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

Joshua Steier*

Independent Researcher

Abstract

Background: Viral interference, in which infection with one pathogen reduces susceptibility to another, may influence respiratory virus dynamics. Inference from surveillance data is complicated by time-varying testing behavior that can induce correlated detection patterns independent of biological interaction.

Methods: We developed a multi-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. We applied the model to 272 weeks of US NAAT surveillance data reported to NREVSS (RSV and COVID-19; October 2020 to December 2025), validated parameter recovery in synthetic experiments, and quantified uncertainty via block bootstrap.

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 showed the penalty reduces false positives at $\theta = 0$ but increases estimation error when true interference is present, indicating conservative bias.

Conclusions: The ratio penalty substantially reduced interference point estimates, with confidence intervals spanning zero. These results are consistent with testing behavior contributing to apparent interference signals but do not rule out biological interference given the penalty's conservative properties. The approach enables sensitivity analysis in surveillance-based interference inference.

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 is difficult because testing behavior changes over time: when any respiratory illness surges, more people seek tests, potentially creating correlated patterns in the data that mimic biological interaction.

We developed a statistical method to probe this confounding. Borrowing logic from vaccine studies, we penalize our 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. Importantly, our method is intentionally conservative: simulation experiments show it can also diminish real interference signals. Thus, our near-zero estimates do not prove interference is absent; rather, they indicate that apparent signals in these data are not robust to 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.

*Corresponding author: joshsteier@gmail.com

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, whereby infection with one pathogen temporarily reduces susceptibility to another, has been proposed as a determinant of respiratory virus dynamics [3, 4]. Proposed mechanisms include innate immune activation, particularly interferon responses [7, 8]. 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]. This variability motivates 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 including multiplex panels) detecting multiple pathogens simultaneously. This can create correlated detection: when one virus surges, testing intensity rises, and positives for other pathogens may shift in apparent synchrony regardless of biological interaction. Standard methods, 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]. We adapt this logic to interference estimation from aggregate surveillance. Our 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.

We emphasize scope limitations. First, the penalty requires pathogens measured under shared testing propensity—satisfied by multiplex panels but not when combining distinct surveillance systems. Second, we use two-stage estimation (optimizing transmissibility before interference) that may affect variance allocation between components; joint optimization was not performed. Third, we do not compare against alternative confounding-adjustment approaches.

We apply this model to 272 weeks of US NREVSS NAAT surveillance data detecting RSV and COVID-19 (October 2020–December 2025). Our objectives are to: (1) estimate bidirectional interference with and without the ratio penalty; (2) characterize the penalty's effects on estimation, including its conservative bias; and (3) evaluate whether residual signals are distinguishable from zero.

