Online Appendix

I Robustness and supplementary constructions

Table A2 consolidates the falsification and robustness checks summarized in Sections 8–10, each re-estimating the headline land-crossing-volume gradient under an alternative specification, sample, or outcome transformation. Across the battery the estimate stays between about −0.30% per standard deviation (the triple-difference, its conservative floor) and −0.71% (with tariff and manufacturing-share controls). It is stable in sign and order of magnitude throughout. The subsections that follow give the detail behind the checks Sections 8–10 state in summary form—regional controls, the metro's own tariff shock, seasonality and the pre-trend, the distance-decay bandwidth grid, the administrative air exposure, the synthetic-control case studies, the market-potential gradient, and the job-count constructions—followed by the identifying geography.

I.A Regional controls

The land–air contrast could in principle be regional rather than modal: every high-land-exposure metro sits on the northern tier and every large air destination in the Sun Belt, so a slow southward drift of employment growth could masquerade as a boycott gradient. Forcing identification within region—division×month fixed effects, so that Buffalo is compared with other Mid-Atlantic metros in the same month—retains three-quarters of the land-volume gradient, −0.44% per standard deviation (permutation $p = 0.002$, cluster-robust $p = 0.057$). The proximity gradient does not survive this cut (−0.14%, insignificant), so within-region variation supports the volume measure and we read proximity as corroboration only. A state×month specification absorbs most of the identifying variation outright, exposure being spatially clustered within states; it returns a noisy −0.10% whose confidence interval comfortably contains the baseline, so that test is uninformative rather than contradicting. The air gradient stays nonnegative under every regional control.

Spatial correlation in the northern-tier exposure also bears on inference directly. Because the design is a single national shock interacted with a spatially clustered cross-section, three complementary checks guard against the over-rejection that metro-level clustering can produce. (i) Conley spatial-HAC standard errors (Conley, 1999), allowing arbitrary spatial correlation to a 300 km cutoff and serial correlation to six months, leave the land-volume gradient significant ($p = 0.016$ at 300 km, $p = 0.003$ at 200 km). (ii) Clustering at the state rather than the metro level, which permits arbitrary within-state correlation, gives a cluster-robust $p = 0.005$ and a wild-cluster-bootstrap $p = 0.034$ over 52 state clusters (Cameron et al., 2008); at the coarser nine-division level the wild bootstrap is insignificant ($p = 0.12$, few enough clusters that the test has little power, MacKinnon and Webb, 2018). The division-level sample excludes Puerto Rico's four metros, which belong to no Census division, and on that slightly smaller sample the coefficient is −0.45%. (iii) A spatially-blocked permutation, which permutes exposure at the level of contiguous spatial blocks so that each placebo draw is as spatially clustered as the real exposure. Because the choice of block definition matters, we report the full grid: the level gradient returns $p = 0.0025$ over forty spatial blocks, 0.0045 over twenty, 0.013 over ten, 0.006 with states as blocks, and 0.024 with

census divisions, clearing conventional levels under every partition. The same test applied to the within-metro triple-difference returns $p = 0.10$ to 0.18, against 0.046 under an i.i.d. permutation; the triple-difference is identified off a within-metro contrast but its exposure varies across metros, and that variation is spatially clustered. Finally, because domestic spending displaced from the border could contaminate nearby control metros, dropping low-exposure controls within 150 km of a high-exposure metro leaves the gradient at -0.61%, if anything slightly larger, so any such contamination biases the estimate toward zero.

I.B The metro's own tariff shock

The most-exposed land metros are also among the most tariff-exposed. Detroit is the second-largest crossing metro and the center of the auto-tariff shock, so a land-border decline could in principle be a local manufacturing recession reaching into services. Two checks argue against it. Adding a metro's predetermined tariff exposure and its manufacturing share, each interacted with post, does not attenuate the land coefficient; it moves from -0.57% to -0.71% per standard deviation. Dropping each border metro in turn leaves the estimate in a tight band between -0.59% and -0.55%, all significant, and leave-out-Detroit returns -0.58%, essentially the baseline. The continuous gradient does not depend on the one metro where the tariff confound is strongest, which is itself an advantage of the continuous design over a binary one.

I.C The bilateral exchange rate

Section 9 rules out a bilateral Canadian-dollar move as the driver of the U.S.-specific travel decline; Figure A3 gives the underlying series. The loonie depreciated about 4.8 percent against the U.S. dollar between the 2024 average and January–March 2025, but only 1.9 percent against the euro over the same window, so the onset move was largely dollar-specific rather than a broad Canadian weakening. It was also transitory: the bilateral rate had returned to its 2024 level by the first quarter of 2026, and over the March–December 2025 decline window the average depreciation was only 1.5 percent. At the day-trip real-exchange-rate elasticities of Chandra et al. (2014)—about 0.9 in absolute value when the U.S. dollar is strong, as in 2025—even crediting the full onset depreciation as permanent implies a trip decline in the low single digits, an order of magnitude short of the observed 23.7 percent. The currency also runs the wrong way for persistence: through early 2026 crossings held about 20 percent below their 2024 level while the bilateral rate was back to par, and later in 2025 the loonie weakened against the euro even as Canadians raised their overseas travel. This is consistent with the border-county evidence that real-exchange-rate movements of this size move cross-border retail modestly (Campbell and Lapham, 2004).

I.D Third-country arrivals

Section 9 distinguishes a Canada-specific land withdrawal from a general aversion to U.S. travel using NTTO arrivals to the United States by country of residence. Over the post-onset window (March 2025 through April 2026 against the same months a year earlier), Canadian all-mode arrivals fell 20.2 percent, the third-largest decline among the 38 origins with at least 200,000 annual arrivals

and about six times the overseas mean of 3.4 percent; Mexican and Japanese arrivals rose. The two larger declines, Denmark (−26.3 percent) and Venezuela (−20.4 percent), are origins with U.S.-specific frictions of their own over the same months, so the cross-origin ranking corroborates the U.S.-specificity of the Canadian pullback without adjudicating its motive. The Canada-specificity is a statement about *magnitude* and about the land channel, not about air travel in isolation: a broad but shallow softening of U.S. inbound air travel is visible across origins, and Canada's own air arrivals, down 11.4 percent, sit inside that band rather than above it. The gap between Canada's 20.2 percent all-mode decline and its 11.4 percent air decline is the land day-trip collapse, a margin no overseas origin shares. Because the employment gradient is identified off land-crossing exposure and overseas air origins do not cross land borders, the general air softening cannot contaminate the land estimate; it is instead consistent with the air null of Section 8. The test cannot separate the boycott from other U.S.-specific land-crossing deterrents such as border-enforcement apprehension, which is why Section 9 identifies the redirection rather than a single motive.

