Trade and the Competition for Transport:

How (a Lack of) Competition in the Transportation Industry Affects Regional Trade Outcomes*

C Adam Pfander
University of Colorado, Boulder
Contact: charles.pfander@colorado.edu

Updated: November 18, 2023
Most recent version available here.

Abstract

This paper seeks to understand how market structure for freight transportation affects domestic trade and its outcomes. I modify a Ricardian trade model to incorporate imperfectly-collusive transporters spanning multiple modes. Estimating the model's fundamentals reveals the expected per-mile iceberg trade cost for specific modes, the correlation of trade costs across geographically distinct locations, and the intensity of competition for freight transportation along individual segments of the domestic transit network. Calibrating the model to domestic trade flows, I find that: i) current losses due to non-competitive pricing are substantial, amounting to 5% of baseline welfare; ii) the exercise of freight market power has an outsized impact on exports relative to imports; iii) these impacts are concentrated in rural areas throughout the Southeast and Mountain West, as well as small urban areas in the Midwest and Gulf states; iv) exercise of freight market power exacerbates the impact of mode-specific shocks; and v) non-competitive freight pricing does little to attenuate international trade shocks, due to the concentration of freight market power away from international gates.

*This paper benefited tremendously from the feedback of Anja Gruber, Taylor Jaworski, Wolfgang Keller, Austin Kennedy, Richard Mansfield, and Scott Savage. I also thank the members of the trade, IO, and grad-student brownbags at the University of Colorado, Boulder for their continued input. A special thanks goes to the participants of the graduate student workshop at the Western Economic Conference for their invaluable feedback. I also owe a particular debt of gratitude to my advisor, Jeronimo Carballo, for his constant guidance, support, and patience. I also give thanks to the team behind the R Project for statistical computing (R Core Team, 2022), as well as the authors of the following packages: dplyr (Wickham et al., 2022), magrittr (Bache and Wickham, 2022), data.table (Dowle and Srinivasan, 2023), doParallel (Microsoft and Weston, 2022a), foreach (Microsoft and Weston, 2022b), fixest (Bergé, 2018), DescTools (Andri et mult. al., 2022), stringr (Wickham, 2022), ggmap (Kahle and Wickham, 2013), dqrng (Stubner, 2021), and RPostgres (Wickham et al., 2023). I also acknowledge the creators of and contributors to PostgreSQL.

All remaining errors are my own.

1 Introduction

Transportation is central to trade —how a good moves from its origin to its destination greatly influences its final price. Yet conventional trade models abstract away from the freight-service industry by assuming exogenous transport costs.[1] Recognizing this insufficiency, a raft of recent papers has sought to incorporate more realistic transportation sectors into canonical trade models. Specifically, these recent works model route-choice (Allen and Arkolakis, 2022; Allen and Arkolakis, 2014), mode-choice (Fuchs and Wong, 2023), and port-choice (Brancaccio et al., 2020) to generate trade costs as a function of existing infrastructure. However, the role of transporters' strategic price-setting behavior remains unexplored. Hence, in this paper, I analyse the domestic trade impacts of market power among freight service providers.

I develop a Ricardian trade model, in which each origin-destination combination is served by a representative freight transporter of a particular mode. Multiple modes – namely, Trucking (or, "Road"), Rail, Water, Air, and Multi-Modal – service each origin-destination pairing,[2] and each of these intermediaries may pick among a selection of routes per mode. Transporters need not take the shortest path, but may elect more circuitous routes if they appear more profitable due to unobserved, exogenous cost components. Further, transporters set markups endogenously to maximize profits; the key inputs into the transporter's problem include demand for freight transport between any origin-destination pairing, the costs incurred en-route, and the degree of freight-market competition in their geographic region. Regarding competition among transporters: I do not take a stand a priori on a particular solution to the transporter's pricing game, but adapt Bresnahan (1982)'s method to develop a pricing rule that admits a wide array of competitive outcomes, including the extremes of perfect competition and perfect monopoly.[3] Additionally, to accurately model demand for freight transport, I follow Lind and Ramondo (2023) to incorporate correlation of prices across competing origins, modes, and routes, which permits more realistic substitution patterns across competing trade routes.[4] The model thus provides a tractable framework to both estimate the

[1] This strict exogeneity assumption stems from Samuelson (1954), who first proposed an exogenous iceberg trade cost formulation in order to generate a tractable expression of trade flows. While convenient, this assumption is undoubtedly an oversimplification

[2] I emphasize that the set feasible of modes into any one location is determined by available infrastructure – Water transport is not possible without a navigable Waterway.

[3] A more recent application of a conduct-parameter approach is Miller and Weinberg (2017).

[4] The assumption of perfectly independent trade costs, made originally for analytical convenience (Eaton and Kortum, 2002), implies unrealistic substitution patterns, and thus undermines the reliability of counterfactual simulations

extent of freight market power, and to evaluate how both trade flows and transport prices adjust to exogenous shocks.

The model yields five principal insights, three of which are descriptive. First, I find that current losses due to non-competitive freight pricing are substantial, equaling approximately 5% of baseline welfare. The lion's share of these losses stem from the Road, a finding supported by anecdotal evidence of pervasive bottlenecks throughout the trucking industry, which lend market power to available operators. Second, I document significant geographic heterogeneity in the exercise of freight market power. Potential welfare gains from elimination of freight market power are concentrated in rural areas in the Southeast and Mountain West, as well as small urban areas throughout the Midwest and along the Gulf of Mexico; major international gates stand the least to gain from a perfectly-competitive freight market, as these markets currently exhibit the least concentration of market power. Third, the exercise of freight market power has an outsized influence on exports; imposing perfectly-competitive freight pricing may increase exports by as much as 20% but imports only 7%. This result suggests that managing non-competitive freight pricing offers a new lever to influence the U.S. trade deficit, which is of substantial interest to politicians and policy makers. Together, these results demonstrate how the current exercise of freight market power shapes domestic trade outcomes.

My final two principal results pertain to the simulated impact of exogenous transport shocks. Specifically, my fourth finding demonstrates that shutting down a single mode – perhaps due to a labor strike – sparks dramatic welfare losses; the pricing response among remaining modes exacerbates these losses.[5] This result implies that, in the event of national, mode-specific shocks, we should expect a substantial pricing response from competing modes. Finally, due to the concentration of freight market power away from heavily-trading regions, I report near-perfect pass-through of a negative international trade shock; inland areas along sparse portions of the network see the greatest (relative) losses. This result is important, as it underscores that one of the few potential benefits of a concentrated freight market – namely, imperfect pass-through of exogenous shocks – remains unrealized in the case of international shocks. Coupled with the large welfare losses reported in baseline, as well as the exacerbating effect of mode-specific shocks, I conclude that freight

(Train et al., 1987). See Lind and Ramondo (2023) for a discussion of correlated prices in Ricardian models.

[5] The analysis that yields this result is inspired by the narrowly-avoided Rail strike of 2022, which would have halted all Rail traffic and the vast majority of multi-modal movements.

market power offers little upside at substantial cost.

The natural question, then, is what is the appropriate policy response; the model offers some insight here as well. I document that market power is concentrated along trucking and within rural and small, urban areas. This finding accords with anecdotal evidence of pervasive bottlenecks throughout the trucking industry, granting market power to available operators (Costello and Karickhoff, 2019).[6] Moreover, I demonstrate that multi-modal transport is trucking's closest competitor; expanding the availability of multi-modal transshipment exchanges will offset the market-power of road transport. These interventions should be targeted to remote areas along sparse portions of the domestic transit network, where market power is concentrated.

The model offers highly-granular, route-level insights; it is therefore notable that I estimate these fundamentals using exclusively public data. To elaborate, I populate potential trade routes using detailed topographic information on the domestic transportation network from the National Transportation Atlas Databases (NTAD); I retrieve domestic trade flows in 2017 from the Freight Analysis Framework version 5 (FAF5). Both of these datasets are publicly-available from the Bureau of Transportation Statistics (BTS). However, the FAF5 data are aggregated by origin, destination, and mode; it does not provide route-level trade flows, which are necessary for estimation. I therefore rely on an expectation-maximization (EM) process – a well-known simulated method of moments strategy popular in computer science, which has been used to some extent in the field of economics (e.g. Bonadio, 2021). No existing paper provides such a comprehensive and detailed view of the transportation sector using only public information.[7]

This paper contributes to a growing body of research that explores the impact of freight transportation and infrastructure on trade outcomes. Recent papers along this line of inquiry have analysed the welfare effects of road and/or port congestion (Allen and Arkolakis, 2022; Bonadio, 2021;

[6]The American Trucking Associations points specifically to a severe shortage of qualified drivers causing disruptions throughout the industry. They also outline specific interventions designed to increase the supply of drivers, including: i) increasing recruitment and retention of women among the driving workforce; ii) relaxing strict drug-screening requirements; iii) lowering the minimum age required to drive freight across state lines; iv) increasing time at-home for long-haul truckers; v) increasing available parking stops and decreasing congestion at truck stops along major freight corridors; and vi) relaxing high standards as regards to driving records and criminal histories (ATA, October 25, 2021).

[7]Other papers estimate detailed trade costs using confidential data. Most recently, Fuchs and Wong (2023) utilize the confidential waybill sample to estimate the impact of Rail (as well as port) congestion throughout the domestic transport network; Brancaccio et al. (2020) estimate bargained freight rates using proprietary transaction data from Clarkson's Research. Others utilize public data, but narrow their focus to a subset of modes – for example, Allen and Arkolakis (2022) focus on roads; Donaldson (2018) focus on Rails; Ducruet et al. (2020) focus on shipping networks.

Fuchs and Wong, 2023), containerization (Coşar and Demir, 2018), and most popularly, network development/expansion (Jaworski et al., 2023; Ganapati et al., 2021; Donaldson, 2018; Ducruet et al., 2020; Donaldson and Hornbeck, 2016; Faber, 2014). However, all of these papers analyze only one or two modalities – e.g. roads, Rail, or ports/ships. In contrast, I model all available modes simultaneously. Capturing the multi-faceted nature of the US transit network is important when studying the transporter's pricing decision, as competition weighs both within and across modes. Put differently, the consumer likely does not care *how* a good got to its destination, only that it arrived and it is cheaper than all (quality-adjusted) alternatives. To stay cost-competitive, modes must engage with one another.[8] My primary contribution to this literature is to model competition among freight service providers, capturing price pressures both within and across modes.

The importance of multi-modal transport to domestic trade is highlighted in the parallel work by Fuchs and Wong (2023). In this important paper, the authors build upon a highly-tractable model of traffic flows (Allen and Arkolakis, 2022) to encompass multiple modes and costly transshipment. The upshot of this approach is that the authors may quantify how congestion along any one mode spills into the wider network; importantly, they identify key bottlenecks throughout the domestic transit network and estimate how even slight investments at these choke points create substantial welfare gains throughout the nation. This paper underscores the importance of multi-modal shipping to domestic trade outcomes; however, the authors do not consider the role of competition in setting freight rates. I build upon their framework by incorporating the transporter's strategic pricing decision.

The present paper is most closely related to Brancaccio et al. (2020). There, the authors model the transport market with a matching function; freight rates are set via Nash bargaining. The paper is the first to model how transporters' strategic behavior impacts trade flows; however, their focus is solely on ships in the international freight market. In contrast, I focus on the multi-modal domestic transport network. Further, their matching function approach – which they utilize to explain ships' ballasting choices – is governed by one, global bargaining parameter. In contrast, I evaluate the state of market power within distinct markets throughout the mainland U.S.. My

[8] Throughout the latter half of the 20th Century, domestic freight in the US was dominated by trucking. However, recent technological improvements, which facilitate near-seamless transshipment, have given new life to competing modes; multi-modal movements are now the second-most popular form of transit domestically (Bureau of Transportation Statistics and Federal Highway Administration, U.S. DOT, 2022).

model thus offers a highly-granular assessment of how competition for transport within individual locations affects trade locally, and nationally.

The rest of the paper is organized as follows. Section 2 develops my theoretical model, distinguishing between transport demand – equilibrium trade flows – and supply – a model of transporter's optimal pricing rule. Section 3 details my strategy for estimating the model fundamentals; it presents my data sources and expounds upon the EM algorithm, which exploits the structure of the model to estimate route-level trade costs from aggregated data. Section 4 discusses my estimated parameters, model validation exercises, and robustness checks. Section 5 details my counterfactual simulations. Section 6 concludes.

2 Model

My theoretical model may be described simply as a general equilibrium model for freight transport services. On the demand side, I modify a standard Ricardian trade model to account for: i) multiple modes serving every origin-destination pairing; ii) multiple routes per mode; iii) correlation of prices across origins, modes, and routes; and iv) an endogenously-determined transportation markup set by freight-service firms. The trade flows implied by this model constitute the demand for freight transportation. On the supply-side, I develop an approximately linear pricing rule that expresses these markups as a function of (observable) trade shares and (estimable) correlation parameters. I now detail the model.

2.1 Transportation Demand: Equilibrium Trade Flows

I begin with the familiar Eaton and Kortum (2002) framework, henceforth termed "EK". Trade in a unit-mass of goods Ω occurs amongst a discrete, finite set of locations $\mathcal{S}$; let $i, j \in \mathcal{S}$ denote the ultimate origin, destination. I further distinguish between domestic locations $\mathcal{S}^d \subset \mathcal{S}$ and foreign locations $\mathcal{S}^f \subset \mathcal{S}$, noting that $\mathcal{S}^d \cap \mathcal{S}^f = \emptyset$. Let $e, x \in \mathcal{S}^d$ denote a shipment's domestic origin, destination; in the case of international trade, this will denote the shipment's port of entry, exit.[9]

All domestic locations lie on a graph and are connected by a set of edges. A route between a domestic origin and destination, $e, x \in \mathcal{S}^d$, comprises a set of edges that form an unbroken

[9] Note that, for domestic origins (destinations), the choice of port of entry (exit) is degenerate.

path between e and x. Routes need not be direct, but I do require them to be simple – that is, a route may take a circuitous path on its way from e to x, but no repeating edges.[10] Let $\mathcal{R}_{ex}$ denote the set of all routes linking e to x. I further decompose this set by mode of transport. Let $\mathcal{M} = \{Road, Rail, Water, Air, Multi\text{-}modal\}$ denote the set of available modes and let $\mathcal{R}_{ex}^m \subseteq \mathcal{R}_{ex}$ denote the subset of routes between e and x via mode $m \in \mathcal{M}$.[11] Figure 1 provides an illustrative example of the domestic transit network.

Traversing the network is costly. Let $\tau_{ir(e,x,m)j}$ denote the cost of traversing route $r \in \mathcal{R}_{ex}^m$ between i and j;[12] these transit costs take the standard iceberg formulation. I further assume that each origin enjoys a uniform input cost c_i and employs constant returns to scale technology. Finally, I impose perfect competition in the goods market, such that the price of good ω, produced in i, consumed in j, and taking route r via mode m between the two is given by

$$p_{irj}(\omega) = \frac{c_i \tau_{ijr} \mu_{irj}}{\epsilon_{irj}(\omega)}, \tag{1}$$

where μ_{irj} is an endogenous markup set by the transporter, and $\epsilon_{ijr}(\omega)$ is a composite production and transportation efficiency, which I will elaborate shortly. But first, I want to highlight that prices vary not only by their origin, but also by the route taken to market. Thanks to perfect competition in the market for tradeable goods, the price paid by the consumer is the cheapest available alternative. That is:

$$p_j(\omega) = \min_{i \in \mathcal{S}} \left\{ \min_{e,x \in \mathcal{S}^d} \left\{ \min_{m \in \mathcal{M}} \left\{ \min_{r \in \mathcal{R}_{ex}^m} p_{ijr}(\omega) \right\} \right\} \right\}.$$

As in the traditional EK model, I rely on a probabilistic representation of composite efficiencies to relate trade flows across locations and routes to a handful of easy-to-interpret parameters. Unlike the standard model, I follow McFadden (1984) to allow correlation of efficiencies within distinct destinations, origins, ports of entry/exit, and modes. Explicitly, the joint distribution function for

[10] Note also that routes are direction-dependent: the set of edges forming a path from e to x is distinct from the same path moving from x to e, though their observable characteristics are identical.

[11] It is important to note that the set of modes is exhaustive, such that $\mathcal{R}_{ex} = \cup_{m \in \mathcal{M}} \mathcal{R}_{ex}^m \; \forall \; e, x \in \mathcal{S}^d$. Additionally, with the exception of Multi-modal transport, these paths are mutually exclusive, $\mathcal{R}_{ex}^m \cap \mathcal{R}_{ex}^{m'} = \emptyset \; \forall \; m, m' \in \mathcal{M} \setminus \{Multi\text{-}Modal\}, m \neq m'$, and $\forall \; i, j \in \mathcal{S}$.

[12] Note that knowledge of the route implies knowledge of the domestic origin e, destination x, and mode m. For ease of exposition, I will omit the extraneous subscripts where expedient.

the vector of log efficiencies is given by:

$$F_j\begin{pmatrix}\ln\epsilon(\omega)_1\\ \ln\epsilon(\omega)_2\\ \vdots\end{pmatrix}=\exp\left(-\sum_{i\in\mathcal{S}}\left(\sum_{e,x\in\mathcal{S}^d}\left(\sum_{m\in\mathcal{M}}\left(\sum_{r\in\mathcal{R}^m_{ex}}\exp\left(\frac{-\ln A_i-\theta\ln\epsilon_r(\omega)}{\varphi\rho\zeta}\right)\right)^{\varphi}\right)^{\rho}\right)^{\zeta}\right). \tag{2}$$

In words, I assume a generalized extreme value distribution with location parameter A_i and scale parameter θ. The parameters $\varphi, \rho,$ and $\zeta \in (0,1]$ respectively govern the correlation of efficiencies within distinct nests— respectively, modes, ports of entry/exit, and origins. A value of 1 implies no correlation among these efficiencies within a given nest, while 0 implies perfect correlation. Due to this correlation structure, the interpretation of parameters is somewhat altered from the original EK model. A_i governs absolute advantage, while $\varphi, \rho,$ and $\zeta,$ along with θ jointly determine the prevalence of comparative advantage.[13] An illustration of the nesting structure are presented in Figure 2.

Equation (2) is a generalization of the hallmark Fréchet assumption from the canonical Eaton and Kortum model. This formulation is attractive for two key reasons. First, this logit demand formulation treats transport along different routes as differentiated services. Hence, trade from competing origins, modes, and even different routes serving the same market, function as imperfect substitutes. Second, the nesting structure dispenses with the assumption that efficiency shocks are independent across locations; rather, the model permits a highly flexible correlation structure across destinations, origins, ports of entry/exit, and modes. This correlation structure relaxes the independence of irrelevant alternatives (IIA) assumption inherent in the standard logit model, which implies strict and often unrealistic substitution patterns. In this model, I require IIA to hold within, but not across nests.[14]

I assume the following functional form for trade costs:

$$\tau_{ir(e,x,m)j}=\exp\Big(\underbrace{\mu_{irj}}_{\text{markup}}+\underbrace{\kappa_r}_{\text{deterministic costs}}+\underbrace{u_{irj}}_{\text{unobserved quality}}+\underbrace{\eta_{ie}+\nu_{xj}}_{\text{international trade cost}}+\underbrace{\alpha_{iem}}_{\text{first mile}}+\underbrace{\gamma_{xmj}}_{\text{final mile}}\Big) \tag{3}$$

[13]See Lind and Ramondo (2023) for a complete discussion of how allowing correlation of efficiencies changes the interpretation of the EK model.