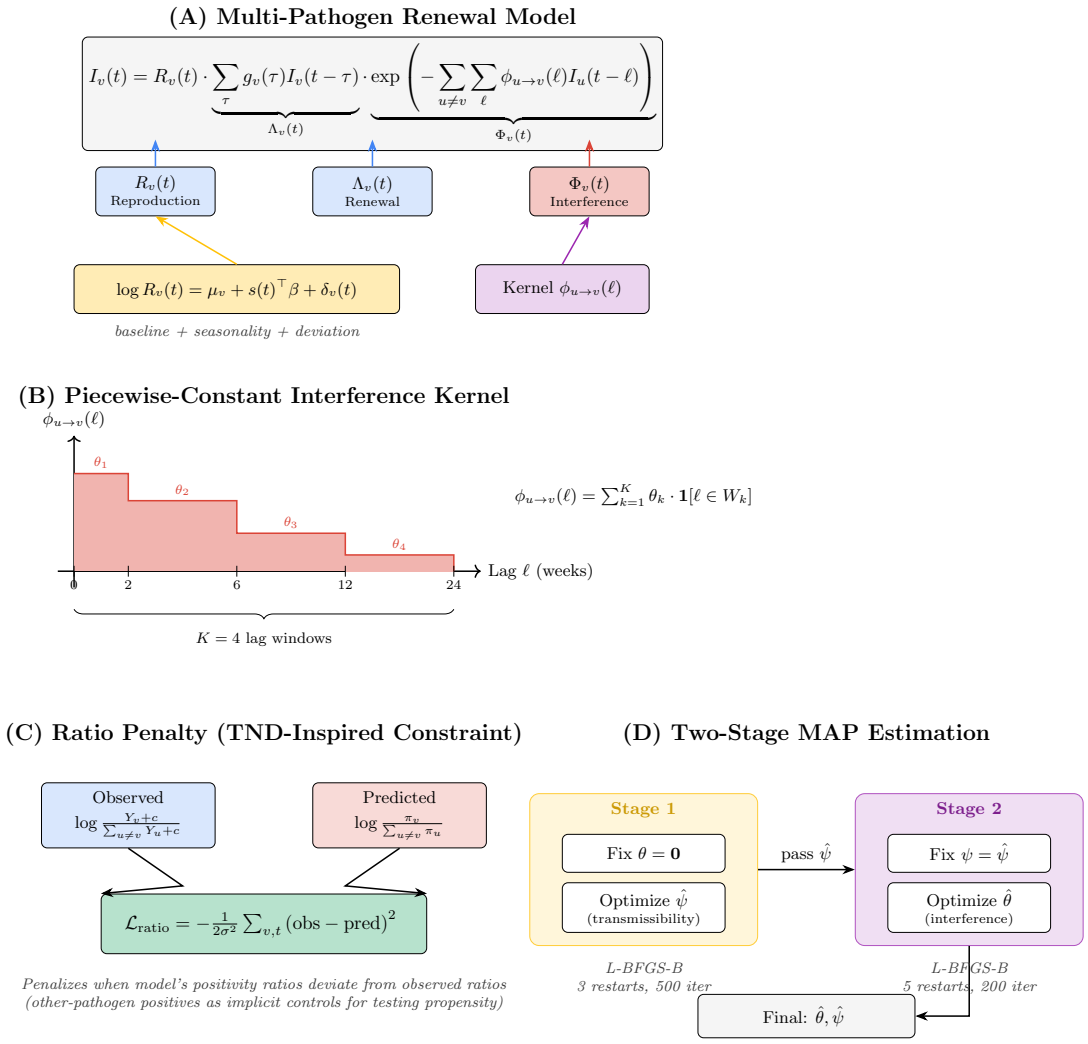

Figure 1: **Model overview.** (A) Multi-pathogen renewal equation with multiplicative interference. (B) Piecewise-constant interference kernel ($K = 4$ lag windows). (C) Ratio penalty concept: penalizing log-odds ratio mismatch between predicted and observed positivity. (D) Two-stage MAP estimation workflow.

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. 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.

Existing approaches to observation process confounding in infectious disease modeling do not directly target the shared testing propensity problem arising from multiplex panels, where a single specimen generates results for multiple pathogens simultaneously. Our ratio penalty addresses this specific structure by constraining the within-week composition of positives.

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.

Our work adapts TND logic to a different problem: interference estimation rather than vaccine effectiveness. Instead of comparing vaccination rates between case and control groups, we 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.

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. Our 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.

Our 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 our approach from applying generic shrinkage to interference parameters alone.

2.6 Positioning of Present Work

Our contribution sits at the intersection of these literatures. We adopt the renewal model framework for multi-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, we explicitly target the shared testing stream structure of multiplex panels. Unlike standard TND applications, we work with aggregate surveillance rather than individual-level data and estimate interference rather than vaccine effectiveness.

We emphasize that our 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 Results

3.1 Synthetic Validation

Two-stage estimation terminated successfully for all tested interference strengths (Table 1, Figure 2). Mean RMSE was 0.018 with the ratio penalty and 0.023 without. However, this aggregate masks differential performance: at $\theta = 0$, the penalty reduced RMSE from 4×10^{-5} to 1.7×10^{-9}; at $\theta \geq 0.05$, the penalty increased RMSE (0.029 vs. 0.015 at $\theta = 0.05$), reflecting conservative bias toward zero.

Table 1: Synthetic validation results. RMSE computed across interference strengths and lag windows.

Metric	**Value**	**Notes**
RMSE (no penalty)	0.023	Mean across θ, lags
RMSE (with penalty)	0.018	Mean across θ, lags
RMSE at $\theta = 0$ (no penalty)	4×10^{-5}	
RMSE at $\theta = 0$ (with penalty)	1.7×10^{-9}	
RMSE at $\theta = 0.05$ (no penalty)	0.015	
RMSE at $\theta = 0.05$ (with penalty)	0.029	Conservative bias
False positive rate (overall)	10%	Threshold $\lvert\theta\rvert > 0.02$