I.E Sectoral composition and the regional confound

Section 7 reports that the metro-level gradient is not tourism-specific. Running the headline specification on each CES supersector in turn makes the point concrete; Table 5 reports the full battery. Unconditionally the gradient is broad: wholesale trade (−0.84% per standard deviation, permutation $p = 0.005$) and financial activities (−0.67%, $p = 0.003$) show the two largest declines, both larger in absolute value than retail trade (−0.36%) and neither reachable by a cross-border day-trip. Government is flat (+0.06%, $p = 0.72$). A local multiplier running out from the visitor economy would concentrate in retail, transport, and other services and leave wholesale and finance alone; it does not, which points to a regional trend common to the Great Lakes industrial corridor.

Census-division-by-month effects remove it. The non-tourism aggregate falls from −0.28% ($p = 0.001$) to −0.12% ($p = 0.33$), while leisure and hospitality holds at −0.45% (permutation $p = 0.003$). We do not interpret individual sector cells within division, for two reasons. Thirteen sectors estimated under two specifications produce wrong-signed significant coefficients at about the rate chance predicts. And the division-by-month control is expensive: it absorbs roughly three-fifths of the identifying variation in exposure, against about one-tenth for a control that groups metros by the border segment they face. Two of the sign reversals it produces—financial activities from −0.67% to +0.14%, professional and business services from −0.12% to +0.45%—are as consistent with that variance loss as with the removal of a confound, and we do not read them as evidence either way. The two facts we rely on are the breadth of the unconditional gradient and the tourism-versus-rest split once regional trends are absorbed.

The two sectors carrying the unconditional gradient are worth characterizing, since they are the clearest evidence that it is regional. Both declines pre-date the episode: on placebo onsets inside the pre-boycott period the wholesale and finance gradients are −0.59% and −0.67%, the latter equal to its own post-onset value, and over 2015–2019 wholesale is −1.03%. Neither is the metro's own tariff shock: a sector-loaded tariff control leaves both slightly more negative, though that control finds no employment gradient in manufacturing and so bounds the tariff channel weakly. The wholesale gradient is also not a goods-trade channel. Holding allocated crossing volume fixed and varying only the freight intensity of the ports a metro draws on, metros facing more truck-intensive crossings

Table 5. The Land-Crossing Gradient Sector by Sector

		Metro + month FE		Metro + division×month FE	
Sector	Share, %[b]	Coef.	Perm. p	Coef.	Perm. p
Treated sector					
Leisure & hospitality	10.5	−0.57	< 0.001	−0.45	0.003
Day-trip visitor channel					
Retail trade	9.3	−0.36	0.003	−0.20	0.092
Transportation, warehousing & utilities	4.6	−0.60	0.100	+0.25	0.479
Other services	3.8	−0.14	0.494	+0.65	0.002
Ambiguous channel					
Wholesale trade	3.9	−0.84	0.004	−0.55	0.046
Mining, logging & construction	4.9	+0.33	0.264	−0.42	0.141
No day-trip channel					
Financial activities	6.2	−0.67	0.002	+0.14	0.495
Information	2.2	−0.28	0.612	−0.44	0.451
Professional & business services	15.8	−0.12	0.620	+0.45	0.072
Education & health services	17.6	−0.36	0.072	−0.70	< 0.001
Manufacturing	6.8	−0.30	0.198	+0.52	0.029
Government	13.8	+0.06	0.720	−0.09	0.602
Memo					
Trade, transportation & utilities[a]	17.9	−0.48	0.001	−0.12	0.338

Notes: Each cell is a separate estimate of the headline land-crossing-volume specification with the named sector's log employment as the dependent variable: coefficient on exposure × post per +1 standard deviation of exposure, in percent, with the exposure-permutation p beside it. Sector panels group supersectors by whether a cross-border day-trip has a direct channel into them. Sector coverage in the published metro CES series varies, so samples range from 311 to 352 metros (12,751 to 14,432 metro-months). Twenty-six cells estimated under two specifications will produce wrong-signed significant coefficients at roughly the rate chance predicts, and we do not interpret individual within-division cells (see text). Cluster-robust p-values and sample sizes per cell are in `ext_D_supersector_incidence.csv`.

[a] Aggregate of wholesale trade, retail trade, and transportation, warehousing & utilities.

[b] The sector's 2024 share of metro nonfarm employment.

show *higher* wholesale employment (+1.0% per standard deviation, $p = 0.0004$), the opposite sign to a freight-exposure mechanism. It is not composition either: exposure is essentially uncorrelated with the wholesale employment share ($r = 0.07$), and the share does not move. At the three-digit level the decline sits in durable-goods wholesaling rather than agents and brokers, where a customs-brokerage channel would appear. It weakens to −0.36% ($p = 0.37$) within division by month. We carry it as a regional gradient of unidentified origin that bears on the level specification and not on the within-metro comparison.

The same ladder applies to the mode contrast, and it survives: the land-minus-air difference is −0.60 percentage points ($t = -4.1$) unconditionally and −0.52 points ($t = -2.3$, permutation $p = 0.007$) within division. Under state-by-month effects the difference is −0.12 points with a standard error of 0.64; with exposure clustered inside states, that specification's confidence interval contains both zero and the baseline estimate and so does not discriminate between them. The pooled single-index gradient, by contrast, falls from −0.36% to −0.02% within division, so the cost of pooling land and air rises rather than falls under the more demanding comparison.

I.F Seasonality and the pre-trend

The metro series are not seasonally adjusted, and the dynamic path on the raw series carries residual seasonality. Deseasonalizing log leisure and hospitality employment metro by metro with an STL decomposition (Cleveland et al., 1990) and re-estimating leaves the land-volume gradient at -0.55% ($p < 0.001$) and proximity at -0.34% ($p = 0.010$), with air still null; the deseasonalized gradient is, if anything, slightly larger than the unadjusted one, so the decline is not an artifact of cold-weather border seasonality. A specification that avoids the decomposition entirely points the same way: adding metro-by-calendar-month fixed effects, which absorb each metro's own seasonal profile without estimating a decomposition, leaves the land-volume gradient at -0.58% (permutation $p < 0.001$) and proximity at -0.36%, with air null. The forty-one-month panel spans only a little over three seasonal cycles, so we read the STL pass as corroboration that the raw and adjusted estimates coincide rather than as a free-standing seasonally adjusted series.

