Merchandise Trade by State *

Kyle Hood, Gabriel Medeiros, Dirk van Duym

U.S. Bureau of Economic Analysis

March 2026

Abstract

Understanding intranational trade flows is essential for analyzing regional interconnectedness, evaluating place-based and trade-related policies, and supporting the construction of sub-national supply-use and input–output systems. Unlike in the case of international trade, however, interstate trade flows are not observed as a by-product of customs enforcement, since shipments across internal jurisdictional boundaries are not subject to duties or tariffs. As a result, information on sub-national trade has been lacking. It must be inferred from surveys collected for other purposes—most notably infrastructure and freight planning—which are not designed to measure production-to-consumption trade flows and embed substantial pass-through activity from wholesaling and warehousing. In this paper, we develop a framework for estimating interstate trade in goods by repurposing infrastructure-oriented shipment data. We factor observed flows into production-related shipments and pass-through activity arising from distribution, storage, and wholesale intermediation, allowing for cross-hauling between regions. Using a set of transparent assumptions and a large-scale numerical solver, we reconcile these components into a single matrix (for each commodity) of production-origin to final-destination trade flows that is internally consistent. The framework provides a policy-relevant representation of interstate goods trade that can be used as an input into broader systems of regional economic analysis.

*Any views expressed in this paper are those of the authors and do not necessarily reflect those of the Bureau of Economic Analysis or the Department of Commerce.

1 Introduction

Understanding the movement of goods within a country is central to analyzing regional economic structure, assessing place-based and trade-related policies, and tracing supply-chain linkages across space. Unlike international trade flows—which are recorded as a by-product of customs enforcement—interstate trade flows in the United States are not directly observed. As a result, researchers and policymakers rely on indirect evidence—primarily shipment data collected for infrastructure and freight-planning purposes—to infer the economic incidence of production-to-consumption flows. However, these data present significant challenges: shipment surveys such as the Commodity Flow Survey (CFS) from the Census Bureau and its extension, the Freight Analysis Framework (FAF) (Bureau of Transportation Statistics, 2026c), capture the movement of goods rather than their economic use, embedding substantial pass-through activity from wholesalers, warehouses, and other distribution intermediaries.

The implications are far-reaching. Because shipments frequently reflect routing through distribution hubs rather than the location of ultimate users, raw flows can be highly misleading. Even the combined CFS–FAF system can show dramatic inconsistencies due to the absence of industry detail for FAF extensions and the survey design of CFS. For example, CFS data record roughly $38 billion in wholesaling shipments of gasoline within the Austin area but only about $300 million in inbound producer shipments—a discrepancy that cannot be reconciled without additional modeling structure or constructed shipments.

These data challenges create obstacles for academic research and applied analysis alike. State-level supply–use tables (SUTs) require a decomposition of shipments into true economic flows—those that connect producers to final or intermediate users—rather than physical movements that reflect distribution logistics. While interstate trade flows are only one component of a full SUT system, they form a critical backbone: without a credible representation of where goods produced in one state are ultimately absorbed, regional accounting identities cannot be closed. This paper therefore focuses exclusively on the estimation of interstate merchandise trade flows, using SUTs only as conceptual motivation.

To address the core measurement problem, we develop a framework that re-purposes infrastructure-oriented shipment data into economically interpretable trade flows. The approach separates observed bilateral shipments into (i) shipments originating from producers, (ii) shipments originating from non-producers—such as wholesalers, warehouses, and retailers—and (iii) the implied economic absorption of goods at each stage of the distribution chain. Using a set of transparent assumptions and a large-scale numerical optimization procedure, the model factors observed flows into production-related and trans-shipment components, allowing for multi-step routing and cross-hauling between regions. The result is a single matrix of production-origin to final-destination trade flows that reconciles observed shipment patterns while remaining internally consistent.

A second contribution is an empirical assessment of these estimated flows across a broad set of Standard Classification of Transported Goods (SCTG) commodities (e.g., coal, gasoline, base metals, metallic ores). We document how the model reshapes the geography of U.S. goods flows by shifting mass away from distribution hubs toward true producer–user relationships. We also present gravity-based diagnostics, comparing producer–to–user flows to simple distance–production relationships. These comparisons reveal substantial heterogeneity: some commodities align well with gravity-type distance decay, while others exhibit pronounced deviations due to clustering of specific uses, multi-stage processing, or idiosyncratic production structures.

Taken together, the framework and results constitute a step toward reliable, policy-relevant measures of interstate goods trade. While preliminary—and not yet validated against independent consumption or output measures already published by the Bureau of Economic Analysis (BEA)—the estimates provide a clearer view of economic flows than raw shipment data and support a wide range of regional applications, from understanding supply-chain vulnerabilities to informing future development of coherent state-level input–output or supply–use systems.

Table 1: BEA–FAF Comparison, 2017: Goods Trade and Double Counting

Category (Millions of dollars)	**BEA**	**FAF**
Imports	2,369,900	2,270,006
Exports	1,546,657	1,555,032
Output / Domestic origin shipments	6,207,105	16,636,779
Residual	10,429,674	
Cross-hauling (direct)	9,632,407	

Notes: "Domestic origin shipments" from FAF substantially exceed domestic economic output because shipments include distribution-chain legs (e.g., producer → warehouse → wholesaler → retailer). "Cross-hauling" entries summarize the magnitude of pass-through shipments when interpreted as producer-to-user trade; residual and direct variants reflect alternative accounting treatments used in the reconciliation of shipment and production totals. Values are in millions of dollars.

2 Background and Data Challenges

The principal U.S. sources for sub-national goods movements are the Commodity Flow Survey (CFS) and the Freight Analysis Framework (FAF). The CFS is fielded every five years as part of the Economic Census and provides establishment-level shipments by commodity, origin, destination, value, weight, shipment mode, and distance (Bureau of Transportation Statistics, 2026a,b). FAF integrates CFS with additional inputs (e.g., foreign trade and out-of-scope sectors such as farms, petroleum/natural gas, and retail) to produce origin–destination (OD) estimates by commodity and mode for states and metro areas (Bureau of Transportation Statistics, 2026c; Federal Highway Administration, 2026). The FAF5 methodology report documents annual benchmarking and adjustments used to reconcile sources across years and modes (United States Department of Transportation, Bureau of Transportation Statistics, 2025).

A central measurement issue for intranational trade is that shipment data record *movements* rather than *economic incidence*. Distribution chains—including wholesalers, warehouses, auxiliaries, and retail—can dominate observed OD flows, obscuring producer-to-user transactions and leading to double counting when distribution legs are interpreted as trade between producers and users. This problem is especially acute for commodities with heavy wholesale and storage activity and for geographies hosting major logistics hubs. Gravity-style diagnostics and model-based reconciliation are commonly used to evaluate whether inferred economic flows behave plausibly with respect to distance and production capacity (Phillips, 2024; Yotov, 2021; Helpman et al., 2007; Park et al., 2008).

To motivate the distinction between economic trade and shipment activity, Table 1 presents a simple comparison of merchandise activity for a benchmark year. The comparison contrasts national-level values from BEA's economic accounts (imports, exports, and domestic output) with FAF-based domestic origin shipments. The gap between domestic output and FAF domestic origin shipments reflects distribution-chain pass-through and cross-hauling rather than additional production.[1]

This comparison underscores why shipment OD matrices cannot be read as economic trade flows without adjustment: a large share of observed movements are distribution legs that must be decomposed into local absorption and trans-shipments before producer-to-user flows can be inferred. In the sections that follow, we describe the data harmonization process, develop a flow model to perform this decomposition using transparent assumptions and a large-scale numerical optimizer, and we evaluate the plausibility of inferred flows using gravity-style diagnostics (Phillips, 2024; Yotov, 2021; Helpman et al., 2007), before examining results for a selection of commodities.

