Appendix A. Data provenance and verification

This appendix records the source, format, and verification of each benchmark; the panel-construction scripts reproduce every file from the sources cited. The bilateral all-commodity panel spans all seven waves; the commodity panel (SCTG two-digit) spans every wave but 1993. No cell in either panel is interpolated, concorded across classification breaks, or otherwise manufactured. Nor does any deflator enter the analysis: no statistic compares dollar levels across waves—rank correlations are ordinal, destination-normalized shares cancel any within-wave price level, and gravity is estimated wave by wave—so the few dollar magnitudes quoted are nominal values for their own benchmark year.

Wave-specific notes. The 1993 extraction reads the 50 state-report PDFs of the 1992 Economic Census transportation program (typeset 1996 output with an intact text layer; no OCR), passes two internal checksums per file—destination sums against the printed Total row, every cell against its printed percent column—and a cell-for-cell external comparison with the NBER transportation archive, with all eleven adjudicated disagreements resolving against the NBER transcription. The 2007 tables are read from the BTS ROSAP archive: State Tables 7 and 8 for all-commodity totals, and the origin-destination-by-commodity workbook (final version, October 2011) for pair-by-commodity detail with cell-level CVs on value, tons, and ton-miles; the two sources agree exactly on all 2,266 comparable cells, and state names govern where the workbook's state codes err (1,442 rows). The 2012, 2017, and 2022 microdata aggregates match published U.S. totals to the dollar; suppressed-geography rows (about 0.1 percent of value) are dropped and logged. Survey scope changes—most prominently the 2007 addition of auxiliary establishments—affect levels, not the destination-normalized shares that carry the headline results.

Distance measures. Great-circle distances connect 2020 state population centroids and are held fixed across waves. Truck travel times aggregate the American Transportation Research Institute (ATRI) county-to-county GPS database to state pairs with county-employment weights; Alaska and Hawaii enter by 20-knot sea legs from Anchorage and Honolulu to gateway ports plus truck legs.

Appendix B. The Rauch classification of SCTG commodities

Table B1 reports the full mapping of two-digit SCTG commodities to Rauch (1999) classes—w for organized-exchange, r for reference-priced, n for differentiated—in the conservative and liberal variants, with the note governing each contested assignment. The mapping is a judgment classification at the two-digit level; both variants are carried through every result in Section 4.1 and Appendix C.

Table B1. SCTG-2 to Rauch class mapping

SCTG	Description	Cons.	Lib.	Note
01	Live animals and fish	r	w	auction markets for livestock argue w under liberal
02	Cereal grains	w	w	organized exchanges (CBOT)

03	Other agricultural products	r	w	produce reference-priced; some exchange-traded
04	Animal feed and products of animal origin nec	r	w	soybean meal etc. exchange-traded under liberal
05	Meat, fish, seafood, and preparations	r	r	reference-priced
06	Milled grain products and bakery	r	n	branded bakery differentiated under liberal
07	Other prepared foodstuffs; fats and oils	r	r	mixed; fats/oils near w
08	Alcoholic beverages	n	n	branded
09	Tobacco products	n	n	branded
10	Monumental or building stone	r	r	
11	Natural sands	r	r	
12	Gravel and crushed stone	r	r	
13	Nonmetallic minerals nec	r	r	
14	Metallic ores and concentrates	w	w	exchange-priced metals content
15	Coal	w	w	
17	Gasoline and aviation turbine fuel	w	w	
18	Fuel oils	w	w	
19	Coal and petroleum products nec	w	w	
20	Basic chemicals	r	w	several basic chemicals exchange-listed under liberal
21	Pharmaceutical products	n	n	
22	Fertilizers	r	w	urea/DAP near-exchange under liberal
23	Chemical products and preparations nec	r	n	formulated products differentiated under liberal
24	Plastics and rubber	r	r	primary forms reference-priced dominate value
25	Logs and other wood in the rough	r	w	lumber futures argue w under liberal
26	Wood products	r	r	
27	Pulp, newsprint, paper, paperboard	r	r	

