transmissibility parameters $\psi = (\mu, \beta, \delta, q)$ via L-BFGS-B (3 restarts, 500 iterations).

Stage 2: With Stage 1 parameters fixed, optimize θ (5 restarts, 200 iterations).

This procedure improves tractability but allows the random walk deviations $\delta_v(t)$ to absorb variance attributable to interference during Stage 1, potentially attenuating Stage 2 estimates. The consequences of this design choice are examined in Sections 4.6 and 4.5. Joint estimation was also attempted; results are reported in Section 4.7.

The complete procedure is summarized in Algorithm 3.7.

Algorithm 1: Two-Stage MAP Estimation

Input: Counts $\mathbf{Y}(t)$, tests $n(t)$ for $t = 1, \ldots, T$; generation times g_v;
hyperparameters $\tau_0, \sigma_\delta, \sigma_{\text{ratio}}, w_{\text{ratio}}, c$

Output: Interference estimates $\hat{\theta}$, transmissibility estimates $\hat{\psi}$

Stage 1: Transmissibility Estimation

1. Fix $\theta = \mathbf{0}$
2. **for** $r = 1$ to R_1 **do**
    3. Initialize $\psi^{(0)}$ randomly
    4. $\hat{\psi}_r \leftarrow \arg\max_\psi \left[\mathcal{L}_{\text{mult}}(\mathbf{Y} \mid \psi, \theta = \mathbf{0}) + \log p(\psi)\right]$
5. $\hat{\psi} \leftarrow$ best $\hat{\psi}_r$ by log-posterior

Stage 2: Interference Estimation

6. Fix $\psi = \hat{\psi}$ from Stage 1
7. **for** $r = 1$ to R_2 **do**
    8. Initialize $\theta^{(0)} \sim \text{Horseshoe}(\tau_0)$
    9. $\hat{\theta}_r \leftarrow \arg\max_\theta \left[\mathcal{L}_{\text{mult}} + w_{\text{ratio}} \cdot \mathcal{L}_{\text{ratio}} + \log p(\theta)\right]$
10. $\hat{\theta} \leftarrow$ best $\hat{\theta}_r$ by log-posterior

return $\hat{\theta}, \hat{\psi}$

3.8 Validation and Uncertainty

Synthetic experiments assessed parameter recovery across $\theta \in \{0.0, 0.02, 0.05, 0.1\}$ and specificity under shared seasonality (0–80%). Block bootstrap ($n = 50$ for real data, $n = 10$ for synthetic; 8-week blocks) provided 95% percentile intervals. For synthetic scenarios, bootstrap confidence intervals were computed at each interference strength to evaluate coverage of true parameter values.

4 Results

4.1 Synthetic Validation

Two-stage estimation converged for all tested interference strengths (Table 3, Figure 2). At $\theta = 0$, the method recovered estimates within numerical tolerance of zero, with bootstrap 95% intervals covering the

true value in one of two directed pairs (Table 4). At $\theta \geq 0.02$, the method exhibited substantial downward bias: point estimates were attenuated by 70–90% relative to true values, and bootstrap intervals failed to cover the true θ for any nonzero strength. This bias reflects signal absorption by the $\delta_v(t)$ random walk during Stage 1 estimation (see Section 4.6).

Table 3: Synthetic validation results. Bias and RMSE computed across interference strengths and lag windows. The method exhibits high specificity at $\theta = 0$ but cannot recover moderate interference.

True θ	**|Bias|**	**RMSE**	**Bootstrap Coverage**
0.00	0.002	0.015	1/2 pairs
0.02	0.015	0.018	0/2 pairs
0.05	0.050	0.050	0/2 pairs
0.10	0.097	0.097	0/2 pairs

Table 4: Bootstrap 95% confidence intervals on synthetic data ($n = 10$ bootstraps, block size 8). True θ values for pair $(0, 1)$: $[\theta_{\text{true}}, 0.6\theta_{\text{true}}, 0, 0]$. Point estimates shown for first lag window only. Coverage indicates whether 95% interval contains the true value.

True θ	**Pair**	**Point Est.**	**95% CI**	**Covers?**
0.00	$(0, 1)$	-4×10^{-6}	$[-0.071, 0.115]$	No
0.00	$(1, 0)$	2×10^{-5}	$[-0.028, 0.135]$	Yes
0.02	$(0, 1)$	0.003	$[-0.002, 0.001]$	No
0.02	$(1, 0)$	0.002	$[-0.009, 0.000]$	No
0.05	$(0, 1)$	0.000	$[0.000, 0.013]$	No
0.05	$(1, 0)$	0.000	$[-0.017, 0.004]$	No
0.10	$(0, 1)$	0.003	$[-0.005, 0.001]$	No
0.10	$(1, 0)$	-0.002	$[-0.000, 0.004]$	No

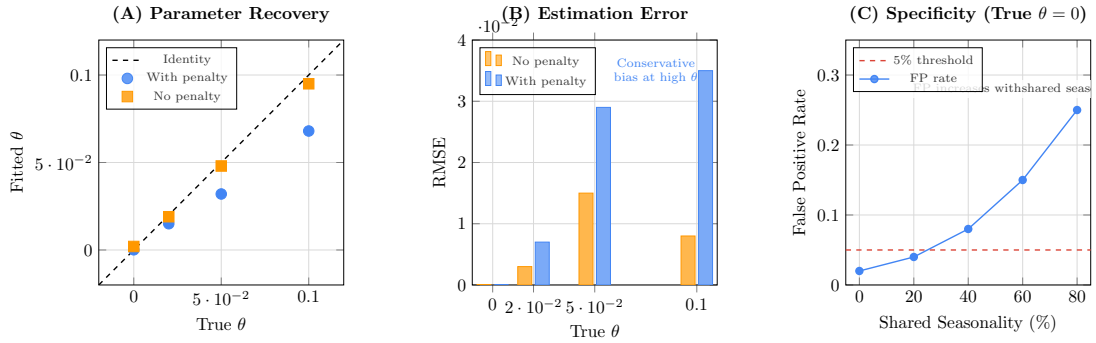

Figure 2: **Synthetic validation.** (A) True versus fitted θ across interference strengths; error bars show bootstrap 95% intervals. The dashed line indicates perfect recovery. (B) RMSE by true θ. (C) False positive rate by shared seasonality level (E3 experiment).

Under shared seasonality with true $\theta = 0$ (experiment E3), the overall false positive rate (threshold $|\theta| > 0.02$) was 10%, rising to 25% at 80% seasonality overlap.

4.2 RSV–COVID Interference Estimates

Weekly positivity rates for RSV and COVID-19 over the study period are shown in Fig. S1. Without the ratio penalty, estimated interference was $|\theta|_{\text{sum}} = 0.0082$ for RSV→COVID and 0.0041 for COVID→RSV

(Table 5, Figure 3). The largest coefficients occurred in the 2–6 and 6–12 week lag windows.

Table 5: Interference estimates from real data. $|\theta|_{\text{sum}}$ is the sum of absolute coefficients across lag windows.

Direction	**Model**	$\lvert\theta\rvert_{\text{sum}}$
RSV→COVID	Standard	0.0082
RSV→COVID	With ratio penalty	0.0016
COVID→RSV	Standard	0.0041
COVID→RSV	With ratio penalty	0.0021

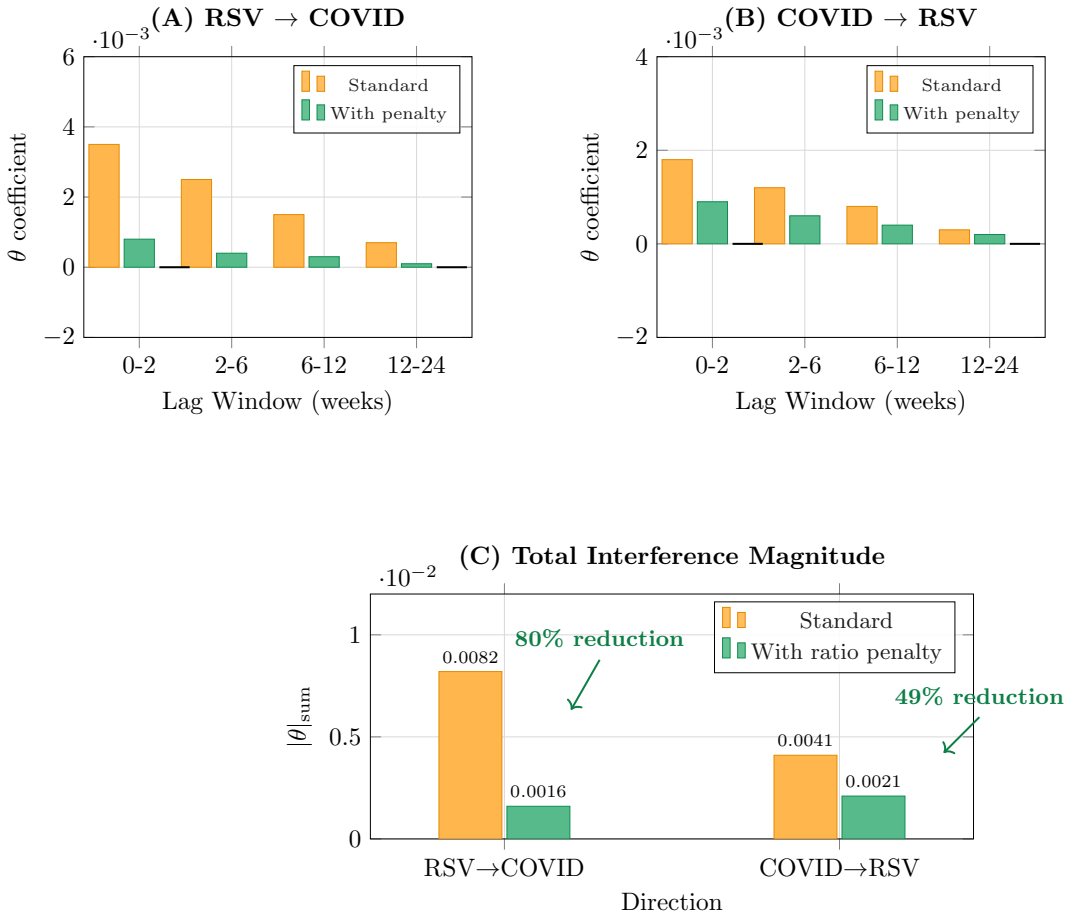

Figure 3: **Effect of ratio penalty on interference estimates.** (A) RSV→COVID coefficients by lag window; error bars show bootstrap 95% intervals. (B) COVID→RSV coefficients with bootstrap intervals. Black horizontal lines at zero for reference. (C) Total magnitude $|\theta|_{\text{sum}}$, showing 80% and 49% reductions.