Two formal checks show why the cross-sectional estimate, not the dynamic path, must carry the result. First, an in-time placebo: re-dating onset to March 2024, a year before the boycott, and re-estimating on the true pre-period produces no exposure gradient on the deseasonalized series (land volume between -0.0012 and -0.0015, permutation $p \approx 0.25$–0.34), so there is no spurious gradient at an arbitrary pre-boycott date.[25] Second, a Rambachan–Roth sensitivity confirms that the dynamic path cannot carry the result on its own. The pre-period exhibits the converge-then-break pattern familiar from import-competing geographies: high-exposure metros sit on a relatively higher, declining path before onset, and the post-onset coefficient breaks downward, reaching about -0.49% by the end of the window, continuing through the February 2026 tariff reset. The average post-onset effect is therefore sensitive to a relative-magnitudes pre-trend violation (Rambachan and Roth, 2023), which is why we read the headline as a descriptive association and the job counts as an upper bound (Section 10). The same fragility appears parametrically: adding an exposure-specific linear trend over the full window, which lets high-exposure metros follow their own declining path throughout, absorbs the post-onset break, leaving a post coefficient near zero ($+0.12\%$, insignificant) with a negative trend term. That is expected given the converge-then-break shape, and it is precisely why identification rests on the designs that do not share the trend: the triple-difference, whose metro-by-month fixed effects absorb any metro-specific trend and which still returns -0.30% (Section 7), and the placebos. Figure A1 displays that path, with the air gradient flat at zero throughout.

The Canadian-side series carry their own lead diagnostics: the Wald tests reported with Figure 1 reject strict flatness of the 2024 leads ($p = 0.03$ for the United States–overseas contrast, $p < 0.001$ for automobile–air), with the drift small relative to the post-onset break and, for the automobile–air contrast, opposite in sign to it.

[25] On the *raw* series the same placebo returns a significant *positive* gradient once the winter months are dropped from the pre-window—the differential summer seasonality of the border metros, not a treatment effect, and exactly the contamination the STL adjustment removes, which is why the placebo is read on the deseasonalized series.

I.G The triple-difference pre-trend record

Section 9 summarizes the pre-trend record of the within-metro triple-difference; this subsection reports it in full. On the deseasonalized series the immediate pre-onset leads are flat: a joint test over the six months before onset does not reject ($p = 0.14$ asymptotic; $p = 0.17$ under design-based inference that recomputes the test across exposure permutations, $p = 0.44$ state-blocked). A joint test over the full 25-lead window does reject ($p = 4.5 \times 10^{-6}$), driven by distant 2023 leads that form a small negative offset converging toward zero. That shape differs from the level design's converge-then-break in the direction that matters: the pre-lead slope is positive (+0.013 percent per month), so extrapolating it across the post window would predict a gradient near +0.20 percent, and trend continuation therefore understates rather than manufactures the post-onset break, which reaches about −0.57% per standard deviation by the end of the window and continues through the tariff reset (Figure A2). A relative-magnitudes sensitivity analysis (Rambachan and Roth, 2023) on this event study is uninformative in both directions and we do not lean on it: unrestricted, the average post-onset effect is not robust to violations of any size (breakdown $\bar{M} = 0$), and the restricted upper bounds are not monotone in $\bar{M}$—a pattern we read as grid noise rather than robustness—so we report no breakdown value. Two further checks bear on the tourism-specific trend concern directly. Replacing the sector-by-month fixed effects with division-by-sector-by-month effects, so the tourism-versus-rest differential trend may vary freely across census divisions, leaves the triple-difference at −0.32% (permutation $p = 0.035$). And the same design run on the southern border, where crossings did not decline, returns −0.11% (permutation $p \approx 0.45$): whatever secular tourism-versus-rest trend border metros share does not produce a gradient where the withdrawal did not occur. The proximity version of the triple-difference is right-signed but insignificant (−0.19 percent, $p = 0.17$), consistent with proximity's weaker showing throughout. The division-by-month comparison of Online Appendix I.A guards against regional confounds but uses the same cross-metro exposure variation.

I.H A five-year pre-boycott window

The estimation panel begins in 2023, allowing a two-year pre-period. The underlying Bureau of Labor Statistics state and area series run back to 2006, and we use 2015–2019—dropping 2020–2022 as pandemic-contaminated—to test the pre-period over a longer horizon.

Over 2015–2019, land-crossing exposure predicts declining employment in essentially every sector. Total nonfarm falls 0.67% per standard deviation (permutation $p < 0.001$); leisure and hospitality falls 0.79%, a trend of about 0.34 percentage points a year, which exceeds the post-onset gradient this paper estimates. The pattern is not driven by particular places or years. Leaving out each of the twenty-five most exposed metros in turn moves the leisure and hospitality slope within [−0.81%, −0.76%], and dropping Buffalo, Detroit, or Bellingham individually moves it by less than 0.01 points. Exposure deciles are monotone, with the most exposed decile growing 1.33 points a year more slowly than the least. Estimated year by year the gradient is steady and still widening at the end of the window (−0.59% in 2015, −0.67% in 2019), so it is not an artifact of the mid-decade commodity cycle.

The measure is a 2024 cross-section applied retroactively, but the geography of Canadian cross-

ing volume is close to static: rebuilding exposure from contemporaneous 2015–2019 volumes, with the same gravity allocation and decay bandwidth, reproduces the gradient (-0.73% for leisure and hospitality), and the two cross-sections correlate at 0.998. Trimming or winsorizing exposure leaves the result intact or strengthens it.

The tourism-minus-other-sectors differential behaves differently, which is why identification rests on the within-metro design rather than the level. Over the same five years the differential is -0.14% ($p = 0.35$), with the sign running against the post-onset estimate, and it is flat under every leave-one-out, trimming and winsorizing variant and on the leakage-free 2023–2024 window. Within the Great Lakes alone it is -0.45% ($p \approx 0.02$); grouping metros by the border segment they face does not absorb this, since a segment-by-month control has no leverage on a trend running inside the segment. In-time placebo windows end in December 2024 throughout, because the Canadian withdrawal is visible in the crossing data from January 2025 and later placebo windows would pick up the treatment. Of the design-based lead tests run at placebo onsets inside the window, one, at mid-2024, rejects flatness; the others do not, which is the second qualification noted in Section 10.