[1]FAF domestic origin shipments encompass many movements that are not final producer-to-user transactions; they include wholesale and distribution legs as well as the movement of used goods United States Department of Transportation, Bureau of Transportation Statistics (2025).

3 Data Preparation

This section documents the preparation of shipment data for eventual estimation of trade flows from origin to final destination. Public documentation for CFS and FAF anchors all definitions, coverage, and benchmarking steps (Bureau of Transportation Statistics, 2026a,b,c; United States Department of Transportation, Bureau of Transportation Statistics, 2025).

3.1 Harmonizing CFS and FAF

The CFS provides establishment-reported shipments by commodity, origin, destination, value, weight, mode, and distance every five years. Microdata on shipments is publicly available and can be used to replicate published estimates. The FAF integrates CFS with additional inputs (including foreign trade data and flows from agriculture, oil and gas extraction, utilities, construction, retail, and service sectors) to produce more complete OD estimates by commodity and mode for states and metropolitan areas (Bureau of Transportation Statistics, 2026c). However, crucially the FAF does not report any industry information for the additions it makes to CFS, and microdata is not available. To harmonize the datasets, we merge the CFS microdata with FAF, and we use FAF's methodology report to impute the industry of the shipments. For example, agricultural products are only shipped by food manufacturing establishments in CFS microdata. FAF adds in shipments by farms, using a variety of data sources including the Census of Agriculture. In this case, it is clear that the FAF additions are producer rather than passthrough shipments. We back out these farm shipments by aggregating CFS microdata at the origin by destination by commodity level and subtracting from reported FAF values at that same level. Then the shipments from producers consist of both the shipments from food manufacturers as well as from farms. But in other cases outside of agriculture, FAF makes additions for multiple reasons: for example, for furniture, FAF adds in shipments by retailers and by moving companies. In our framework, retail shipments need to be classified as passthrough, but moving company shipments need to be excluded entirely (the moving company is simply providing a service–the furniture itself is not current production). This reclassification process is performed for each commodity in the FAF (apart from waste/scrap, which is assumed to be zero value). This ensures consistency in the integration of non-CFS sectors into the OD flow matrices.

3.2 Summary

The data preparation steps produce harmonized OD matrices by commodity, flow type (production or passthrough) and export/import/domestic status. These preparations are grounded in the public documentation of CFS and FAF coverage, integration, and benchmarking, and set the stage for the modeling framework in Section 4 and for results that can be compared against gravity-style expectations (Bureau of Transportation Statistics, 2026a,b,c; United States Department of Transportation, Bureau of Transportation Statistics, 2025).

4 Modeling Framework

This section presents a flow model that repurposes shipment data to estimate economically meaningful interstate trade in goods. The framework links shipments originating at producers with shipments originating at non-producers (e.g., wholesalers, warehouses, retailers) to infer producer-to-user flows, while allowing for multi-step routing and cross-hauling (the same commodity being shipped in both directions along a trade route).

4.1 Objective: Producer-to-User Flows

Let c index commodities, and let i, j, k index regions (states or FAF/CFS geographies). Denote observed shipment values by

$$S_{ij}^{c} \geq 0,$$

which represent movements of commodity c from origin i to destination j recorded in shipment datasets. Our objective is to recover an economically interpretable matrix of producer-to-user flows,

$$T_{ij}^c \geq 0,$$

that connects production locations i to final- or intermediate-using locations j (absorption), net of trans-shipment legs. Shipment data record movements rather than economic incidence; thus, S_{ij}^c can include distribution-chain legs and may substantially exceed producer-to-user incidence in hub-heavy networks.

4.2 Data Inputs and Classification

For each commodity c, we prepare the following matrices from the harmonized CFS/FAF workflow (Section 3):

- P^c: producer-origin *domestic* shipments (origin–destination matrix).
- P_m^c: *import-origin* shipments entering the domestic network (OD matrix).
- T^c: target *domestic* forwarding/cross-haul matrix to be matched by the model.
- T_x^c: target *export-routing* matrix to be matched by the model.

FAF/CFS definitions guide commodity coverage and shipment-type interpretation; targets T^c and T_x^c embody the observed routing patterns that the model will reproduce after decomposing local absorption and forwarding (Bureau of Transportation Statistics, 2026a,c; United States Department of Transportation, Bureau of Transportation Statistics, 2025).

4.3 Two-Step Estimation Strategy

We implement a two-step procedure, using a unified notation that controls *how much* mass is absorbed or forwarded at each destination and *where* forwarded mass is routed. Step 1 estimates absorption/forwarding propensities via constrained optimization; Step 2 performs gravity-gradient refinement of P (producer-origin) and then re-solves Step 1.

Routing matrices and propensities. For each commodity, define *row-normalized* routing matrices

$$S_{ij} \;=\; \frac{T_{ij}}{\sum_k T_{ik}}, \qquad S_{ij}^x \;=\; \frac{(T_x)_{ij}}{\sum_k (T_x)_{ik}},$$

so that each origin's forwarded mass is distributed across destinations according to observed targets. Let N be the number of regions. Define diagonal propensity matrices

$$Q \;=\; \mathrm{diag}(q_1, \ldots, q_N), \qquad Q_x \;=\; \mathrm{diag}(q_1^x, \ldots, q_N^x),$$

where $q_i \in [0,1]$ controls the *domestic forwarding* share from destination i (the complement $1 - q_i$ is local absorption), and $q_i^x \in [0,1]$ controls the *export diversion* share from destination i (the complement $1 - q_i^x$ is the mass remaining available for domestic absorption/forwarding). For clarity, commodity subscripts are suppressed, but Q and Q_x are estimated separately for each commodity c.

Step 1: Constrained optimization (absorption and forwarding). Let R_t and R_t^m be *state* (diagonal) matrices holding the mass available at destinations at recursion step t for producer-origin and import-origin stocks, respectively. Given (Q, Q_x) and (S, S^x), flows evolve as:

$$X_t^x \;=\; R_t \, Q_x \, S^x,$$

$$R_t^{\mathrm{rem}} \;=\; R_t \, (I - Q_x),$$

$$X_t^P \;=\; R_t^{\text{rem}}\, Q\, S, \qquad X_t^{Pm} \;=\; R_t^m\, Q\, S,$$

with state updates

$$R_{t+1} \;=\; \operatorname{diag}\!\big(\mathbf{1}^\top X_t^P\big), \qquad R_{t+1}^m \;=\; \operatorname{diag}\!\big(\mathbf{1}^\top X_t^{Pm}\big).$$

Model-implied domestic and export totals are

$$X_{\text{tot}} \;=\; \sum_t \big(X_t^P + X_t^{Pm}\big), \qquad X_{\text{tot}}^x \;=\; \sum_t X_t^x.$$

Parameters (Q, Q_x) are estimated by solving

$$\min_{Q,Q_x} \;\big\|X_{\text{tot}} - T\big\|_F^2 \;+\; \big\|X_{\text{tot}}^x - T_x\big\|_F^2 \quad \text{s.t. } 0 \le q_i \le 1,\;\; 0 \le q_i^x \le 1,$$

using L-BFGS-B with JAX-based gradients. Economic decomposition is immediate: local absorption at step t is

$$E_t^P \;=\; R_t^{\text{rem}}\,(I - Q), \qquad E_t^{Pm} \;=\; R_t^m\,(I - Q),$$

and forwarded quantities are X_t^P (domestic forwarding from producer-origin mass), X_t^{Pm} (domestic forwarding from import-origin mass), and X_t^x (export forwarding).