28	Paper or paperboard articles	r	n	converted articles differentiated under liberal
29	Printed products	n	n	
30	Textiles, leather, and articles	n	n	
31	Nonmetallic mineral products	r	r	cement/glass reference-priced
32	Base metal in primary or semifinished form	r	w	LME metals argue w under liberal
33	Articles of base metal	n	n	
34	Machinery	n	n	
35	Electronic and other electrical equipment	n	n	
36	Motorized and other vehicles	n	n	
37	Transportation equipment nec	n	n	
38	Precision instruments and apparatus	n	n	
39	Furniture, mattresses, lamps	n	n	
40	Miscellaneous manufactured products	n	n	
41	Waste and scrap	r	w	scrap grades near-exchange under liberal
43	Mixed freight	drop	drop	unclassifiable

Notes: w = organized exchange; r = reference-priced; n = differentiated (Rauch, 1999). SCTG 43 (mixed freight) and commodity-unknown cells are excluded from all class-level results.

Appendix C.

Additional results

Dynamic estimates including the diagonal. Recomputing Section 3.3 on shares of total (rather than interstate) sourcing, so that the diagonal enters both numerator and support, moves the fixed-effects estimates by a few hundredths (LSDV .04, bias-corrected .21) and leaves the noise-robust estimators unchanged to reported precision.

Link margins under the grouped-code harmonization. Retaining the PUF disclosure groupings under the coarsest common partition (ten commodity groups) raises measured survival—count .95, value .9985 to .9989 per transition, births at most half a percent of value—because coarser cells churn less; the pure-code design of Section 3.4 is the conservative choice.

The Rauch gradient under the liberal mapping and within well-measured commodities. The liberal variant narrows the gradient without disturbing its order (value-weighted survival .93

to .95, .99, and .997 across 2017→2022). Restricting to the commodities the survey measures well—the medium and well-covered reliability buckets—raises exchange-class survival to .89 over five years and .92 over ten and leaves differentiated-class survival at .998; the gradient is not a measurement artifact.

Gravity robustness.

Rerunning the wave-by-wave PPML in the great-circle metric: with tonnage as the dependent variable, the elasticity runs from −1.76 (1993) to −1.61 (2022)—uniformly steeper than value, as expected for weight-intensive shipments, and no flatter at the end of the period than at its start. Dropping Alaska, Hawaii, and the District of Columbia reproduces the headline path to within a few hundredths (−1.03 in 1993, −1.06 in 2022). Retaining the diagonal under the internal-distance convention $d_{ii} = (2/3)\sqrt{(area_i/\pi)}$ (Redding and Venables, 2004), with land areas from Census TIGERweb, steepens the level to −1.38 through −1.48 while preserving the same mild 2012 hump. No variant moves the conclusion: the impedance elasticity of interstate trade does not flatten.

Tonnage counterparts and class-level gravity. In tons rather than dollars, wave-on-wave rank correlations run .93 to .96 and the 30-year correlation is .89; tons-weighted link survival is .92 to .96 across the microdata transitions—lower than the value-weighted rates, as the Rauch gradient predicts, since tonnage overweights the exchange-traded commodities that churn. Class-level gravity completes the picture: across the six commodity waves (1997–2022) the distance elasticity of differentiated goods stays between −0.75 and −0.84, reference-priced commodities steepen from −0.92 to −1.17, and the small exchange-traded class steepens from −0.51 to −1.59 with wide standard errors. No class flattens, so the stability of Table 4 is not an artifact of composition.

Corridor events behind Section 5. Table C1 lists the documented cost events read against the recovered friction: the corridors each touched, what a cost-tracking friction would record, and what the recovered friction shows

.

Table C1. Corridor events and the recovered friction

Event	Window	Corridors	A cost-tracking friction would show	The recovered friction shows
Arizona-Nevada single-crossing corridor	All six transitions	Arizona-Nevada	Stability: the road did not change	Reversal of direction at every transition, swinging by factors of two to four
Union Pacific-Southern Pacific service meltdown, 1997–98	1997 (collected mid-crisis)	Gulf and Southwest corridors	A visible rise in 1997	The all-pair average; if anything below unaffected controls
Hoover Dam truck ban and its 2010 reversal	1997→2002; 2007→2012	Arizona-Nevada	A rise at the ban, a fall at the reversal	No systematic mark; implied changes scatter in sign
Heartland and National Gateway clearances	2007→2012; 2012→2017	Mid-Atlantic-Midwest corridors	A fall on the cleared corridors	No systematic mark; sign scatters corridor to corridor

Notes: Events with documented direction, location, and timing, read against the friction φ of equation (2); Section 5 reports each test. All statements concern movement only and hold across $\sigma \in \{4, 6, 8\}$. Five-year benchmarks are coarse instruments for single events, so null results are read as limits on interpretation rather than as measured zeros.