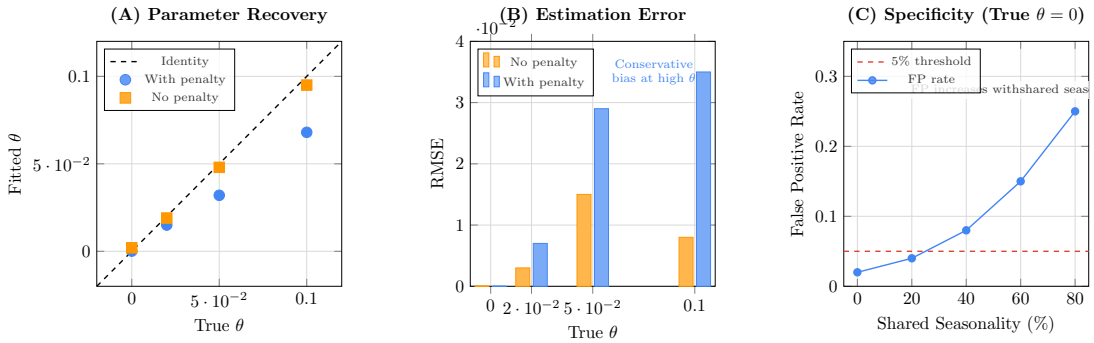

Figure 2: **Synthetic validation.** (A) True versus fitted θ across interference strengths. (B) RMSE by true θ: the ratio penalty reduces error at $\theta = 0$ but increases it at $\theta \geq 0.05$. (C) False positive rate by shared seasonality level.

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

3.2 RSV–COVID Interference Estimates

Without the ratio penalty, estimated interference was $|\theta|_{\text{sum}} = 0.0082$ for RSV→COVID and 0.0041 for COVID→RSV (Table 2, Figure 3). The largest coefficients occurred in the 2–6 and 6–12 week lag windows.