I.I Sensitivity to the distance-decay bandwidth

The two land exposures embed a distance-decay bandwidth, $\tau = 150$ km for the gravity allocation of crossing volume and $\tau = 200$ km for the proximity decay. Table A4 rebuilds each exposure across a grid of τ and re-estimates the headline gradient. The estimate is stable in sign, magnitude, and significance throughout: the volume gradient stays between -0.57% and -0.46% per standard deviation and the proximity gradient between -0.46% and -0.23%, all significant at the five percent level, so the result is not an artifact of the bandwidth choice. The volume gradient peaks near $\tau \approx 150$–200 km; the proximity gradient tightens as the window widens, because a longer decay length pulls more inland metros into the gradient.

I.J The air null under an administrative exposure

The air-destination share is a survey object, and classical measurement error in an exposure attenuates its coefficient toward zero, so the air null could in principle be noise in the regressor. We rebuild the exposure from administrative T-100 segment data covering every carrier, Canadian ones included: each metro's 2024 Canada-to-U.S. passengers at its airports, with each airport assigned to the metropolitan or micropolitan area whose boundary contains it (full-resolution TIGER/Line 2023 polygons; airports outside every such area carry 0.01 percent of Canadian traffic). Canada-to-U.S. passengers fell from 15.9 to 14.4 million between 2024 and 2025, or 9.3 percent, consistent with the I-94 arrivals decline once T-100's broader coverage is taken into account, since it counts every passenger on the segment rather than Canadian residents alone. The employment gradient on the administrative exposure is $+0.15\%$ per standard deviation (cluster-robust $p = 0.15$, permutation $p = 0.31$); on a Canadian-share-of-arrivals version, $+0.17\%$ (permutation $p = 0.24$; the cluster-robust p of 0.047 lies on the positive side, the opposite of a loss); and the pooled land–air difference re-estimated with the administrative measure is -0.77%, $t = -4.3$. The null also survives its support. Estimated only on the 46 metros with positive survey exposure it is $+0.20\%$ and insignificant; winsorizing Las Vegas to the second-highest share, or dropping it, leaves $+0.09\%$ and $+0.08\%$;

in raw units the gradient is +0.73% per ten percentage points of destination share, insignificant throughout. Nor is the null an artifact of averaging over a non-winter post window that would miss a delayed snowbird response: restricting the post period to the first full snowbird season, November 2025 through April 2026, leaves the air gradient statistically indistinguishable from zero on both the survey and the administrative measures (survey +0.26%, permutation $p = 0.22$; administrative Canadian-share −0.02%, $p = 0.92$).

The same data explain the null. The Canadian decline was somewhat deeper at the most exposed decile of metros, 9.6 percent against 6.4 percent elsewhere, with no statistically significant exposure gradient in the decline itself, and Canadian passengers were only about 3 percent of arrivals even at the most exposed destinations, so the withdrawal removed a few tenths of a percent of their arrivals—a dose an order of magnitude too small to move metropolitan employment. The most air-exposed destinations in fact added leisure and hospitality jobs in 2025 while the most land-exposed metros shed them; Figure 4 displays the airport-level pattern.

I.K Synthetic-control case studies

Synthetic controls for the four most day-trip-dependent metros give a mixed but mechanism-consistent picture rather than uniform corroboration of the level gradient. Bellingham, which depends heavily on Vancouver day-trippers and sits far from replacement demand, diverges clearly below its synthetic (−2.7 percent post-onset against a pre-period fit of 0.012 RMSE). The higher-market-potential Buffalo and Burlington show no gap; Buffalo in fact sits slightly above its synthetic. Detroit's exact pre-period fit (RMSE near zero across 335 donors) marks its series as overfit, so we do not read its gap as informative. The synthetic evidence therefore supports the market-potential mechanism of Section 8—the loss concentrates where replacement demand is thin—not the level gradient metro by metro.

I.L The market-potential gradient

Section 8 reports that the land decline concentrates where the target's spending is a large share of local demand and thins where local demand is deep. We make that margin continuous by interacting land-crossing exposure with a domestic market-potential index, the gravity sum of other U.S. metros' employment within short-haul reach, in the market-access tradition of Donaldson and Hornbeck (2016). The gradient is −0.84% per standard deviation in isolated border corners with little domestic demand to draw on and a flat −0.02% where market potential is high, the interaction significant at $p < 0.001$. Market potential proxies jointly for a low visitor share of local demand and for the depth of replacement spending, so the interaction identifies their combined gradient rather than isolating either channel; we read it as a joint incidence gradient, not as a causal estimate of replaceability. It is the continuous form of the discrete land-versus-air contrast, separating Bellingham, a remote Vancouver-dependent crossing, from Detroit and Buffalo, large metros embedded in dense domestic markets.

I.M Job-count constructions

The bracket of Section 11 crosses two constructions with the three gradients of Section 10. The first construction extrapolates a per-standard-deviation gradient: multiplying each metro's exposure by the gradient and by its 2024 baseline leisure and hospitality employment, summed over metros with positive land-crossing exposure. This is a linear extrapolation across the exposure distribution and is sensitive to that functional form, so we read it as an order of magnitude rather than a point estimate. The second is independent of the functional form: instrumenting each metro's realized lost crossings with its predetermined exposure (first-stage $F \approx 570$ in levels) gives -0.75 log points of leisure and hospitality employment per million lost distance-allocated crossings at the metro-level gradient, the same units as the exposure construction, and cumulating this over metros gives roughly 12,500 (95% confidence interval 1,500–23,500). Because that cumulation is linear in the coefficient, it rescales in proportion across the ladder.

Gradient	Per SD	Extrapolated	Instrumented
Within-metro triple-difference (flat pre-period)	−0.30%	16,400	6,500
Leisure and hospitality within division × month	−0.45%	24,600	9,800
Metro-level gradient (upper bound)	−0.57%	31,600	12,500

Set against the 9.7 million foregone crossings, these aggregates bracket the intensity between roughly one job per six hundred and one per fifteen hundred lost trips. The land-crossing series observes Canadian land visits only; the air channel is not separately identified at this frequency. Because the instrument is the same predetermined exposure that drives the reduced form, the instrumented construction is a reparameterization into visits units rather than independent evidence, and its confidence interval is sampling uncertainty conditional on the identification caveats of Section 10.