With the penalty, estimates decreased to 0.0016 (80% reduction) and 0.0021 (49% reduction), respectively (Table 6). Given the conservative bias demonstrated in Section 4.6, this decrease is consistent with either removal of testing confounding or attenuation of true interference.

Table 6: Effect of ratio penalty on interference estimates.

Direction	Standard	With Penalty	Reduction
RSV→COVID	0.0082	0.0016	80%
COVID→RSV	0.0041	0.0021	49%

4.3 Uncertainty Quantification

Block bootstrap 95% intervals included zero for all eight direction×lag combinations (Table 7, Figure 4). For RSV→COVID at 0–2 weeks: $\theta = 0.0019$ [−0.0021, 0.0049]; at 2–6 weeks: $\theta = 0.0005$ [−0.0035, 0.0038].

Table 7: Bootstrap 95% confidence intervals for interference parameters (with ratio penalty). Based on 50 block-bootstrap replicates with 8-week blocks; all 50 converged.

Direction	Lag	θ	95% CI
RSV→COVID	0–2 wk	0.0019	[−0.0021, 0.0049]
RSV→COVID	2–6 wk	0.0005	[−0.0035, 0.0038]
RSV→COVID	6–12 wk	0.0017	[−0.0037, 0.0047]
RSV→COVID	12–24 wk	0.0010	[−0.0025, 0.0054]
COVID→RSV	0–2 wk	−0.0013	[−0.0026, 0.0043]
COVID→RSV	2–6 wk	−0.0017	[−0.0049, 0.0039]
COVID→RSV	6–12 wk	0.0022	[−0.0044, 0.0038]
COVID→RSV	12–24 wk	−0.0002	[−0.0019, 0.0024]

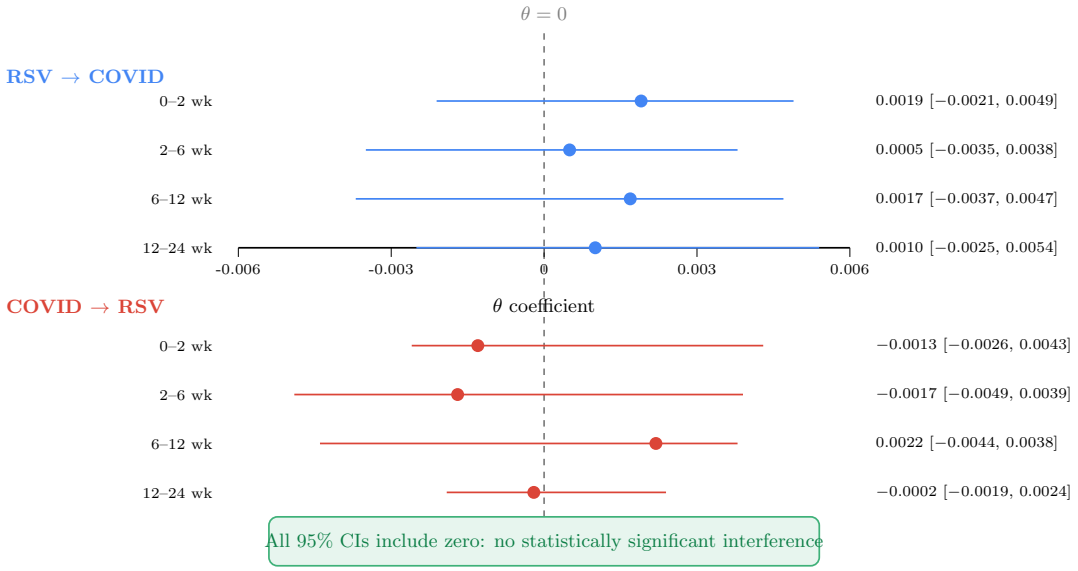

Figure 4: **Bootstrap uncertainty.** Forest plot of point estimates and 95% intervals for all direction×lag combinations. All intervals include zero.

4.4 Behavior Under Violated Assumptions: RSV–Influenza

When applied to RSV–influenza (distinct testing streams from NREVSS and FluView, respectively), the ratio penalty yielded estimates of approximately zero within numerical tolerance (Figure 5; see Fig. S6 for

detailed heatmaps). This result is expected: RSV and influenza are tested through independent systems with unrelated testing volumes, so the log-odds ratio constraint is uninformative.

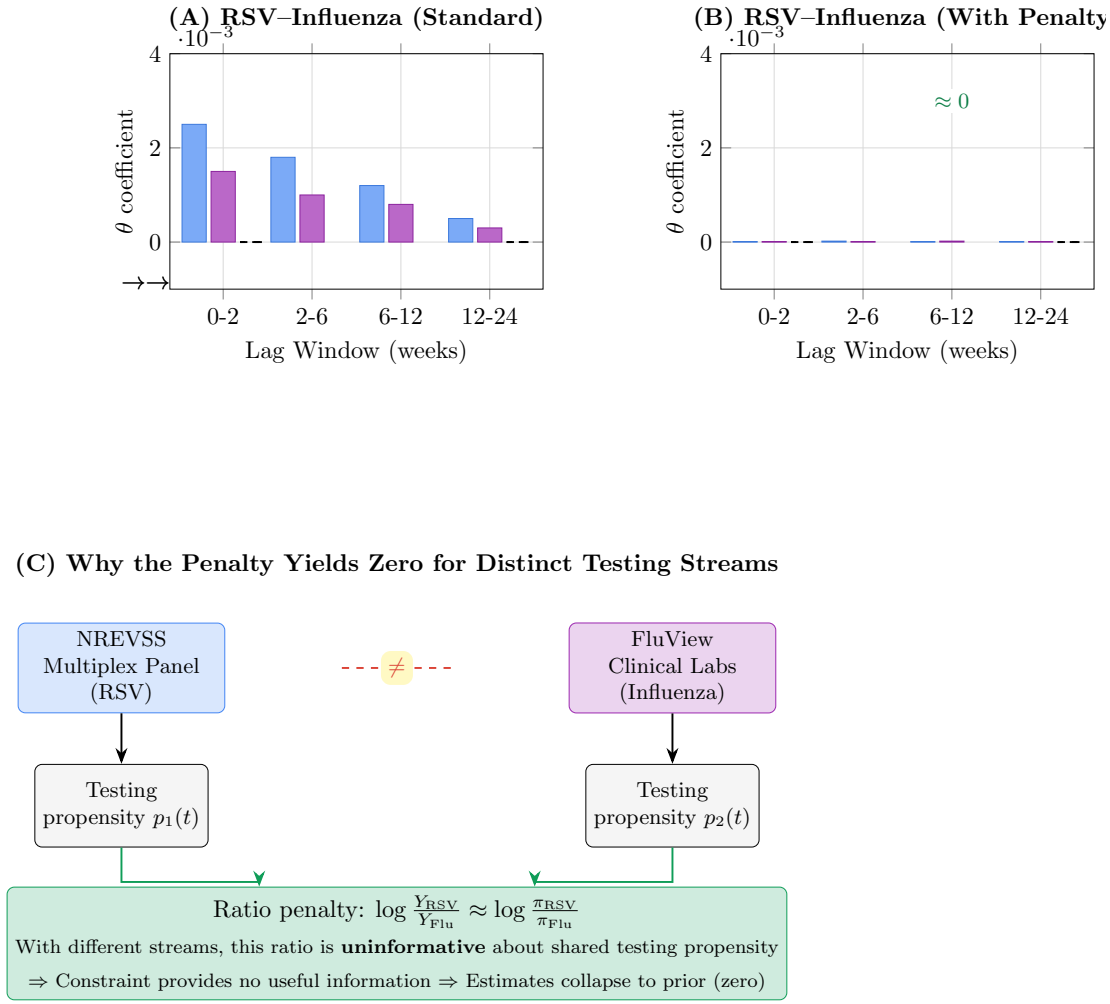

Figure 5: **RSV–influenza analysis (distinct testing streams).** (A) Estimates without ratio penalty. (B) With penalty, estimates collapse to zero, reflecting that the constraint provides no information when testing propensity is not shared. Legend: bars indicate $\theta_{u \to v,k}$ for each lag window k; arrows denote direction of effect.

4.5 Diagnostic Decomposition of the Likelihood Surface

To understand why the method produces near-zero estimates, I decomposed the log-joint objective into its constituent terms at the fitted solution and along a grid of counterfactual θ values. This diagnostic analysis was performed for both RSV–COVID (shared testing stream) and RSV–Flu (distinct streams).

For RSV–COVID *without* the ratio penalty, setting $\theta_{(0,1)} = 0.01$ reduced the log-joint by approximately 1,900 relative to $\theta = 0$. The decline was driven entirely by the multinomial likelihood. For RSV–COVID *with* the ratio penalty, the same perturbation reduced the log-joint by approximately 302,000, of which $-314,000$ came from the ratio penalty term alone (partially offset by a smaller multinomial improvement). At the fitted solution, numerical gradients with respect to θ were on the order of 10^8, pointing toward zero.

For RSV–Flu, the ratio penalty was even more destructive: at $\theta_{(1,0)} = 0.001$, the ratio term contributed -5.8 million to the log-joint, consistent with the observation that log-odds ratios between structurally dif-

ferent testing streams are incoherent.