Step 2: Gravity-gradient refinement (production anchors). When Step 1 leaves residual mismatch, we refine P using gradient information filtered by gravity plausibility, then re-estimate (Q, Q_x). Define the loss

$$\mathcal{L}(P) \;=\; \min_{Q,Q_x} \;\big\|X_{\text{tot}} - T\big\|_F^2 \;+\; \big\|X_{\text{tot}}^x - T_x\big\|_F^2,$$

and compute row-constrained shadow values

$$G \;=\; \nabla_P\, \mathcal{L}(P) \quad \text{subject to fixed row sums of } P.$$

Construct gravity anchors

$$\Gamma_{ij} \;\propto\; \frac{\text{prod}_i}{d_{ij}^{\delta}},$$

where prod_i proxies origin production mass, d_{ij} is distance, and δ is a commodity-specific decay parameter, consistent with gravity-style diagnostics for domestic flows (Phillips, 2024; Yotov, 2021; Helpman et al., 2007). Column-normalize Γ and update P only in beneficial directions (negative gradient signal weighted by gravity), with adaptive step size η_t:

$$P'_{ij} \;=\; P_{ij} - \eta_t\, \Gamma_{ij}\, G_{ij},$$

followed by row normalization to preserve origin totals:

$$\sum_j P'_{ij} \;=\; \sum_j P_{ij} \quad \forall\, i.$$

After each outer iteration, warm-start (Q, Q_x) and re-solve Step 1. Backtracking reduces η_t when $\mathcal{L}$ increases. The procedure is production-conserving, improves fit, and prioritizes reallocations that are empirically plausible under gravity intuition (Phillips, 2024; Yotov, 2021).

4.4 Treatment of Flow Types

We separate estimation by flow type:

- **Domestic-to-domestic** flows: full absorption/forwarding decomposition applies using (Q, S).
- **Import-related** flows: foreign-origin shipments enter at points of entry (P_m) and are decomposed via (Q, S) to final domestic users.

- **Domestic-to-export** flows: shipments destined for ports of exit are forwarded until reaching export nodes via (Q_x, S^x); absorption at export nodes is zero and mass is routed to exit points.

This triage reflects FAF's distinction between domestic and foreign-related movements and acknowledges port-centric routing observed in shipment data (Bureau of Transportation Statistics, 2026c; United States Department of Transportation, Bureau of Transportation Statistics, 2025).

4.5 Key Assumptions and Identification

The framework rests on transparent assumptions:

1. **Routing proportionality from observed targets:** S and S^x are fixed, row-normalized from T and T_x, and determine *where* forwarded mass goes; Q and Q_x determine *how much* is forwarded vs. absorbed.

2. **Stationarity within benchmark year:** Q and Q_x are destination- and commodity-specific but time-invariant within the benchmark year.

3. **Limited forwarding depth:** Recursive forwarding continues for a finite number of steps (e.g., up to four), reflecting typical distribution-chain lengths.

4. **Gravity-consistent production adjustments:** When observed non-producer shipments cannot be explained by producer-origin inflows, small, production-conserving, gravity-consistent adjustments to P are permitted to restore feasibility.

Identification exploits variation in (P, P_m) (sources), (T, T_x) (targets), the network structure in (S, S^x), and commodity heterogeneity. Gravity intuition aids plausibility checks because economically meaningful flows tend to decline with distance and scale with origin/destination mass (Phillips, 2024; Yotov, 2021).

4.6 Practical Implementation

Step 1 is implemented as a large-scale constrained least squares for each commodity c, with (Q, Q_x) bounded in $[0, 1]$ and (S, S^x) row-normalized from (T, T_x). Step 2 applies modest gravity-style perturbations to P under row-sum constraints, then re-optimizes Step 1 with warm starts. The entire procedure is conducted for each commodity.

4.7 Identifying Inconsistent or Incomplete Flows

We flag OD pairs where reported non-producer shipments (e.g., wholesaling, warehousing, retail auxiliaries) substantially exceed feasible inbound shipments for the same commodity and geography. These flags are motivated by the known limitations of CFS (four representative weeks with scaling to annual totals) and FAF's integration of non-CFS sectors without industry detail (Bureau of Transportation Statistics, 2026a,b,c). In practice, we:

- Compute inbound producer-origin shipments for each destination and commodity from harmonized CFS/FAF data as described above.

- Compare these to outbound non-producer shipments originating in the same destination and commodity; large discrepancies are flagged for reconciliation (Bureau of Transportation Statistics, 2026a,c).

- Use FAF methodological notes to interpret whether discrepancies plausibly arise from benchmarking fills, suppressed cell treatments, or inclusion of out-of-scope sectors (United States Department of Transportation, Bureau of Transportation Statistics, 2025).

4.8 Constructing Supplemental Producer Shipments

When the linear reconciliation fails due to survey under-coverage or timing mismatches, we construct *small* supplemental producer shipments to restore feasibility. These constructed shipments are commodity- and geography-specific, limited in magnitude relative to observed flows, and treated as inputs to the optimizer rather than final outputs. The justification follows the CFS process of benchmarking to Economic Census data on production, but not benchmarking to data on shipment destination, because no other data besides the CFS exists. Our process is as follows:

1. For flagged OD pairs, we add minimal producer-origin mass sufficient to explain observed non-producer movements once local absorption and forwarding are applied.
2. We cap adjustments using commodity-level controls, consistent units, and documented FAF base-year totals to avoid over-correction.
3. Sensitivity checks later assess the influence of these additions on inferred producer-to-user flows.

5 Results

This section presents inferred interstate producer-to-user flows for a set of benchmark commodities, highlights how these flows differ from raw shipment origin–destination (OD) patterns, and interprets the spatial structure that emerges once distribution-chain pass-through is decomposed. We focus on domestic-only flows for clarity, although we also have results for import-related and export-bound legs. All commodity definitions and OD units follow the public CFS/FAF documentation (Bureau of Transportation Statistics, 2026a,c).

5.1 Overview of Inferred Domestic Flows

The flow model reallocates mass away from logistics hubs toward destinations with substantial local absorption, yielding producer-to-user matrices that are more consistent with economic incidence than raw shipments. This is expected because shipment datasets record movements (including wholesaling and warehousing) rather than final use, and FAF extends CFS with non-CFS sectors without industry detail (Bureau of Transportation Statistics, 2026a,c). As a result, the inferred flows typically show:

- Greater concentration near producing states for bulky raw materials (e.g., ores, coal) and short radii of dispersion relative to finished goods.[2]
- Redistribution away from wholesale-intensive metro areas toward manufacturing or final-demand states for refined petroleum and intermediate inputs.
- Port-oriented routing for export-bound domestic shipments and inland absorption for domestic-only flows.

5.2 Commodity-Specific Flow Patterns

We illustrate the reallocation using a selection of representative commodities. The maps show FAF zones, defined as metro areas (broken out by state, e.g. Chicago is split between Illinois and Indiana), rest-of-state, or the entire state for states without large metro areas. There are 132 zones in the U.S. Alaska and Hawaii are estimated but not mapped, for clarity. The largest producing zones are mapped with their shipment links drawn until production shipments reach 90 percent of total commodity domestic production. Colors are arbitrarily assigned to zones. Circle size of producers is proportional to production, and line thickness is proportional to shipments. Note that within-zone shipments are common and included in determining circle sizes.

[2]This aligns with FAF commodity grouping and OD coverage that distinguish raw materials from manufactured goods (Bureau of Transportation Statistics, 2026c).

Agricultural commodities. For cereal grains and other agricultural goods, inferred flows show short-haul dispersion from producer states to nearby processing and consumption regions, with absorption patterns consistent with perishable handling and rail/truck dominance documented in CFS/FAF commodity definitions (Bureau of Transportation Statistics, 2026a,c).