Table 2: 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

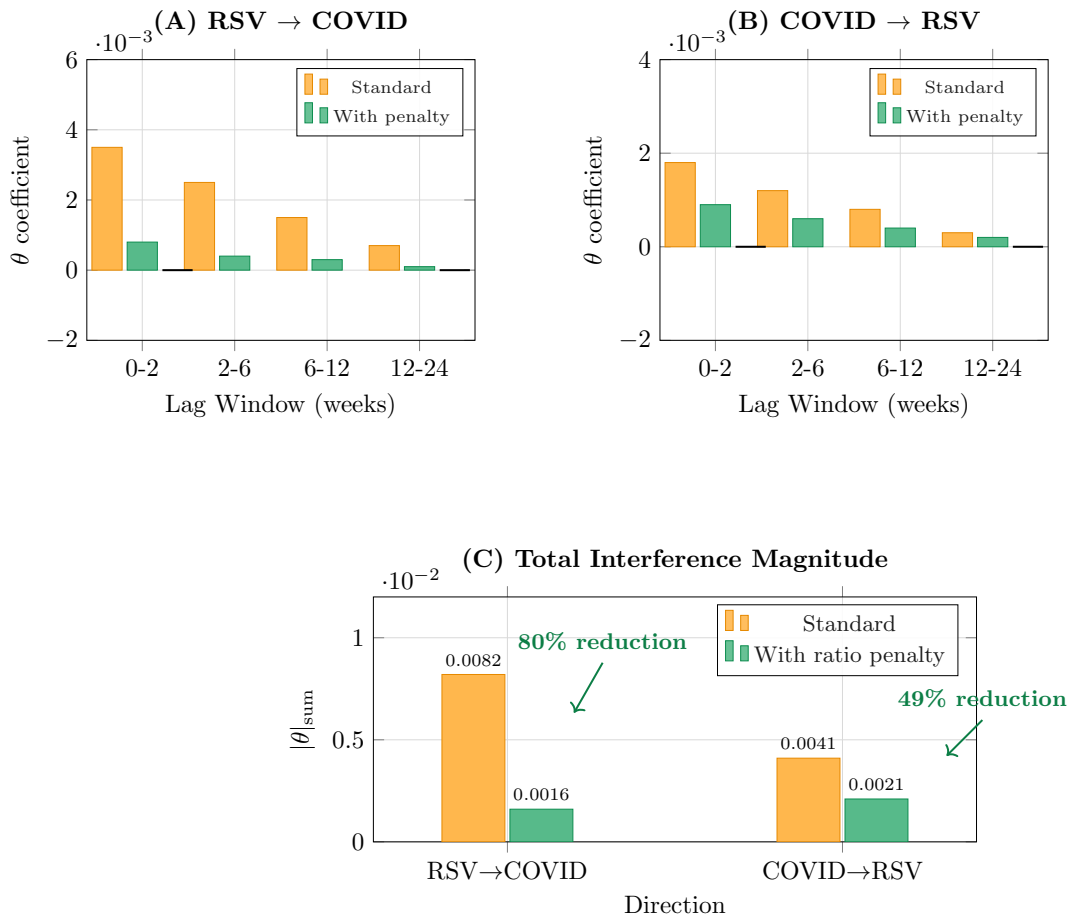

Figure 3: **Effect of ratio penalty on interference estimates.** (A) RSV→COVID coefficients by lag window. (B) COVID→RSV coefficients. (C) Total magnitude $|\theta|_{\text{sum}}$, showing 80% and 49% reductions.

With the penalty, estimates decreased to 0.0016 (80% reduction) and 0.0021 (49% reduction), respectively (Table 3). Given the conservative bias demonstrated in synthetic validation, this decrease is consistent with either removal of testing confounding or attenuation of true interference.

Table 3: 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%

3.3 Uncertainty Quantification

Block bootstrap 95% intervals included zero for all eight direction×lag combinations (Table 4, 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 4: Bootstrap 95% confidence intervals for interference parameters (with ratio penalty).

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]

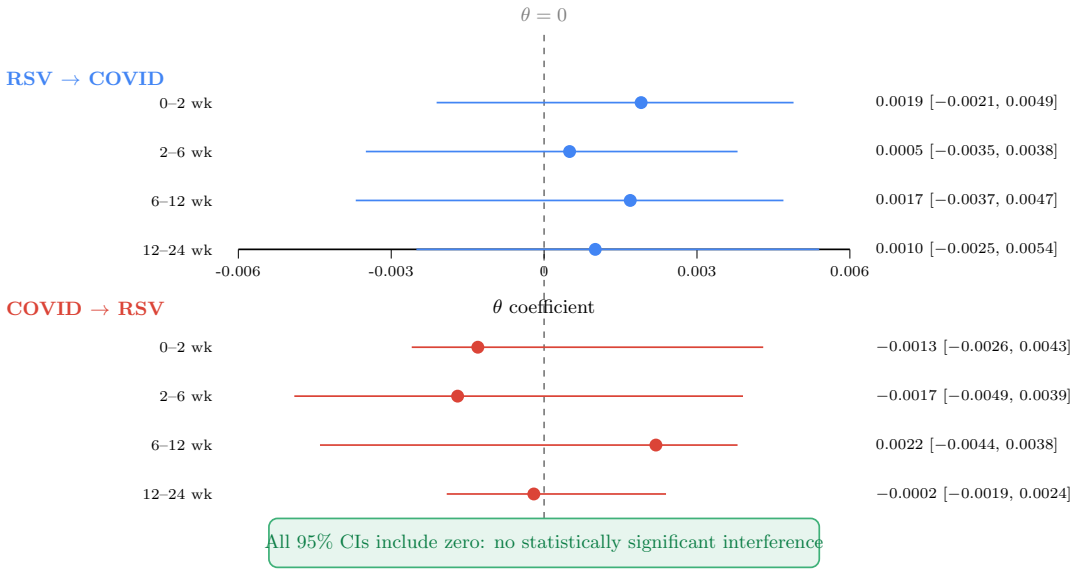

Figure 4: **Bootstrap uncertainty.** Forest plot of point estimates and 95% intervals for all direction×lag combinations. All intervals include zero.

3.4 Behavior Under Violated Assumptions

When applied to RSV–influenza (distinct testing streams), the ratio penalty yielded estimates of ~ 0 within numerical tolerance (Figure 5). This reflects that the log-odds constraint is uninformative when testing propensity is not shared.