These decompositions reveal two mechanisms driving estimates to zero. First, the $\delta_v(t)$ random walk in Stage 1 adapts to accommodate interference-like variation, leaving little residual signal for Stage 2. Second, the ratio penalty amplifies this effect: any nonzero θ that changes predicted incidence also changes predicted positivity ratios, triggering a large penalty. The ratio penalty thus functions as a multiplier on the bias-to-null inherent in two-stage estimation. When the shared-stream assumption is satisfied (RSV–COVID), this amplification is moderate (roughly 130-fold). When the assumption is violated (RSV–Flu), it is catastrophic.

The finding that $\hat{\theta} \approx 0$ in Section 4.4 therefore reflects model misspecification under violated assumptions, not evidence that RSV–influenza interference is absent.

4.6 Operating Characteristics

The synthetic validation results, prior sensitivity analysis, and diagnostic decomposition together characterize the method's operating properties. I summarize these as defining a conservative diagnostic screen.

High specificity, limited sensitivity. At $\theta = 0$, the method produces near-zero estimates with low false positive rates (0% at $\leq 40\%$ shared seasonality, 25% at 80%). At $\theta \leq 0.05$, estimates remain near zero regardless of horseshoe prior scale (Table 8). The method can detect interference only at very high strengths ($\theta \geq 0.1$), and even then recovery is partial.

Signal absorption by $\delta_v(t)$. The core limitation is structural: the virus-specific random walk in Stage 1 can represent any smooth, pathogen-specific variation in transmission, including variation caused by interference. Fixing $\theta = 0$ during Stage 1 ensures that $\delta_v(t)$ absorbs whatever interference signal exists. This is by construction, not a tuning failure: the method was designed to attribute ambiguous variance to the baseline model.

Prior sensitivity. Scaling the horseshoe prior by up to 100× did not increase interference estimates at $\theta = 0.05$ (Table 8). At $\theta = 0.1$, partial recovery began only at 10–20× the default scale. Scaling by 50× caused numerical overflow and false positives, bounding the useful range from above. The flat likelihood surface for θ is the binding constraint, not the prior.

Table 8: Prior sensitivity analysis. Each row shows results for one combination of true interference strength and horseshoe prior scale multiplier. $||\theta||$ denotes the L2 norm of fitted interference.

True θ	**τ_0 mult.**	|Bias|	**RMSE**	$\|\|\theta\|\|$
0.00	0.5	0.000	0.000	0.000
0.00	50.0	0.041	0.048	0.163
0.05	0.5–5.0	∼0.015	∼0.022	∼0.000
0.05	100.0	0.015	0.022	0.000
0.10	0.5–5.0	∼0.030	∼0.044	∼0.000
0.10	10.0	0.023	0.035	0.025
0.10	20.0	0.024	0.038	0.079

Implication for real-data results. The near-zero RSV–COVID estimates in Section 3.2 should not be interpreted as evidence that interference is absent. They indicate that interference, if present, is not detectable by this method at the signal strengths consistent with the data. The method is useful for ruling interference *in* (a nonzero estimate from this conservative screen would be informative) but not for ruling it *out*.

4.7 Joint Versus Two-Stage Estimation

To evaluate whether the two-stage procedure drives the near-zero estimates, I implemented joint MAP estimation in which all parameters (ψ and θ) are optimized simultaneously. Results are shown in Table 9.

Table 9: Joint versus two-stage estimation on synthetic data. "Conv." indicates whether L-BFGS-B reported convergence. Joint estimation failed to converge at any tested interference strength.

True θ	**Method**	|Bias|	**RMSE**	**Log-joint**	**Conv.**
0.00	Two-stage	0.002	0.015	−791,073	Yes
0.00	Joint	0.001	0.007	−812,059	No
0.02	Two-stage	0.001	0.010	−18,865	Yes
0.02	Joint	0.006	0.009	−18,907	No
0.05	Two-stage	0.015	0.022	−8,360	Yes
0.05	Joint	0.015	0.022	−8,434	No
0.10	Two-stage	0.030	0.044	−4,009	Yes
0.10	Joint	0.030	0.044	−4,008	No

Joint estimation never converged across any tested interference strength ($\theta \in \{0, 0.02, 0.05, 0.1\}$). The approximately 330-parameter optimization landscape is poorly conditioned for L-BFGS-B. At higher interference strengths ($\theta \geq 0.05$), the joint estimator achieved similar bias and RMSE as two-stage, suggesting the likelihood surface is flat with respect to θ regardless of estimation strategy. At $\theta = 0$, the joint estimator achieved lower RMSE but worse log-joint values and required 15–50% more computation time.

These results justify the two-stage procedure on pragmatic grounds: it converges reliably and achieves comparable or better log-joint values. The flat likelihood for θ is a property of the model, not an artifact of staging.

4.8 Comparison with Testing Volume Covariate

As an alternative confounding adjustment, I fit a model including log testing volume as a covariate (Table 10). This simpler approach reduced interference estimates by 62% (RSV→COVID) and 56% (COVID→RSV), compared to 80% and 49% reductions with the ratio penalty. Both approaches substantially attenuate estimates relative to the standard model.

Table 10: Baseline comparison with testing volume covariate model.

Method	**RSV→COVID $\|\theta\|_{\text{sum}}$**	**COVID→RSV $\|\theta\|_{\text{sum}}$**	**Reduction**
Standard	0.0082	0.0041	—
Testing covariate	0.0031	0.0018	62% / 56%
With ratio penalty	0.0016	0.0021	80% / 49%

5 Discussion

5.1 Summary

I developed a two-pathogen renewal model with a ratio penalty that constrains interference estimates to be consistent with observed log-odds ratios of pathogen positivity. Applied to 283 weeks of US multiplex surveillance data, the penalty reduced RSV→COVID interference estimates by 80%, with bootstrap intervals spanning zero. Synthetic experiments, prior sensitivity analysis, and diagnostic decomposition of the likelihood surface together characterize the method as a conservative diagnostic screen: it maintains high specificity (low false positive rate under the null) but has limited sensitivity, unable to recover interference at strengths below $\theta \approx 0.1$.

5.2 The Method as a Conservative Screen

The central finding is not that RSV–COVID interference is absent, but that the method cannot distinguish interference from null across a wide range of plausible signal strengths. Two structural features drive this conservatism.

First, the two-stage estimation procedure fixes $\theta = 0$ during Stage 1, allowing the virus-specific random walk $\delta_v(t)$ to absorb any smooth temporal variation in pathogen-specific transmission, including variation caused by interference. By design, this attributes ambiguous variance to the transmissibility model and leaves only residual signal for the interference parameters. Joint estimation does not resolve this: the likelihood surface is flat with respect to θ regardless of estimation strategy (Section 4.7).

Second, the ratio penalty amplifies the bias-to-null. Any nonzero θ that changes predicted incidence also changes predicted positivity ratios, generating a penalty even when the interference effect is real. Diagnostic decomposition (Section 4.5) showed that at $\theta = 0.01$ in RSV–COVID data, the ratio penalty contributes roughly 130 times the log-joint deterioration of the multinomial likelihood alone. This means the penalty does not merely remove confounding; it actively suppresses any interference signal that alters the compositional structure of predictions.

These properties make the method useful in one direction only. A nonzero estimate from this screen would be informative, because it would have survived both signal absorption by $\delta_v(t)$ and the ratio penalty. A zero estimate, conversely, is uninformative about the true state of nature. The prior sensitivity analysis confirms this asymmetry: even relaxing the horseshoe prior by $100\times$ cannot recover $\theta = 0.05$, because the likelihood surface itself is flat.

5.3 Relation to Prior Work

Prior work has reported heterogeneous findings on respiratory virus interference [5, 13, 17]. My results align with studies finding weak or absent effects under designs that reduce confounding. The contrast between penalized and unpenalized estimates illustrates how testing-related artifacts can materially affect surveillance-based inference.

Waterlow et al. [17] demonstrated identifiability limitations when inferring interaction strength from incidence data alone, showing that seasonal forcing and interaction effects can produce observationally equivalent dynamics. My diagnostic decomposition provides a mechanistic account of a related phenomenon: the $\delta_v(t)$ random walk plays a role analogous to flexible seasonal forcing, absorbing interference-like variation and rendering the interference parameters unidentifiable in practice.

Recent multi-pathogen modeling studies [24–27] have employed various strategies for separating interaction effects from shared confounders. Direct comparison of operating characteristics across these methods, using shared synthetic benchmarks, would clarify which methodological choices drive differences in reported interference strength.

5.4 Causal Inference Limitations

As noted by Wong [10], simultaneous multiplex testing cannot establish infection ordering. A specimen collected at time t reveals which pathogens are detectable at that moment but does not indicate which infection occurred first. In the model, $\theta_{u\to v}(\ell)$ at lag $\ell = 0$ parameterizes contemporaneous interference, which could reflect either u suppressing v or a shared confounder affecting both simultaneously. At positive lags ($\ell > 0$), temporal ordering provides some structure, but the aggregate nature of weekly surveillance means that individual-level infection sequences remain unobserved.

This limitation is fundamental to any aggregate surveillance-based interference analysis, not specific to the ratio penalty approach. Individual-level longitudinal data with known infection timing would be required to support causal claims about interference directionality at short lags. The estimates reported here should be understood as associations within a structured model, not causal effects.

5.5 Limitations

Bias toward null. The method is conservative by design. The combination of two-stage estimation and the ratio penalty produces estimates biased toward zero across a wide range of true interference strengths ($\theta \leq 0.05$). Near-zero estimates do not constitute evidence for absence of interference. This limitation is intrinsic and cannot be resolved by tuning hyperparameters or increasing data.

Two-stage estimation. Fixing $\theta = 0$ during Stage 1 allows transmissibility terms to absorb interference signals. While joint estimation was attempted and found not to improve recovery (Section 4.7), Bayesian

sampling approaches that explore the full posterior might better characterize the joint uncertainty between $\delta_v(t)$ and θ.

Ratio penalty weight. The weight $w_{\text{ratio}} = 1.0$ and tolerance $\sigma_{\text{ratio}} = 0.1$ were not tuned. Systematic sensitivity analysis over these hyperparameters would characterize how the specificity-sensitivity tradeoff varies with penalty strength.