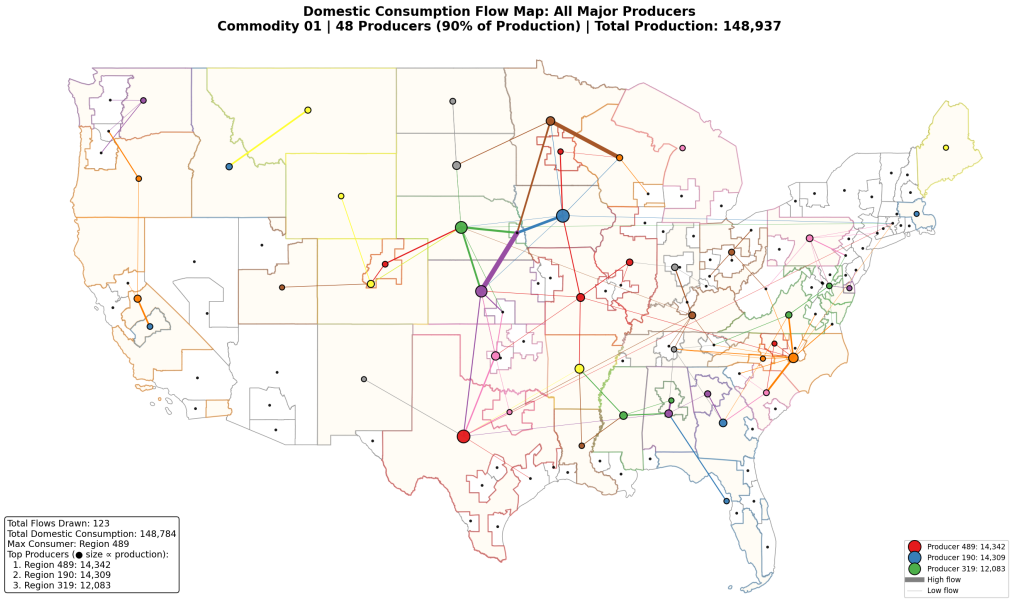


Figure 1: Live animals and fish

Stone, gravel, and building materials. Bulk, low-value-per-weight commodities display pronounced local absorption near production sites, with minimal long-distance forwarding; this conforms to FAF's treatment of heavy, low value-to-weight flows that are typically trucked over short distances (Bureau of Transportation Statistics, 2026c).

Metallic ores and coal. Primary raw materials exhibit corridor-like patterns from extraction states to concentrated user industries (e.g., smelters, utilities). The inferred flows reduce apparent wholesaler hub dominance and restore producer-to-user linkages.

Refined petroleum (e.g., gasoline). For refined products, the inference substantially rebalances flows away from wholesaling hubs and toward refinery-to-market corridors, reflecting multi-step distribution with final absorption in consumer-heavy states.

Pharmaceuticals and base metals. High-value manufactured goods show broader dispersion and multi-node forwarding before absorption. The inferred maps retain some hub roles (e.g., air and specialized distribution centers) but reassign final use toward manufacturing clusters and population-dense markets, consistent with CFS/FAF commodity distinctions and modes (Bureau of Transportation Statistics, 2026a,c).

5.3 Differences Between Observed Shipments and Inferred Economic Flows

A summary comparison highlights how the model changes the geography of flows:

1. **Hub de-emphasis.** OD mass in logistics hubs drops when distribution legs are decomposed, addressing the double-counting inherent in shipment-based OD matrices.