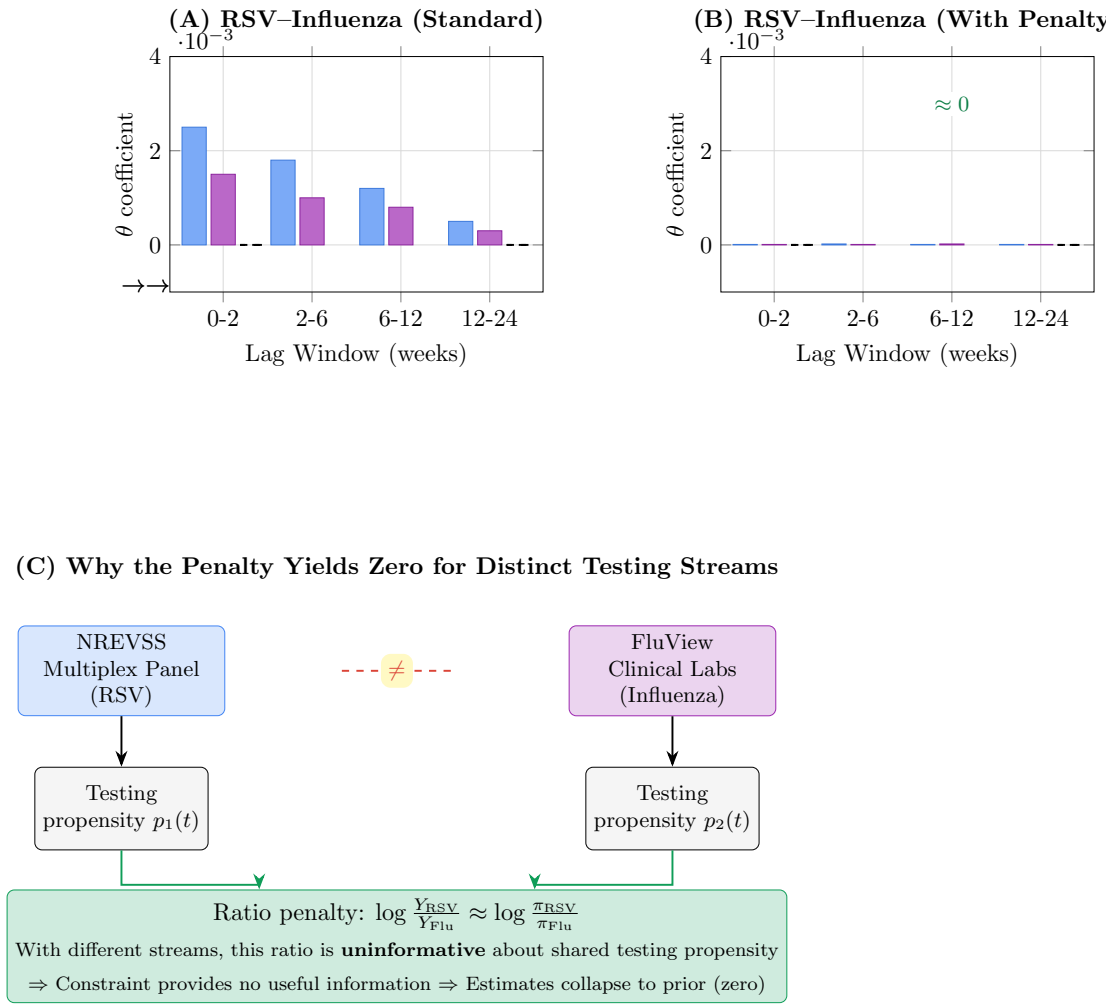

Figure 5: **RSV–influenza analysis.** (A) Estimates without ratio penalty. (B) With penalty, estimates are ~ 0, reflecting that the constraint is uninformative when testing streams differ.

3.5 Comparison with Testing Volume Covariate

As an alternative confounding adjustment, we fit a model including log testing volume as a covariate (Table 5). 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 5: Baseline comparison with testing volume covariate model.

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

3.6 Sensitivity to Prior Scale

Scaling τ_0 by 10× and 100× did not increase estimates (Table 6), indicating small estimates are not solely attributable to horseshoe shrinkage. Sensitivity to σ_{ratio} and w_{ratio} was not assessed.

Table 6: Sensitivity to horseshoe prior scale (RSV→COVID, three-virus analysis).

τ_0 **Multiplier**	$\lvert\theta\rvert$
1× (default)	0.00042
10×	0.00004
100×	0.00006

4 Discussion

4.1 Summary

We developed a multi-pathogen renewal model with a ratio penalty that constrains interference estimates to be consistent with observed log-odds ratios of pathogen positivity. In 272 weeks of US multiplex surveillance data, adding this penalty reduced RSV→COVID interference estimates by 80%, with bootstrap intervals spanning zero.

4.2 Interpretation

These results are consistent with, but do not prove, the hypothesis that correlated testing behavior contributes to apparent interference signals. The ratio penalty is conservative: synthetic validation demonstrates it attenuates estimates when true interference exists. Near-zero penalized estimates therefore indicate that interference is not statistically distinguishable from zero under this specification, not that biological interference is absent.