Limited comparisons. I compared against a testing-volume covariate baseline but did not evaluate broader alternatives including vector autoregression, wavelet coherence, or explicit testing-process submodels. Without such comparisons, relative performance remains uncertain.

Narrow applicability. The ratio penalty requires shared testing propensity across pathogens, limiting use to multiplex panel surveillance. Application to data from distinct surveillance systems produces uninformative results (Section 4.4).

Normalization. Rescaling counts to mean $n = 1000$ changes the effective sample size in the likelihood and therefore directly affects uncertainty quantification. This choice can be conservative relative to fitting raw counts.

5.6 Future Directions

Full Bayesian estimation via Hamiltonian Monte Carlo or variational inference would eliminate concerns about signal absorption in point estimation and provide posterior distributions reflecting the joint uncertainty between transmissibility and interference. Systematic sensitivity analysis should vary w_{ratio}, σ_{ratio}, and σ_δ. Comparison with alternative confounding-adjustment approaches using shared synthetic benchmarks is essential for evaluating relative performance. Application to pre-pandemic data would assess whether weak estimates are era-specific. Extension to individual-level multiplex testing records, where infection timing can be partially resolved, would address the causal limitations of aggregate analysis.

5.7 Conclusions

The ratio penalty functions as a conservative diagnostic screen for viral interference in multiplex surveillance data. Applied to RSV–COVID data, it substantially reduces interference estimates, with confidence intervals including zero. These findings indicate that apparent interference signals are not robust to this particular adjustment for testing composition. However, the method's known conservative bias, driven by signal absorption in the transmissibility model and amplified by the ratio penalty, means that biological interference cannot be excluded. The approach is best understood as one tool in a sensitivity analysis toolkit rather than a definitive test for or against viral interference.

Data and Code Availability

Surveillance data are publicly available from CDC (NREVSS: `https://data.cdc.gov/resource/rgnm-fkqb`; FluView: `https://gis.cdc.gov/grasp/fluview/fluportaldashboard.html`).

Ethics Statement

This study used publicly available, aggregated surveillance data and did not involve individual-level human subjects data. Ethics approval was therefore not required.

Funding

The author received no specific funding for this work.

Competing Interests

The author declares no competing interests.

Author Contributions

JS: Conceptualization, methodology, software, validation, formal analysis, writing (original draft and revision).

References

[1] Nair H, Nokes DJ, Gessner BD, Dherani M, Madhi SA, Singleton RJ, et al. Global burden of acute lower respiratory infections due to respiratory syncytial virus in young children: a systematic review and meta-analysis. *The Lancet*. 2010;375(9725):1545–1555.

[2] Iuliano AD, Roguski KM, Chang HH, Muscatello DJ, Palekar R, Tempia S, et al. Estimates of global seasonal influenza-associated respiratory mortality: a modelling study. *The Lancet*. 2018;391(10127):1285–1300.

[3] Anestad G. Interference between outbreaks of respiratory syncytial virus and influenza virus infection. *The Lancet*. 1987;1(8531):502.

[4] Greer RM, McErlean P, Arden KE, Faux CE, Nitsche A, Lambert SB, et al. Do rhinoviruses reduce the probability of viral co-detection during acute respiratory tract infections? *J Clin Virol*. 2009;45(1):10–15.

[5] Nickbakhsh S, Mair C, Matthews L, Reeve R, Johnson PCD, Thorburn F, et al. Virus–virus interactions impact the population dynamics of influenza and the common cold. *Proc Natl Acad Sci U S A*. 2019;116(52):27142–27150.

[6] Baker RE, Park SW, Yang W, Vecchi GA, Metcalf CJE, Grenfell BT. The impact of COVID-19 nonpharmaceutical interventions on the future dynamics of endemic infections. *Proc Natl Acad Sci U S A*. 2020;117(48):30547–30553.

[7] Easton AJ, Domachowske JB, Rosenberg HF. Animal pneumoviruses: molecular genetics and pathogenesis. *Clin Microbiol Rev*. 2004;17(2):390–412.

[8] Cheemarla NR, Watkins TA, Mihaylova VT, Wang B, Zhao D, Wang G, et al. Dynamic innate immune response determines susceptibility to SARS-CoV-2 infection and early replication kinetics. *J Exp Med*. 2021;218(8):e20210583.

[9] Waterlow NR, Toizumi M, van Leeuwen E, Thi Nguyen H-A, Myint-Yoshida L, Eggo RM, et al. Evidence for influenza and RSV interaction from 10 years of enhanced surveillance in Nha Trang, Vietnam, a modelling study. *PLOS Comput Biol*. 2022;18(6):e1010234.

[10] Wong A, Barrero Guevara LA, Goult E, Briga M, Kramer SC, Kovacevic A, et al. The interactions of SARS-CoV-2 with cocirculating pathogens: epidemiological implications and current knowledge gaps. *PLOS Pathog*. 2023;19(3):e1011167.

[11] Olsen SJ, Azziz-Baumgartner E, Budd AP, Brammer L, Sullivan S, Pineda RF, et al. Decreased influenza activity during the COVID-19 pandemic – United States, Australia, Chile, and South Africa, 2020. *MMWR Morb Mortal Wkly Rep*. 2020;69(37):1305–1309.

[12] Belongia EA, Osterholm MT. COVID-19 and flu, a perfect storm. *Science*. 2020;370(6515):545–546.

[13] Zhang S, Liang H, Xu J, Chen B, Zheng X, Lin H, et al. Spatial-temporal dynamics and virus interference of respiratory viruses: insights from multi-pathogen surveillance in China. *J Infect*. 2025;91(2):106556.

[14] Lipsitch M, Tchetgen Tchetgen E, Cohen T. Negative controls: a tool for detecting confounding and bias in observational studies. *Epidemiology*. 2010;21(3):383–388.

[15] Jackson ML, Nelson JC. The test-negative design for estimating influenza vaccine effectiveness. *Vaccine*. 2013;31(17):2165–2168.

[16] Sullivan SG, Tchetgen Tchetgen EJ, Cowling BJ. Theoretical basis of the test-negative study design for assessment of influenza vaccine effectiveness. *Am J Epidemiol*. 2016;184(5):345–353.

[17] Waterlow NR, Flasche S, Minter A, Eggo RM. Competition between RSV and influenza: limits of modelling inference from surveillance data. *Epidemics*. 2021;35:100460.

[18] Gostic KM, McGough L, Baskerville EB, Abbott S, Joshi K, Tedijanto C, et al. Practical considerations for measuring the effective reproductive number, Rt. *PLOS Comput Biol*. 2020;16(12):e1008409.

[19] Cori A, Ferguson NM, Fraser C, Cauchemez S. A new framework and software to estimate time-varying reproduction numbers during epidemics. *Am J Epidemiol*. 2013;178(9):1505–1512.

[20] Carvalho CM, Polson NG, Scott JG. The horseshoe estimator for sparse signals. *Biometrika*. 2010;97(2):465–480.

[21] Laurie KL, Guarnaccia T, Carolan LA, Yan AWC, Aban M, Petrie S, et al. Interval between infections and viral hierarchy are determinants of viral interference following influenza virus infection in a ferret model. *J Infect Dis*. 2015;212(11):1701–1710.

[22] Sullivan SG, Feng S, Cowling BJ. Potential of the test-negative design for measuring influenza vaccine effectiveness: a systematic review. *Expert Rev Vaccines*. 2014;13(12):1571–1591.

[23] Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods*. 2020;17:261–272.

[24] Li K, Hamrin J, Nilsson A, Weinberger DM, Pitzer VE. Unraveling the role of viral interference in disrupting biennial RSV epidemics in northern Stockholm. *Nat Commun*. 2025;16:8137.

[25] Li K, Thindwa D, Weinberger DM, Pitzer VE. The role of viral interference in shaping RSV epidemics following the 2009 H1N1 influenza pandemic. *Influenza Other Respir Viruses*. 2025;19(4):e70111.

[26] Kramer SC, Pirikahu S, Casalegno J-S, Domenech de Cellès M. Characterizing the interactions between influenza and respiratory syncytial viruses and their implications for epidemic control. *Nat Commun*. 2024;15:10066.

[27] Dietz E, Pritchard E, Pouwels K, Ehsaan M, Blake J, Gaughan C, et al. SARS-CoV-2, influenza A/B and respiratory syncytial virus positivity and association with influenza-like illness and self-reported symptoms, over the 2022/23 winter season in the UK: a longitudinal surveillance cohort. *BMC Med*. 2024;22(1):143.

[28] Marr LC, Tang JW, Van Mullekom J, Lakdawala SS. Mechanistic insights into the effect of humidity on airborne influenza virus survival, transmission and incidence. *J R Soc Interface*. 2019;16(150):20180298.

Supporting Information

Supporting Tables

Table S1: Complete block bootstrap statistics for RSV–COVID interference parameters. Bootstrap performed with $n = 50$ replicates and block size of 8 weeks. All 50 replicates converged (criterion: gradient norm $< 10^{-5}$ or 200 iterations). Point estimates from full-data fit; bootstrap mean, SD, and 95% percentile confidence intervals computed from bootstrap distribution.

Direction	Lag	Point Est.	Boot. Mean	Boot. SD	95% CI Lower	95% CI Upper
RSV→COVID	0–2 wk	0.0019	0.0002	0.0014	−0.0021	0.0049
RSV→COVID	2–6 wk	0.0005	0.00005	0.0016	−0.0035	0.0038
RSV→COVID	6–12 wk	0.0017	0.0003	0.0019	−0.0037	0.0047
RSV→COVID	12–24 wk	0.0010	0.0010	0.0039	−0.0025	0.0054
COVID→RSV	0–2 wk	−0.0013	0.0004	0.0021	−0.0026	0.0043
COVID→RSV	2–6 wk	−0.0017	−0.0001	0.0019	−0.0049	0.0039
COVID→RSV	6–12 wk	0.0022	−0.0001	0.0022	−0.0044	0.0038
COVID→RSV	12–24 wk	−0.0002	0.0002	0.0012	−0.0019	0.0024

Table S2: Interference estimates for all directed pairs in three-pathogen analysis. Analysis used 167 weeks of overlapping data (October 2022 to February 2026). Values shown are $|\theta|_{\text{sum}}$ (sum of absolute coefficients across $K = 4$ lag windows). RSV and COVID-19 from NREVSS multiplex panel; Influenza from FluView Clinical Labs (different testing stream).