Appendix D. Persistence estimation under measurement error

This appendix gives the fuller derivations, reliability calculations, and robustness checks behind the measurement-error correction of Section 3.3: the attenuation and within-transformation algebra, the reliability implied by the 2007 coefficients of variation, the Nickell bias, the errors-in-variables correction, the long-difference and cumulation arguments, and the calibrated simulation.

D.1 Specification and the estimator path

The dependent variable is the destination-normalized off-diagonal value share: the fraction of destination d's out-of-state purchases sourced from origin o in benchmark wave t. The balanced complete-case support is 1,719 off-diagonal pairs observed in all seven waves. The partial-adjustment specification is

$$s_{od}^{t} \;=\; \rho s_{od}^{t-1} \;+\; \gamma_{o}^{t} \;+\; \delta_{d}^{t}\,[+\,\mu_{od}] \;+\; \varepsilon_{od}^{t}, \qquad \text{(D1)}$$

with origin-wave effects γ and destination-wave effects δ absorbing supply and demand shifters, and, in fixed-effects variants, a pair effect μ. A five-year coefficient converts to an annual rate as $\rho^{1/5}$. The estimator path (main-text Table 5) runs from pooled OLS with a five-year ρ of .92 (annual .983), to the pair-fixed-effects (LSDV) estimate of −.03, to the Nickell-corrected within estimate of .135. Read naively, the collapse looks like an absence of persistence. It is the opposite: it is what a within estimator does to a variable measured with sampling error. The rest of this appendix shows that a single highly persistent process, measured with CFS-sized error, generates the entire spread.

D.2 Classical attenuation and the within transformation

The survey never reports a pair's true share s; it reports an estimate s* = s + η, where η is sampling error. Under classical measurement error (η uncorrelated with s), regressing s*(t) on s*(t−1) attenuates the slope toward zero by the reliability ratio

$$\lambda \;=\; \frac{\mathrm{Var}(s)}{\mathrm{Var}(s*)} \;=\; \frac{\mathrm{Var}(s)}{\mathrm{Var}(s)+\mathrm{Var}(\eta)}, \qquad \operatorname{plim}\hat{\rho} \;=\; \lambda\cdot\rho. \qquad \text{(D2)}$$

There is no single reliability ratio. λ depends on which part of the variation an estimator exploits. Pairs differ enormously in how much they trade, so the cross-pair variation in shares is large relative to its sampling error and its reliability is high, $\lambda_{\text{between}} \approx .96$. A pair's own share barely moves across waves, so its wave-to-wave changes are dominated by η and their reliability is low, $\lambda_{\text{within}} \approx .39$. Both are computed from the 2007 cellwise coefficients of variation (Section D.3); the ratio of cross-pair to within-pair variation in the shares is about 17 to 1.

Pooled OLS draws on the high-reliability cross-pair variation and is therefore only mildly attenuated: .92 against a true value near unity. The within (pair-fixed-effects) estimator discards the cross-pair variation by construction and is left with the low-reliability within-pair variation. With $\lambda_{\text{within}} \approx .39$, its probability limit is at most $.39\cdot\rho \approx .39$ before any dynamic-panel bias—so a within coefficient at or below .15 is exactly what a highly persistent series measured this way must produce. The within estimator regresses noise on noise.