I.N The withdrawal in dollars: margins, a spending-side benchmark, and modes

Section 4 reports the headline dollar margins from National Travel Survey aggregates (Statistics Canada Table 24-10-0070): a C$2.95 billion decline over 2025Q2–Q4 (about US$2.1 billion), driven entirely by the extensive margin, two-thirds leisure with family-visit spending up, and a BEA mirror agreeing to within about 20 percent. This subsection adds construction detail and the benchmark calculations built on those figures. By purpose, leisure accounts for C$1.95 billion of the decline (66 percent) and travel for other purposes for C$1.11 billion (38 percent), while spending on visits to friends and relatives rose C$0.12 billion. Within the leisure decline the leisure-overnight cell alone is C$1.92 billion. The same purpose ordering holds on visit counts, where the other-purpose decline (38 percent) exceeds leisure's (27 percent) and family visits fall least (13 percent)—a pattern Section 9 reads as constraining but not discriminating between the candidate motives. The same-day cells contribute only about C$0.34 billion of the national dollar decline, which is not a tension with the paper's border-metro incidence but the reason for it: day-trip dollars are small in the aggregate and concentrated in a handful of small border economies, while the far larger leisure-overnight dollars are spread across the whole country, which is why the national total is modest and the measurable employment imprint is local.

The spending decline supplies a benchmark for the job counts of Section 11 from independent

data. Accommodation, restaurant-and-bar, and recreation spending—the categories that meet local leisure and hospitality labor—are 68 percent of Canadian spending in the United States. Applying a 30 to 40 percent labor share of leisure and hospitality revenue and the $25,000–$35,000 annual wage of Section 11 to the annualized local-consumption decline implies roughly 17,000 to 31,000 national leisure-and-hospitality job-equivalents. That an aggregate spending-side calculation, from survey data and with no employment estimate in it, lands on the same range as the exposure-gradient constructions is a consistency check on both; and because the spending figure is national while the gradient counts land-border metros, the overlap leaves little room for additional air-metro losses, consistent with the air null of Section 8.

Allocating the dollar decline across travel modes with the frontier counts (Table 24-10-0053) and the survey's duration-specific spend per visit (assuming spend per visit within a duration class is common across modes, which understates the longer-haul air trips and so bounds the air share from below), the automobile margins carry about four-fifths of the dollar decline—69 percent overnight automobile, 11 percent same-day automobile—and air at most about a fifth. The mode split of the dollars matches the mode split of the measured employment incidence. Outputs and construction in `ext_D_nts_decomposition.csv` and `ext_D_nts_boe.csv`; survey aggregates are quoted without sampling inference.

I.O The identifying geography

Figure A5 maps the predetermined exposure that identifies the estimates: Canadian land-crossing volume on the northern tier, Canadian air-destination share in the Sun Belt, and the Mexican land-crossing volume that anchors the southern-border placebo of Section 9.

Table A1. Summary Statistics

	Mean	SD	10th	Median	90th
Leisure & hospitality employment (000s)	41.8	89.9	5.0	14.6	102.2
Total nonfarm employment (000s)	391.3	867.5	48.8	127.5	1,064.9
Canadian air-destination share (%)	0.2	1.3	0.0	0.0	0.2
Distance to nearest Canada crossing (km)	853	650	175	711	1,695
Distance to nearest Mexico crossing (km)	1,601	811	451	1,644	2,626

Notes: Metro cross-section, 390 metropolitan areas, 15,990 metro-months (forty-one months, January 2023 through May 2026). Distributional moments are computed over the 352 metros with centroid coordinates that form the continuous-exposure estimation cross-section; employment is the 2024 metro mean. Ten metros constitute the legacy binary land-border set and 46 carry a positive Canadian air-destination share; 15 metros lie within 100 km of a Canada land crossing and 11 within 100 km of a Mexico crossing. The border-county panel covers 649 counties in the thirteen border states plus Alaska. *Sources:* BLS Current Employment Statistics; U.S. NTTO Survey of International Air Travelers; Bureau of Transportation Statistics Border Crossing data; author's calculations.