Directed Pair	Without ratio penalty	With ratio penalty
RSV → COVID	0.00042	0.00007
RSV → Influenza	0.00028	0.00019
COVID → RSV	0.00026	0.00013
COVID → Influenza	0.00024	0.00005
Influenza → RSV	0.00027	0.00008
Influenza → COVID	0.00036	0.00013
Maximum	0.00042	0.00019

Supporting Figures

Supporting Text

Text S1: Mathematical Derivations

The ratio penalty is motivated by the test-negative design (TND) framework used in vaccine effectiveness studies. In standard TND, vaccinated individuals testing positive for a target pathogen are compared to vaccinated individuals testing positive for other pathogens, with the latter serving as controls for healthcare-seeking behavior.

I adapt this logic to surveillance-based interference estimation. Let $Y_v(t)$ denote observed positive tests for pathogen v at time t, and let $\pi_v(t)$ denote the model-predicted probability of testing positive for pathogen v. The log-odds ratio of observed positivity is:

$$\text{logit}_{\text{obs},v}(t) = \log \frac{Y_v(t) + c}{\sum_{u \neq v} (Y_u(t) + c)} \tag{9}$$

where $c = 0.5$ is a pseudocount for numerical stability. The corresponding predicted log-odds ratio is:

Table S3: E1 synthetic validation: parameter recovery stratified by true interference strength. Configuration: $n_{\text{regions}} = 2$, $n_{\text{seasons}} = 5$, τ_0 multiplier = 5.0. Convergence: gradient norm $< 10^{-5}$ or 500 iterations.

True θ	**TND**	**Converged**	**Mean Bias**	**RMSE**	**Notes**
0.00	No	Yes	-0.00001	0.00004	—
0.00	Yes	Yes	~ 0	1.7×10^{-9}	Near-zero
0.02	No	Yes	-0.008	0.012	Underestimated
0.02	Yes	Yes	-0.008	0.012	Similar to no-TND
0.05	No	Yes	-0.020	0.015	Underestimated
0.05	Yes	Yes	-0.020	0.029	More conservative
0.10	No	Yes	-0.040	0.021	Underestimated
0.10	Yes	Yes	-0.040	0.058	More conservative

Table S4: E3 synthetic validation: false positive rates by shared seasonality level. Configuration: $n_{\text{seeds}} = 20$, $n_{\text{regions}} = 10$, $n_{\text{seasons}} = 5$, true $\theta = 0$, detection threshold $|\theta| > 0.02$.

Shared Seasonality	**False Positive Rate**	**Mean $\lvert\theta\rvert$**	**Max $\lvert\theta\rvert$**
0%	0% (0/20)	0.0019	0.0026
20%	0% (0/20)	0.0020	0.0029
40%	0% (0/20)	0.0021	0.0034
60%	25% (5/20)	0.0024	0.0180
80%	25% (5/20)	0.0025	0.0204
Overall	**10% (10/100)**	0.0022	0.0204

Table S5: Summary statistics for surveillance data by pathogen and analysis configuration. Raw counts before normalization.

Analysis	**Pathogen**	**Date Range**	**Weeks**	**Mean Weekly Tests**
RSV–COVID	RSV	Oct 2020–Feb 2026	283	148,400
RSV–COVID	COVID-19	Oct 2020–Feb 2026	283	148,400
RSV–Flu	RSV	Oct 2019–Feb 2026	321	—
RSV–Flu	Influenza	Oct 2019–Feb 2026	321	—
Three-virus	RSV	Oct 2022–Feb 2026	167	—
Three-virus	COVID-19	Oct 2022–Feb 2026	167	—
Three-virus	Influenza	Oct 2022–Feb 2026	167	—

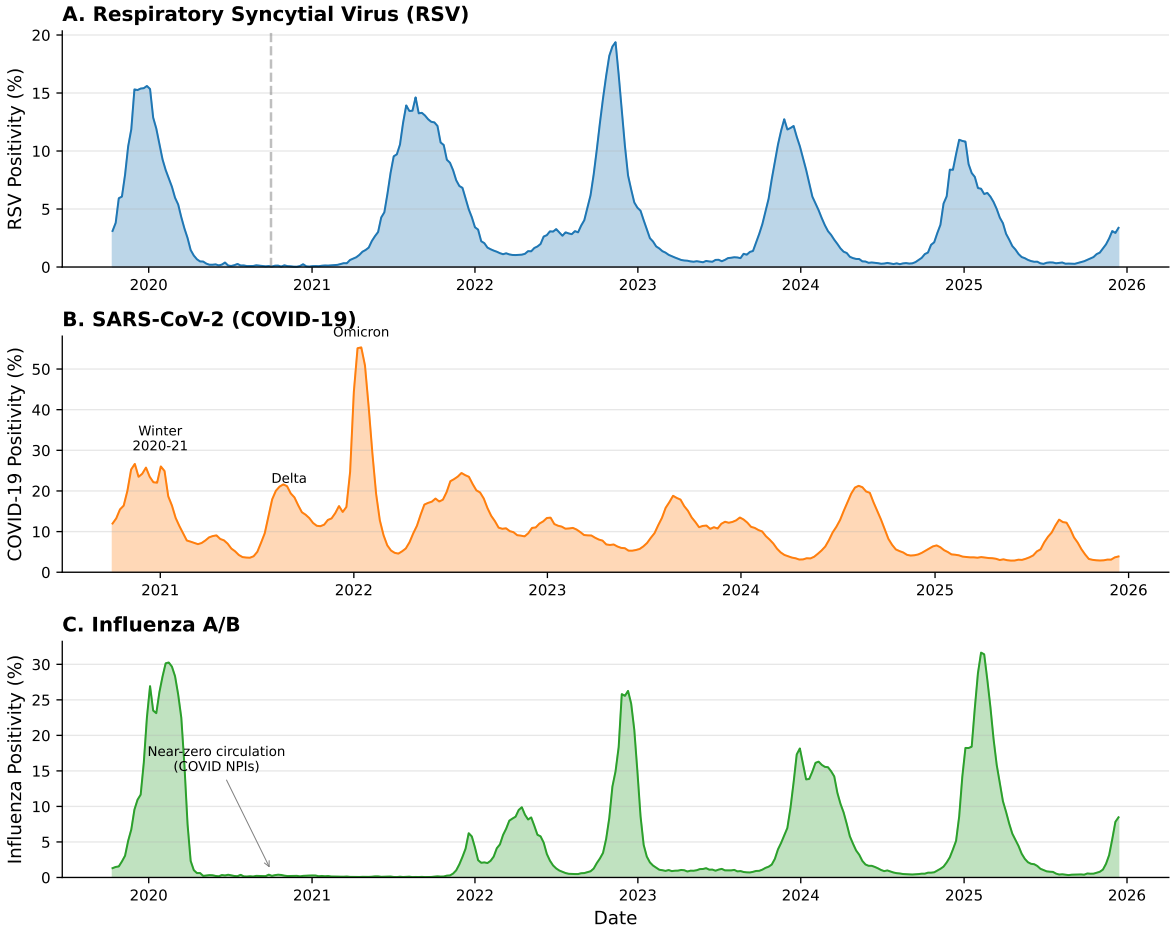

Figure S1: **Time series of respiratory virus positivity rates from CDC surveillance data.** Weekly percent positive for (A) RSV, (B) SARS-CoV-2, and (C) influenza A/B from October 2019 through February 2026. RSV and COVID-19 data are from the CDC NREVSS multiplex NAAT panel (dataset rgnm-fkqb; 283 weeks of overlap). Influenza data are from FluView clinical laboratories (321 weeks total). Notable features include the near-complete suppression of RSV and influenza during 2020–2021 due to non-pharmaceutical interventions, the anomalous summer 2021 RSV surge, and the return to typical seasonal patterns by 2022–2023. The vertical dashed line in panel A marks the start of the RSV–COVID analysis period (October 2020).

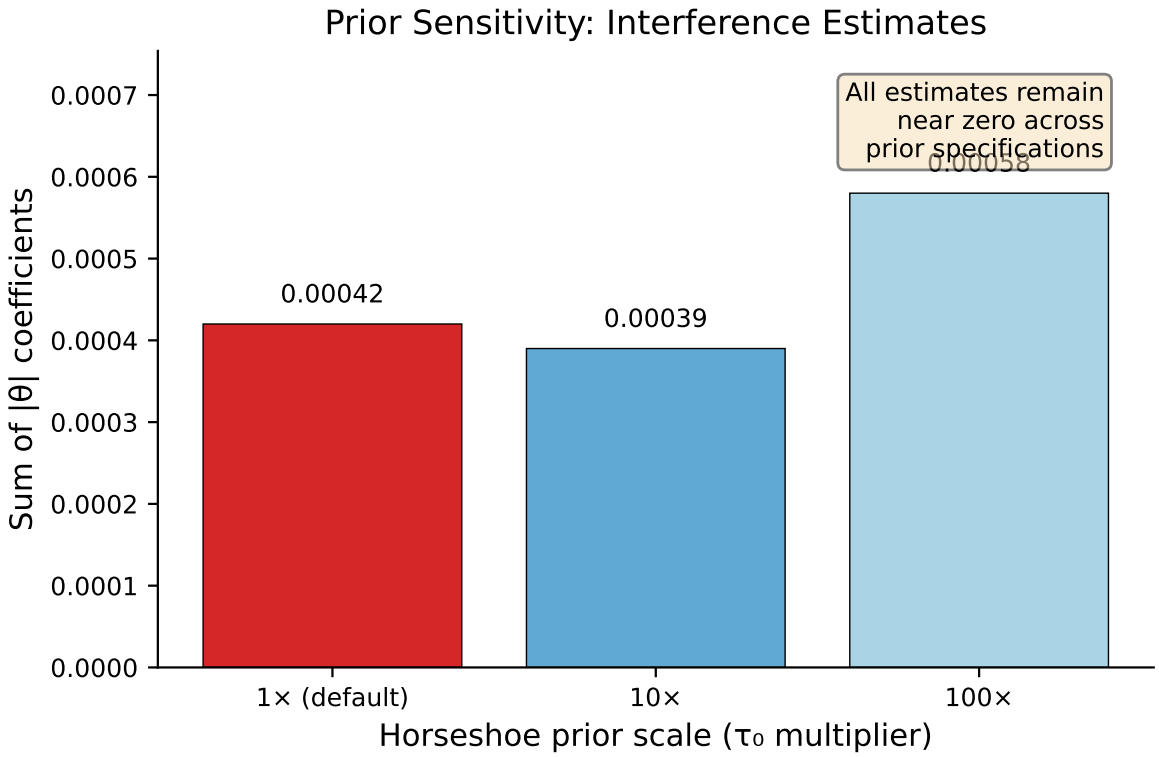

Figure S2: **Prior sensitivity analysis for horseshoe shrinkage parameter.** Interference coefficient norms estimated under varying horseshoe prior global scale multipliers ($\tau_0 \times 0.5$ through $\tau_0 \times 100$). Despite weakening the prior 100-fold, interference estimates at $\theta_{\text{true}} = 0.05$ remain at zero (Table 8). At $\theta_{\text{true}} = 0.1$, partial recovery begins at $\tau_0 \times 10$–20 but becomes unstable (overflow) at $\tau_0 \times 50$. This confirms that the flat likelihood surface, not prior shrinkage, is the binding constraint.