[14]To illustrate the potential pitfalls of IIA, consider the following, stylized example. Imagine that a transporter introduces a new Rail route between an arbitrary origin and destination. Naturally, this new route will siphon traffic away from existing trade routes. IIA requires that the new route will draw from other routes in proportion to their trade shares, such that the ratio of trade shares among existing routes remains constant. However, it seems far more likely that the new Rail line will disproportionately draw trade away from other Rail routes, and have less of an impact on trade via other modes. The nested structure of my model captures this more realistic substitution pattern.

I now detail each component. Transporters set μ_{irj} to maximize profit, based on the demand and competition for transport services in geographically distinct markets (see Section 2.2 for detail). The deterministic component, κ_r, reflects the costs incurred per-mile while moving along a route – practically, this includes the cost of fuel, labor, and maintenance, and is taken as exogenous. The parameter u_{irj} is a structural error term and reflects unobserved route quality. Following Berry (1994), I assume $u_r \sim N(0, v)$, and is drawn independently across routes. The parameters η_{ie} and ν_{xj} capture costs incurred along the international leg of the journey (if any). Explicitly, η_{ie} presents the cost of travelling from an international origin i to a domestic port of entry e, while ν_{xj} is the cost of travelling from a domestic port of exit x to a foreign destination j. Finally, the parameters α_{iem} and γ_{xmj} capture a wide array of otherwise unobserved location-by-mode-specific costs. These include the cost of entering/exiting a domestic location via mode m, and reflect, e.g., port, yard, and/or inner-city congestion. These parameters also capture stand-alone, mode-specific costs, such as loading times, product composition, network-wide congestion, and/or reliability.

The distributional assumption on composite efficiencies, coupled with perfect competition in the market for tradeable goods, yields a nested logit[15] formulation for trade shares within each market j:[16]

$$\pi_{r|iexmj} = \exp\big(-\theta[\mu_{irj} + \kappa_r + u_{irj}]/(\varphi\rho\zeta) - J_{iexmj}\big) \tag{4}$$

$$\pi_{m|iexj} = \exp\big(-\theta[\alpha_{iem} + \gamma_{mxj}]/(\rho\zeta) + \varphi J_{iexmj} - V_{iexj}\big) \tag{5}$$

$$\pi_{ex|ij} = \exp\big(-\theta[\eta_{ie} + \nu_{xj}]/\zeta + \rho V_{iexj} - I_{ij}\big) \tag{6}$$

$$\pi_{i|j} = \exp\big(-\ln A_i - \theta \ln c_i + \zeta I_{ij} - Q_j\big) \tag{7}$$

where

$$J_{iexmj} = \ln \sum_{r \in \mathcal{R}^m_{ex}} \exp\big(-\theta[\mu_{irj} + \kappa_r + u_{irj}]/(\varphi\rho\zeta)\big) \tag{8}$$

$$V_{iexj} = \ln \sum_{m \in \mathcal{M}} \exp\big(-\theta[\alpha_{iem} + \gamma_{jxm}]/(\rho\zeta) + \varphi J_{iexmj}\big) \tag{9}$$

$$I_{ij} = \ln \sum_{e,x \in \mathcal{S}^d} \exp\big(-\theta[\eta_{ie} + \nu_{xj}]/\zeta + \rho V_{iexj}\big) \tag{10}$$

[15]Note that if any one correlation parameter equals 1, then that nest effectively dissolves and is absorbed into the higher nest; if $\varphi = \rho = \zeta = 1$, the model collapses to a standard EK trade model with trade elasticity θ (McFadden, 1984; Eaton and Kortum, 2002).

[16]See Train et al. (1987) for detailed derivations.

$$Q_j = \ln \sum_{i \in \mathcal{S}} \exp\big(- \ln A_i - \theta \ln c_i + \zeta I_{ij} \big). \tag{11}$$

Combining these conditional probabilities yields the following log-linear expression for the trade share along any one route:

$$\begin{aligned}
\pi_{ir(e,x,m)|j} =& \pi_{r|iexmj} \pi_{m|iexmj} \pi_{ex|ij} \pi_{i|j} \\
=& \exp\Big(- \ln A_i - \theta \big[\ln c_i + (\mu_{irj} + \kappa_r + u_{irj})/(\varphi\rho\zeta) + (\alpha_{iem} + \gamma_{xmj})/(\rho\zeta) + (\eta_{ie} + \nu_{xj})/\zeta \big] \\
& - (1-\varphi) J_{iexmj} - (1-\rho) V_{iexj} - (1-\zeta) I_{ij} - Q_j \Big).
\end{aligned}$$

Some algebra yields:[17]

$$\begin{aligned}
\pi_{ir(e,x,m)|j} =& \exp\Big(- \ln A_i - \theta \big[\ln c_i + \mu_{irj} + \kappa_r + u_{irj} + \alpha_{iem} + \gamma_{xmj} + \eta_{ie} + \nu_{xj} \big] \\
& + (1 - \varphi\rho\zeta) \ln \pi_{r|iexmj} + (1-\rho\zeta) \ln \pi_{m|iexj} + (1-\zeta) \ln \pi_{ex|ij} - Q_j \Big).
\end{aligned} \tag{12}$$

Finally, consumers enjoy CES utility with elasticity parameter $\sigma \in (1, \theta + 1]$. Explicitly, consumers demand a quantity of each good, $q(\omega)$ to maximize $U = \left(\int_\Omega q(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}}$ subject to a budget constraint that aggregates to a trade balance condition. That is, the sum total value of imports must equal the sum total value of exports in location j. The assumption of CES utility implies a CES price index in location j (see Appendix A.2 for detail). Further, I adhere to the structure laid out in Eaton and Kortum (2002) and assume a Cobb-Douglas aggregate of input costs, with labor having a constant share β. These conditions yield expressions for input costs, prices, and wages:

$$c_i = w_i^\beta p_i^{1-\beta} \tag{13}$$

$$p_j = \Gamma \left(\frac{\theta + 1 - \sigma}{\theta} \right)^{\frac{1}{1-\sigma}} Q_j^{-\frac{1}{\theta}} \tag{14}$$

$$L_i w_i = \sum_{k \in \mathcal{S}} \sum_{e,x \in \mathcal{S}} \sum_{r \in \mathcal{R}_{ex}} \pi_{ir|k} L_k w_k \tag{15}$$

where w_i, L_i are the wage, population in location i, and Γ denotes the gamma function. Taken together, Equations (4) - (15) characterize a modified Ricardian model of trade, allowing for a multiplicity of trade routes between locations, imperfect modal substitution, and imperfect competition for freight transport services.

[17] see Appendix A.1 for detail.

2.2 Transportation Supply: Setting Markups

Given appropriate data, estimating Equation (12) is straightforward. However, nationally-representative data on costs and/or prices in the freight industry are not publicly available (see Section 3.1 for detail), meaning route-level markups μ_{irj} are unknown ex-ante. I thus impose additional structure on the transporter's pricing decision, the goal of which is to parameterize μ_{irj} such that: i) markups are a function of observable data; and ii) I do not assume a particular solution to the transporter's pricing game, but permit a wide array of potential equilibria, including the extremes of perfect competition and perfect monopoly. Research in empirical IO provides such a framework (Bresnahan, 1982; Miller and Weinberg, 2017); I now detail the structure that yields this result.

A representative transporter of mode m serves an origin-destination pairing ij. Given the scale of the transporter's problem, it is convenient to characterize its total profit in matrix notation. Accordingly, define $\boldsymbol{\pi}_{imj}^{\top} = \begin{bmatrix} \pi_{ir_1|j} & \pi_{ir_2|j} & \dots \end{bmatrix}$, $\boldsymbol{\mu}_{imj}^{\top} = \begin{bmatrix} \mu_{ir_1j} & \mu_{ir_2j} & \dots \end{bmatrix}$, and $\boldsymbol{\tau}_{imj}^{\top} = \begin{bmatrix} \tau_{ir_2j} & \tau_{ir_2j} & \dots \end{bmatrix}$, the vectors of market shares, markups, and iceberg trade costs along all routes carrying freight from i to j via m. Total profits are thus:

$$\Pi_{imj} = B_j \boldsymbol{\pi}_{imj}^{\top} \left[(\boldsymbol{\mu}_{imj} - \mathbf{1}) \circ \boldsymbol{\tau}_{imj} \right]$$

where B_j is the total spending of destination j, which the transporter treats as exogenous, and $\circ$ denotes the element-wise matrix product. It follows that the transporter's marginal profit with respect to its vector of markups $\boldsymbol{\mu}_{imj}$ is given by:

$$\begin{aligned} \nabla\Pi_{imj} &= \frac{\partial \Pi_{imj}(\pi_{ir_1|j}, \pi_{ir_2|j}, \dots)}{\partial(\mu_{ir_1j}, \mu_{ir_2j}, \dots)} \\ &= \left(\frac{\partial(\ln \mu_{ir_1j}, \ln \mu_{ir_2j}, \dots)}{\partial(\mu_{ir_1j}, \mu_{ir_2j}, \dots)} \right) \left(\frac{\partial \Pi_{imj}(\pi_{ir_1|j}, \pi_{ir_2|j}, \dots)}{\partial(\ln \mu_{ir_1j}, \ln \mu_{ir_2j}, \dots)} \right) \\ &= B_j \left(\mathbf{DS}_{imj}^{\top} (\boldsymbol{\mu}_{imj} - \mathbf{1}) + \lambda_i \boldsymbol{\pi}_{imj} \right) \circ \frac{\boldsymbol{\tau}_{imj}}{\boldsymbol{\mu}_{imj}} \end{aligned}$$

where $\mathbf{DS}^{\top} = \begin{bmatrix} \frac{\partial \pi_{ir_1|j}}{\partial(\ln \mu_{ir_1j}, \ln \mu_{ir_2j}, \dots)} & \frac{\partial \pi_{ir_2|j}}{\partial(\ln \mu_{ir_1j}, \ln \mu_{ir_2j}, \dots)} & \dots \end{bmatrix}$ is the Jacobian matrix of trade shares with respect to the vector of log markups, while $\lambda_i \in [0, 1]$ captures the transporter's pricing conduct. This λ_i parameter reflects the level of freight market competition for transport services out of i. $\lambda_i = 0$ corresponds to Bertrand competition: freight markets are perfectly competitive and markups as a percent of marginal cost go to 1. In contrast, λ_i parameters different from zero point to deviations from Bertrand competition, suggesting exercise of freight market power; at the extreme, $\lambda_i = 1$ corresponds to a perfect monopoly (oligopoly) in market i. This parameterization of the transporter's profit-maximizing behavior thus permits a wide variety of solution concepts.

Manipulating the transporter's first-order condition yields the pricing rule:

$$\begin{aligned}
\mathbf{0} =& B_j \left(\mathbf{DS}_{imj}^{\top} (\boldsymbol{\mu}_{imj} - \mathbf{1}) + \lambda_i \boldsymbol{\pi}_{imj} \right) \circ \frac{\boldsymbol{\tau}_{imj}}{\boldsymbol{\mu}_{imj}} \\
\boldsymbol{\mu}_{imj} =& \mathbf{1} - \lambda_i \big(\mathbf{DS}_{imj}^{\top}\big)^{-1} \boldsymbol{\pi}_{imj} \\
\ln \boldsymbol{\mu}_{imj} \approx & - \lambda_i \big(\mathbf{DS}_{imj}^{\top}\big)^{-1} \boldsymbol{\pi}_{imj}
\end{aligned}$$

where this last step follows from a first-order Taylor series expansion about the perfectly-competitive equilibrium.[18] It will prove expedient to define:

$$\boldsymbol{\delta}_{imj} = \begin{bmatrix} \delta_{ir_1 j} \\ \delta_{ir_2 j} \\ \vdots \end{bmatrix} = -\theta \big(\mathbf{DS}_{imj}^{\top}\big)^{-1} \boldsymbol{\pi}_{imj} \tag{16}$$

Thanks to the logit-demand structure, the $\mathbf{DS}_{imj}^{\top}$ matrix may be easily calculated from the (observable) trade shares and correlation parameters; see Appendix A.3 for detail. Markups along an individual route are given by

$$\ln \mu_{irj} \approx (\lambda_i / \theta) \delta_{irj}. \tag{17}$$

I thus arrive at a simple, linear expression for markups that encompasses a wide array of solution concepts. Importantly, the linear markup input δ_{irj} may be easily calculated from observable data.

3 Empirical Strategy

Having laid out the theory, I turn my attention to estimating the model fundamentals. In this section, I demonstrate how the structure developed through Section 2 yields a simple, linear estimating equation that identifies: i) expected trade costs per mile, by mode; ii) the current severity of freight-market competition within and across distinct modes; and iii) the structural correlation parameters ρ and ζ. I also detail the data used to estimate the model. Notably, my model charac-

[18] To be clear, the perfectly-competitive outcome corresponds to $\mu_{imj} = 1$.

terizes trade along individual routes, while the data are aggregated by origin, destination, port of entry/exit (if any), and mode; I thus develop and detail a method to estimate model fundamentals from aggregated data.

3.1 Data

I utilize domestic trade data from the Freight Analysis Framework, version 5 (FAF5), compiled by the Bureau of Transportation Statistics (BTS) and Federal Highway Administration (FHA). This dataset lists total annual trade volumes in 2017 disaggregated by mode among major metropolitan areas in the U.S. and the following international regions: Canada, Mexico, Rest of Americas, Europe, Africa, Southwest and Central Asia, Eastern Asia, Southeast Asia, and Oceania.[19] The FAF also encompasses large, rural areas in each state. Figure 3 displays the domestic regions. I restrict my analysis of mode and route choice to the mainland U.S. due to the significant natural and political boundaries facing Hawaii and Alaska. For the same reasons, I do not model mode or route choice for the international portions of a journey. I thus analyse transportation among 129 domestic locations, as well as the eight international regions. The data are aggregated by origin, port of entry (for imports), port of exit (for exports), destination, and mode.[20] I restrict my analysis to five domestic modes: Road, Rail, Water, Air, and Multi-Modal; I do not consider freight movements via pipeline or with an unknown domestic mode. I also remove trade flows valued at less than $1 million total in 2017, to ensure enough variation in each origin-destination pairing to reliably identify the parameters of interest. In summary, the FAF provides total annual domestic trade flows in 2017, disaggregated by origin, destination, port of entry/exit (if any), and mode.

Second, I utilize detailed geographic information on the domestic transit network from the National Transportation Atlas Database (NTAD). The NTAD provides highly granular detail on domestic roadways, Rail lines, navigable Waterways, Airports, and intermodal exchanges. I assume that travel may occur directly between any Airport. Along the roadway, I limit my attention to arterial roads and interstates, as these carry the most interstate commerce. From these disparate modal networks, I create a graphical representation of a single, multi-modal transport network.

[19] Unfortunately, the FAF does not offer more granular international geographies.

[20] It is important to note that the Air freight included in the FAF5 encompasses shipments that are serviced by both road and Air. The reason for this aggregation is that Air typically cannot perform final-mile services, and thus relies on trucks to complete the trip. This aggregation is not a problem when estimating my model, as I account for final-mile costs via destination $\times$ mode fixed effects.

Importantly, exchanges between modes may only occur at designated intermodal exchanges. Figure 1 provides a simplified example of my geography; Figure 4 displays a map of my network (excluding routes via international Waters).

3.2 Estimating Model Parameters

I begin with my ideal regression specification. I parameterize the deterministic component of trade costs as follows:

$$\kappa_{r(m)} = Miles_r \beta_m, \tag{18}$$

where $Miles_r$ is the length of the route in thousands of miles, and β_m captures the expected cost per-mile along a particular mode. Plugging Equations (17) and (18) into Equation (12) yields the following log-linear expression:

$$\begin{aligned}\pi_{ir(e,x,m)|j} = \exp\Big(& -\ln A_i - \lambda_i \delta_{irj} - \theta \big[\ln c_i + Miles_r \beta_m + u_{irj} + \eta_{ie} + \nu_{xj} + \alpha_{iem} + \gamma_{xmj} \big] \\ & + (1 - \varphi\rho\zeta) \ln \pi_{r|iexmj} + (1 - \rho\zeta) \ln \pi_{m|iexj} + (1 - \zeta) \ln \pi_{ex|ij} - Q_j \Big). \end{aligned} \tag{19}$$

Given appropriate data, this expression provides a convenient, log-linear framework to estimate the model parameters.

3.2.1 Identification

Before detailing my estimation procedure, I must emphasize a few points regarding identification. First, because the log conditional probabilities $\ln \pi_{r|iexmj}$, $\ln \pi_{m|iexj}$, and $\ln \pi_{ex|ij}$ are included to identify the correlation coefficients, the cost and conduct parameters, β_m and λ_i, are identified by observed variation in trade shares across origins and destinations; similarly, ρ and ζ are respectively identified by observed variation in conditional trade shares across modes, domestic ports of entry. The parameter φ, however, cannot be estimated, as I do not observed trade flows at the route-level in my data.[21] I thus fix this parameter to 1 at the outset, implying zero correlation of transport efficiencies within origin, mode pAirs. In Section 4.2, I show how the results change with lower values of φ (i.e. non-zero correlations among routes of the same origin, mode).

[21] I estimate my route-level model using a simulated method of moments strategy; see Section 3.2.2 for detail.

Second, Equation (19) makes evident that all cost coefficients are identified up to scale; I cannot separately identify θ. This scaling is not strictly a problem in the context of my model, but requires that I rely on estimates of θ from the literature when calibrating wages and performing counterfactual analysis. Henceforth, I assume a standard trade elasticity of 4 (Simonovska and Waugh, 2014).[22]

Third, there is the natural issue of endogeneity of the markup term δ_{irj}. This regressor is endogenous for two reasons: i) the pricing rule laid out in Equations (16) and (17) makes clear that δ_{irj} is structurally correlated with u_{irj}, and ii) the linear approximation of markups introduces further error into the regression equation, which will necessarily be correlated with δ_{irj}. I overcome these identification problems with an instrumental variables strategy.