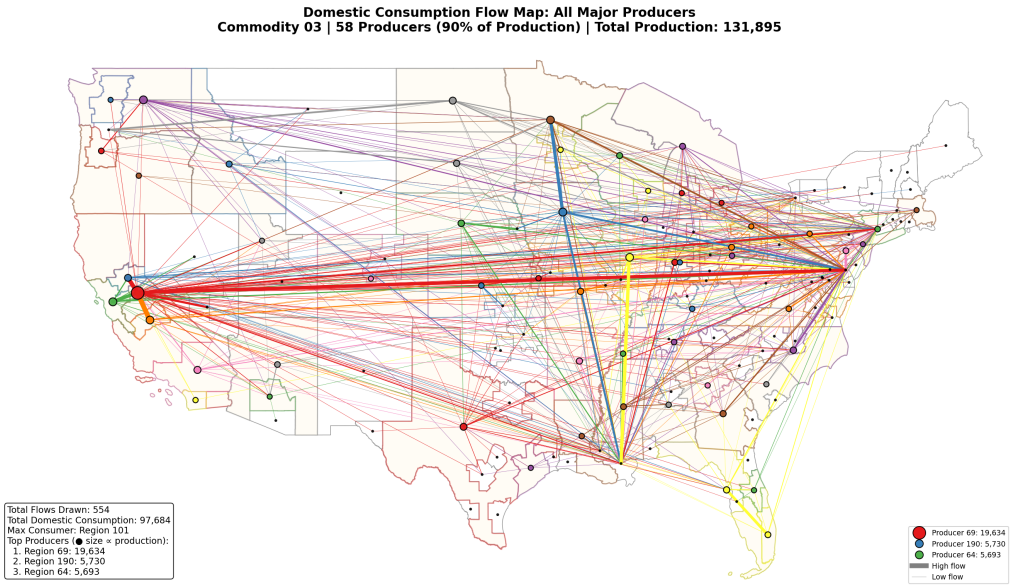


Figure 2: Fruits, vegetables, soybeans, and other non-grains and non-animal feed

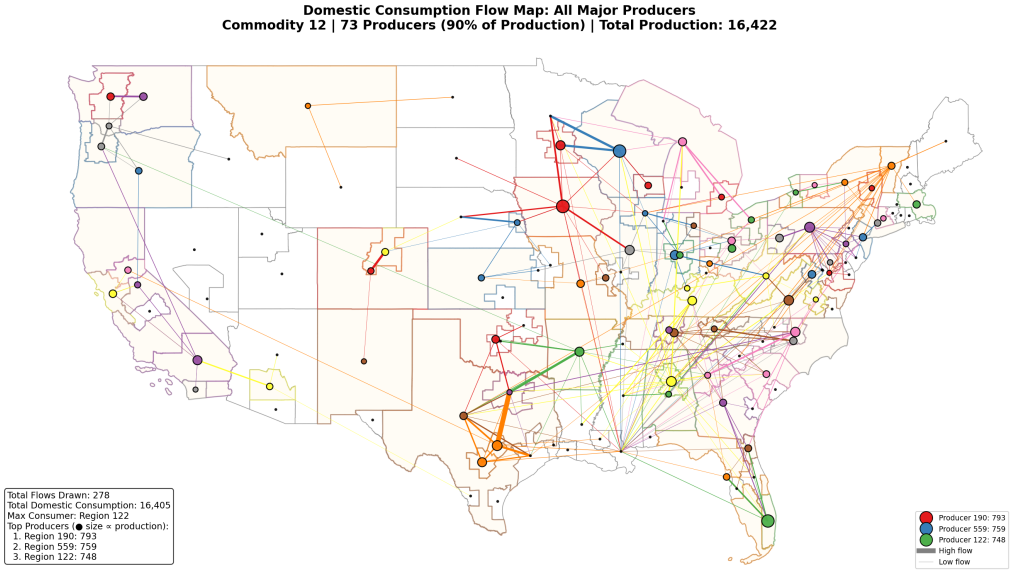


Figure 3: Gravel and crushed stone

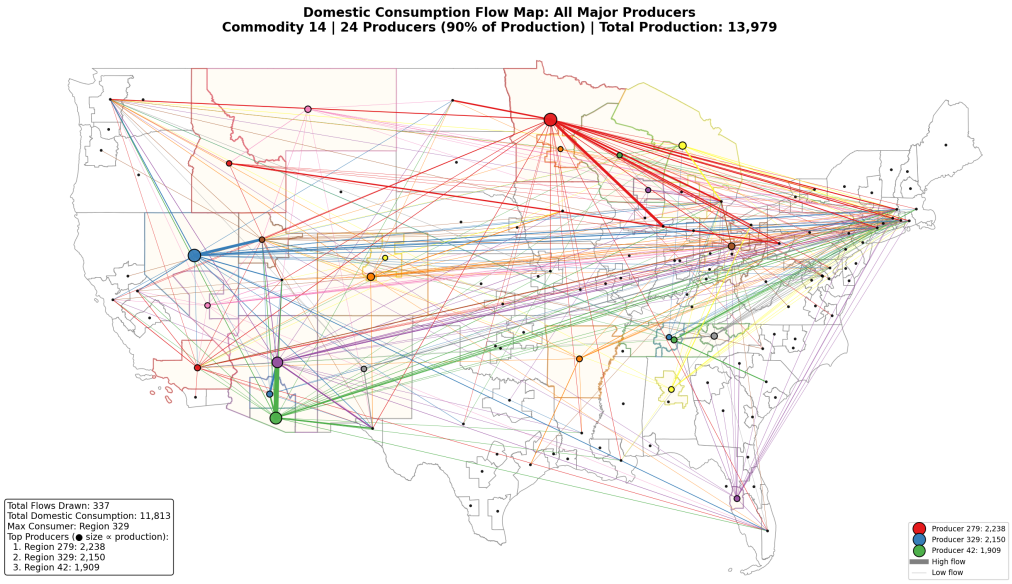


Figure 4: Metallic ores

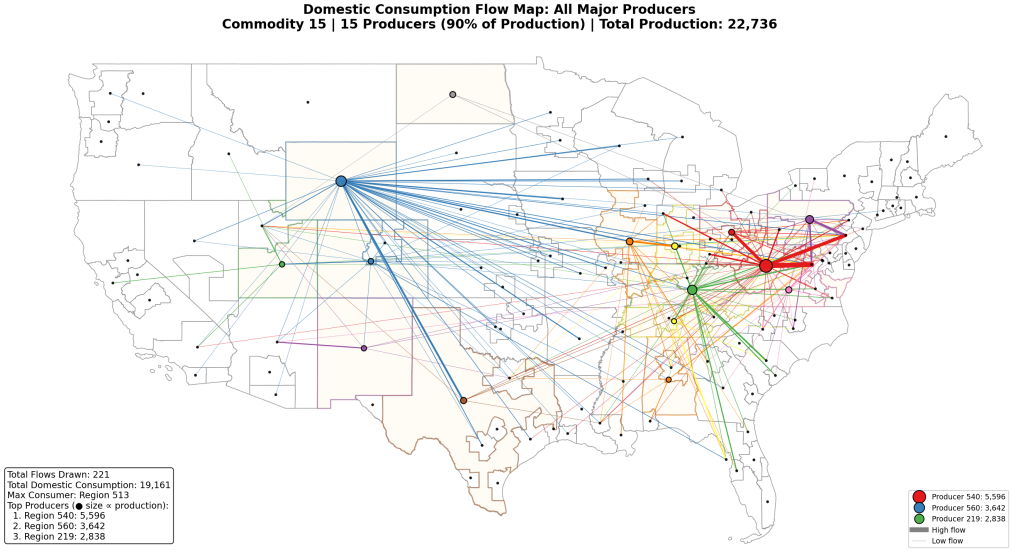


Figure 5: Coal

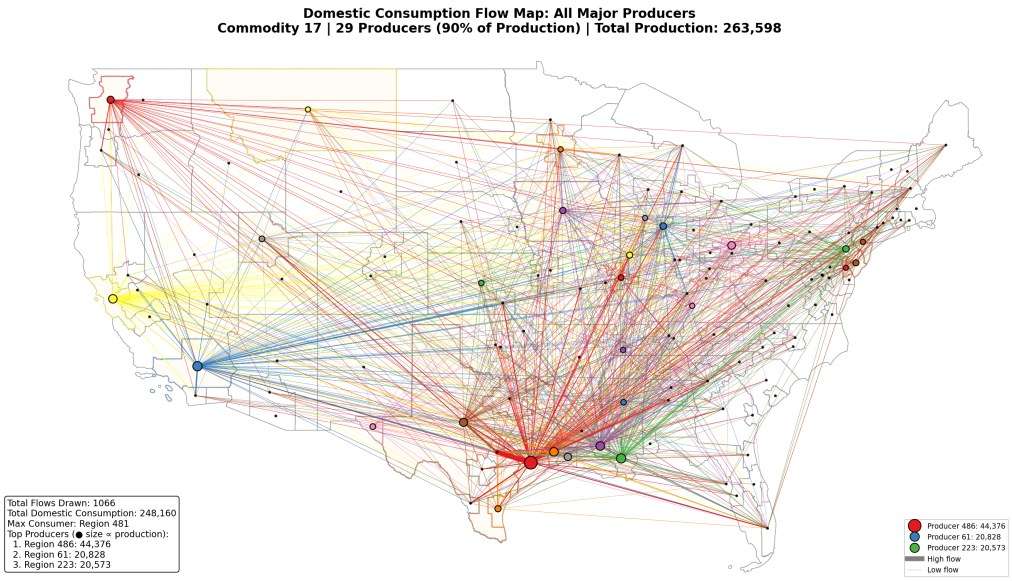


Figure 6: Gasoline

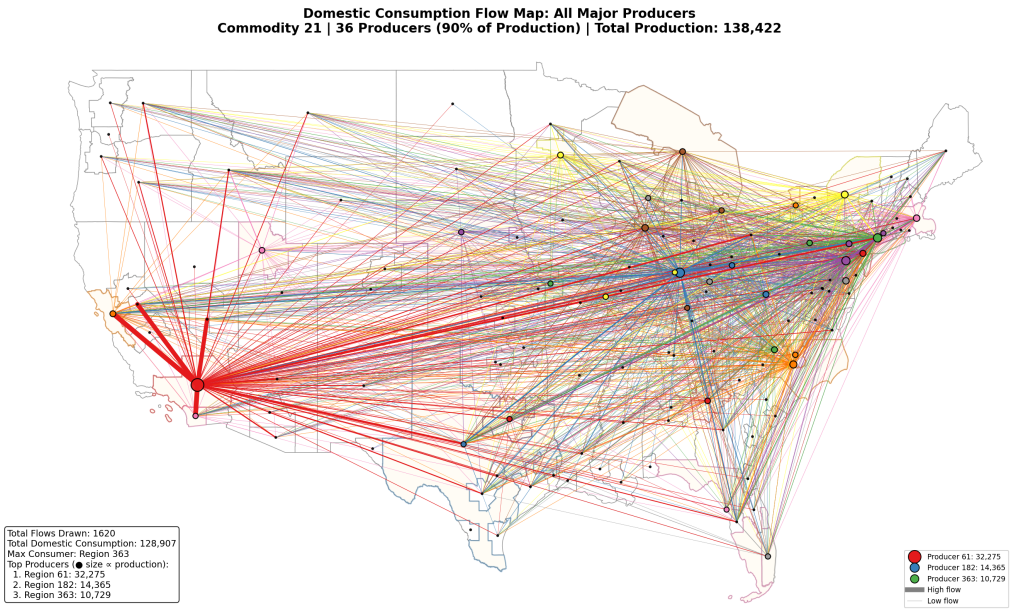


Figure 7: Pharmaceutical products

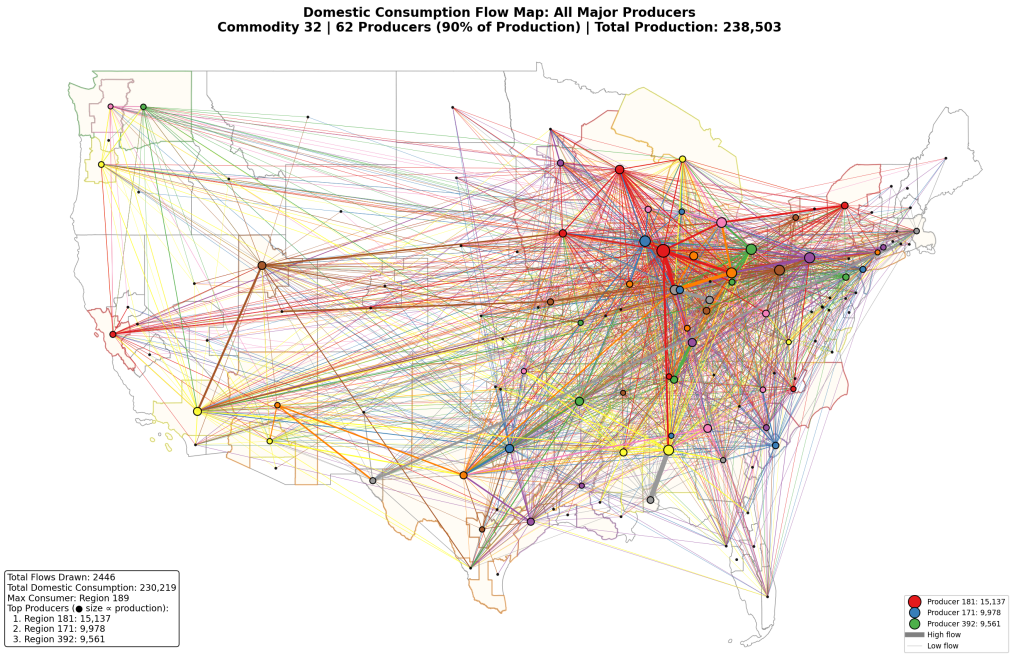


Figure 8: Base metals

2. **Producer adjacency.** Producer-adjacent destinations gain share in the inferred matrix, especially for heavy raw materials whose shipments are costly over distance.

3. **Port specialization.** Export nodes retain strong roles for outbound legs but contribute less to domestic absorption once producer-to-user flows are isolated.

5.4 Interpretation of Spatial Patterns

The inferred patterns align with standard gravity intuition: flows decline with distance and scale with origin production and destination demand, with stronger distance decay for bulk raw materials than for high-value manufactured goods. This is consistent with subnational gravity evidence and within-country gravity variation documented in recent work (Phillips, 2024; Yotov, 2021). Two qualitative regularities are visible in the maps:

- **Distance decay heterogeneity.** Commodities differ markedly in decay rates, with raw materials showing tighter spatial radii than refined or high-value goods (Phillips, 2024).

- **Clustered absorption.** Downstream industrial clusters and large consumer markets exhibit concentrated absorption, even when shipment routing is multi-step, which comports with structural gravity expectations for domestic flows (Yotov, 2021).

5.5 Gravity-Style Distance Relationships

To evaluate whether the inferred flows exhibit gravity-like behavior, we relate bilateral *producer-to-user* flows to bilateral distances. For each commodity c, let $U_{ij}^c \geq 0$ denote the model-implied *producer-to-user* flow from origin i to destination j. Using these bilateral flows, define

$$\text{Prod}_i^c = \sum_j U_{ij}^c \qquad \text{and} \qquad \text{Use}_j^c = \sum_i U_{ij}^c,$$

i.e., total *production* in origin i and total *use* in destination j are obtained directly by summing U_{ij}^c over destinations and origins, respectively.

Let $d_{ij} > 0$ denote the bilateral distance between i and j. We construct a commodity-specific gravity benchmark

$$G_{ij}^{c} \;=\; \frac{\mathrm{Prod}_{i}^{c}\,\mathrm{Use}_{j}^{c}}{d_{ij}^{2}},$$

which captures the standard multiplicative size term (origin production × destination use) divided by squared distance.