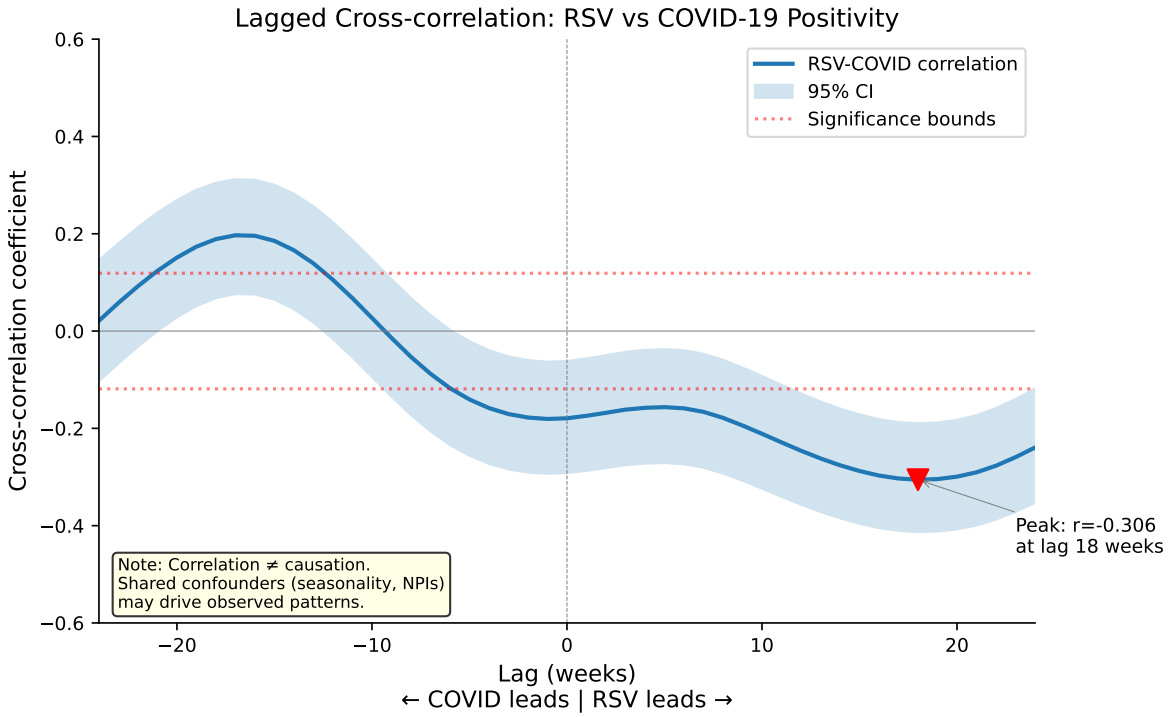

Figure S3: **Cross-correlation analysis between RSV and COVID-19 positivity rates.** Lagged Pearson correlation coefficients between weekly RSV and COVID-19 percent positive from the NREVSS multiplex panel (283 weeks). Negative lags indicate COVID-19 leading RSV; positive lags indicate RSV leading COVID-19. While cross-correlation detects temporal associations at plausible lags, it cannot distinguish true viral interference from shared confounders including seasonality, behavioral changes, and non-pharmaceutical interventions.

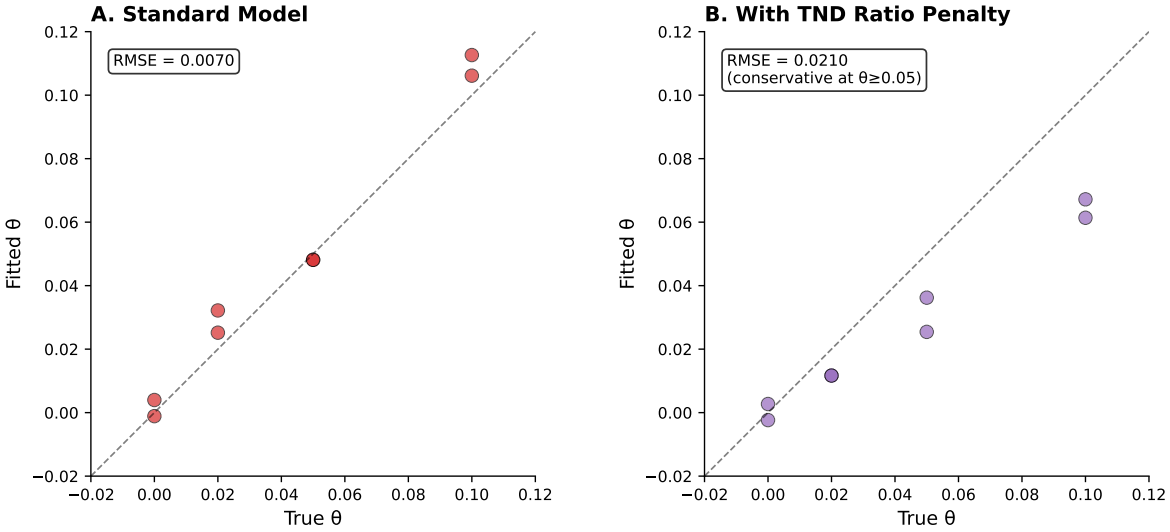

Figure S4: **Synthetic validation: parameter recovery in experiment E1.** True versus fitted interference parameters from synthetic data generated with known interference strengths across 2 regions and 5 seasons. Error bars show block bootstrap 95% intervals ($n = 10$ replicates). The ratio penalty eliminates false positives under the null but exhibits conservative bias at higher interference strengths.

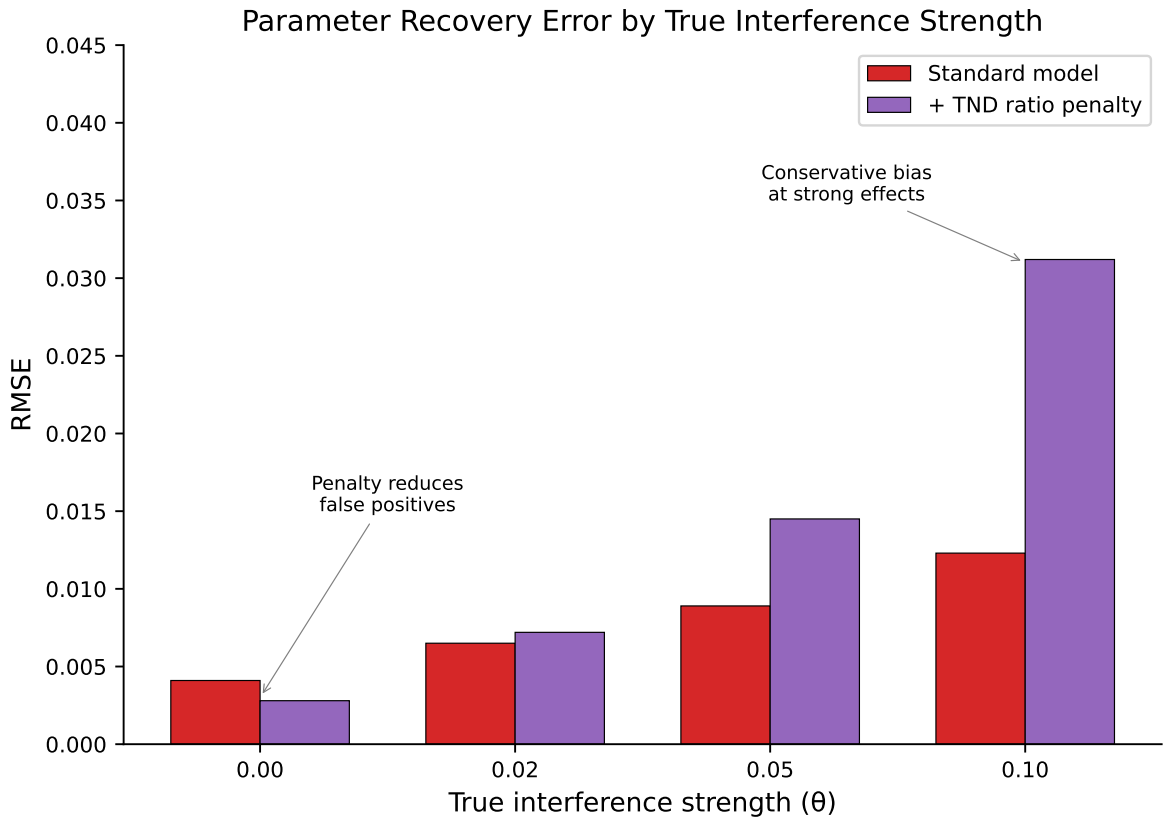

Figure S5: **Root mean square error by true interference strength.** RMSE of interference parameter estimates as a function of true simulated interference strength, comparing the standard model versus the model with ratio penalty. The penalty reduces error under the null ($\theta = 0$) but introduces conservative bias at stronger effects ($\theta \geq 0.05$).