To instrument the endogenous markup com, I propose the cost of competing transport services, which, in this context, is a log-linear function of the distance of routes along competing modes. Explicitly, I utilize the following instruments:

$$\boldsymbol{Z}_{ir(e,x,m)j} = \begin{bmatrix} \sum_{r' \in \mathcal{R}_{ex}^{m'}} \mathbb{1}\big(k(r) = k(r')\big) Miles_{r'} & \ldots \end{bmatrix} \tag{20}$$

for $m' \neq m$, and where $\mathbb{1}()$ is the indicator function and $k(r)$ denotes the relative ordering of the route. In words, I utilize the mileage along competing modes carrying freight from i to j of the same relative ordering as r – so, the shortest path on Rail from e to x is instrumented by the shortest path on Road, Water, Air, and Multi-Modal making the same trip; the second-shortest path on Rail is instrumented by the second-shortest path on Road, Water, Air, and Multi-Modal; while the shortest path on Water is instrumented by the shortest path on Road, Rail, Air, and Multi-modal.[23]

For my IV estimates to be valid, these instruments must satisfy the exclusion restriction and relevance assumptions. Regarding the former: the exclusion restriction is satisfied so long as routed miles along alternate modes are independent of the error term for a given trade route. This strict exogeneity holds by assumption in my theoretical context, and is valid in reality if trade flows today are independent of the trade flows that influenced the planning and development of

[22]Note that a well-known result in the trade literature is that varying the trade elasticity will scale the estimated parameters accordingly (Anderson and van Wincoop, 2003).

[23]It should be noted that I have four instruments for every mode, leading to an over-identified model. I include all available instruments in my empirical specification as I am agnostic about which mode is likely to be most relevant for the pricing of another mode.

the disparate modal networks decades (or in some cases, centuries) ago. Moreover, many modal networks were built for reasons distinct from the efficient movement of goods: transportation of people, as well as national defense resources, are a major motivation underlying many transit networks, particularly the interstate highway system. Hence, my independence assumption seems reasonable. The relevance assumption holds so long as freight firms compete with one another. This assumption holds so long as distinct modes compete, which seems reasonable a priori. However, absent any price data from the transport industry, it is difficult to verify. The relevance of my proposed instruments may therefore be suspect – however, I may empirically evaluate the relevance of the instrument via the first-stage F-statistic (detailed in Panel D of Table 2). Reassuringly, all estimated F-statistics suggest that my instruments are relevant.

3.2.2 Estimation

Given route-level data on domestic trade as well as knowledge of the profit input δ_{irj}, it would be straight-forward to estimate my model parameters via 2SLS. However, this cannot be achieved with publicly-available data. As made evident by Equation (16) as well as Appendix A.3, δ_{irj} is a nonlinear function of φ and ρ, which I aim to estimate in the same equation. Second, and much more problematic, I do not observe route-level trade shares; rather, my data only reports aggregate trade flows for an origin, destination, port of entry/exit (where relevant), and mode. I thus rely on a simulated method of moments strategy to overcome these data shortcomings.

As a first step, I populate the set $\mathcal{R}_{ij}^{m}$ by mapping successively longer paths between e and x along each mode.[24] I also remove the first and last 25 miles from each journey, as these are the most-expensive parts of the trip and should be captured by the α_{iem} and γ_{jxm} parameters.[25] Theoretically, I could repeat this process almost infinitely; practically, I consider only the 5 shortest paths for each mode. I have two main reasons for stopping at this threshold: i) routes outside of this set seem sufficiently circuitous to be cost-prohibitive, and therefore add little information to my analysis, and ii) computation time grows with each additional route, becoming excessive after 5.[26] This mapping process generates a set of 5 distinct routes – as well as the corresponding mileage

[24] I require these paths to be substantially different, so that the changes to the distance are meaningful. Table 1 demonstrates the average evolution of mileage across these paths.

[25] This truncation, as well as the fact that I cannot route within a region, means that I hold $\kappa_r = 0$ for all routes that start and end in the same region.

[26] Estimated using a custom-built desktop, whose computing power (RAM, threads, and graphics card) exceeds

– for domestic origin, destination, and mode combination.

Table 1 describes the evolution of average mileage across these subsequently longer routes for each mode. To construct this table, I scale the distance of each routing by the shortest distance for every domestic origin, destination, and mode pairing. I then average these scaled distances across origin-destination combinations to find the average, scaled distance for every routing. Columns (1) and (2) demonstrate that, on average, the second-shortest routing via Road is approximately 4.4% longer than the shortest path on the Road; the distance of the third-shortest path is 7.4% long than the shortest path, and 2.8% larger than the second-shortest path; and so on. Notably, for Roads, Rails, and Multi-modal, the marginal increase in mileage along each route generally shrinks as the number of routes increases; I attribute this trend to the relative density of these networks: as routes become more and more circuitous, the marginal increase in mileage becomes less and less. Likewise, the marginal increase in Air distance with each route is consistently small and stable; this trend owes to the fact that I allow a straight-line path between all freight Airports, making the Air network the most dense by far. Finally, the increases in routed distance over Water are somewhat erratic, reflecting the relative sparsity of the Water network. The last key takeaway from this table is the sizeable difference between the shortest path and the fifth-shortest path: for Roads and for Multi-modal, the fifth-shortest path is approximately 11-12% longer than the shortest path on average; this number hovers around 15% for the Rail and Water networks. For Air, the fifth-shortest path is only 1.3% greater than the shortest trip on average, again reflecting the unique ability of Air to travel in a straight line, anywhere.

The routed mileage identifies the per-mile cost parameters, β_m. The question now becomes how to estimate route-level costs (and markups), from aggregated data. Towards this end, I employ an Expectation-Maximization (EM) algorithm. The algorithm proceeds accordingly:

1. Take an initial guess at the model parameters, $\lambda_i^{(0)}, \beta_m^{(0)}, \zeta^{(0)}, \rho^{(0)}$, and $v^{(0)}$. Keep in mind that I have already fixed the parameters φ and θ.

2. (Expectation Step) Simulate $\mathbb{E}\left[\pi_{r|iexmj} \middle| \lambda_i^{(0)}, \beta_m^{(0)}, \zeta^{(0)}, \rho^{(0)}, v^{(0)}, \varphi, \theta\right]$ according to Equation (4) under the assumption that $u_{irj} \sim N(0, v^{(0)})$. Denote this conditional probability $\pi^{(0)}_{r|iexmj}$.

most commercially-available desktops and is on-par with low- to mid-range cloud solutions; exact specifications are available upon request. In Section 4.2, I demonstrate the sensitivity of my estimates to different numbers of potential routings – in general, my estimates remain stable.

3. Calculate $\pi^{(0)}_{ir|j} = \pi^{(0)}_{r|iexmj}\pi_{m|iexj}\pi_{ex|ij}\pi_{i|j}$. Note that only the first of these conditional probabilities is simulated; the rest are observed in the data.

4. Calculate the markup term $\delta^{(0)}_{irj}$ as described by Equation (16) and Appendix A.3.

5. Estimate the first stage via OLS:

$$\delta^{(0)}_{irj} = \boldsymbol{Z}_{irj}\zeta_m + Miles_r\phi_m + \vartheta_{iem} + \vartheta_{jxm} + \varepsilon_{irj}$$

where ϑ_{iem} and ϑ_{jxm} are fixed effects. Denote the fitted values from this regression $\hat{\delta}^{(0)}_{irj}$.

6. Define $y^{(0)}_{irj} = \ln \pi^{(0)}_{ir|j} - \left(1 - \varphi\rho^{(0)}\zeta^{(0)}\right) \ln \pi^{(0)}_{r|iexmj}$.

7. (Maximization Step) Estimate the second stage via OLS:

$$\begin{aligned} y^{(0)}_{irj} =& - \theta Miles_r\beta_m - \lambda_i\hat{\delta}^{(0)}_{irj} + (1-\rho\zeta)\ln \pi_{m|iexmj} + (1-\zeta)\ln \pi_{ex|ij} - \theta u_{irj} + \\ & \underbrace{\text{Origin} \times \text{Port} \times \text{Mode FEs}}_{-\theta(\nu_{xj}+\gamma_{xjm})-Q_j} + \underbrace{\text{Destination} \times \text{Port} \times \text{Mode FEs}}_{-\ln A_i - \theta(\ln c_i + \eta_{ie} + \alpha_{iem})} \end{aligned} \tag{21}$$

Denote the new set of parameter estimates $\lambda^{(1)}_i, \beta^{(1)}_m, \zeta^{(1)}, \rho^{(1)},$ and $v^{(1)}$.[27]

8. Return to step 2, and repeat the process with the new parameter estimates. Continue until the estimates converge.

Under standard assumptions required for a linear regression model, the parameter estimates of the EM process converge to the (local) maximum-likelihood estimates. See Appendix A.4 for further detail.

4 Descriptive Results

The end result of the EM process just laid out in Section 3.2 are estimated model fundamentals. Namely, these encompass per-mile iceberg cost coefficients β_m, the correlation parameters ρ and ζ, as well as the conduct parameters λ_i, which in turn inform route-level markups μ_{irj}. Table 2 summarizes the results of my main specification; additional robustness checks and model validation exercises are discussed in Section 4.2.

[27] It should be noted that, while I do not immediately estimate $v^{(1)}$ via OLS, I can estimate $v^{(1)} = \text{Var}(\hat{u}_{irj})$.

4.1 Primary Specification

I begin discussion of my reduced-form estimates with Panel A of Table 2. This panel reports the parameter estimates that inform geographic costs in my trade model. I breakout per-mile costs by mode of transport, or – in the case of multi-modal – the cost of a single transshipment. As these coefficients inform iceberg trade costs, interpretation is somewhat involved.[28] However, their magnitude and relative ordering align with similar estimates in the literature (Allen and Arkolakis, 2014; Donaldson, 2018; Allen and Arkolakis, 2022). Importantly, Road is estimated as the most expensive mode per-mile, followed by Rail and Water, which have similar magnitudes. Air is estimated as cheapest mode per-mile, which may prove initially surprising, given Air's high fuel intensity and maintenance requirements. However, bear in mind that these estimated coefficients exclude first- and final-mile costs, as well as any costs accrued at a destination or origins.[29] The final reported coefficient in Panel A is rarely estimated: the cost of transshipment across modes. I estimate that the cost of a single transshipment accounts for approximately 1.1% of the final value of a good. Note that this effect compounds across multiple exchanges.

I now shift focus to the estimated correlation coefficients listed in Panel B of Table 2. Note that these are inverse correlation coefficients: higher values of ρ, ζ imply lower correlations among composite efficiencies among modes, origins. In particular, the real correlation may be calculated as 1 minus the square of the estimated coefficient (Heiss, 2002). My estimate of the modal correlation parameter ρ thus implies that, on average, transportation efficiencies – and hence, prices – across domestic modes enjoy a correlation of $1 - 0.596^2 = 0.645$; production efficiencies across domestic locations report a correlation of $1 - 0.379^2 = 0.856$. Both of these estimates are statistically distinct from 1, which implies zero correlation. These substantial correlation estimates underscore the importance of allowing correlation in my model. In Section 4.2, I explore the consequences of imposing zero correlation, as is common the conventional EK setting.

Panel C summarizes the distribution of estimated non-competitive markups $\hat{\mu}_{irj} = \hat{\lambda}_i \delta_{irj}$. I focus on the distribution of markups – as opposed to the distribution of conduct parameters λ_i – due to ease of interpretability.[30] The table demonstrates that, across all domestic trade routes,

[28]Interpreting the first reported coefficient: 1,000 additional miles on road increases iceberg trade costs, τ_{irj}, by $\exp(0.197) \approx 1.218$, or about 21.8%. Interpretation of the remaining coefficients is analogous.

[29]These route-invariant costs are estimated via fixed-effects in my regression specification. See Section 3 for detail.

[30]While the extreme cases of $\lambda_i = 1$ and $\lambda_i = 0$ are easily interpretable as perfect monopoly and perfect competition,

markups average 3% above marginal cost. The distribution is slightly skew right; the highest estimated markup is approximately 13% above marginal cost.[31] These estimated markups are fairly low, suggesting that, in general, freight markets are highly competitive. However, due to the necessary centrality of the transport industry, even a small exercise of market power could have a profound influence on national welfare.[32]

4.2 Model Validation & Robustness Checks

In Table 3, I evaluate my reduced-form estimates relative to those implied by simpler modelling approaches. Column (1) of this table re-states my reduced-form estimates under my main specification: estimating freight market power, allowing correlation of prices across both modes and geographies, and encompassing 5 distinct routings per-mode along every domestic origin-destination pairing.[33] Column (2) re-estimates my geographic cost and correlation coefficients, imposing perfectly-competitive freight markets (i.e. $\lambda_i = 0 \, \forall i$). Comparison of these first two columns reveals that assuming perfect freight competition (as is common in the trade literature) results in very minor changes to estimated fundamentals. This stability is unsurprising in light of my finding that freight markets are, generally, near-perfectly competitive. Column (3) further restricts the model by assuming zero correlation of prices across markets; these estimates correspond to the parameters implied by a conventional EK framework. My cost estimates remain fairly stable: estimated Road and Air coefficients increase slightly, while the coefficients on Rail and Water drop. The largest change comes from Multi-Modal transshipment, whose coefficient drops by nearly half. Finally, in Column (4), I re-run the model using only one routing per-mode; substitution across parallel trade routes in this model is impossible.[34] This change causes the estimated coefficient on Road to again increase, while the coefficients on Rail and Water decreases slightly. Overall, this exercise reveals that simpler models induce very limited bias in estimated fundamental parameters.[35]

---

intermediate values are more obscure. I thus focus discussion of my reduced-form results on the distribution of markups, as opposed to the conduct parameters themselves. A histogram of the estimated λ_i parameters is presented in Figure 5.

[31] I will analyze the geographic distribution of these markups – or rather, their consequences – in Section 5.

[32] Indeed, I demonstrate in Section 5.1 that even these moderate markups produce aggregate welfare losses equal to approximately 5% of total perfectly-competitive welfare.

[33] The numbers reported in this column are identical to those reported in Table 2.

[34] This highly constrained model most closely resembles the routing logic utilized in Allen and Arkolakis (2014).

[35] I should emphasize that, while my complex modelling framework does not lead to drastic differences in estimated fundamentals, the welfare consequences of freight market power are still profound (see Section 5).

I also run the model through a battery of robustness checks, summarized in Tables 4 - 6. Fist, I re-estimate my primary specification allowing for non-zero correlation of prices across parallel trade routes, which I achieve by setting the route-correlation parameter φ to values less than 1. Results of this exercise are displayed in Table 4. Column (1) re-states my main results; each subsequent column displays analogous estimates assuming an intermediate value of φ. Across these alternate specifications, there is no discernable difference in my estimated coefficients. This stability stems from the fact that my parameters are identified by observed variation in trade shares at the origin-destination level, which is independent of the simulated route-level variation (see Section 3.2 for detail). Hence, Table 4 confirms that my model is insensitive to the values of the route-level correlation parameter φ.

In my next battery of robustness checks, I re-estimate my primary specification with amended or different data; results of these checks are reported in Table 5. As before, Column (1) re-states my main results. Column (2) restricts the set of commodities, from which I generate aggregate trade flows, to those that are easily transported by any mode.[36] Results of this exercise reveal generally similar parameter estimates; notably, the distribution of markups is slightly higher than those reported in baseline. These results suggest that I may under-estimate market power in my primary specification. Next, in Column (3) I re-estimate the model excluding Air and Water, which account for a small share of domestic trade and do not enjoy wide geographic representation. Reassuringly, my estimates are largely the same. The distribution of markups is again slightly larger relative to my main specification; however, as I elaborate in Section 5.1, Air and Water report relatively little exercise of market power, and I expect the distribution of markups to shift right when removing two relatively competitive modes. In Column (4), I re-estimate the model using only urban areas. Reported cost coefficients are slightly smaller, but of a similar magnitude, while the distribution of markups compresses slightly. Again, this result accords with expectation, based on the geographic distribution of markups (see Section 5.1 for detail). Finally, in Columns (5) and (6) I address external validity concerns by re-estimating the data in 2012 and in 2022.[37]

[36] I achieve this by restricting my trade data to the subset of commodities that report at least \$10 million in total volume within the U.S. across all modes in 2017. These commodities include: cereal grains and seed, agricultural products (excluding grains, feed, and forage), alcoholic beverages, coal, crude petroleum, other coal and petroleum products, raw chemicals, other chemical products and preparations, plastics and rubber, base metals in a semi-finished form or basic shape, machinery, motorized and other vehicles, and transport equipment not elsewhere defined.

[37] I elect these years to create an even interval about my estimation year, 2017. Moreover, these are years, in which the Commodity Flow Survey (CFS) is conducted. The FAF data I use is based on the CFS, and is most accurate in

Reassuringly, the estimated cost and correlation parameters are similar. Notably, markups are declining over time– there is a steady, approximately linear downward trend in the distribution of markups as I move from 2012 to 2022. This decline is striking, but further investigation is beyond the scope of this paper. Overall, these results demonstrate the stability of my estimates to different data-cleaning decisions.

Finally, I analyse the impact of different starting values in my EM algorithm; results of this exercise are reported in Table 6. In my main specification, I elect $\beta_m = \mathbf{1}$ as my starting value, which implies that each mode has the same cost per-mile at the outset of the EM process; results stemming from these starting values are reported in Column (1). In the subsequent Columns, I change my starting values to equal 1 for a particular mode, and zero for the remaining; hence, in these alternate specifications, one mode is particularly costly to traverse, while the others are free (outside of the first- and final-mile). Reassuringly, the model converges to the exact same parameter estimates.

5 Counterfactual Simulations

Having estimated the fundamental parameters, I calibrate the model to domestic trade flows in 2017; I then run three counterfactual simulations, designed to highlight the role of freight transport in impacting domestic trade. The first such exercise investigates which regions and modes are most or least impacted by imperfect competition in the market for freight transport. Explicitly, I compare calibrated welfare under a counterfactual, perfectly-competitive pricing regime to calibrated welfare under the status-quo; this juxtaposition reveals where and along what modes current non-competitive losses are concentrated. My second counterfactual evaluates the welfare consequences of removing an entire mode of transport from the domestic transit network: specifically, Rail and (consequently) multi-modal.[38] This simulation is inspired by the proposed Rail strike of 2022, which was narrowly avoided at the eleventh hour by federal intervention, and to some controversy (Shepardson and Baertlein, November 29, 2022). My model sheds light on the potential impacts of

survey years.

[38]Why multi-modal as well? The lion's share of domestic multi-modal movement's encompass the Rail network (AAR, October, 2023); hence, removing Rail also vastly limits the scope for multi-modal travel. For the purpose of my counterfactual simulation, I make the simplifying assumption that, if Rail shuts down, multi-modal movements disappear as well.

such a strike, how trade would have moved to competing modes, and how freight market power may amplify or attenuate downstream welfare losses to consumers. My final counterfactual simulates an international trade shock: specifically, a sharp increase in trade costs at international gates, meant to emulate supply chain bottlenecks sparked by the Covid-19 pandemic. This counterfactual reveals how strategic pricing across a geographically dispersed, multi-modal network exacerbates or mitigates an exogenous, geographically concentrated trade shock. Together, these three scenarios highlight how the transport sector influences domestic trade outcomes.