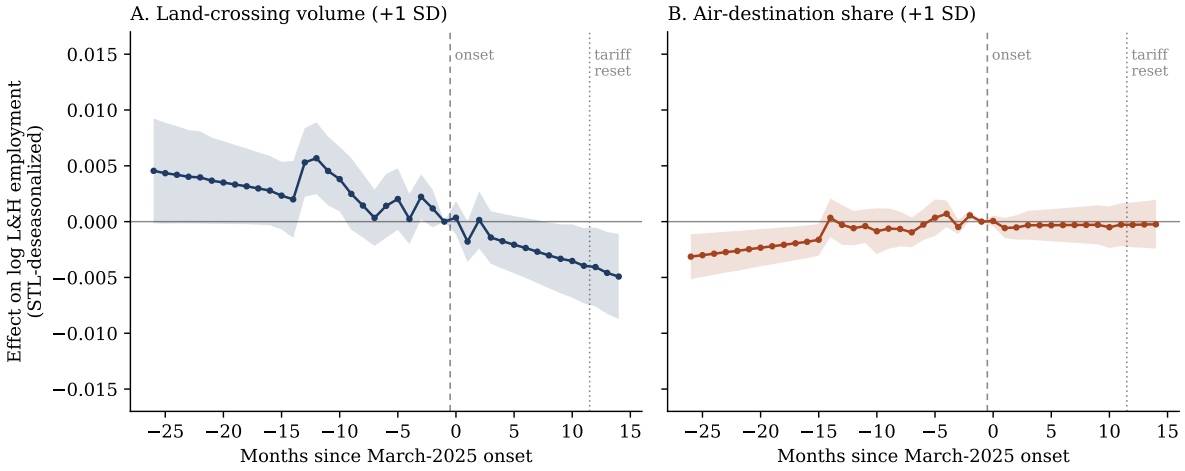

Figure A1. Event-Study Gradients on the Deseasonalized Series

Notes: Event-time coefficients on exposure interacted with month-since-onset dummies (reference February 2025, $t = -1$), estimated on STL-deseasonalized log metropolitan leisure and hospitality employment with metro and month-of-sample fixed effects; shaded bands are cluster-robust 95% intervals (metro clusters). Panel A is the land-crossing-volume exposure: high-exposure metros sit on a higher, declining path before onset, converge by the reference month, and break downward afterward, declining to about −0.49% per standard deviation at the end of the window, continuing through and beyond the February 2026 tariff reset (dotted line). Panel B is the air-destination exposure, flat at zero after onset. The paper does not rest on this path: a Rambachan–Roth sensitivity shows the average post-onset coefficient is not robust to relative-magnitudes pre-trend violations of any size (breakdown $\bar{M} = 0$), which is why we read the headline as a descriptive association supported by the triple-difference, the division-by-month comparison, the in-time placebo, the southern-border placebo, and the Canadian-side redirection (Section 9). *Sources:* BLS Current Employment Statistics; Bureau of Transportation Statistics; author's calculations.

Table A2. Robustness of the Land-Crossing-Volume Gradient

Specification	Land-volume gradient	Inference / note
Baseline (Table 2)	−0.57%	perm. $p = 0.001$; CRV1, WCR (Rademacher/Webb), perm. agree
metro wild-cluster restricted bootstrap	−0.57%	Rademacher $p = 0.0002$; Webb $p = 0.0002$ (352 clusters)
+ own-tariff exposure × post, mfg. share × post	−0.71%	confound moves estimate away from 0
Normalized exposure (volume per 2024 L&H worker, logs)	−0.62%	$p < 0.001$; not merely "large crossing nearby"
Pooled land + air exposure (single index)	−0.36%	perm. $p = 0.015$; land loss averaged against air null
Leave-one-border-metro-out (range over 18)	−0.59% to −0.55%	all significant
Leave-out-Detroit (enters once, as MSA)	−0.58%	essentially the baseline
Division×month fixed effects	−0.44%	perm. $p = 0.002$; proximity does not survive
State×month fixed effects	−0.10%	ns; SE quadruples, CI contains baseline
Metro × calendar-month fixed effects	−0.58%	$p < 0.001$; proximity −0.36%
Exposure × linear trend (full window)	+0.12%	ns; trend absorbs converge-then-break (see triple-diff)
Conley spatial-HAC SE (300 km, 6-mo. lag)	−0.57%	$p = 0.016$; $p = 0.003$ at 200 km
Wild-cluster bootstrap, state clusters	−0.57%	$p = 0.034$ (52 clusters); division $p = 0.12$ (9, excl. PR)
Spatially-blocked permutation (full block grid)	−0.57%	$p = 0.0025$ (40 blocks), 0.0045 (20), 0.013 (10), 0.006 (states), 0.024 (divisions)
Drop controls < 150 km from treated (SUTVA)	−0.61%	$p < 0.001$; contamination biases toward 0
STL-deseasonalized outcome	−0.55%	$p < 0.001$ (proximity −0.34%, $p = 0.010$)
Triple-difference (tourism vs. rest, within metro-month)	−0.30%	perm. $p = 0.046$; CRV1 0.072, WCR 0.077
pre-onset leads, design-based flatness test	—	fail to reject flat: perm. $p = 0.17$ (i.i.d.), 0.44 (state-blocked)
division×sector×month fixed effects	−0.32%	perm. $p = 0.035$; tourism trend free by division
southern-border (Mexico) triple-difference	−0.11%	ns (perm. $p \approx 0.45$); no gradient absent the withdrawal
Honest-DiD, sign-restricted late horizon	not reported	bounds non-monotone in $\bar{M}$ and within 4×10^{-4} of zero throughout; treated as grid noise. Unrestricted avg. post $\bar{M} = 0$
Post-reset persistence (split at 2026m2, panel to 2026m7)	−0.22% / −0.52%	boycott window / post-reset, division×sector×month spec. (Table 4); perm. $p = 0.014$, blocked 0.058 on post-reset. Pre-period-fixed seasonal; the STL deseasonalizer overstates the post-reset deepening at the sample endpoint
Onset date Dec 2024–Apr 2025 (deseasonalized)	−0.51% to −0.56%	stable; Jan. onset −0.53%, perm $p < 0.001$
Placebo onset (March 2024, deseasonalized)	−0.12% to −0.15%	ns (perm. $p \approx 0.25$–0.34); small, insignificant pre-period gradient; consistent with modest pre-onset

Table A3. Persistence Split: Robustness

	Boycott window		Post-reset		p, difference
North, within metro, i.i.d. permutation	−0.25	(0.073)	−0.39	(0.057)	0.156
North, within metro, state-blocked	−0.25	(0.196)	−0.42	(0.161)	0.367
North, division×month FE (preferred)	−0.22	(0.124)	−0.52	(0.014)	0.007
state-blocked				(0.058)	0.062
South (Mexico) placebo, within metro	−0.12	(0.399)	−0.22	(0.276)	0.324

Notes: Post-period split of the within-metro tourism-minus-nontourism differential at the February 2026 reset, panel through July 2026, pre-period-fixed seasonal factors throughout; coefficients in percent per standard deviation of land-crossing exposure with permutation p-values in parentheses (2,000 replications, both split coefficients jointly per draw; state-blocked draws reassign state-mean exposure across states). The Mexico row applies the identical construction to the 28 southern land ports on the same metros. The final column tests post-reset minus boycott-window. Companion to Table 4.
Sources: Bureau of Transportation Statistics; BLS Current Employment Statistics; author's calculations.

Table A4. Land-Gradient Sensitivity to the Distance-Decay Bandwidth τ

	Land-crossing volume		Border proximity	
τ (km)	Coef. (%)	CRV1 p	Coef. (%)	CRV1 p
50	−0.46	0.002	−0.23	0.027
100	−0.53	< 0.001	−0.29	0.016
150 (volume)	−0.57	< 0.001	−0.34	0.009
200 (prox.)	−0.57	< 0.001	−0.39	0.004
250	−0.55	< 0.001	−0.43	0.002
300	−0.52	< 0.001	−0.46	0.001

Notes: Each cell re-estimates the headline two-way (metro and month-of-sample) fixed-effects gradient on log metropolitan leisure and hospitality employment for a one-standard-deviation increase in the stated land exposure interacted with the post-March-2025 indicator, after rebuilding that exposure at the given bandwidth τ and z-scoring it over the estimation cross-section (n = 14,432 metro-months). Coefficients are in percent of baseline leisure and hospitality employment. The headline specifications use $\tau = 150$ km for land-crossing volume and $\tau = 200$ km for border proximity (marked). CRV1 is the cluster-robust p-value (metro clusters); at the headline bandwidths the permutation p-values are 0.001 (volume) and 0.009 (proximity). All estimates are observational.
Sources: BLS Current Employment Statistics; Bureau of Transportation Statistics; Census cartographic boundaries; author's calculations.