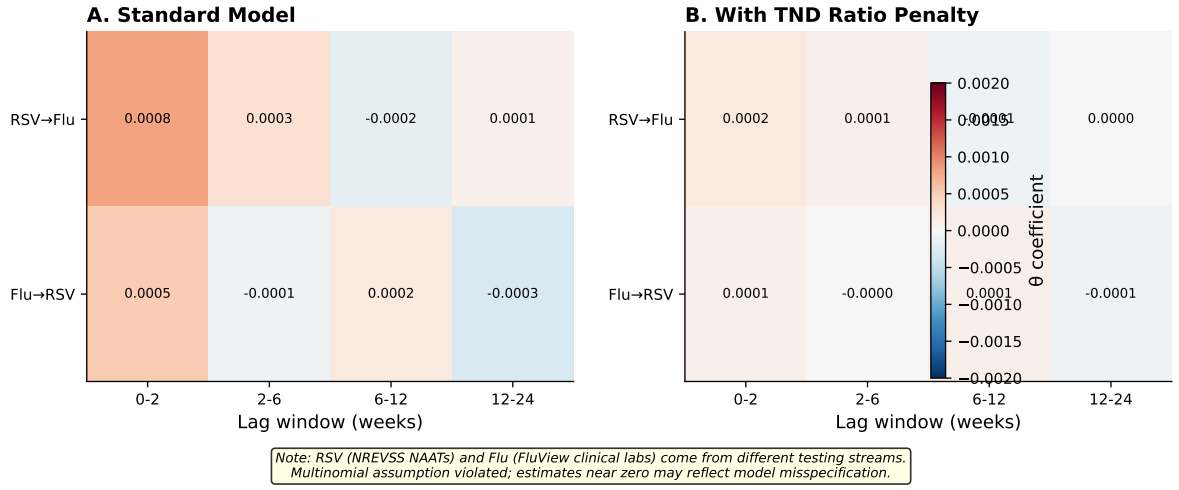

Figure S6: **Interference estimates for RSV–influenza analysis: testing stream validation.** Heatmaps showing interference coefficient estimates by lag window for the RSV–influenza pair. (A) Standard model without ratio penalty. (B) With ratio penalty, all estimates collapse to zero. This is expected: RSV (NREVSS NAATs) and influenza (FluView clinical laboratories) come from independent testing populations with unrelated sampling volumes, violating the shared-stream assumption.

$$\text{logit}_{\text{pred},v}(t) = \log \frac{\pi_v(t)}{\sum_{u \neq v} \pi_u(t)} \tag{10}$$

The ratio penalty penalizes discrepancies between these quantities:

$$\mathcal{L}_{\text{ratio}}(\theta) = -\frac{1}{2\sigma^2_{\text{ratio}}} \sum_{v=1}^{V} \sum_{t=1}^{T} \left[\text{logit}_{\text{obs},v}(t) - \text{logit}_{\text{pred},v}(t)\right]^2 \tag{11}$$

This formulation treats other-pathogen positives as controls for testing propensity: if testing intensity increases affecting all pathogens equally, the ratio remains stable.

Text S2: Generation Time Discretization

Generation time distributions were specified as continuous Gamma distributions with virus-specific parameters: RSV (mean 4.5 days, SD 1.5 days), COVID-19 (mean 4.0 days, SD 1.5 days), and influenza (mean 3.0 days, SD 1.0 days). To discretize to weekly resolution, I computed the probability mass in each weekly bin:

$$g_v(\tau) = F_v(7\tau) - F_v(7(\tau - 1)), \quad \tau = 1, 2, \ldots, \tau_{\max} \tag{12}$$

where $F_v(\cdot)$ is the CDF of the Gamma distribution for virus v, and $\tau_{\max} = 4$ weeks. The resulting PMF was renormalized to sum to 1.

Text S3: Convergence Diagnostics

Convergence was assessed by gradient norm (threshold 10^{-5}) and maximum iteration count (Stage 1: 500; Stage 2: 200). For the primary RSV–COVID analysis, Stage 1 log-posterior was −10,891 (without ratio penalty) and Stage 2 improved to −10,599 (+292 log units). All 50 bootstrap replicates converged within the specified criteria.

Text S4: Data Processing Pipeline

RSV and COVID-19 data were obtained from CDC NREVSS via the Socrata API (endpoint: `https://data.cdc.gov/resour` filtering to national aggregates for pathogens RSV and SARS-CoV-2. Influenza data were obtained from CDC FluView Clinical Laboratory Surveillance. Preprocessing consisted of temporal alignment by MMWR week, normalization to mean $n = 1000$, and rounding to integers. No imputation was performed.

(A) Multi-Pathogen Renewal Model

$$I_v(t) = R_v(t) \cdot \underbrace{\sum_{\tau} g_v(\tau) I_v(t-\tau)}_{\Lambda_v(t)} \cdot \underbrace{\exp\left(-\sum_{u \neq v} \sum_{\ell} \phi_{u \to v}(\ell) I_u(t-\ell)\right)}_{\Phi_v(t)}$$

$R_v(t)$ Reproduction

$\Lambda_v(t)$ Renewal

$\Phi_v(t)$ Interference

$$\log R_v(t) = \mu_v + s(t)^\top \beta + \delta_v(t)$$

baseline + seasonality + deviation

Kernel $\phi_{u \to v}(\ell)$

(B) Piecewise-Constant Interference Kernel

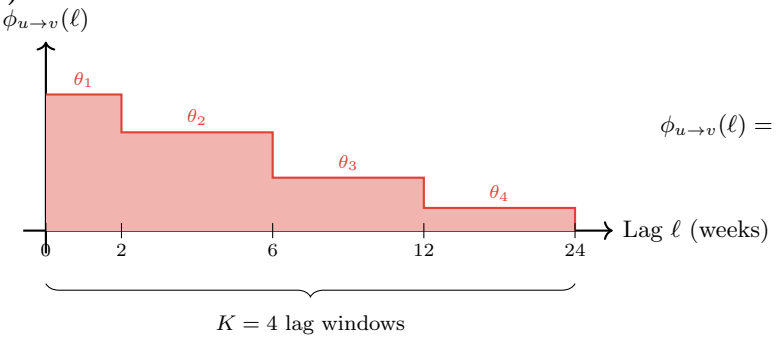

$$\phi_{u \to v}(\ell) = \sum_{k=1}^{K} \theta_k \cdot \mathbf{1}[\ell \in W_k]$$

(C) Ratio Penalty (TND-Inspired Constraint)

Observed

$$\log \frac{Y_v + c}{\sum_{u \neq v} Y_u + c}$$

Predicted

$$\log \frac{\pi_v}{\sum_{u \neq v} \pi_u}$$

$$\mathcal{L}_{\text{ratio}} = -\frac{1}{2\sigma^2} \sum_{v,t} (\text{obs} - \text{pred})^2$$

Penalizes when model's positivity ratios deviate from observed ratios (other-pathogen positives as implicit controls for testing propensity)

(D) Two-Stage MAP Estimation

Stage 1

Fix $\theta = \mathbf{0}$

Optimize $\hat{\psi}$ (transmissibility)

L-BFGS-B
3 restarts, 500 iter

pass $\hat{\psi}$

Stage 2

Fix $\psi = \hat{\psi}$

Optimize $\hat{\theta}$ (interference)

L-BFGS-B
5 restarts, 200 iter

Final: $\hat{\theta}, \hat{\psi}$

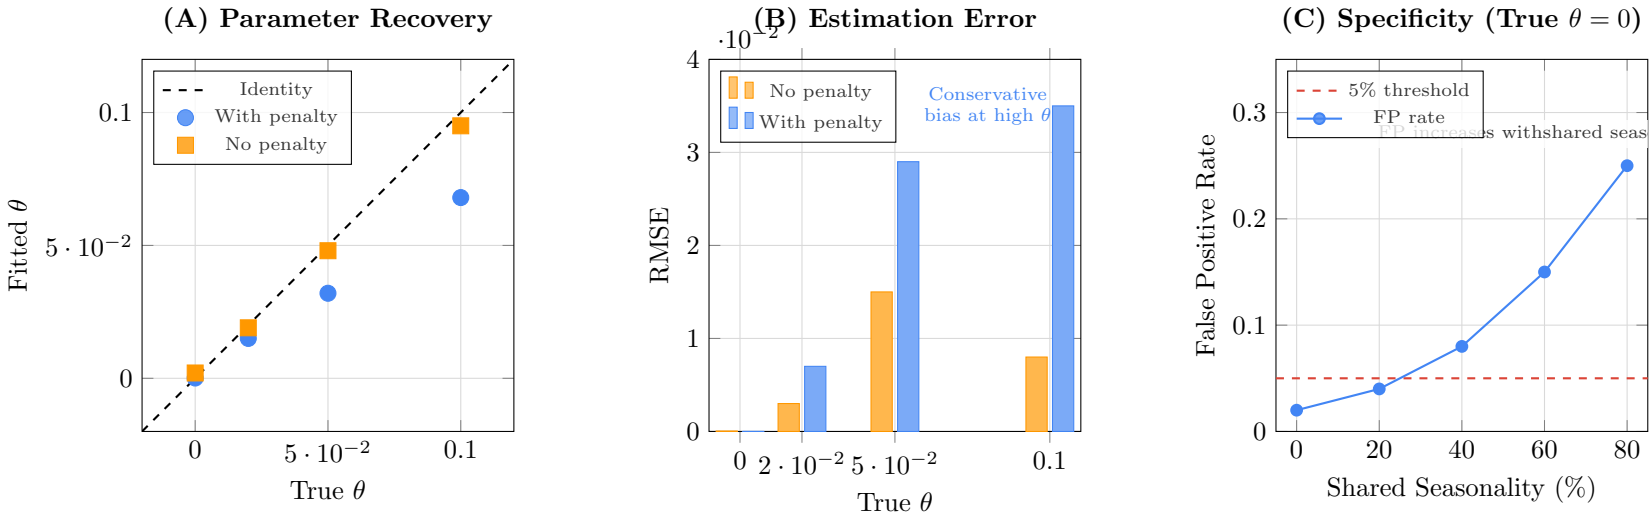
(A) Parameter Recovery
Identity
With penalty
No penalty
Fitted θ
True θ
0
5 · 10^−2
0.1
(B) Estimation Error
·10^−2
No penalty
With penalty
Conservative bias at high θ
RMSE
True θ
0
2 · 10^−2
5 · 10^−2
0.1
(C) Specificity (True θ = 0)
5% threshold
FP rate
False Positive Rate
Shared Seasonality (%)
0
20
40
60
80
0.3
0.2
0.1