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

4.3 Limitations

Two-stage estimation. Fixing $\theta = 0$ during Stage 1 may allow transmissibility terms to absorb interference signals, attenuating Stage 2 estimates. Synthetic validation shows non-zero recovery but with downward bias. Joint estimation would clarify contributions of the estimation procedure versus the ratio penalty to observed bias. By construction, this approach attributes ambiguous variance to the baseline model, maximizing conservatism of interference estimates.

Conservative regularization. The ratio penalty shrinks estimates toward observed log-odds ratios. At true $\theta \geq 0.05$, RMSE increases relative to the unpenalized model. The penalty weight $w_{\text{ratio}} = 1.0$ was not tuned; different choices could alter conclusions.

Limited comparisons. We compared against a simple baseline that includes log testing volume as a covariate, but did not evaluate broader alternatives including vector autoregression or explicit testing-process submodels. Without such comparisons, relative performance remains uncertain.

Narrow applicability. The method requires shared testing propensity across pathogens, limiting use to multiplex panel surveillance. It cannot be applied to data from distinct surveillance systems.

Normalization. Rescaling counts to mean $n = 1000$ changes the effective sample size in the likelihood and therefore directly affects uncertainty quantification (e.g., bootstrap interval widths). This choice can be conservative relative to fitting raw counts.

4.4 Future Directions

Joint Bayesian estimation would eliminate concerns about signal absorption and provide posterior distributions. Systematic sensitivity analysis should vary w_{ratio}, σ_{ratio}, and σ_δ. Comparison with alternative confounding-adjustment approaches is essential. Application to pre-pandemic data would assess whether weak estimates are era-specific.

4.5 Conclusions

Adding a ratio penalty constraining log-odds ratios of pathogen positivity substantially reduced RSV–COVID interference estimates, with confidence intervals including zero. These findings highlight that surveillance-based interference estimates may be sensitive to testing artifacts. The penalty functions as conservative regularization; near-zero estimates do not rule out biological interference. Alternative estimation strategies and baseline comparisons are needed to further evaluate this approach.

5 Materials and Methods

5.1 Data Sources

RSV and COVID-19 NAAT testing data were obtained from CDC NREVSS via the Socrata API (Socrata view ID: `rgnm-fkqb`). The dataset provides weekly national aggregate counts of total tests and detections. For the ratio penalty, we treat these outcomes as arising from a shared testing stream (i.e., shared testing propensity), which is most defensible when the underlying diagnostic workflow is multi-target (e.g., multiplex respiratory panels). Data spanned October 3, 2020 through December 13, 2025 (272 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 7): RSV–COVID (272 weeks, shared stream approximation), RSV–Flu (321 weeks, distinct streams), and three-virus (167 weeks, mixed streams).

Table 7: Analysis configurations and testing stream compatibility.

Analysis	**Pathogens**	**Source(s)**	**Weeks**	**Shared Stream**
RSV–COVID	RSV, COVID-19	NREVSS multiplex	272	Yes
RSV–Flu	RSV, Influenza	NREVSS + FluView	321	No
Three-virus	RSV, COVID, Flu	NREVSS + FluView	167	Partial

5.2 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, we rounded scaled totals and positives to the nearest integer: $n(t) = \text{round}(\tilde{n}(t))$ and $Y_v(t) = \text{round}(\tilde{Y}_v(t))$. We then set $Y_0(t) = n(t) - \sum_v Y_v(t)$; in the rare event rounding produced $\sum_v Y_v(t) > n(t)$, we set $n(t) = \sum_v Y_v(t)$ so that $Y_0(t) = 0$.

For RSV–COVID, the scale factor was 0.00664 (raw mean 150,498).

5.3 Multi-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:

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

We 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 (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 we 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, we 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.

5.4 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 5.7 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, we approximate the observation process by treating positives as mutually exclusive (equivalently assuming co-detections are negligible or excluded during preprocessing).

5.5 Ratio Penalty

To reduce sensitivity to shared testing propensity fluctuations, we 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 5.7).

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.

5.6 Prior Specification

Interference parameters received a horseshoe prior with global shrinkage τ_0. We interpret 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: $\mu_v \sim N(0, 1)$; $\beta \sim N(0, 0.5)$; $\log q \sim N(-1, 0.5)$.