It is important to emphasize that my model is agnostic about the source of market power: I cannot differentiate between non-competitive pricing stemming from collusion or from barriers to entry. However, with the notable exception of Rail, freight markets are typically saturated with multiple service providers,[39] while fixed-costs in the industry are exceptionally high.[40] It thus appears likely a-priori that barriers to entry, coupled with inelastic demand for freight transport along specific modes, are the culprit for exercise of freight market-power; accordingly, I direct my policy discussion towards lowering barriers to entry, as opposed to anti-trust enforcement.

I now detail the results of my counterfactual simulations.

5.1 Simulating a Perfectly Competitive Freight Market

For this counterfactual exercise, I set (log) markups to zero across all domestic trade routes; I then utilize Equations (4) - (17) to simulate trade flows and prices with this perfectly-competitive freight market. Comparing calibrated welfare under the status-quo and counterfactual pricing regimes reveals current losses due to non-competitive freight pricing. I further decompose these aggregate effects by mode and geography in Table 7 and Figure 6, respectively. The former exploits my logit structure to simulate substitution patterns across competing modes; the latter reveals

[39] In 2021, approximately 600 thousand for-hire interstate trucking companies operated in the United States (FMCSA, 2021); approximately 543 Air cargo carriers (**Airfreight**), 37 U.S.-flagged maritime carriers (MARAD, September, 2022), and just 7 Class-1 freight Railroads operated in the U.S. concurrently.

[40] Trucking has the lowest entry cost: a new Semi truck runs between \$100 and \$200 thousand, with an additional \$50 thousand for a new trailer on average; this estimate excludes (substantial) insurance and licensing costs, as well as the cost of any sophisticated trailer – e.g., refrigeration, logging, HAZMAT, etc.. (Harris, 2021). The remaining modes enjoy much higher fixed costs: a new diesel locomotive costs between \$1.5 to \$2 million on average (World Wide Rails, 2023), while a new Rail car runs between \$100 and \$150 thousand (Blaze, 2019); a new cargo plane averages between \$49 and \$77 million, but can reach over \$400 million, depending on capacity and technology (Walsh, 2023); and a newly-constructed container ship costs between \$14 and \$300 million, again depending on capacity and technology (UNCTAD, 2010). In addition, while roads, Waterways, and Airways are (generally) free-to-use, Rail companies must create and maintain their own infrastructure.

what locations enjoy the greatest/least concentrations of freight market power, and thus stand the most/least to gain from an erosion of that pricing power. I now discuss these results in detail.

I find that, as a result of removing freight market power, real welfare in the U.S. increases by approximately 5.04%; Panel A of Table 7 details total welfare changes accruing to each mode and overall. Column (1) of the table lists the baseline share of real trade volumes; Column (4) details the percentage change in this trade share brought about by the imposition of perfect freight competition. Columns (2) and (3) decompose this aggregate change into a modal substitution and income effect. The former effect is calculated by simulating the change in trade volumes, holding constant income in each region;[41] the latter effect is calculated by simulating the change in aggregate trade volumes to each region, holding constant the distribution of trade across modes.[42] The sum of these effects equals the total change in trade volume.

Analysis of Panel A, Column (2) reveals that, holding real incomes in each region constant, elimination of freight markups shifts trade away from multi-modal transport and onto the road. This result may be initially surprising, as it suggests road transport currently enjoys the *highest* markups, despite having by-far the highest market concentration and barriers to entry. One potential explanation for this perhaps counter-intuitive outcome is the unique ability of truck transport to complete the first- and final-mile. This supreme routing flexibility limits the scope for cross-modal competition: no alternative save multi-modal can offer comparable services. Indeed, Column (2) displays that the vast majority of this new trucking volume comes from multi-modal; very little crosses over from Rail, Water, or Air.[43] In summary, Column (2) displays the imperfect substitution of trade across modes: price decreases that mainly affect the Road draw trade primarily from multi-modal.

I now shift attention to Panel A, Columns (3) and (4). The former isolates the income effect: the change in welfare brought about by making certain regions wealthier via a reduction in trade costs. This column aligns with expectation: all modes see an increase in real trade volumes, and the relative distribution of these increases roughly corresponds to the baseline trade allocation

[41] If markups were uniformly distributed across modes, then the substitution effect would simply equal the change trade costs multiplied by the baseline trade share of each mode.

[42] If markups were uniformly distributed across geography, then the income effect would simply equal the change in income multiplied by the baseline trade share.

[43] A more complete evaluation of this "final-mile hypothesis" requires more refined geographic data, which is not publicly-available at a national scale.

displayed in Column (1).[44] Finally, Column (4) presents the aggregate change in welfare accruing to each mode. Notably, the income effect dominates for all modes except multi-modal. Hence, the lion's share of welfare improvements accruing to each mode stems from these income gains following the reduction in trade costs.

Panels B and C of Table 7 repeat much the same exercise, focusing instead on imports and exports; the outcomes reported in these panels largely mirror the effects on domestic welfare. Focusing first on Column (2), roadways see the largest relative benefit holding income in each region constant. Among imports, these newfound roadway flows again draw primarily from multi-modal; among exports, this substitution effect pulls primarily from Rail. In general, Rail, Air, and Water see much higher substitution effects among these international movements than among domestic trade, a trend likely attributable to their higher baseline trade volumes (see Column (1)). Column (3) demonstrates that, for both imports and exports, trucking is again the largest benefactor from the income effect; however, among exports, all modes report substantial increase in volume due to increasing incomes. Finally, in Column (4), I report aggregate changes: roadways again see the largest change in total imports and exports, driven largely the income effect reported in Column (3). Total imports increase by 7.52%, while exports climb by nearly 20%. This outsized effect on exports is striking, and suggests that managing freight market power is a powerful policy tool to influence the international trade balance.

Finally, in Figure 6, I breakout the geographic incidence of removing non-competitive freight rates. Panel A reports real welfare impacts: generally, the largest welfare gains accrue to rural areas in the Southeast and Mountain West, as well as to small urban areas in the Midwest and along the Gulf of Mexico. Strikingly, major international gates, including Los Angeles/Long Beach, CA, San Francisco/Oakland, CA, New York, NY, Chicago, IL, Baltimore, MD, Norfolk, VA, Atlanta, GA, and Ft. Lauderdale, FL, are among the *least* affected areas. While this result aligns with expectation – high trade volumes should attract more shippers, thereby reducing market power – it suggests that the aggregate impact of removing freight market power will be muted relative to the local impact in rural regions. Indeed, the Panel A of Figure 6 reveals that local impacts range from 2 to 11% of baseline welfare.

[44] Note that this relative distribution does not perfectly correspond to Column (1) due to geographic heterogeneity: as just discussed, the largest gains from this counterfactual are concentrated on the roadway; areas that trade more heavily in the road will see the largest income gains. Figure 6 provides a more detailed geographic analysis.

Panels B and C of Figure 6 report the geographic incidence among imports and exports, respectively. I begin with imports, where the largest gains accrue mostly to the coasts, as well as the northern Midwest and Great Plains. Regions up the middle of the country – including Lake Charles, LA, Austin, TX, Laredo, TX, East St. Louis, IL, Rural Georgia, Rural Indiana, most regions in Appalachia, and all of Michigan – experience modest increases or even *declines* in imports. Notably, these areas also saw substantial welfare gains due to elimination of markups; this pattern suggests that, thanks to the alleviation of freight market power, these regions import relatively less from abroad as domestic goods become cheaper. Switching focus to Panel C, increases in real exports accrue primarily to the regions bordering the Gulf of Mexico, as well as rural areas up the east coast and in the Northwest. There is not a discernable geographic relationship between welfare changes and export increases.

In summary, my first simulation explores the welfare consequences of removing freight market power. Most of these welfare gains accrue to the Road, and shift domestic trade away primarily from multi-modal, and to a lesser extent, Rail. This substitution pattern suggests that market power is currently concentrated along the roadway–an initially counter-intuitive finding that is somewhat supported by anecdotal evidence of a pervasive shortage of qualified truck drivers. Effects on international trade are similar, but larger. Most strikingly, real U.S. exports increase by nearly 20%, relative to only a 7.5% increase in real imports. Geographically, gains from the simulation are concentrated in rural and small, urban areas throughout the Southeast, Midwest, and Gulf states; major international gates see the smallest welfare impacts. Notably, areas that see large welfare gains report the smallest gains or even modest declines in imports as a result of removing freight market power. This latter trend is likely the result of domestic freight becoming cheaper, leading to a lower reliance on imports.

5.2 Simulating a Rail Closure

The counterfactual detailed in this Section is primarily policy-focused. Explicitly, I simulate the effect of removing Rail and multi-modal transport from the domestic transit network – this simulation emulates the narrowly-avoided Rail strike of 2022. To highlight the role of freight-market concentration in determining domestic welfare, I simulate the shock twice: once under status-quo freight pricing and once imposing perfect competition. Comparison of calibrated welfare under these two scenarios reveals how the downstream impacts of this exogenous shock are amplified/attenuated by market power in the freight industry.

Table 8 reports the impacts to domestic welfare as a result of removing Rail and multi-modal. Panel A reports simulated welfare changes under status-quo freight pricing; Panel B reports the same under perfect freight market competition. Comparison of these results reveals a key conclusion: removal of the Railway (and consequently, multi-modal) leads to substantial welfare losses, regardless of freight market competition; however, losses are *worse* thanks to imperfectly-competitive freight pricing. Real welfare declines by 4.31% under the status-quo, and only 2.86% under perfect competition. The reason for this substantial amplification stems from the exercise of market power: as displayed by Column (2), the vast majority of freight volumes shift onto the roadway, where freight market power is concentrated. Freight-service providers may then take advantage of the drop in cross-modal competition and raise prices. Column (3) reports welfare losses stemming from the effective change in income as a result of removing the Railway. Note again that this income effect is concentrated in trucking. However, Panel (B) demonstrates that, absent an endogenous pricing response on behalf of freight carriers, welfare losses stemming from this income drop amount to only 2.77%. Hence, of the 4.23% drop reported in Panel A, approximately 1.46 percentage points – nearly 35% of the total change – is attributable to price increases stemming from a decline in modal competition. This income effect accounts for the majority of the aggregate changes reported in Column (4).

Tables 9 and 10 repeat this exercise for real imports and exports. These results are very similar to the patterns exhibited for domestic welfare: losses are substantial, and made worse by the exercise of freight market power. Specifically, real imports decline by about 6.03% under perfect competition, and by 8% under status-quo pricing; real exports decline by approximately 1.23% under perfect freight market competition, and by 4.17% under the status-quo. Note that the change in real export volumes is much larger than the change in real imports, a result that is emblematic of the outsized effect freight market power exerts over exports, as previously discussed. The table demonstrates that the vast majority of international freight shifts to the road for the domestic portion of its journey; however, in contrast to domestic welfare, a substantial share shifts to the Air, especially among imports. Given the centrality of Air and Water transport to international transactions, this result is not surprising.

Finally, Figure 7 breaks out these real welfare changes by geography; Panel A reports results under status-quo freight pricing, and Panel B reports under perfect competition. Overall, losses are concentrated in the Southeast, northern Midwest, and Great Plains. Urban areas along the coasts also report moderate losses. Notably, a handful of areas report welfare *gains* – these areas are served more readily by the road, Air, and Water networks, and thus see an increase in total trade volumes as trade shifts off the Rail network. Comparison of Panels A and B reveals that exercise of freight market power exacerbates welfare losses among rural and small urban areas throughout the U.S.. This result is not surprising, as these areas report the highest concentrations of market power ex-ante. Geographic trends for imports and exports are similar, and reported in Figures 8 and 9.[45]

In summary, shutting down the Rail system would have a substantial, negative impact on domestic welfare; these losses are made worse by the exercise of market power along competing modes. I report similar impacts for imports and exports; most strikingly, freight market power has an outsized impact on exports– real export losses nearly quadruple under status-quo freight pricing relative to perfect freight market competition. Turning to the geographic incidence, losses are concentrated in the Southeast, northern Midwest, and Great Plains. Some areas, which are more heavily served by the road, see *gains*. Exercise of freight market power exacerbates losses to rural and small urban areas throughout the country. Speaking broadly, these results emphasize that, should any single mode shut down, we should expect a meaningful pricing response along competing modes, most especially trucking; this pricing response will most-impact rural and small urban areas.

5.3 Simulating an International Trade Shock

My final counterfactual explores how freight transport influences a geographically-concentrated international trade shock. Specifically, I simulate the welfare consequences of a 36% increase in trade costs along international routes, meant to emulate price increases witnessed in the transport sector during the Covid-19 pandemic.[46] This shock is geographically concentrated in the sense that

[45] These results make clear that inland areas in the Great Plains are by-far the worst impacted among international trade flows.

[46] The magnitude of this shock comes from recent analysis by the St. Louis Federal Reserve (Leibovici, May 11, 2022).

only international gates feel the immediate impact; inland locations are not directly affected, but as I will demonstrate, suffer substantial losses due to equilibrium effects. I again simulate this shock twice: once under status-quo freight pricing, and once under perfect freight market competition. Contrasting the changes in simulated welfare under these two pricing scenarios reveals the role of strategic freight pricing in determining domestic trade outcomes. More broadly, this counterfactual reveals how the freight industry influences the domestic welfare consequences of an international trade shock.

Table 11 reports real welfare changes from this simulation. Again, Panel A reports welfare impacts when the freight industry is allowed to exercise market power; Panel B reports welfare impacts under perfect freight market competition. In this case, we see that market power *attenuates* welfare changes, but very, very slightly. This result accords with expectation: market power implies imperfect pass-through of exogenous cost shocks; however, as demonstrated previously, freight market power is concentrated away from international gates. Hence, non-competitive freight pricing can have only a limited impact on international trade shocks. As displayed by Column (2), modal substitution effects are limited in magnitude, and concentrated along road and multi-modal, while Air sees the largest (though still small) gains. Switching focus to Column (3), income effects are all negative and concentrated along the roadway. In total, domestic welfare declines by approximately 7.11% under status-quo freight pricing, and 7.22% under perfect competition.

Unsurprisingly, this shock has a much larger impact on international trade. Tables 12 and 13 report simulated changes in real imports, exports under the status-quo and perfect freight market competition. As reported in Panel B of each table, imports decline by approximately 17.97%, while exports decline by about 22.59% in the absence of freight market power. Attenuation rates due to non-competitive pricing are again minuscule, or even slightly negative: as reported in Panel A of each table, under the status-quo, imports decline by 17.84%, while exports decline by 22.67%. Focusing on Column (2), we see very little substitution across modes, implying a near-uniform shock across modes; this outcome is not surprising, as it is rare for an international gate to specialize in any, one mode. Column (3) reports substantial income effects; in-line with the uniform effect across modes, the relative distribution of these effects roughly corresponds to each mode's baseline trade share: Road sees the biggest declines, followed by multi-modal, Air, Rail, and finally, Water.

The geographic distribution of the welfare impacts stemming from this shock are presented in

Figure 10. Panel A reports real welfare changes under the status-quo, while Panel B reports welfare changes under perfect competition. Unsurprisingly, the two panels are nearly identical: rural Maine, Georgia, and Alabama, as well as East St. Louis, IL and Omaha, NE, report the only observable attenuation of welfare losses. Moreover, the figure reveals that Washington, Wyoming, Wisconsin, and West Virginia are the worst-impacted states in either pricing scenario: welfare losses in each of these states amount to at least 15%, and can climb as high as 20%. This geographic concentration is telling, as – with the exception of Washington – none of these states contain major international gates. Indeed, most international gates report small to moderate impacts. This finding suggests that rural, inland areas are worst impacted by an international trade shock. One explanation for this result is that major gates lie centrally in the domestic transit network, and are thus able to easily pick among cheaper suppliers/customers in response to more expensive foreign imports/exports. However, these increased costs are detrimental to rural areas along sparse portions of the network, who cannot easily adjust to consumption. Figures 11 and 12 breakout these effects for imports and exports and reveal largely the same pattern – the one, notable exception being for exports, where Florida reports the largest declines.

In summary, a sharp rise in international transport prices, similar to what the U.S. witnessed during the Covid-19 pandemic, creates substantial domestic welfare losses. These losses equal approximately 7% of baseline welfare, and are even more pronounced among imports and exports. Due to the concentration of freight market power away from international gates, this shock reports near-perfect pass through. Moreover, losses are generally concentrated within inland areas along relatively sparse portions of the transit network; those locations along relatively dense portions of network more easily adjust production or consumption in response to the shock.

6 Conclusion

This paper quantifies the impact of the transport sector in setting domestic trade outcomes. I modify a standard Ricardian trade model to account for route-choice, mode-choice, port-choice, correlations across these disparate nests, and finally, imperfect competition for freight services in geographically distinct markets. The result of these extensions is a highly plausible model of the transportation industry, as well as its effect on domestic trade. The trade costs implied by my model reflect imperfect competition for freight services in distinct locations; imperfect substitution across origins, modes, and even parallel trade routes; and finally, available infrastructure. This framework allows me to evaluate the domestic welfare and trade incidence of a number of transportation shocks: i) removal of freight market power; ii) the complete removal of Rail and (consequently) multi-modal transport; and iii) a sharp spike in international freight costs, similar to those witnessed during the Covid-19 pandemic. The model thus offers a comprehensive view of how freight markets influence domestic trade outcomes.

I estimate fundamentals and calibrate the model to domestic trade flows in 2017. The estimation yields the following insights: i) an estimate of per-mile trade costs along each mode; ii) the degree of correlation in prices across modes, origins; and iii) location-specific conduct parameters, which, in-turn, influence route-level markups throughout the domestic U.S.. The latter calibration yields 5 key insights: i) current losses due to non-competitive pricing are substantial, amounting to roughly 5% of baseline welfare; ii) the exercise of freight market power has an outsized impact on exports– elimination of freight market power could increase national exports by nearly 20%; iii) these impacts are concentrated in rural areas throughout the Southeast and Mountain West, as well as small urban areas in the Midwest and Gulf states; iv) exercise of freight market power exacerbates the impact of mode-specific shocks; and v) non-competitive freight pricing does does little to attenuate international trade shocks, due to the concentration of freight market power away form international gates.

These results yield a number of insights relevant to policy-markers. First, losses due to non-competitive freight pricing are substantial, and perhaps surprisingly, concentrated in the trucking industry. Anecdotal evidence suggests that trucking in particular suffers from a shortage of qualified drivers, granting market power to existing operators. With the exception of multi-modal, competition from other modes is insufficient to solve this problem, as trucking offers the unique ability to fulfill the first- and final-mile. Hence, addressing the driver shortage is an obvious and immediate policy intervention, which should counter the majority of welfare losses due to non-competitive freight rates; further expansion of multi-modal terminals will also deflate freight market power. Additionally, given the concentration of freight market power in rural and remote urban areas, such an intervention should help reduce economic inequality across the U.S.. Hence, action to reduce barriers to entry in trucking should achieve large impacts on domestic welfare.