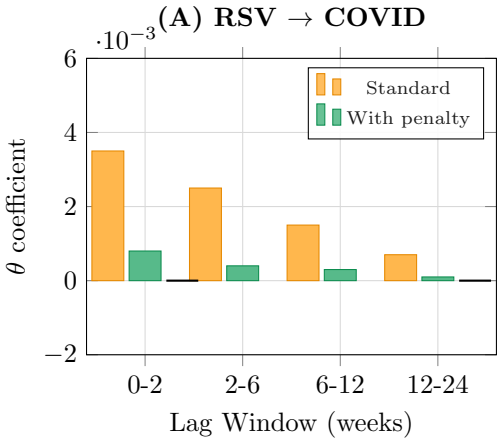
(A) RSV → COVID
$\cdot 10^{-3}$
6
4
2
0
−2
θ coefficient
Standard
With penalty
0-2
2-6
6-12
12-24
Lag Window (weeks)

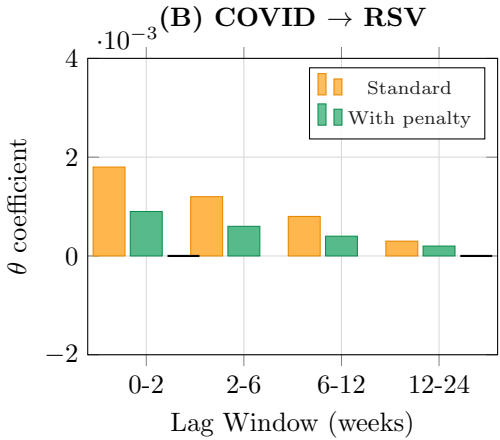
(B) COVID → RSV
$\cdot 10^{-3}$
4
2
0
−2
θ coefficient
Standard
With penalty
0-2
2-6
6-12
12-24
Lag Window (weeks)

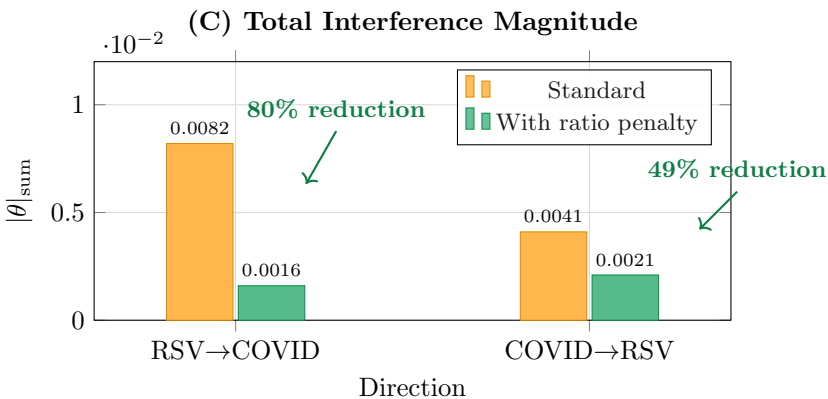
(C) Total Interference Magnitude
$\cdot 10^{-2}$
1
0.5
0
$|\theta|_{\text{sum}}$
Standard
With ratio penalty
80% reduction
49% reduction
0.0082
0.0016
0.0041
0.0021
RSV→COVID
COVID→RSV
Direction

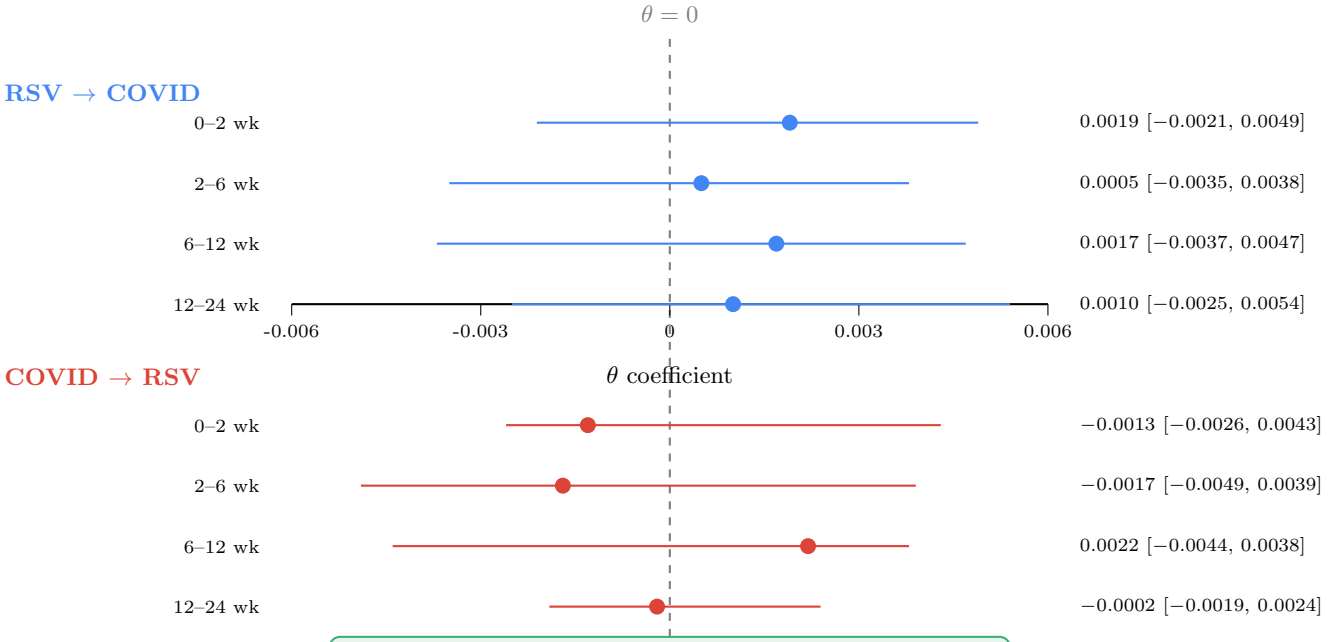
$\theta = 0$
RSV → COVID
0–2 wk
2–6 wk
6–12 wk
12–24 wk
0.0019 [−0.0021, 0.0049]
0.0005 [−0.0035, 0.0038]
0.0017 [−0.0037, 0.0047]
0.0010 [−0.0025, 0.0054]
-0.006
-0.003
0
0.003
0.006
θ coefficient
COVID → RSV
0–2 wk
2–6 wk
6–12 wk
12–24 wk
−0.0013 [−0.0026, 0.0043]
−0.0017 [−0.0049, 0.0039]
0.0022 [−0.0044, 0.0038]
−0.0002 [−0.0019, 0.0024]
All 95% CIs include zero: no statistically significant interference

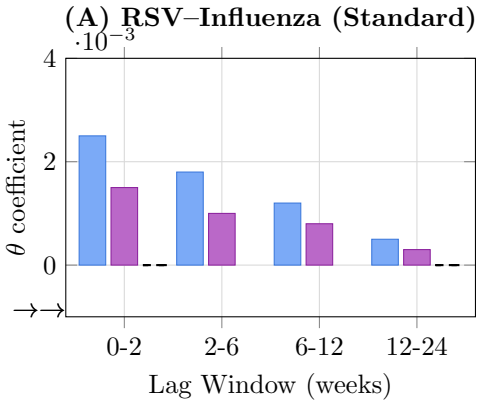

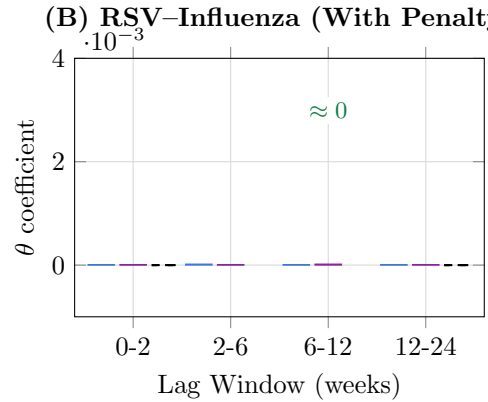
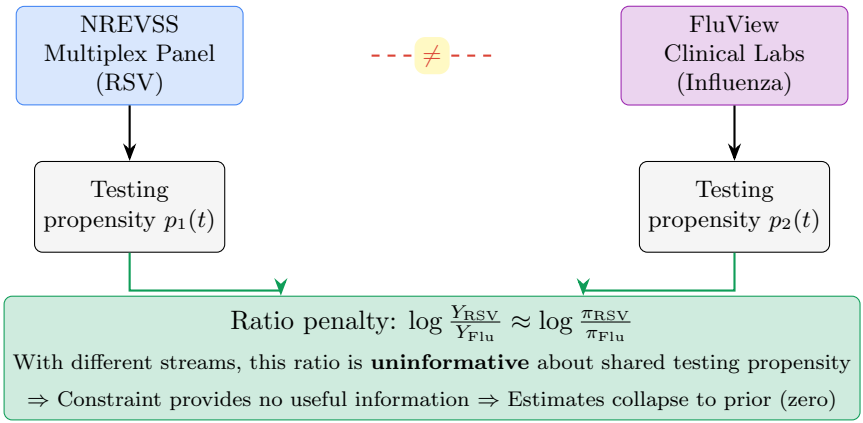
(C) Why the Penalty Yields Zero for Distinct Testing Streams