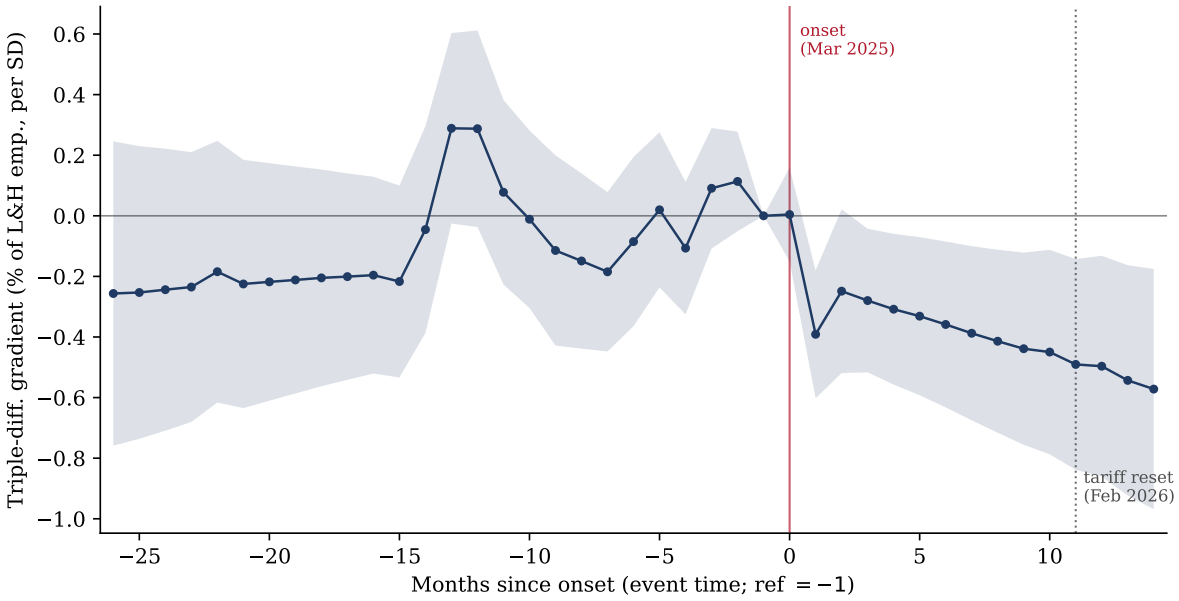

Figure A2. Triple-Difference Event Study on the Deseasonalized Series

Notes: Event-time coefficients on Canadian land-crossing exposure interacted with a leisure-and-hospitality indicator and month-since-onset dummies (reference February 2025, $t = -1$), from the within-metro triple-difference of Section 7, with metro-by-sector, metro-by-month, and sector-by-month fixed effects, estimated on STL-deseasonalized log metropolitan employment; shaded bands are cluster-robust 95% intervals (metro clusters). The immediate pre-onset leads are flat—a joint Wald test over the five lead coefficients before the reference month ($t = -6$ through $t = -2$) does not reject flatness ($p = 0.14$)—and the tourism-differential gradient breaks downward after onset (solid line, March 2025), reaching about −0.57% per standard deviation by the end of the window (numerically coincident with, but a different object from, the −0.57% level gradient of Table 2) and continuing through the February 2026 tariff reset (dotted line). A joint test over the full 25-lead window rejects flatness ($p = 4.5 \times 10^{-6}$), driven by the distant 2023 leads visible at the left of the panel: a small negative offset converging toward zero, whose positive slope (+0.013 percent per month) is the wrong sign to inflate a negative post-onset break—extrapolated forward it would predict a positive gradient. The single-coefficient version of this design is the −0.30% triple-difference of Section 7, −0.29% on the deseasonalized series. *Sources:* BLS Current Employment Statistics; Bureau of Transportation Statistics; author's calculations.

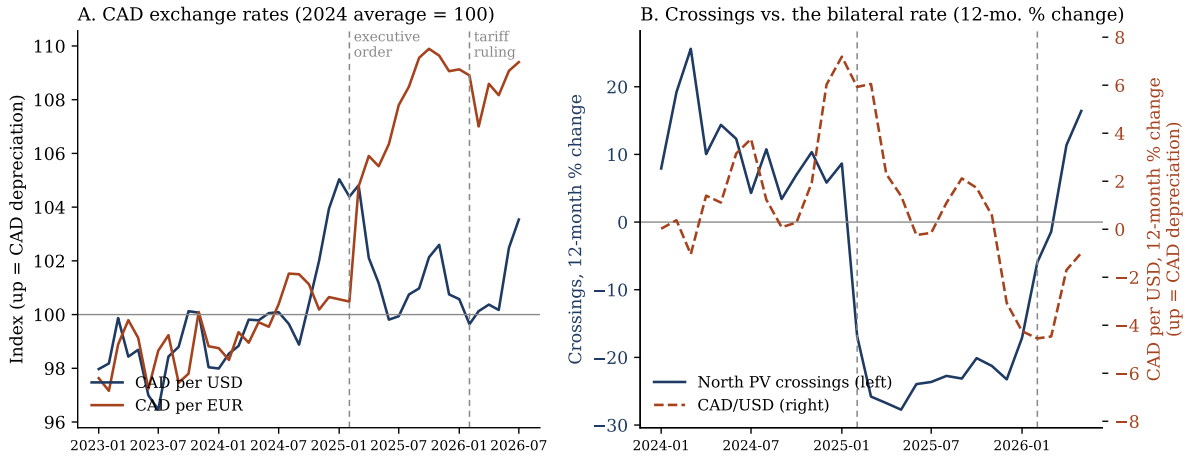

Figure A3. The Canadian Dollar Cannot Account for the Travel Decline

Notes: Panel A plots the monthly Canadian dollar against the U.S. dollar and against the euro, each indexed to its 2024 average = 100 (Bank of Canada daily rates from the Valet API, collapsed to monthly means), with vertical lines at the February 2025 executive order and the February 2026 tariff ruling. The bilateral depreciation against the U.S. dollar peaks near 5% at onset, is larger than the contemporaneous move against the euro, and returns to its 2024 level by early 2026. Panel B overlays the twelve-month change in northern personal-vehicle crossings on the twelve-month change in the bilateral rate (the crossing series is on the left axis, the exchange rate on the much narrower right axis); crossings fall by roughly a quarter at onset, and the positive twelve-month changes in early 2026 reflect the collapsed 2025 base—in levels, crossings remain about 20% below 2024 (Section 4.C) after the currency has recovered. *Sources:* Bank of Canada; Bureau of Transportation Statistics; author's calculations.