Any policy intervention designed to deflate freight market power has important ramifications for the national trade deficit, which is of key interest to many politicians and policy makers. Because non-competitive freight pricing exerts an outsized influence on exports, it offers a heretofore unexplored lever to manage international trade flows. Importantly, it is a purely domestic policy, and does not require multilateral negotiation and enforcement.

Finally, the results bulwark arguments in favor of aggressive policy interventions in response to transport shocks. Explicitly, the closure of the Railway would lead to sizeable welfare losses, made worse by the exercise of market power from competing modes; an international trade shock is born almost exclusively by consumers – freight market power is concentrated away from international gates, and I estimate near-perfect pass through. Hence, heavy-handed policy intervention is necessary to avoid the worst welfare impacts of transport shocks. It should also be noted that, due to its centrality, transportation carries consequences far beyond trade flows. To name a few, the composition of freight carries environmental, health and safety, and equity concerns which are beyond the scope of this paper. A promising avenue for future research may apply the methods developed in this paper to these outcomes.

References

AAR. (October, 2023). *AAR Intermodal Fact Sheet*. https://www.aar.org/issue/freight-rail-intermodal/

Allen, T., & Arkolakis, C. (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics*, *129*(3), 1085–1139.

Allen, T., & Arkolakis, C. (2022). The welfare effects of transportation infrastructure improvements. *Review of Economic Studies*, *89*(6), 2911–2957.

Anderson, J. E., & van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, *93*(1), 170–192. https://doi.org/10.1257/000282803321455214

Andri et mult. al., S. (2022). *DescTools: Tools for descriptive statistics* [R package version 0.99.47]. https://cran.r-project.org/package=DescTools

ATA. (October 25, 2021). *Driver Shortage Update 2021* (tech. rep.). American Trucking Associations, Inc. https://www.trucking.org/sites/default/files/2021-10/ATA%20Driver%20Shortage%20Report%202021%20Executive%20Summary.FINAL_.pdf

Bache, S. M., & Wickham, H. (2022). *Magrittr: A forward-pipe operator for r* [R package version 2.0.3]. https://CRAN.R-project.org/package=magrittr

Bergé, L. (2018). Efficient estimation of maximum likelihood models with multiple fixed-effects: The R package FENmlm. *CREA Discussion Papers*, (13).

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, *25*(2), 242–262.

Blaze, J. (2019). *Railcar economics are as complex as the movement of freight*. Freight Waves. https://www.freightwaves.com/news/economics-of-railcars-are-complex

Bonadio, B. (2021). *Ports vs. roads: Infrastructure, market access and regional outcomes* (tech. rep.). Working Paper.

Brancaccio, G., Kalouptsidi, M., & Papageorgiou, T. (2020). Geography, transportation, and endogenous trade costs. *Econometrica*, *88*(2), 657–691.

Bresnahan, T. F. (1982). The oligopoly solution concept is identified. *Economics Letters*, *10*(1–2), 87–92.

Bureau of Transportation Statistics and Federal Highway Administration, U.S. DOT. (2022). *Freight analysis framework version 5 (faf5) regional database* [Accessed September 06, 2022]. https://faf.ornl.gov/faf5/data/download_files/FAF5.3.zip.

Coşar, A. K., & Demir, B. (2018). Shipping inside the box: Containerization and trade. *Journal of International Economics*, *114*(100), 331–345. https://doi.org/10.1016/j.jinteco.2018.07

Costello, B., & Karickhoff, A. (2019). *Truck driver shortage analysis* (tech. rep.). American Trucking Associations.

Donaldson, D. (2018). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review*, *108*(4–5), 899–934.

Donaldson, D., & Hornbeck, R. (2016). Railroads and american economic growth: A "market access" approach. *The Quarterly Journal of Economics*, *131*(2), 699–858.

Dowle, M., & Srinivasan, A. (2023). *Data.table: Extension of 'data.frame'* [R package version 1.14.8]. https://CRAN.R-project.org/package=data.table

Ducruet, C., Juhász, R., Nagy, D. K., & Steinwender, C. (2020). *All aboard: The effects of port development* (Working Paper No. 28148). National Bureau of Economic Research. https://doi.org/10.3386/w28148

Eaton, J., & Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, *70*(5), 1741–1779.

Faber, B. (2014). Trade integration, market size, and industrialization: Evidence from china's national trunk highway system. *Review of Economic Studies*, *81*(3), 1046–1070.

FMCSA. (2021). *Pocket guide to large truck and bus statistics* (tech. rep.). Federal Motor Carrier Safety Administration, U.S. Department of Transportation.

Fuchs, S., & Wong, W. F. (2023). *Multimodal transport networks* (Federal Reserve Bank of Atlanta Working Paper # 2022-13).

Ganapati, S., Wong, W. F., & Ziv, O. (2021). *Entrepot: Hubs, scale, and trade costs* (tech. rep.). National Bureau of Economic Research.

Harris, S. (2021). *How much does a semi truck cost?* Freight Waves Ratings. https://ratings.freightwaves.com/how-much-does-a-semi-truck-cost/

Heiss, F. (2002). Structural choice analysis with nested logit models. *Stata Journal*, *2*(3), 227–252. https://EconPapers.repec.org/RePEc:tsj:stataj:v:2:y:2002:i:3:p:227-252

Jaworski, T., Kitchens, C., & Nigai, S. (2023). Highways and globalization. *International Economic Review*, *64*(4), 1615–1648. https://doi.org/https://doi.org/10.1111/iere.12640

Kahle, D., & Wickham, H. (2013). Ggmap: Spatial visualization with ggplot2. *The R Journal*, *5*(1), 144–161. https://journal.r-project.org/archive/2013-1/kahle-wickham.pdf

Leibovici, F. (May 11, 2022). International shipping costs during and after covid-19. *Timely Topics (St. Louis Federal Reserve podcast)*. https://www.stlouisfed.org/timely-topics/is-covid-19-sole-cause-of-shipping-disruptions-delays

Lind, N., & Ramondo, N. (2023). Trade with correlation. *American Economic Review*, *113*(2), 317–53. https://doi.org/10.1257/aer.20190781

MARAD. (September, 2022). *Contact information for u.s.-flag carriers* (tech. rep.). Maritime Administration, Office of Cargo Commercial Sealift. %7Bhttps://www.maritime.dot.gov/sites/marad.dot.gov/files/2022-09/U.S.-flag%20Contact%20List%20Sept%202022.pdf%7D

McFadden, D. L. (1984). Econometric analysis of qualitative response models. In Z. Griliches & M. Intriligator (Eds.), *Handbook of econometrics*. Elsevier Science Publishers.

Microsoft, & Weston, S. (2022a). *Doparallel: Foreach parallel adaptor for the 'parallel' package* [R package version 1.0.17]. https://CRAN.R-project.org/package=doParallel

Microsoft, & Weston, S. (2022b). *Foreach: Provides foreach looping construct* [R package version 1.5.2]. https://CRAN.R-project.org/package=foreach

Miller, N. H., & Weinberg, M. C. (2017). Undesrtanding the price effects of the millercoors joint venture. *Econometrica*, *85*(6), 1763–1791. Retrieved April 28, 2023, from http://www.jstor.org/stable/44955182

R Core Team. (2022). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. https://www.R-project.org/

Samuelson, P. A. (1954). The transfer problem and transport costs, ii: Analysis of effects of trade impediments. *The Economic Journal*, *64*(254), 264–289. Retrieved May 30, 2023, from http://www.jstor.org/stable/2226834

Shepardson, D., & Baertlein, L. (November 29, 2022). U.S. House to Vote to Block Rail Strike Despite Labor Objections. *Reuters*.

Simonovska, I., & Waugh, M. E. (2014). The elasticity of trade: Estimates and evidence. *Journal of International Economics*, *92*(1), 34–50. https://doi.org/https://doi.org/10.1016/j.jinteco.2013.10.001

Stubner, R. (2021). *Dqrng: Fast pseudo random number generators* [R package version 0.3.0]. https://CRAN.R-project.org/package=dqrng

Train, K. E., McFadden, D. L., & Ben-Akiva, M. (1987). The demand for local telephone service: A fully discrete choice model of residential calling patterns and service choices. *The RAND Journal of Economics*, *18*(1), 109–123.

UNCTAD. (2010). *Review of maritime transport (2010)* (tech. rep.). United Nations Conference on Trade and Development.

Walsh, S. (2023). *How much does a cargo plane cost?* Pilot Passion. https://pilotpassion.com/how-much-does-a-cargo-plane-cost/

Wickham, H. (2022). *Stringr: Simple, consistent wrappers for common string operations* [R package version 1.5.0]. https://CRAN.R-project.org/package=stringr

Wickham, H., François, R., Henry, L., & Müller, K. (2022). *Dplyr: A grammar of data manipulation* [R package version 1.0.10]. https://CRAN.R-project.org/package=dplyr

Wickham, H., Ooms, J., & Müller, K. (2023). *Rpostgres: Rcpp interface to postgresql* [R package version 1.4.5]. https://CRAN.R-project.org/package=RPostgres

World Wide Rails. (2023). *How much do locomotives cost?* https://worldwiderails.com/how-much-do-locomotives-cost/

Wu, C. J. (1983). On the convergence properties of the em algorithm. *The Annals of Statistics*, *11*(1), 95–103.

7 Figures & Tables

Figure 1: A Simple Network

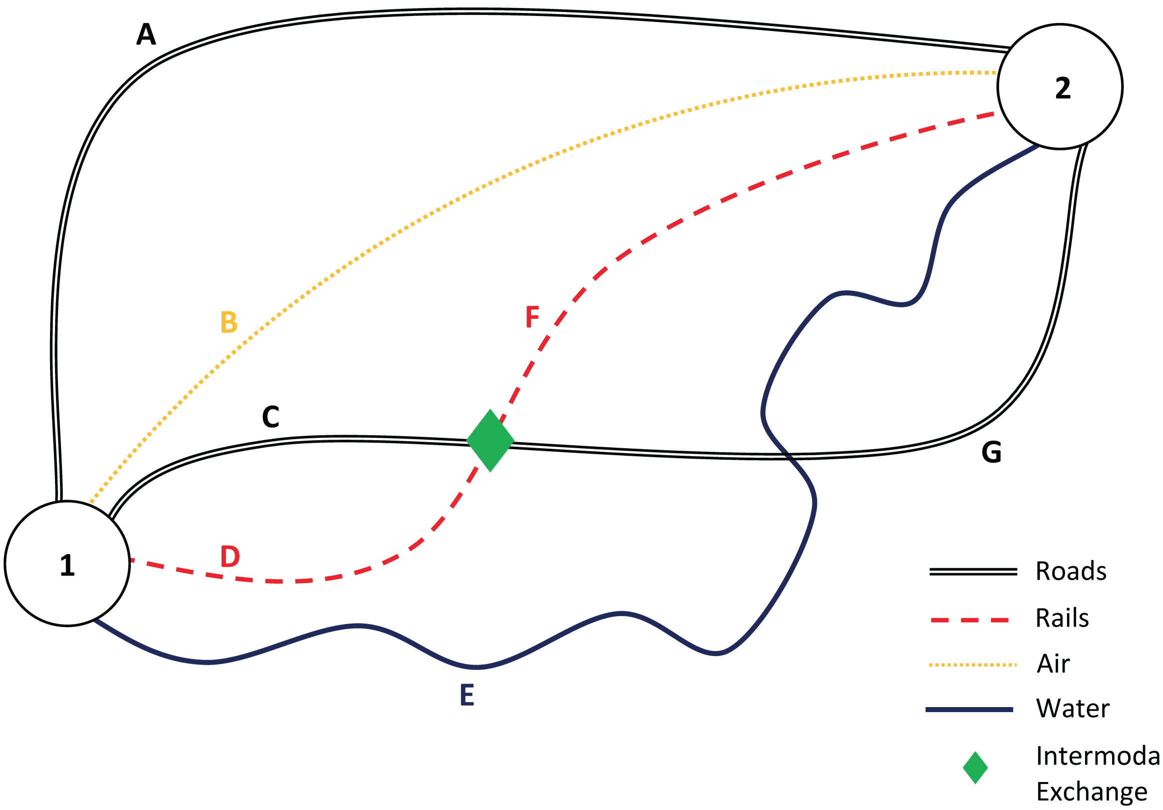

Notes: The above figure displays a simplified transport network to illustrate the basic logic of my setup. In this case, $\mathcal{S}^d = \{1, 2\}$, and the sets of simple (non-repeating) routes along each mode are given by: $\mathcal{R}_{ex}^{Road} = \{A, CG\}$, $\mathcal{R}_{ex}^{Rail} = \{DF\}$, $\mathcal{R}_{ex}^{Air} = \{B\}$, $\mathcal{R}_{ex}^{Water} = \{E\}$, and $\mathcal{R}_{ex}^{Multi\text{-}modal} = \{A, CF, CG, DF, DG\}$. An important point regarding multi-modal: intermodal transshipment may only occur at designated intermodal exchanges (e.g. the green diamond); hence, even though the road network crosses a waterway, an exchange is not possible at this point. Additionally, multi-modal routes need not utilize the intermodal exchange. The defining feature of multi-modal is the *ability* to move across multiple networks, but not a requirement.

Figure 2: Details of the Nesting Structure

(a) Foreign Destinations

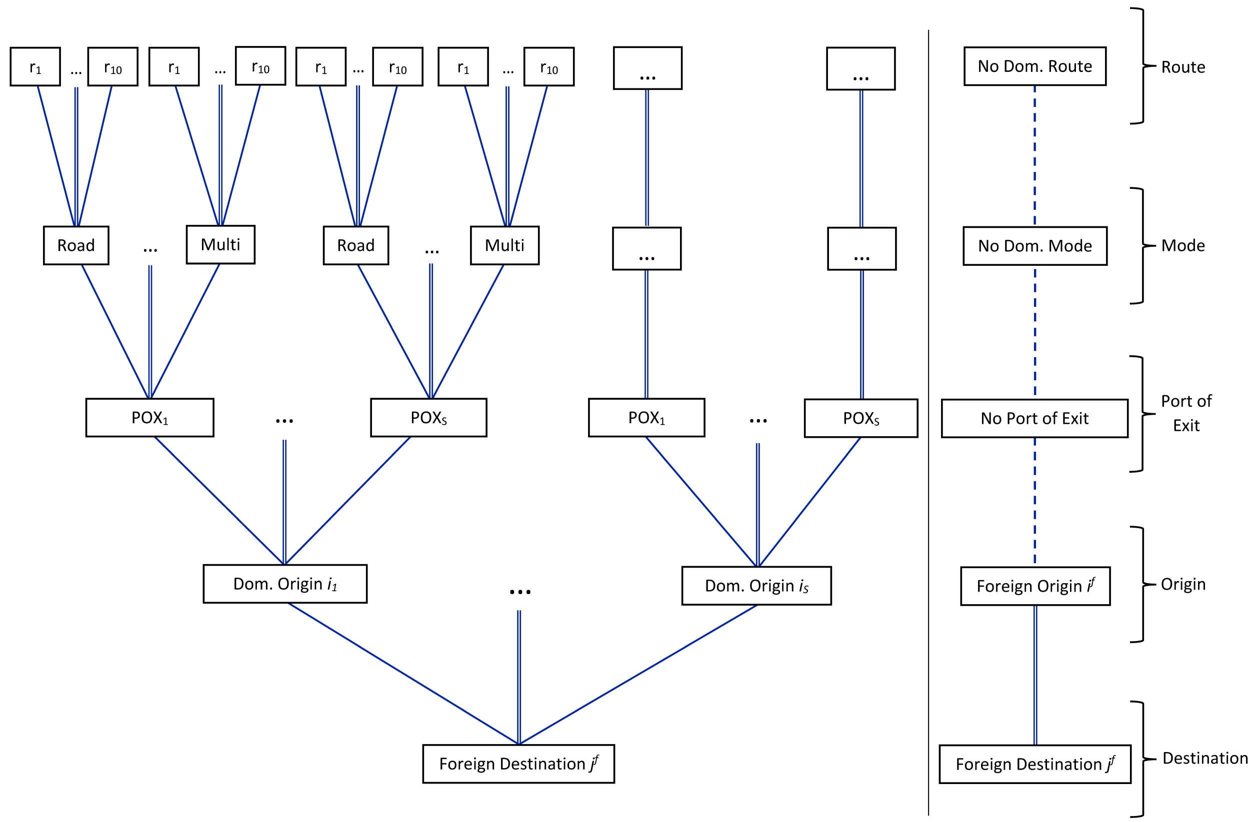

Continued on next page...

(b) **Domestic Destinations**

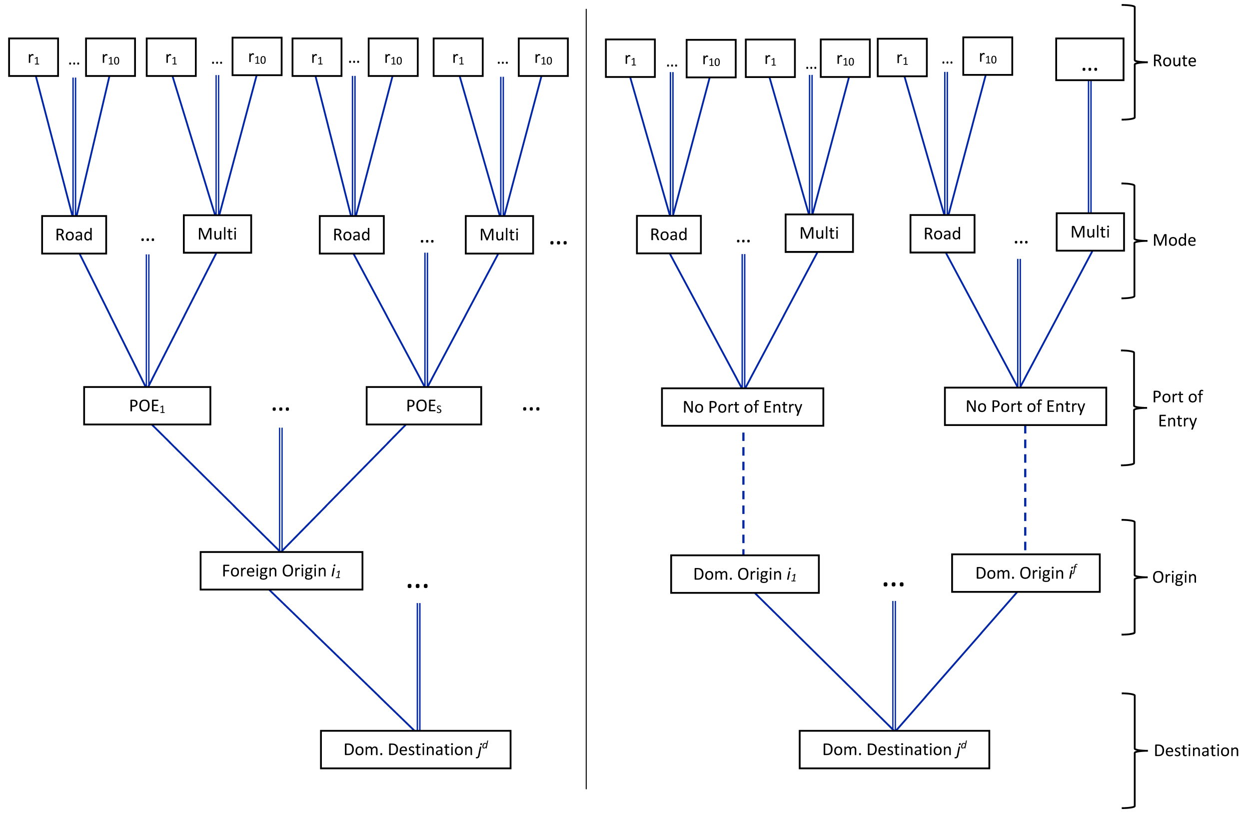

Notes: Panel A presents the nesting structure for foreign destinations; panel B presents the nesting structure for domestic destinations. Multiple options are presented as a compound line; degenerate branches are represented by a dashed line. Note that if a foreign destination elects a foreign origin, the tree becomes completely degenerate thereafter. Similarly, if a domestic location elects a domestic origin, the choice of port of entry/exit becomes degenerate (as there is no choice).