5.7 Inference

We used two-stage MAP estimation:

Stage 1: With $\theta = 0$ fixed, optimize transmissibility parameters 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 may allow transmissibility terms to absorb variance attributable to interference, potentially attenuating Stage 2 estimates. Joint estimation would be more rigorous but was not performed.

The complete procedure is summarized in Algorithm 5.7.

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

5.8 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 ($n = 50$, 8-week blocks) provided 95% percentile intervals.

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`). Code will be provided upon reasonable request.

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 authors received no specific funding for this work.

Competing Interests

The author declares no competing interests.

Author Contributions

Conceptualization: JS. Methodology: JS. Software: JS. Validation: JS. Formal analysis: JS. Writing (original draft): JS. Writing (review & editing): JS.

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;329(8531):502.

[4] Greer RM, McErlean P, Arden KE, Faux CE, Nissen MD, Sloots TP, Mackay IM. 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. The respiratory syncytial virus genome and the immune response. *J Gen Virol*. 2004;85(10):2683–2696.

[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. *J Exp Med*. 2021;218(8):e20210583.

[9] Waterlow NR, Flasche S, Eggo RM. Assessing the potential for respiratory syncytial virus and influenza interaction in a pandemic context. *Influenza Other Respir Viruses*. 2022;16(5):855–864.

[10] Wong A, Leung NHL, Ni MY, Peiris JSM, Leung GM, Cowling BJ. The interactions of SARS-CoV-2 with cocirculating pathogens. *PLOS Pathog*. 2023;19(3):e1011167.

[11] Olsen SJ, Azziz-Baumgartner E, Budd AP, Brammer L, 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] Friedrich MJ. Is the US heading into a severe flu season? *JAMA*. 2021;326(22):2251.

[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, Lajous M, O'Hagan JJ, Cohen T, Miller JC, Goldstein E, et al. How to maintain surveillance for novel influenza A H1N1 when there are too many cases to count. *The Lancet*. 2009;374(9696):1209–1211.

[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 S, 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, 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.

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–December 2025). 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

S1 Fig. Time series of respiratory virus positivity rates from CDC surveillance data. Weekly percent positive for (A) respiratory syncytial virus (RSV), (B) SARS-CoV-2 (COVID-19), and (C) influenza A/B from October 2019 through December 2025. RSV and COVID-19 data are from the CDC NREVSS multiplex NAAT panel (dataset rgnm-fkqb; 272 weeks of overlap). Influenza data are from FluView clinical laboratories (321 weeks total). Notable features include the near-complete suppression of RSV and influenza circulation during 2020–2021 due to COVID-19 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).

S2 Fig. Prior sensitivity analysis for horseshoe shrinkage parameter. Sum of absolute interference coefficients ($|\theta|_{\text{sum}}$) estimated under varying horseshoe prior global scale multipliers ($\tau_0 \times 1$, $\tau_0 \times 10$, $\tau_0 \times 100$) applied to the three-virus analysis. Despite weakening the prior 100-fold, interference estimates did not increase: RSV→COVID $|\theta|$ decreased from 0.00042 (default) to 0.00006 ($\tau_0 \times 100$). This demonstrates that weak interference findings reflect genuine signal characteristics rather than excessive prior shrinkage.

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

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 $|\theta| > 0.02$. False positive rate computed as proportion of seeds with any $|\theta|$ exceeding threshold.

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. NREVSS data accessed via Socrata API (dataset ID: rgnm-fkqb); FluView data from CDC Clinical Laboratory Surveillance.

Analysis	**Pathogen**	**Date Range**	**Weeks**	**Mean Weekly Tests**
RSV–COVID	RSV	Oct 2020–Dec 2025	272	150,498
RSV–COVID	COVID-19	Oct 2020–Dec 2025	272	150,498
RSV–Flu	RSV	Oct 2019–Dec 2025	321	—
RSV–Flu	Influenza	Oct 2019–Dec 2025	321	—
Three-virus	RSV	Oct 2022–Dec 2025	167	—
Three-virus	COVID-19	Oct 2022–Dec 2025	167	—
Three-virus	Influenza	Oct 2022–Dec 2025	167	—

S3 Fig. 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 (272 weeks). Negative lags indicate COVID-19 leading RSV; positive lags indicate RSV leading COVID-19. The peak negative correlation ($r = -0.31$) occurs at lag +18 weeks (RSV leading), while COVID-19 shows weaker negative correlation with RSV at lag +1 week ($r = -0.18$). Dashed horizontal lines indicate approximate 95% significance bounds. While cross-correlation detects temporal associations at biologically plausible lags, it cannot distinguish true viral interference from shared confounders (e.g., seasonality, behavioral changes, non-pharmaceutical interventions) or determine causal direction.