To assess whether inferred flows exhibit gravity-like behavior, we compare U_{ij}^{c} and G_{ij}^{c} in log–log space. An approximately linear association indicates consistency with a simple gravity benchmark; systematic deviations (e.g., clusters, multiple bands, or nonlinear regimes) may reflect commodity-specific processing chains, modal segmentation, or concentrated industrial geography.

The gravity–flow relationships shown in Figures 9–11 illustrate how well the simple gravity formulation captures spatial variation in shipments for several representative commodities. Cereal grains (Figure 9) exhibit one of the strongest monotonic relationships, with a Spearman correlation of $\rho = 0.939$ and a tight cloud of points around the one-to-one line. This suggests a structured and distance-sensitive nature of bulk grain movements, where flows tend to follow well-established production and consumption corridors.

Coal (Figure 10) shows a weaker Pearson association ($r = 0.636$), but still maintains a clear rank order match between predicted and observed flows. We posit that the distinct clusters may relate to different grades of coal (anthracite, bituminous, etc.) being used for different purposes (electricity generation vs. industrial usage), although the limited number of source regions may also play a role.

Pharmaceuticals (Figure 11) provide a contrasting example: despite a high linear correlation ($r = 0.769$), the scatter is more diffuse, especially at lower flow volumes. This behavior is consistent with the commodity's more complex and multimodal supply chain, where high-value shipments may deviate from purely distance-weighted patterns.

Table 2 reinforces this interpretation across all commodities, highlighting how gravity performance varies substantially by sector. Bulk, low-value, and spatially concentrated goods—such as grains, ores, and fuel oils—tend to show the tightest alignment between predicted and observed flows, whereas more specialized or higher-value commodities display greater dispersion around the gravity benchmark.

Table 2: Gravity–Flow Correlations by Commodity

Description	Pearson r	Spearman ρ
Live animals and fish	0.881	0.545
Cereal grains and seeds	0.774	0.939
Agricultural products (excl. grains & feed)	0.584	0.856
Animal feed, eggs, honey, animal products n.e.c.	0.685	0.860
Other prepared foodstuffs, fats, and oils	0.772	0.839
Monumental or building stone	0.559	0.919
Natural sands	0.651	0.719
Gravel and crushed stone	0.640	0.524
Metallic ores and concentrates	0.684	0.830
Coal	0.636	0.505
Gasoline and aviation turbine fuel	0.308	0.954
Fuel oils (incl. diesel, kerosene, etc.)	0.572	0.953
Pharmaceuticals	0.769	0.914
Fertilizers	0.540	0.957
Logs and other wood in the rough	0.658	0.600
Base metals	0.489	0.860