Figure 3: FAF Regions

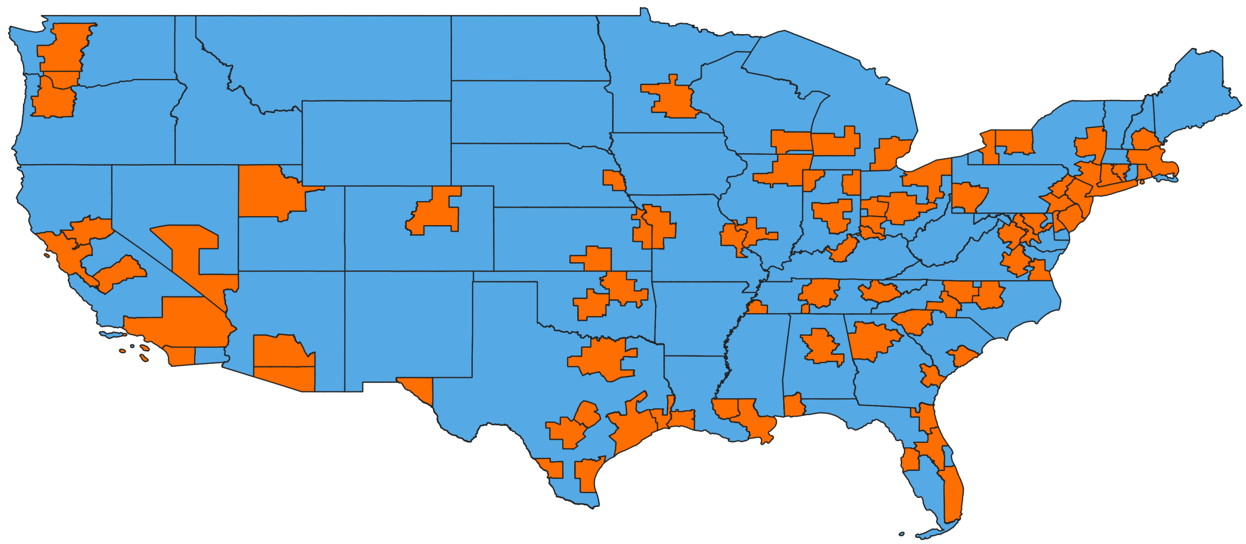

Notes: Black lines denote boundaries between regions. Bright orange areas denote major metropolitan areas. Light blue areas encompass large, rural areas. Note that there is at least one large, rural swath per state; rural areas are distinguished by their state.

Figure 4: National Transport Network

(a) **Roads & Intermodal**

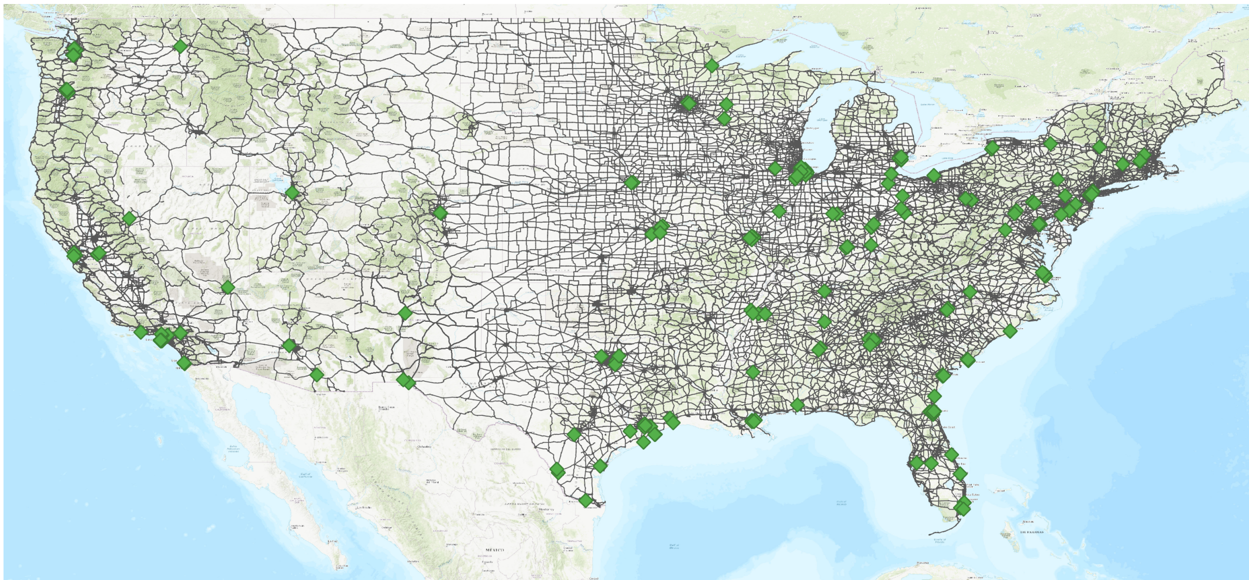

(b) **Rail, Water, Air, & Intermodal**

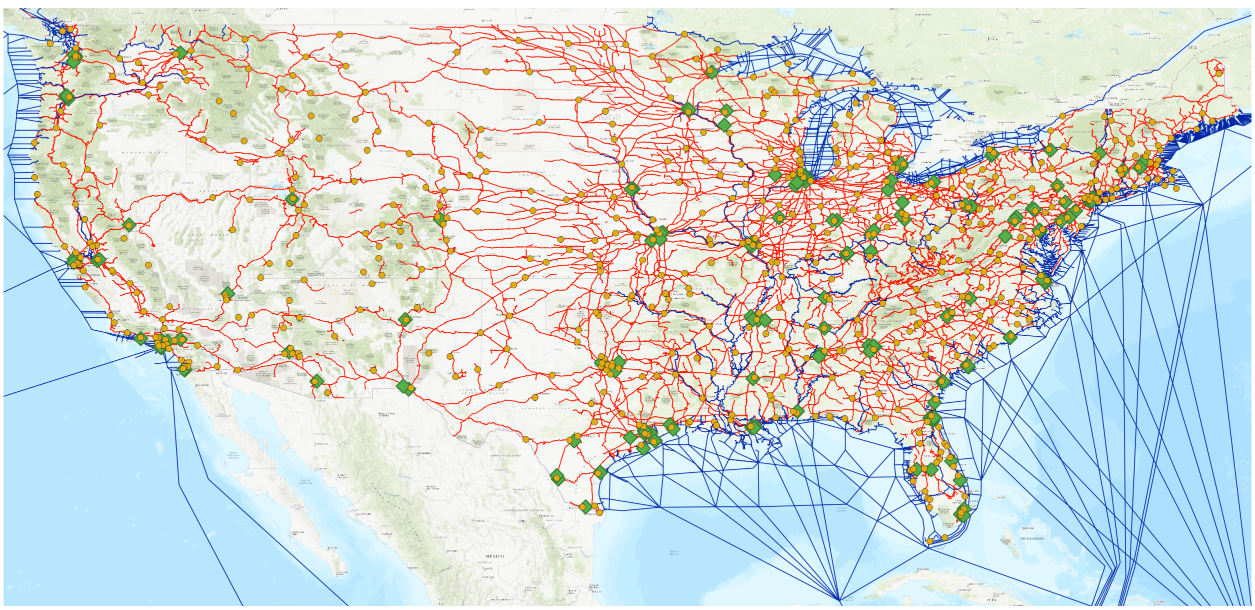

Notes: This figure displays the domestic transport network in the mainland U.S.. Roads are depicted in dark grey, rails in red, navigable waterways in dark blue, airports accepting freight as golden-yellow points, and intermodal exchanges as green diamonds. Panels are split to make the separate networks clearer. Multi-model transshipment may only occur at designated intermodal exchanges. I assume that travel may occur in a straight line between airports. Background map provided by ESRI World Topographical map.

Table 1: Evolution of Routed Mileage by Mode

	Road		Rail		Water		Air		Multi-modal	
Route	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1	1	–	1	–	1	–	1	–	1	–
2	1.044	*+4.4%*	1.055	*+5.5%*	1.042	*+4.2%*	1.002	*+0.2%*	1.043	*+4.3%*
3	1.074	*+2.8%*	1.095	*+3.8%*	1.066	*+2.3%*	1.005	*+0.3%*	1.069	*+2.5%*
4	1.099	*+2.3%*	1.122	*+2.4%*	1.135	*+6.5%*	1.009	*+0.4%*	1.093	*+2.2%*
5	1.119	*+1.8%*	1.154	*+2.9%*	1.153	*+1.6%*	1.013	*+0.4%*	1.114	*+1.9%*

Notes: This table displays the average evolution of mileage along each route across all origin-destination pairs by mode. The odd-numbered columns display the average, scaled mileage per route; the shortest path takes a value of 1 for all origin-destination pairs. The even-numbered columns display the incremental percentage change in this average, scaled mileage.

Table 2: Estimated Fundamentals

A. Geographic Costs $\hat{\beta}_m$

per-Mile Costs				Multi-Modal Exchange
Road	Rail	Water	Air	
0.200***	0.125***	0.111*	0.085***	0.020***
(0.018)	(0.02)	(0.046)	(0.013)	(0.003)

B. Inverse Correlation Coefficients

Route Correlation φ	Mode Correlation $\hat{\rho}$	Port Correlation $\hat{\zeta}$
1	0.543***	0.374***
	(0.013)	(0.009)

C. Markups $\hat{\mu}_{irj}$

Min	P25	P50	P75	Max	Mean	Std. Dev.
1.000	1.022*	1.030*	1.043**	1.130***	1.035***	0.021***
	(0.013)	(0.018)	(0.020)	(0.033)	(0.010)	(0.004)

D. Model Summary

N	Adj. R^2	First-Stage F-stat
382,085	0.921	37.825

Regression specification utilizes 5 routings per mode and are run with Origin $\times$ POE$\times$ Mode and Destination $\times$ POX $\times$ Mode FEs. Bootstrapped standard errors utilizing 500 iterations presented in parentheses. ***, **, and * respectively denote 0.1%, 1%, and 5% significance. Note that, in Panels B and C, I report significance tests relative to the null of 1. All parameter estimates assume a trade elasticity of 4.

Figure 5: Histogram of Estimated Conduct Parameters

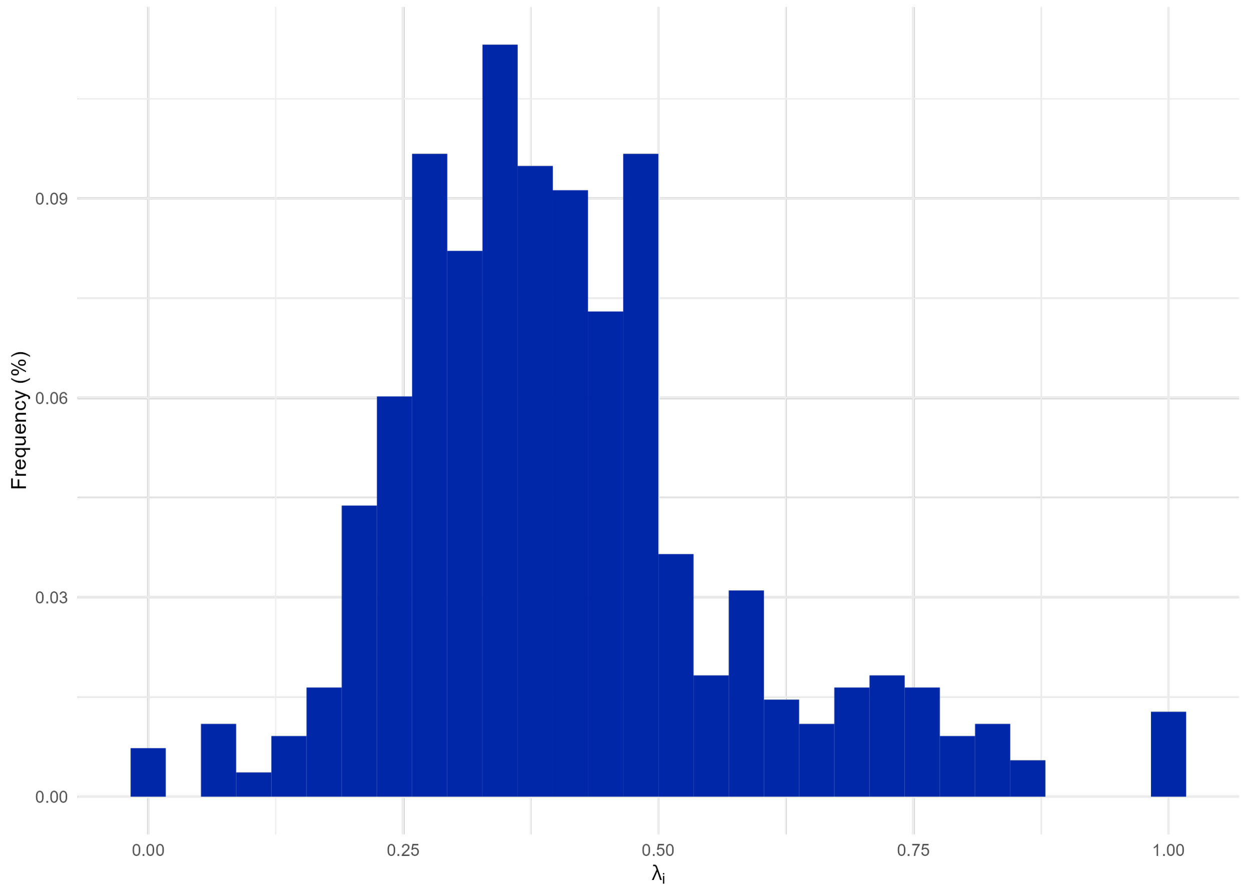

Notes: This figure displays the distribution of my estimated conduct parameters. All estimates assume a trade elasticity of 4, and utilize my favored specification, which encompasses 5 potential routings per origin-destination-mode combination. Note that all estimates are normalized to span the $[0, 1]$ range.

Table 3: Parameter Estimates under Stricter Modelling Assumptions

	Main Specification	No Markups	No Correlation	Single Route
	(1)	(2)	(3)	(4)
A. Cost Coefficients				
Road	0.200	0.198	0.215	0.226
Rail	0.125	0.122	0.099	0.076
Water	0.111	0.108	0.086	0.091
Air	0.085	0.083	0.113	0.113
Multi	0.020	0.019	0.009	0.011
B. Inverse Correlation Parameters				
Route φ	1.000	1.000	1.000	1.000
Mode $\hat{\rho}$	0.543	0.542	1.000	1.000
Origin $\hat{\zeta}$	0.374	0.376	1.000	1.000
C. Distribution of Markups				
Min	1.000	–	–	–
P25	1.022	–	–	–
P50	1.030	–	–	–
P75	1.043	–	–	–
Max	1.130	–	–	–
Mean	1.035	–	–	–
s.d.	0.021	–	–	–
D. Model Summary				
Observations	382,085	382,085	382,085	76,417
Adj. R^2	0.921	0.920	0.843	0.833
1st Stage F-stat	37.823	–	–	–

Notes: This table displays my reduced-form estimates under simpler modelling specifications. Column (1) restates the results of my primary specification; Column (2) shows my cost and correlation estimates, assuming perfect freight-market competition; Column (3) further restricts the model by assuming zero correlation of prices across geographies, modes; and finally, Column (4) imposes a single route, as opposed to 5.

Table 4: Sensitivity to Non-Zero Route-Correlation

	Main Specification	φ 0.75	φ 0.5	φ 0.25
	(1)	(2)	(3)	(4)
A. Cost Coefficients				
Road	0.200	0.200	0.200	0.200
Rail	0.125	0.125	0.125	0.125
Water	0.111	0.111	0.111	0.111
Air	0.085	0.085	0.085	0.085
Multi	0.020	0.020	0.020	0.020
B. Inverse Correlation Parameters				
φ	1.000	0.750	0.500	0.250
$\hat{\rho}$	0.543	0.543	0.543	0.543
$\hat{\zeta}$	0.374	0.374	0.374	0.374
C. Distribution of Markups				
Min	1.000	1.000	1.000	1.000
P25	1.022	1.022	1.022	1.022
P50	1.030	1.030	1.030	1.030
P75	1.043	1.043	1.043	1.043
Max	1.130	1.130	1.130	1.130
Mean	1.035	1.035	1.035	1.035
s.d.	0.021	0.021	0.021	0.021
D. Model Summary				
Observations	382,085	382,085	382,085	382,085
Adj. R^2	0.921	0.921	0.921	0.921
1st Stage F-stat	37.823	37.823	37.823	37.824

Notes: This table displays my reduced-form estimates under different values of the inverse route-correlation parameter φ. Column (1) re-states my main results; Columns (2) - (4) display my results assuming a value of φ equal to 0.75, 0.5, and 0.25, respectively.