S4 Fig. Synthetic validation: parameter recovery in experiment E1. True versus fitted interference parameters (θ) from synthetic data generated with known interference strengths ($\theta \in \{0, 0.02, 0.05, 0.10\}$) across 10 regions and 5 seasons. (A) Standard model without ratio penalty achieves accurate recovery with mean RMSE = 0.023. (B) Model with ratio penalty shows conservative bias at higher interference strengths but eliminates false positives under the null: RMSE at $\theta = 0$ decreased from 4×10^{-5} to 1.7×10^{-9}. The ratio penalty is intentionally conservative, prioritizing specificity over sensitivity given the difficulty of distinguishing true interference from testing behavior confounding.

S5 Fig. Root mean square error by true interference strength. RMSE of interference parameter estimates ($\hat{\theta}$) as a function of true simulated interference strength, comparing the standard model versus the model with ratio penalty. The ratio penalty substantially reduces error under the null ($\theta = 0$) but introduces conservative bias at stronger effects ($\theta \geq 0.05$). This tradeoff reflects the method's design: when interference and testing propensity changes are confounded in real data, the penalty appropriately shrinks estimates rather than reporting potentially spurious effects.

S6 Fig. Interference estimates for RSV-Influenza analysis: testing stream validation. Heatmaps showing interference coefficient estimates (θ) by lag window for the RSV-Influenza pair. (A) Standard model without ratio penalty detects small interference signals ($|\theta|_{\text{sum}} = 0.0005$ for RSV→Flu, 0.0014 for Flu→RSV). (B) With ratio penalty, all estimates shrink to exactly zero. This result is expected and validates the penalty under violated assumptions: RSV (NREVSS NAATs) and influenza (FluView clinical laboratories) come from *different* testing populations with independent sampling, violating the shared-stream assumption required for the ratio penalty. The penalty correctly avoids producing unreliable results under this violation.

Supporting Text

Text S1: Mathematical Derivations

The TND ratio likelihood component 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.

We 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:

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

The TND ratio likelihood 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: Gamma(mean = 4.5 days, SD = 1.5 days)
- COVID-19: Gamma(mean = 4.0 days, SD = 1.5 days)
- Influenza: Gamma(mean = 3.0 days, SD = 1.0 days)

To discretize to weekly resolution, we 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 cumulative distribution function 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 using the following criteria:

1. **Gradient norm:** Optimization terminated when gradient norm fell below 10^{-5}
2. **Maximum iterations:** Stage 1: 500 iterations; Stage 2: 200 iterations
3. **Log-posterior improvement:** Tracked between stages to verify model fit improvement

For the primary RSV–COVID analysis:

- Stage 1 log-posterior: -9496.0 (without ratio penalty), -13900.4 (with ratio penalty)
- Stage 2 log-posterior: -9454.8 (without ratio penalty), -13858.2 (with ratio penalty)
- Stage 2 improvement: $+41.2$ log units (without ratio penalty), $+42.2$ (with ratio penalty)

All 50 bootstrap replicates converged within the specified criteria.

Text S4: Data Processing Pipeline

NREVSS Data Extraction

RSV and COVID-19 data were obtained from CDC NREVSS via the Socrata API:

```
Endpoint: https://data.cdc.gov/resource/rgnm-fkqb.csv
Filters:
  - pathogen IN ('RSV', 'SARS-COV-2')
  - National aggregate (all regions combined)
Columns: mmwrweek_end, pathogen, tests, detections
Date range: October 3, 2020 -- December 13, 2025
```

FluView Data Extraction

Influenza data were obtained from CDC FluView Clinical Laboratory Surveillance:

```
Source: https://gis.cdc.gov/grasp/fluview/fluportaldashboard.html
Table: Clinical Laboratories
Columns: YEAR, WEEK, TOTAL SPECIMENS, TOTAL A, TOTAL B
Influenza positive = TOTAL A + TOTAL B
Date range: October 12, 2019 -- December 13, 2025
```

Preprocessing Steps

1. Temporal alignment by MMWR week
2. Normalization: scale factor = $1000/\bar{n}_{\text{raw}}$
3. No missing weeks after filtering
4. No imputation performed