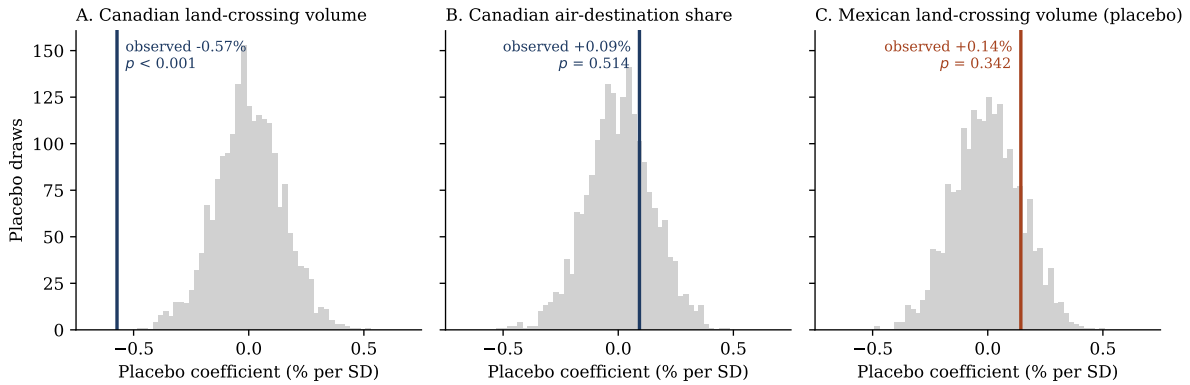

Figure A4. Randomization Distributions for the Exposure Gradients

Notes: Each panel reassigns the static metro-level exposure at random across the 352-metro cross-section 2,000 times, re-estimates the two-way fixed-effects gradient of Table 2 under each placebo assignment, and plots the resulting distribution of placebo coefficients; the vertical line marks the coefficient at the actual assignment. The Canadian land-crossing coefficient lies outside every placebo draw ($p < 0.001$); the Canadian air-destination and Mexican land-crossing coefficients sit inside their null distributions ($p = 0.51$ and $p = 0.34$). These draws are independent of those behind Tables 2 and 3, so the p-values differ from the tabled ones only by simulation noise. *Sources:* BLS Current Employment Statistics; Bureau of Transportation Statistics; U.S. NTTO; author's calculations.

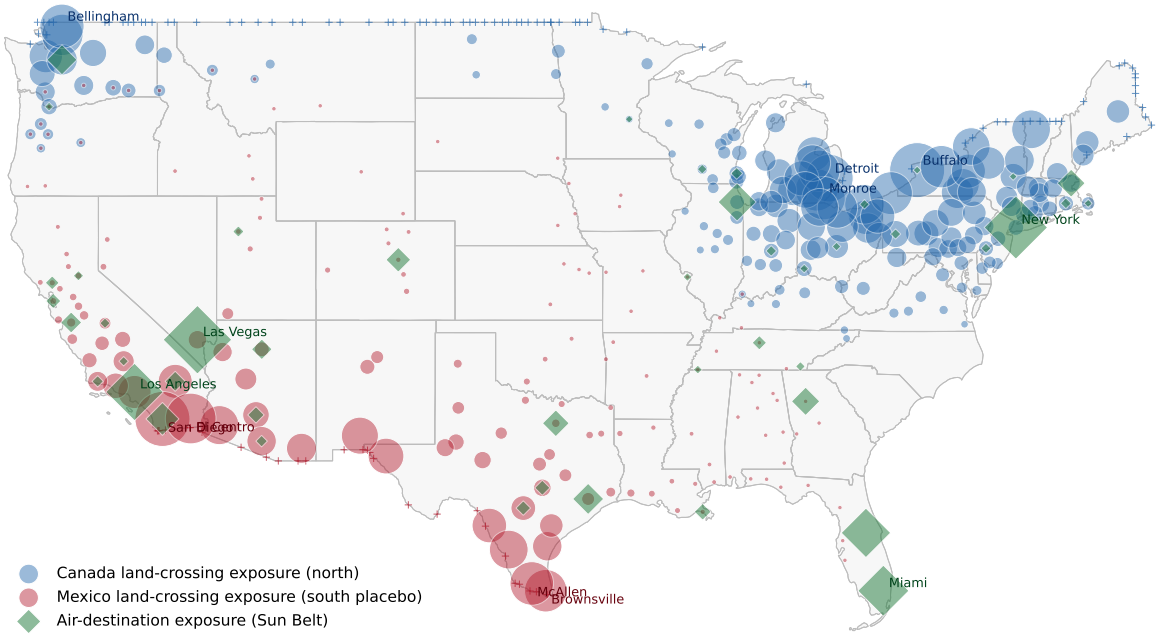

Figure A5. The Identifying Geography on Both Borders

Notes: Each marker is a metropolitan area in the upper part of the exposure distribution (above the 55th percentile of the relevant measure; low-exposure metros are omitted for legibility). Circle area scales with land-crossing exposure, the 2024 inbound personal-vehicle volume gravity-allocated over all Canadian (blue) or Mexican (red) ports of entry with a 150-km decay; green diamonds are Canadian air-destination exposure (NTTO Survey of International Air Travelers). Crosses mark the 89 Canadian and 28 Mexican land ports of entry. The southern-border land exposure provides the placebo of Section 9; the air-destination margin provides the mode contrast of Section 8. *Sources:* Bureau of Transportation Statistics; U.S. NTTO; Census cartographic boundaries.