Table 5: Additional Robustness Checks

	Main Specification	Restricted Commodities	Excluding Water & Air	Urban Areas Only	Alternate Years 2012	Alternate Years 2022
	(1)	(2)	(3)	(4)	(5)	(6)
A. Cost Coefficients						
Road	0.200	0.189	0.202	0.159	0.194	0.200
Rail	0.125	0.107	0.130	0.093	0.126	0.116
Water	0.111	0.059	–	0.079	0.114	0.092
Air	0.085	0.062	–	0.070	0.093	0.090
Multi	0.020	0.006	0.021	0.017	0.007	0.020
B. Inverse Correlation Parameters						
φ	1.000	1.000	1.000	1.000	1.000	1.000
$\hat{\rho}$	0.543	0.643	0.512	0.481	0.625	0.500
$\hat{\zeta}$	0.374	0.399	0.410	0.348	0.393	0.378
C. Distribution of Markups						
Min	1.000	1.000	1.000	1.000	1.000	1.000
P25	1.022	1.033	1.022	1.015	1.026	1.015
P50	1.030	1.047	1.035	1.024	1.035	1.024
P75	1.043	1.067	1.053	1.036	1.049	1.038
Max	1.130	1.165	1.164	1.125	1.123	1.121
Mean	1.035	1.053	1.042	1.029	1.040	1.030
s.d.	0.021	0.031	0.030	0.021	0.022	0.023
D. Model Summary						
Observations	382,085	277,6100	301,520	200,665	359,275	400,140
Adj. R^2	0.921	0.915	0.904	0.935	0.818	0.855
1st Stage F-stat	37.823	22.355	32.167	34.664	19.365	35.512

Notes: This table displays the robustness of my estimates to alternate specifications. Column (1) re-states my main results; Column (2) restricts the set of commodities to those that are frequently carried by all modes; Column (3) excludes Water and Air routes; Column (4) omits rural areas; Columns (5) and (6) re-estimate my primary specification using trade data from 2012 and 2022, respectively.

Table 6: Robustness to Alternate Starting Values

	Main	Mode Receiving Weight				
	Specification	Road	Rail	Water	Air	Multi
	(1)	(2)	(3)	(4)	(5)	(6)
A. Cost Coefficients						
Road	0.200	0.200	0.200	0.200	0.200	0.200
Rail	0.125	0.125	0.125	0.125	0.125	0.125
Water	0.111	0.111	0.111	0.111	0.111	0.111
Air	0.085	0.085	0.085	0.085	0.085	0.085
Multi	0.020	0.020	0.020	0.020	0.020	0.020
B. Inverse Correlation Parameters						
φ	1.000	1.000	1.000	1.000	1.000	1.000
$\hat{\rho}$	0.543	0.543	0.543	0.543	0.543	0.543
$\hat{\zeta}$	0.374	0.374	0.374	0.374	0.374	0.374
C. Distribution of Markups						
Min	1.000	1.000	1.000	1.000	1.000	1.000
P25	1.022	1.022	1.022	1.022	1.022	1.022
P50	1.030	1.030	1.030	1.030	1.030	1.030
P75	1.043	1.043	1.043	1.043	1.043	1.043
Max	1.130	1.130	1.130	1.130	1.130	1.130
Mean	1.035	1.035	1.035	1.035	1.035	1.035
s.d.	0.021	0.021	0.021	0.021	0.021	0.021
D. Model Summary						
Observations	382,085	382,085	382,085	382,085	382,085	382,085
Adj. R^2	0.921	0.921	0.921	0.921	0.921	0.921
1st Stage F-stat	37.823	37.824	37.824	37.823	37.823	37.824

Notes: This table displays the robustness of my estimates to alternate starting values. Explicitly, I parameterize the starting values to take a value of 1 for a specific mode and zero for the remaining. The mode receiving all the weight (initially) is specified in each column.

Table 7: Change in Real Welfare, Imports, and Exports under Perfect Freight Competition

	Baseline Trade Share (%)	Subs. Effect (ppt)	Income Effect (ppt)	Total Change (ppt)
	(1)	(2)	(3)	(4)
A. National Welfare				
Road	78.84	1.85	4.14	5.99
Rail	3.79	-0.17	0.12	-0.04
Water	1.61	0.16	0.13	0.28
Air	1.95	-0.18	0.14	-0.05
Multi	13.82	-1.65	0.52	-1.13
Total	100.00	0.00	5.04	5.04
B. Real Imports				
Road	60.72	3.37	5.22	8.59
Rail	9.53	-0.88	0.15	-0.72
Water	3.35	-0.25	0.16	-0.09
Air	10.27	-0.78	0.81	0.03
Multi	16.13	-1.46	1.18	-0.28
Total	100.00	0.00	7.52	7.52
C. Real Exports				
Road	67.03	1.96	13.59	15.54
Rail	11.39	-1.51	1.58	0.07
Water	8.20	0.11	2.39	2.50
Air	6.33	-0.14	0.99	0.85
Multi	7.05	-0.42	1.29	0.87
Total	100.00	0.00	19.84	19.84

Notes: This table displays changes in real welfare accruing to each mode stemming from elimination of non-competitive pricing in the transport sector. Column (1) displays each mode's share of national trade in baseline, and is included for reference; Column (2) lists the change in trade volumes across modes, holding incomes constant; Column (3) lists the change in trade volumes resulting from a change in income, holding mode, route, and port choices constant; Column (4) = Column (2) + Column (3) lists the total change in transport volume (as a percent of national welfare) along each mode. The simulated changes are calculated using 5 potential routings for each mode-origin-destination combination.

Figure 6: Geographic Incidence from Elimination of Non-Competitive Markups

(a) **Welfare**

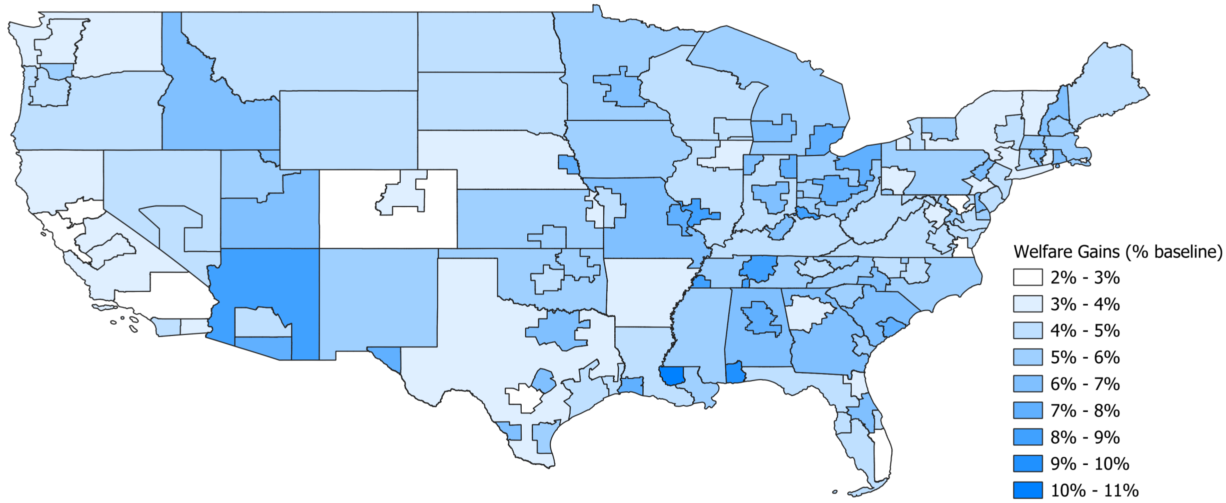

(b) **Imports**

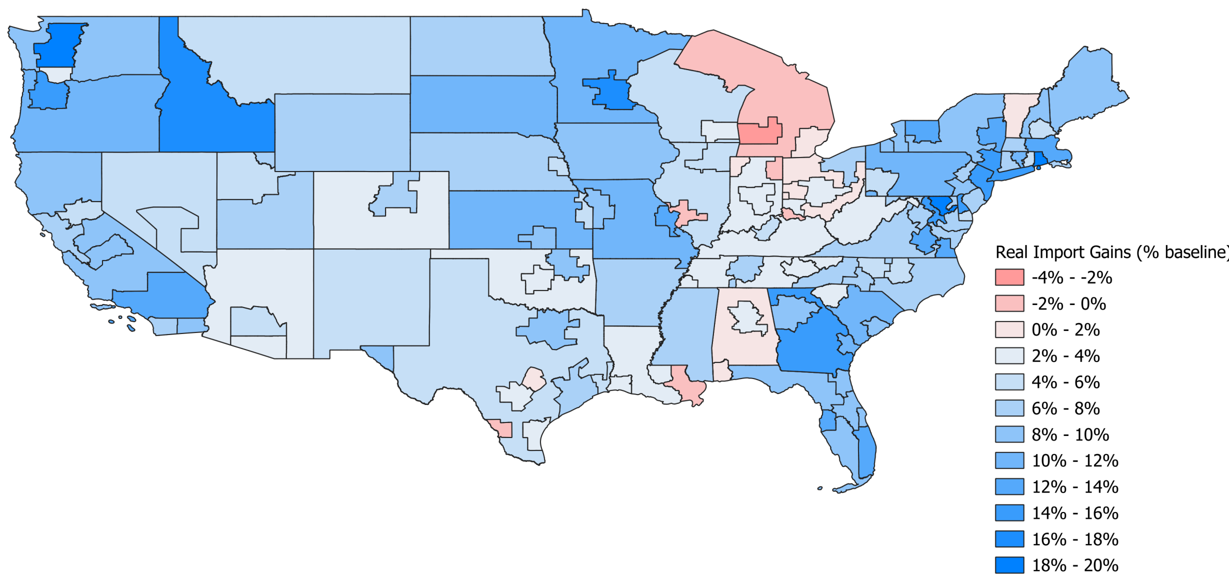

Continued on next page...

(c) **Exports**

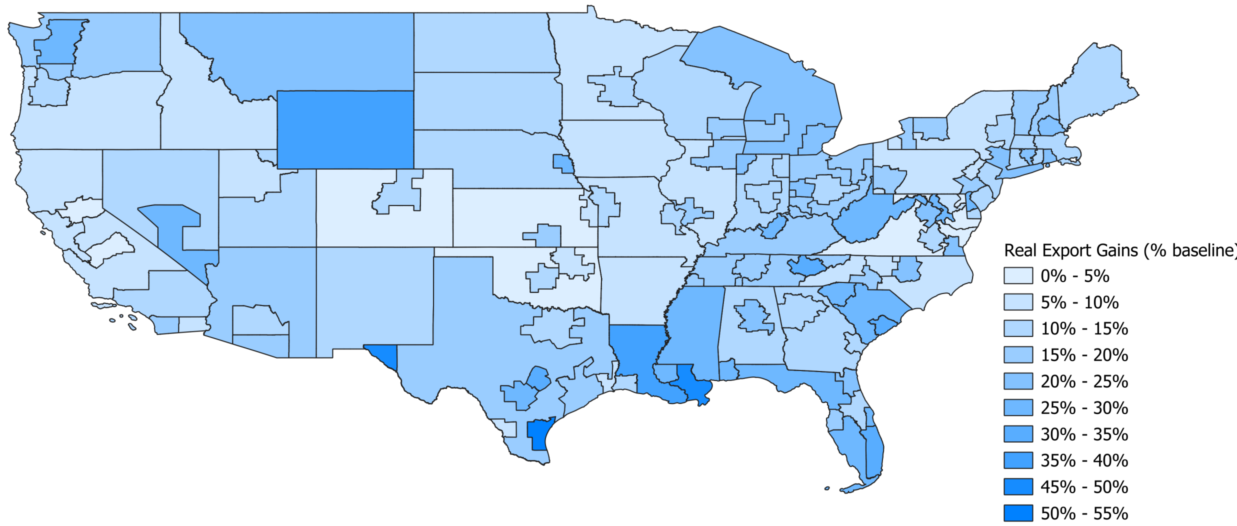

Notes: This figure presents potential change in real welfare, imports, and exports from imposing perfect competition in the transportation sector. These changes are calculated as the percentage increase in real welfare/ imports/ exports when eliminating markups over real welfare/ imports/ exports under the status quo. Darker colors signify greater changes– red denotes a decline in value, while blue denotes a gain.

Table 8: Change in Real Welfare under a Rail Strike

	Baseline Trade Share (%)	Subs. Effect (ppt)	Income Effect (ppt)	Total Change (ppt)
	(1)	(2)	(3)	(4)
A. Status-Quo Freight Pricing				
Road	78.84	15.91	-3.90	12.01
Rail	3.79	-3.79	0.00	-3.79
Water	1.61	0.34	-0.02	0.32
Air	1.95	1.27	-0.30	0.97
Multi	13.82	-13.82	0.00	-13.82
Total	100.00	-0.08	-4.23	-4.31
B. Perfect Freight Market Competition				
Road	80.75	14.43	-2.50	11.93
Rail	3.57	-3.57	0.00	-3.57
Water	1.80	0.32	-0.04	0.28
Air	1.81	0.79	-0.22	0.57
Multi	12.07	-12.07	0.00	-12.07
Total	100.00	-0.09	-2.77	-2.86

Notes: This table displays national welfare changes stemming from removal of rail and multi-modal transport. Column (1) displays each mode's share of national trade in baseline, and is included for reference; Column (2) lists the change in trade volumes across modes, holding incomes constant– note that, in this case, due to imperfect modal substitution and the elimination of certain modes, Column (2) need not sum to 0; Column (3) lists the change in trade volumes resulting from a change in income, holding mode, route, and port choices constant; Column (4) = Column (2) + Column (3) lists the total change in transport volume (as a percent of national welfare) along each mode. The simulated changes are calculated using 5 potential routings for each mode-origin-destination combination.

Table 9: Change in Real Imports under a Rail Strike

	Baseline Trade Share (%)	Subs. Effect (ppt)	Income Effect (ppt)	Total Change (ppt)
	(1)	(2)	(3)	(4)
A. Status-Quo Freight Pricing				
Road	60.72	18.58	-6.13	12.46
Rail	9.53	-9.53	0.00	-9.53
Water	3.35	1.51	-0.43	1.08
Air	10.27	5.57	-1.44	4.13
Multi	16.13	-16.13	0.00	-16.13
Total	100.00	-0.01	-7.99	-8.00
B. Perfect Freight Market Competition				
Road	64.46	17.52	-4.25	13.28
Rail	8.19	-8.19	0.00	-8.19
Water	3.02	1.30	-0.52	0.78
Air	9.57	4.11	-1.25	2.86
Multi	14.75	-14.75	0.00	-14.75
Total	100.00	-0.01	-6.02	-6.03

Notes: This table displays aggregate real import changes stemming from removal of rail and multi-modal transport. Column (1) displays each mode's share of national trade in baseline, and is included for reference; Column (2) lists the change in trade volumes across modes, holding incomes constant– note that, in this case, due to imperfect modal substitution and the elimination of certain modes, Column (2) need not sum to 0; Column (3) lists the change in trade volumes resulting from a change in income, holding mode, route, and port choices constant; Column (4) = Column (2) + Column (3) lists the total change in transport volume (as a percent of national welfare) along each mode .Panel A lists the change under status-quo freight pricing, while Panel B lists the change under perfect freight market competition. The simulated changes are calculated using 5 potential routings for each mode-origin-destination combination.

Table 10: Change in Real Exports under a Rail Strike

	Baseline Trade Share (%)	Subs. Effect (ppt)	Income Effect (ppt)	Total Change (ppt)
	(1)	(2)	(3)	(4)
A. Status-Quo Freight Pricing				
Road	67.03	15.83	-2.97	12.87
Rail	11.39	-11.39	0.00	-11.39
Water	8.20	0.71	-0.42	0.30
Air	6.33	1.89	-0.79	1.10
Multi	7.05	-7.05	0.00	-7.05
Total	100.00	0.00	-4.17	-4.17
B. Perfect Freight Market Competition				
Road	68.90	14.21	-0.35	13.86
Rail	9.56	-9.56	0.00	-9.56
Water	8.93	0.52	-0.29	0.22
Air	5.99	1.44	-0.58	0.86
Multi	6.61	-6.61	0.00	-6.61
Total	100.00	0.00	-1.23	-1.23

Notes: This table displays aggregate real export changes stemming from removal of rail and multi-modal transport. Column (1) displays each mode's share of national trade in baseline, and is included for reference; Column (2) lists the change in trade volumes across modes, holding incomes constant– note that, in this case, due to imperfect modal substitution and the elimination of certain modes, Column (2) need not sum to 0; Column (3) lists the change in trade volumes resulting from a change in income, holding mode, route, and port choices constant; Column (4) = Column (2) + Column (3) lists the total change in transport volume (as a percent of national welfare) along each mode. Panel A lists the change under status-quo freight pricing, while Panel B lists the change under perfect freight market competition. The simulated changes are calculated using 5 potential routings for each mode-origin-destination combination.

Figure 7: Geographic Welfare Incidence of a Rail Strike

(a) Status-Quo Freight Pricing

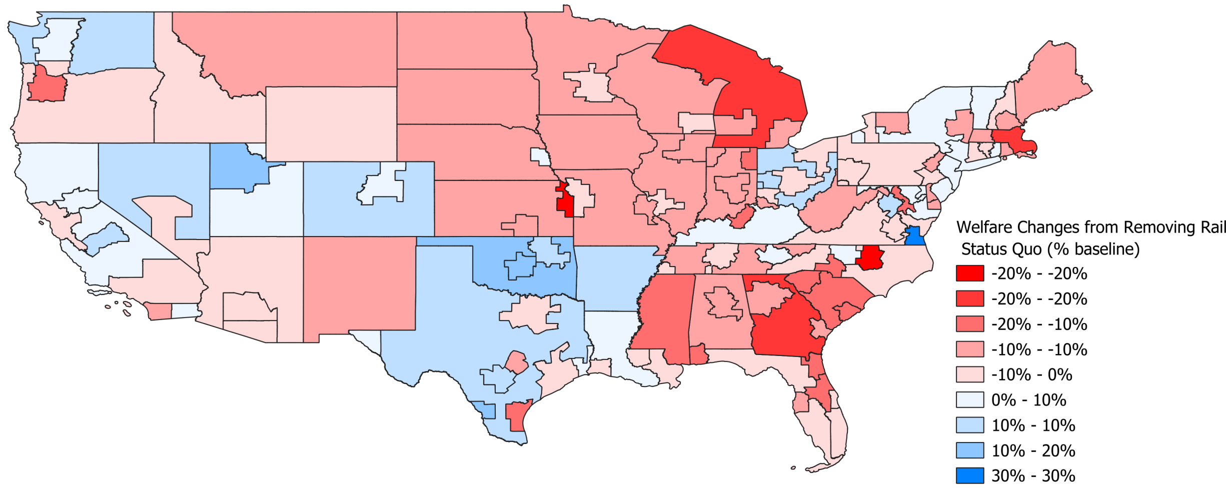

(b) Perfect Freight Market Competition

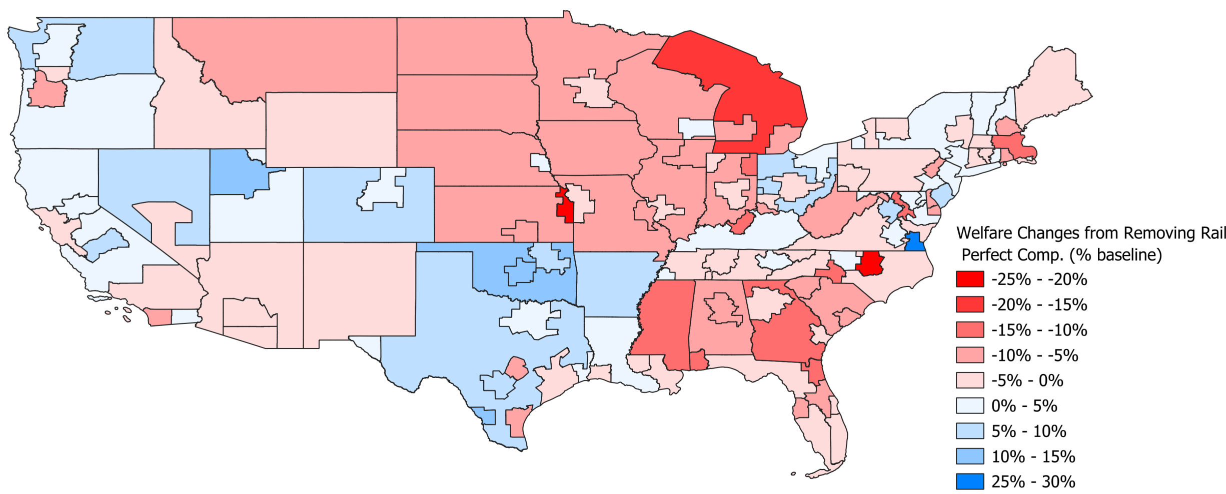

Notes: This figure presents percent changes in real welfare as a result of removing rail and multi-modal transport. Panel A reports welfare changes under the status-quo, while Panel B presents welfare changes under the perfect freight market competition. Darker colors signify greater changes– red denotes a decline in value, while blue denotes a gain.