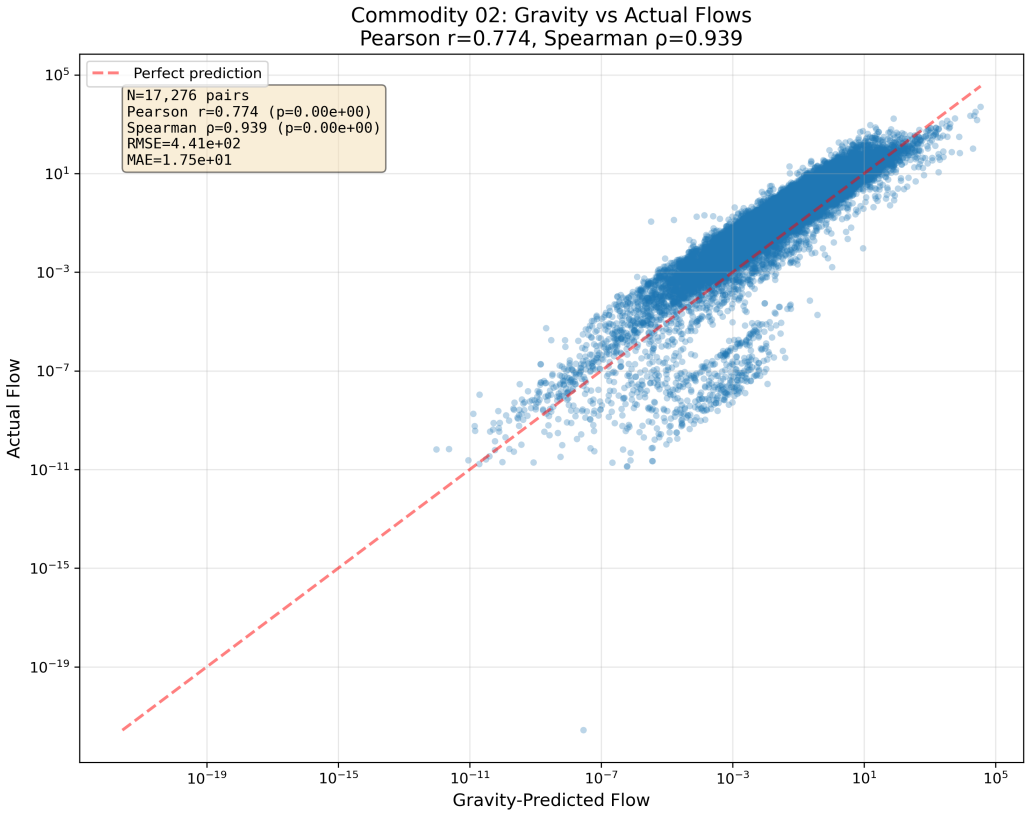


Figure 9: Cereal grains and seeds, gravity

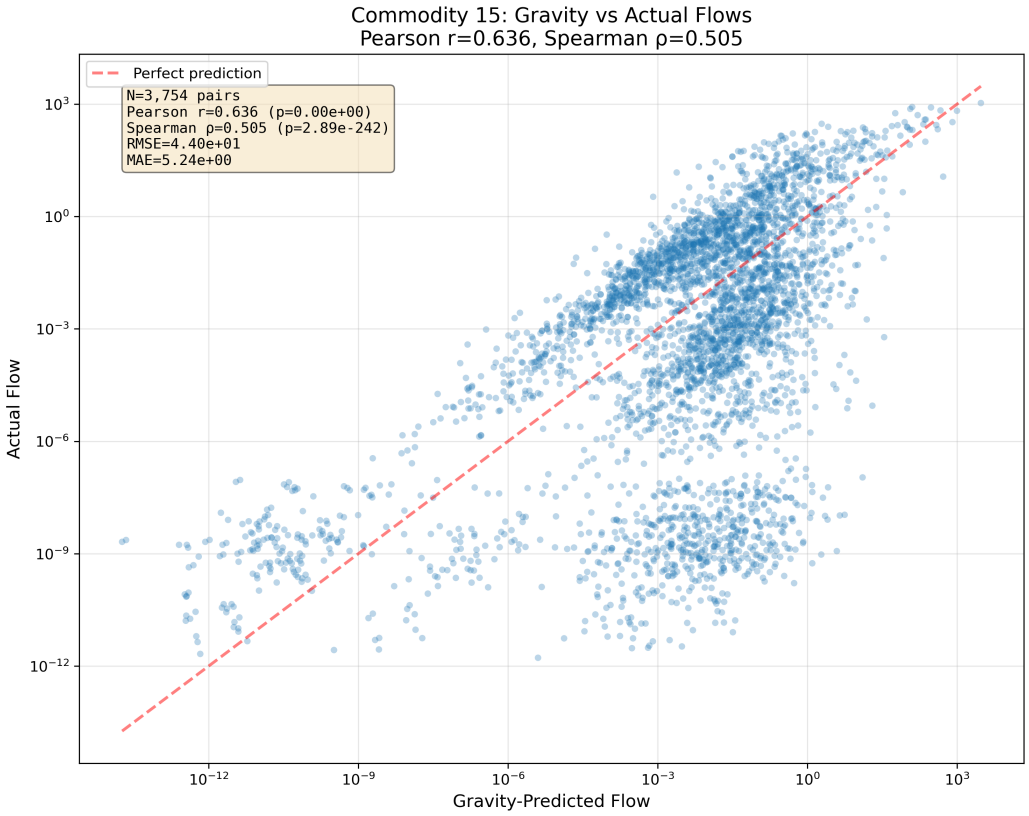


Figure 10: Coal, gravity

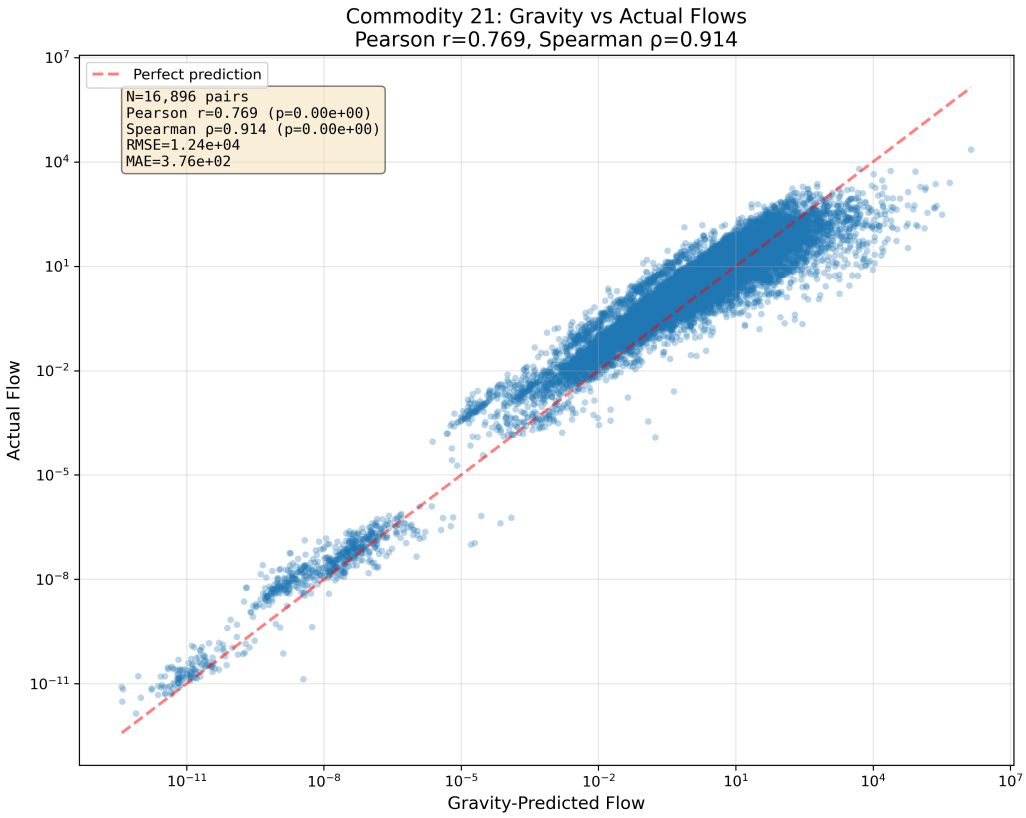


Figure 11: Pharmaceuticals, gravity

6 Discussion

The results highlight several important features of interstate merchandise trade after shipment data are decomposed into economic flows. This section discusses the implications of these findings for understanding U.S. supply chains, the advantages of using inferred producer-to-user flows over raw shipment data, and how the resulting estimates may serve as inputs into broader regional economic analysis.