D.3 Reliability from the 2007 coefficients of variation

The 2007 benchmark is the only wave whose public tables carry a cellwise sampling coefficient of variation (CV) for each origin-destination value. The CV gives the sampling standard error of each cell relative to its level, hence the sampling variance Var(η) of each share after destination normalization. Reliability then follows from (D2) as $\lambda = 1 - \text{mean}[\text{Var}(\eta)] / \text{Var}(s^*)$, computed separately over the cross-pair distribution ($\lambda_{\text{between}} \approx .96$) and over within-pair wave-to-wave changes ($\lambda_{\text{within}} \approx .39$). Sampling variance implied by the CVs alone accounts for 36 percent of the observed variance of share changes across the 2007→2012 transition; if the 2012 microdata aggregates carry comparable sampling error, the accounting reaches roughly three-quarters.

D.4 The Nickell bias

Even without measurement error, the within estimator of a dynamic panel is biased downward in short panels (Nickell, 1981): demeaning induces a mechanical negative correlation between the transformed lagged share and the transformed error, of order $-(1+\rho)/(T-1)$. With T = 7 waves (six transitions) this is about $-(1+\rho)/6$, roughly −.33 for ρ near unity. The Nickell correction moves the within estimate from −.03 to .135—still far below the truth, because it addresses the dynamic-panel bias but not the measurement-error attenuation of D.2, which is the larger effect here. The two biases compound: the within estimator is first attenuated by λ_{within} and then displaced by the Nickell term.

D.5 The errors-in-variables correction

Where the CVs pin the sampling variance directly, the attenuation can be undone rather than merely bounded. The disattenuated coefficient is

$$\hat{\rho}_{\text{corrected}} = \frac{\hat{\rho}_{\text{uncorrected}}}{\lambda}. \qquad \text{(D3)}$$

On the 2007→2012 transition, for which the 2007 CVs give the reliability of the lagged share, the uncorrected destination-effects coefficient is .964, the CV-implied reliability is $\lambda = .952$, and the corrected coefficient is 1.01, with a cluster-bootstrap 95 percent interval of .96 to 1.07—statistically indistinguishable from a unit root.

D.6 The long-difference estimator

Long differences exploit the same asymmetry from the other side. Over 30 years the persistent signal accumulates while transitory sampling error does not, so the reliability of the 30-year-lagged share far exceeds that of a five-year-lagged one. Regressing the 2022 share on the 1993 share with destination effects gives a 30-year coefficient of .87 (s.e. .04); since the 30-year coefficient is ρ^6, this implies a five-year ρ of .977 and an annual .9954, with a cluster-bootstrap 95 percent interval of .992 to .998. The long-difference and CV-corrected estimates agree, and both sit essentially at a unit root.

D.7 The cumulation bound

A complementary argument needs no estimation. If wave-on-wave changes in the shares were independent draws, the reallocation accumulated over the six transitions would grow as the square root of their number: $\sqrt{6} \approx 2.4$ times a single transition's reallocation. The observed 30-year reallocation is 1.2 times a single transition's (main text, Section 3.1). Observed cumulation is half what independence would produce, so most of what any single transition records as reallocation reverses in the next—transitory noise around a fixed structure—and every transition-level persistence metric is therefore a lower bound.

D.8 The calibrated simulation

To confirm that one data-generating process reproduces the whole estimator path, I simulate panels in which the truth is known. Each synthetic panel sets a five-year persistence of $\rho = .975$ around pair-specific targets drawn from the data, with innovations sized to the CV-implied signal variance and measurement error drawn cell by cell from the 2007 CVs; I then run every estimator on each panel. The synthetic estimates reproduce the observed pattern—pooled OLS at .96, the within estimator collapsed to .13, the long difference at an annual-equivalent .994. Inflating the measurement noise to the level the observed within coefficient implies moves the within estimator to .07 and the long difference to .991, bracketing every observed entry. A single highly persistent process, measured with CFS-sized sampling error, generates the full spread from −.03 to near unity.

D.9 What the coefficients can and cannot separate

Two readings remain observationally near-equivalent at benchmark frequency: genuinely slow adjustment of the shares (ρ near 1) and permanent pair heterogeneity with transitory deviations. Both fix the geography; neither is the fast-adjusting flow the standard model assumes. What the data rule out is the third possibility—wave-on-wave changes that cumulate—and that is ruled out not by the dynamic coefficients but by the cumulation bound of D.7. Every remaining bias attenuates, so where correction is impossible the estimates are lower bounds on persistence.