6.1 Implications for Supply-Chain Structure

The inferred flows provide a clearer and more interpretable picture of U.S. supply chains than raw shipment origin–destination matrices. Shipment data inherently emphasize logistics hubs, port-adjacent zones, and wholesaling centers because they record movements rather than uses. By contrast, the inferred matrices emphasize where goods are *absorbed*, shifting mass toward producer-adjacent destinations for bulk goods and toward manufacturing or population centers for refined and high-value goods. These patterns offer insight into: which states act as true producers and true users of each commodity, the extent to which apparent multi-state movements reflect economic linkages rather than logistics systems, and commodity-specific supply-chain breadth, from highly localized raw material flows to widespread distribution of pharmaceuticals and manufactured intermediates.

6.2 Advantages Over Raw Shipment Data

Using shipment data directly as a proxy for trade can lead to substantial misinterpretation. FAF domestic origin shipments often far exceed economic output for comparable categories. The flow model addresses this by factoring these movements into local absorption versus onward forwarding, reducing double counting and aligning flows more closely with economic interpretation.

Several analytical gains follow:

1. **Reduction of false hubs.** States with large wholesaling sectors but limited consumption or production capacity no longer appear as anomalously large traders.
2. **Improved spatial coherence.** For many commodities, the inferred flows exhibit smoother distance decay and fewer routing artifacts than shipment ODs.
3. **Commodity differentiation.** Distinct spatial structures emerge for raw materials, intermediates, and final goods.

6.3 Applications in Economic and Policy Analysis

The inferred producer-to-user flows can support a range of regional analyses. First, they provide a foundation for understanding how supply disruptions propagate geographically, especially for commodities where flows are corridor-like or concentrated in a few production regions. Second, they can inform state-level analyses of exposure to commodity shocks or energy markets, depending on the commodity's geographic pattern. Finally, the flows can serve as components in the development or refinement of regional input–output or supply–use systems, although such broader systems lie beyond the scope of this paper.

6.4 Limitations

Despite the improvements over raw shipment matrices, the current implementation faces several limitations:

- **Dependence on survey data.** CFS's sampling design—four representative weeks scaled to annual totals—implies that some shipments will be missed, and FAF's out-of-scope expansions may require substantial attribution (Bureau of Transportation Statistics, 2026a; United States Department of Transportation, Bureau of Transportation Statistics, 2025).

- **Minor reliance on constructed shipments.** Reconciling data gaps requires supplementing producer-origin flows in limited cases where wholesaling volumes dwarf observed inbound production.
- **Exclusion of services and retail trade.** This strategy for goods shipment does not yield any insights on the increasingly large share of the economy in services, and what interstate trade looks like there.

6.5 Summary

Overall, the inferred flows produce a more economically meaningful representation of interstate trade in goods than raw shipment data. They highlight producer-to-user relationships, correct for logistics-driven distortions, and uncover commodity-specific spatial structures that align with well-documented freight patterns (Bureau of Transportation Statistics, 2026a,c). These results provide a foundation for applications in regional modeling and supply-chain analysis, while also motivating future extensions to additional flow types and validation strategies.

7 Conclusion

This paper develops a framework for estimating interstate trade in goods by repurposing shipment data into economically meaningful producer-to-user flows. Because shipment datasets record the movement of freight rather than its economic use, raw origin–destination matrices can substantially misstate the geography of trade. They embed wholesaling, warehousing, and other distribution-chain legs, and they incorporate benchmarking and out-of-scope flows that inflate shipment totals relative to economic output, as documented in the CFS and FAF public sources (Bureau of Transportation Statistics, 2026a,b,c; United States Department of Transportation, Bureau of Transportation Statistics, 2025). The resulting patterns may overemphasize logistics hubs and understate the importance of producer-adjacent or consumption-heavy destinations.

The modeling approach introduced here addresses these challenges by decomposing shipments into local absorption and forwarded movements using transparent assumptions about distribution behavior, and by applying a numerical optimization procedure to reconcile producer-origin and non-producer-origin shipments. The framework allows for multi-step routing and cross-hauling and yields a matrix of inferred interstate trade flows that better reflects economic incidence. The resulting flows reassign mass away from distribution intermediaries and toward destinations where goods are ultimately used, revealing commodity-specific spatial structures—from localized raw-material corridors to broad national distribution of manufactured goods.

The inferred flows have several practical implications. They produce a clearer view of regional supply chains, highlight economic linkages that are obscured in shipment data, and provide a foundation for analyzing how shocks may propagate across states for particular commodities. They can also serve as inputs to broader regional accounting systems, including state-level supply–use or input–output frameworks, though the development and reconciliation of such systems lies beyond the scope of this paper.

Much remains to be done. Future work will include formal validation of inferred flows against independent measures of production and use, refinement of producer-versus-intermediary attribution in FAF-extended commodities, and an extension to services trade. The development of time series and incorporation of more granular industry information would further enhance the usefulness of these flows for policy and research applications. Continued improvements in freight data collection, integration, and transparency will also play an essential role.

Despite these remaining challenges, the approach presented here demonstrates that it is feasible to transform shipment-based OD data into a coherent and economically interpretable picture of interstate trade in goods. The framework provides a flexible, data-driven foundation for future regional economic analysis and opens a path toward more comprehensive measurement of domestic trade flows in the United States.

References

Bureau of Transportation Statistics (2026a). Commodity flow survey (cfs). Web page. Primary establishment-level shipment survey; includes origins, destinations, value, weight, and mode.

Bureau of Transportation Statistics (2026b). Commodity flow survey — product page. Web page. Overview and dashboards for 2012 and 2017 CFS.

Bureau of Transportation Statistics (2026c). Freight analysis framework (faf). Web page. FAF5.7.1 data and documentation; OD flows by commodity, mode, value/weight/ton-miles.

Federal Highway Administration (2026). Freight analysis framework (faf) — freight management and operations. Web page. FAF resources: downloads, maps, user guides, model network.

Helpman, E., Melitz, M., and Rubinstein, Y. (2007). Estimating trade flows: Trading partners and trading volumes. Technical Report NBER Working Paper No. 12927, National Bureau of Economic Research, Cambridge, MA. Two-stage gravity framework with extensive and intensive margins; treatment of zero flows and firm selection.

Park, J., Gordon, P., II, J. E. M., and Richardson, H. W. (2008). A two-step approach to estimating state-to-state commodity trade flows. *The Annals of Regional Science*, 43(4):1033–1072. Combines CFS-based flows with adjustment and a Fratar-style update to construct state trade tables (NIEMO).

Phillips, P. (2024). An examination of gravity relationships at the subnational level. Technical Report Working Paper 2024–05–A, U.S. International Trade Commission, Office of Economics, Washington, DC. Subnational gravity estimation across 132 U.S. regions and 42 sectors; distance elasticities vary by sector and mode.

United States Department of Transportation, Bureau of Transportation Statistics (2025). Development of the freight analysis framework version 5 (faf5) annual estimates. Technical report, Bureau of Transportation Statistics, Washington, DC. Methodology report documenting data sources, assumptions, benchmarking, and annual/preliminary estimates.

Yotov, Y. V. (2021). The variation of gravity within countries (or 15 reasons why gravity should be estimated with domestic trade flows). Technical Report CESifo Working Paper No. 9057, CESifo, Munich. Survey arguing for inclusion of domestic flows in structural gravity; implications for identification and policy analysis.