Figure 8: Geographic Import Incidence of a Rail Strike

(a) **Status-Quo Freight Pricing**

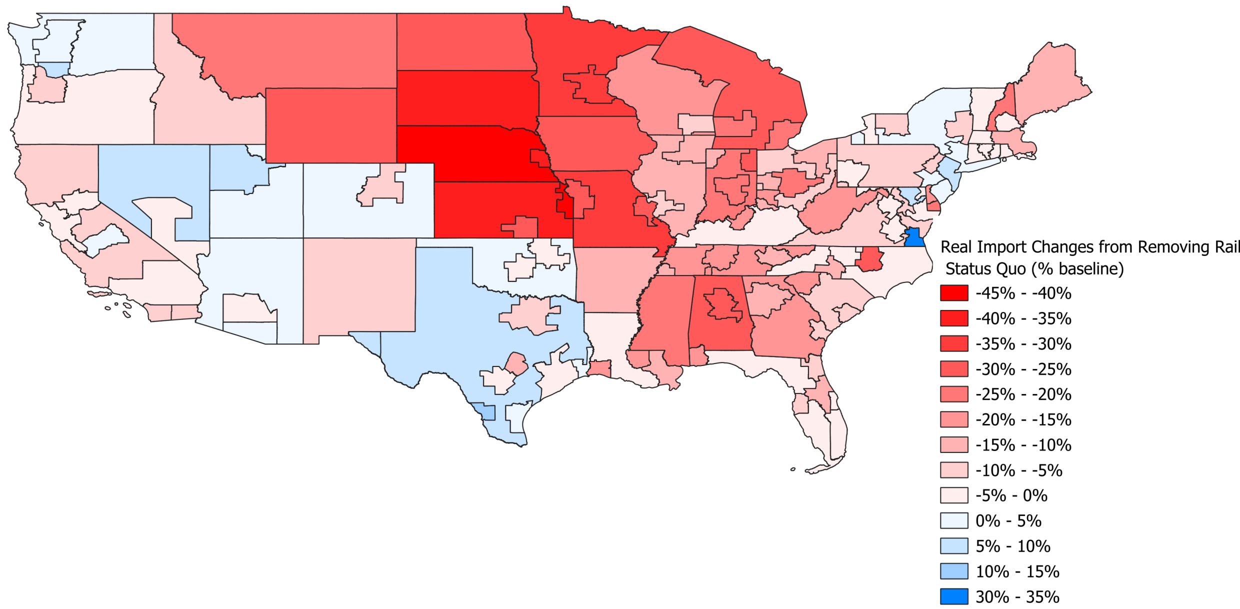

(b) **Perfect Freight Market Competition**

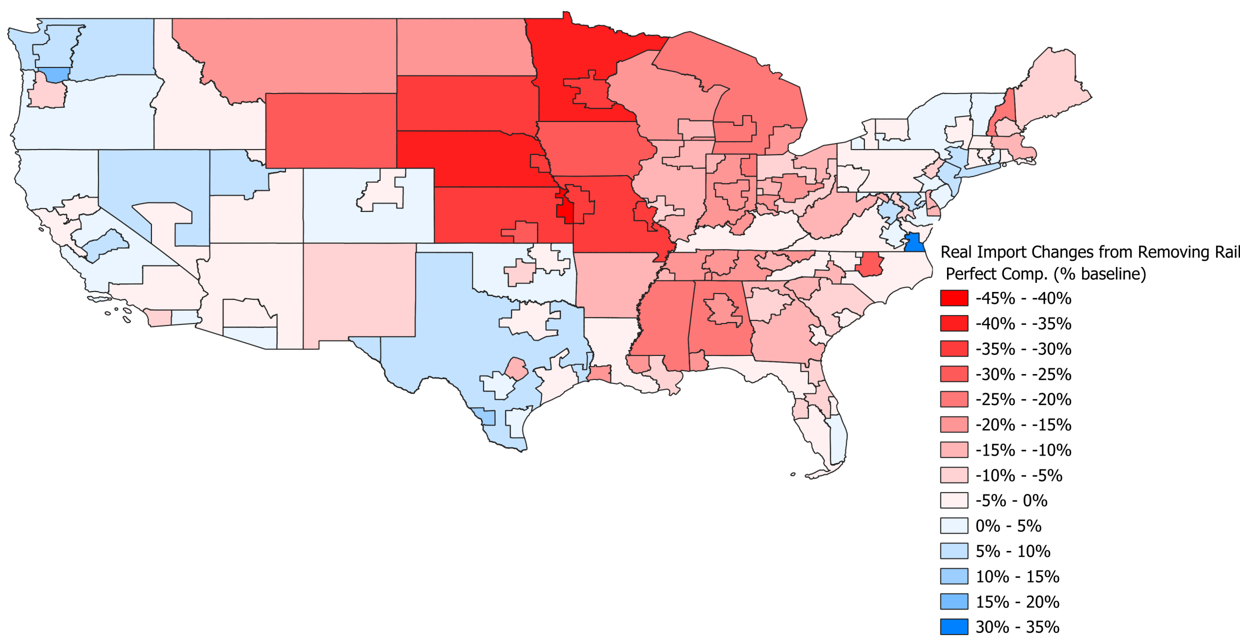

Notes: This figure presents percent changes in real imports as a result of removing rail and multi-modal transport. Panel A reports import changes under the status-quo, while Panel B presents import changes under the perfect freight market competition. Darker colors signify greater changes– red denotes a decline in value, while blue denotes a gain.

Figure 9: Geographic Export Incidence of a Rail Strike

(a) **Status-Quo Freight Pricing**

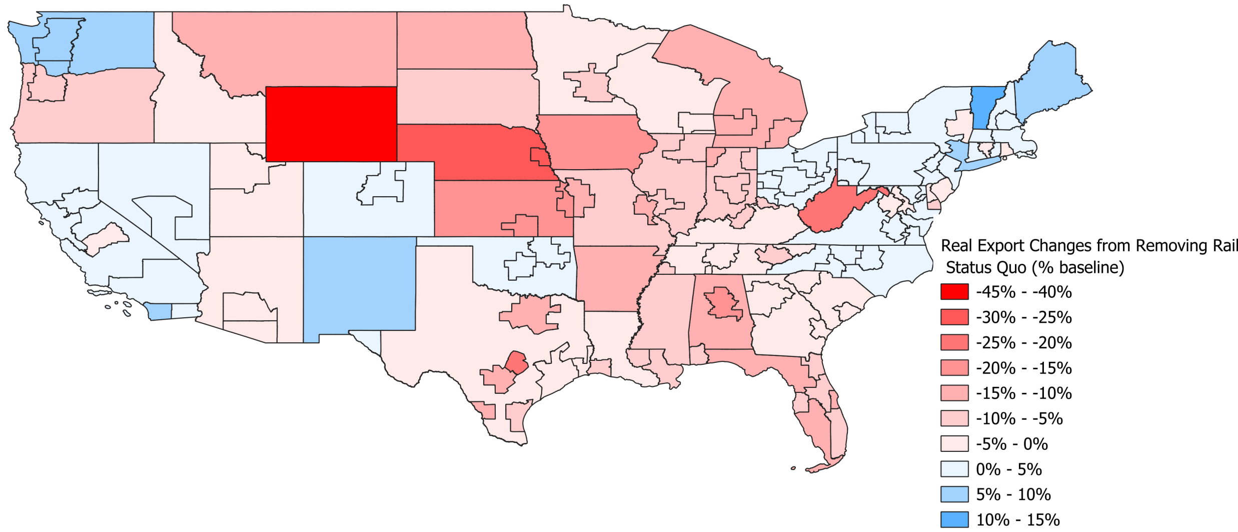

(b) **Perfect Freight Market Competition**

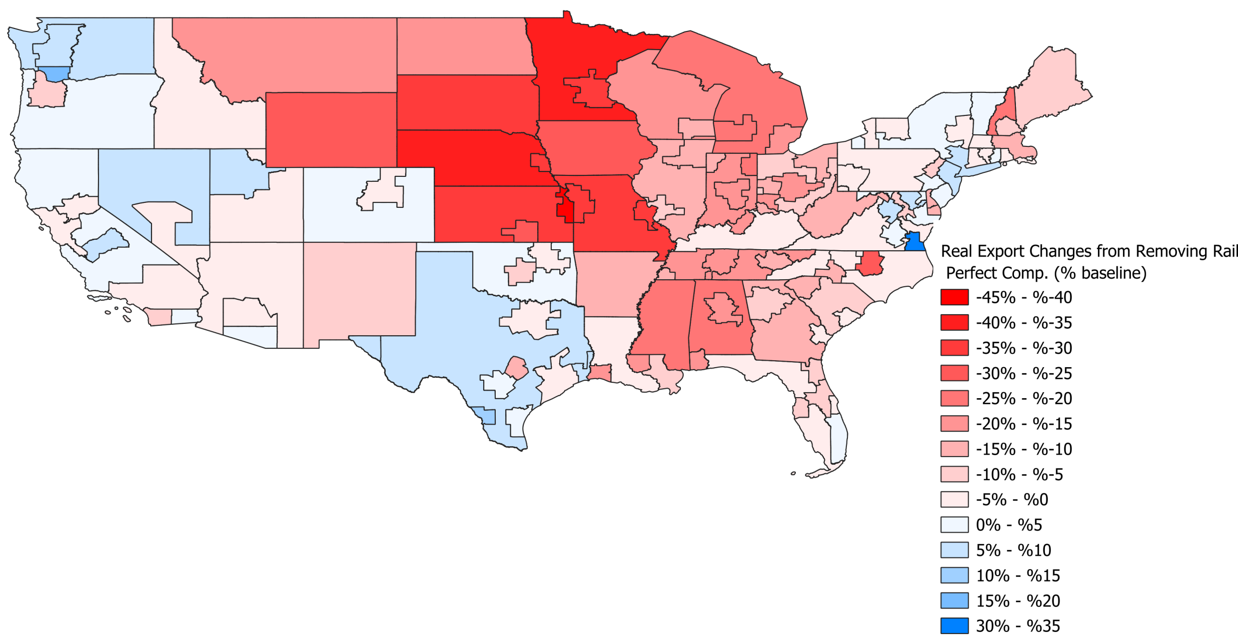

Notes: This figure presents percent changes in real exports as a result of removing rail and multi-modal transport. Panel A reports export changes under the status-quo, while Panel B presents export changes under the perfect freight market competition. Darker colors signify greater changes– red denotes a decline in value, while blue denotes a gain.

Table 11: Change in Real Welfare from an International Trade Shock

	Baseline Trade Share (%)	Subs. Effect (ppt)	Income Effect (ppt)	Total Change (ppt)
	(1)	(2)	(3)	(4)
A. Status-Quo Freight Pricing				
Road	78.84	-0.02	-5.33	-5.35
Rail	3.79	0.03	-0.41	-0.38
Water	1.61	0.02	-0.13	-0.10
Air	1.95	0.18	-0.28	-0.09
Multi	13.82	-0.22	-0.97	-1.19
Total	100.00	0.00	-7.11	-7.11
B. Perfect Freight Market Competition				
Road	80.75	-0.05	-5.56	-5.61
Rail	3.57	0.04	-0.39	-0.35
Water	1.80	0.02	-0.13	-0.11
Air	1.81	0.22	-0.28	-0.05
Multi	12.07	-0.24	-0.87	-1.10
Total	100.00	0.00	-7.22	-7.22

Notes: This table displays national welfare changes stemming from an international trade shock meant to emulate the jump in prices parked by the Covid-19 pandemic. Column (1) displays each mode's share of national trade in baseline, and is included for reference; Column (2) lists the change in trade volumes across modes, holding incomes constant; Column (3) lists the change in trade volumes resulting from a change in income, holding mode, route, and port choices constant; Column (4) = Column (2) + Column (3) lists the total change in transport volume (as a percent of national welfare) along each mode. The simulated changes are calculated using 5 potential routings for each mode-origin-destination combination.

Table 12: Change in Real Imports from an International Trade Shock

	Baseline Trade Share (%)	Subs. Effect (ppt)	Income Effect (ppt)	Total Change (ppt)
	(1)	(2)	(3)	(4)
A. Status-Quo Freight Pricing				
Road	60.72	-0.15	-10.55	-10.70
Rail	9.53	0.27	-1.87	-1.61
Water	3.35	0.17	-0.66	-0.49
Air	10.27	1.50	-2.04	-0.54
Multi	16.13	-1.79	-2.72	-4.51
Total	100.00	0.00	-17.84	-17.84
B. Perfect Freight Market Competition				
Road	64.46	-0.42	-11.08	-11.50
Rail	8.19	0.32	-1.63	-1.31
Water	3.02	0.19	-0.63	-0.44
Air	9.57	1.80	-2.02	-0.22
Multi	14.75	-1.89	-2.56	-4.46
Total	100.00	0.00	-17.92	-17.92

Notes: This table displays aggregate real import changes stemming from an international trade shock meant to emulate the jump in prices parked by the Covid-19 pandemic. Column (1) displays each mode's share of national trade in baseline, and is included for reference; Column (2) lists the change in trade volumes across modes, holding incomes constant; Column (3) lists the change in trade volumes resulting from a change in income, holding mode, route, and port choices constant; Column (4) = Column (2) + Column (3) lists the total change in transport volume (as a percent of national welfare) along each mode. The simulated changes are calculated using 5 potential routings for each mode-origin-destination combination.

Table 13: Change in Real Exports from an International Trade Shock

	Baseline Trade Share (%)	Subs. Effect (ppt)	Income Effect (ppt)	Total Change (ppt)
	(1)	(2)	(3)	(4)
A. Status-Quo Freight Pricing				
Road	67.03	-0.73	-14.78	-15.52
Rail	11.39	0.00	-2.61	-2.61
Water	8.20	0.36	-2.04	-1.68
Air	6.33	-0.19	-1.50	-1.69
Multi	7.05	0.57	-1.74	-1.16
Total	100.00	0.00	-22.67	-22.67
B. Perfect Freight Market Competition				
Road	68.90	-0.87	-15.12	-15.99
Rail	9.56	0.17	-2.23	-2.06
Water	8.93	0.32	-2.19	-1.87
Air	5.99	-0.22	-1.42	-1.63
Multi	6.61	0.60	-1.64	-1.04
Total	100.00	0.00	-22.59	-22.59

Notes: This table displays aggregate real export changes stemming from an international trade shock meant to emulate the jump in prices parked by the Covid-19 pandemic. Column (1) displays each mode's share of national trade in baseline, and is included for reference; Column (2) lists the change in trade volumes across modes, holding incomes constant; Column (3) lists the change in trade volumes resulting from a change in income, holding mode, route, and port choices constant; Column (4) = Column (2) + Column (3) lists the total change in transport volume (as a percent of national welfare) along each mode. The simulated changes are calculated using 5 potential routings for each mode-origin-destination combination.

Figure 10: Geographic Welfare Incidence of an International Trade Shock

(a) **Status-Quo Freight Pricing**

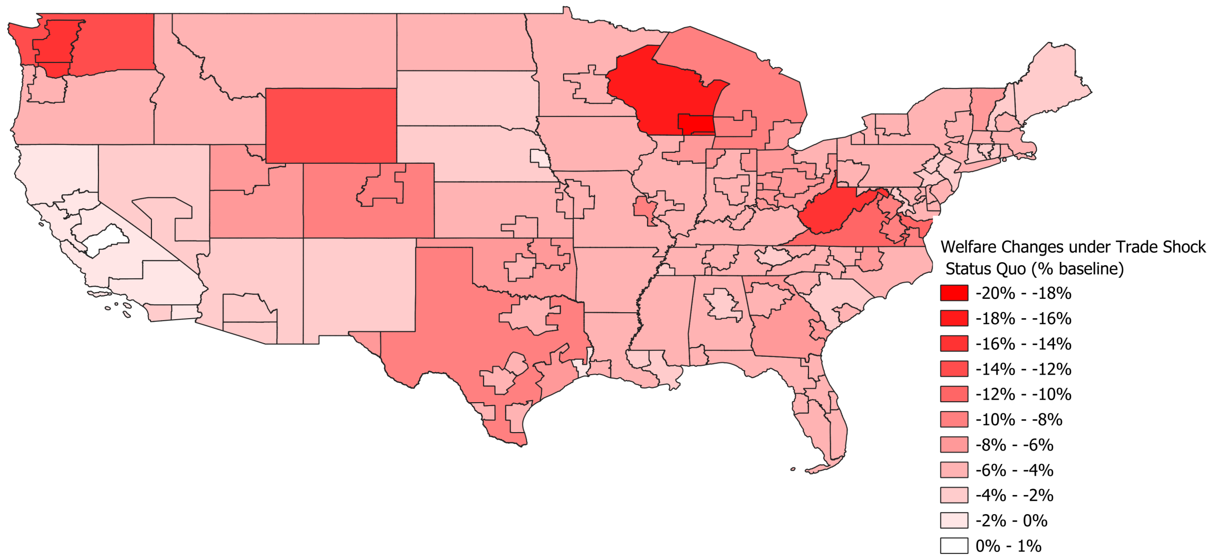

(b) **Perfect Freight Market Competition**

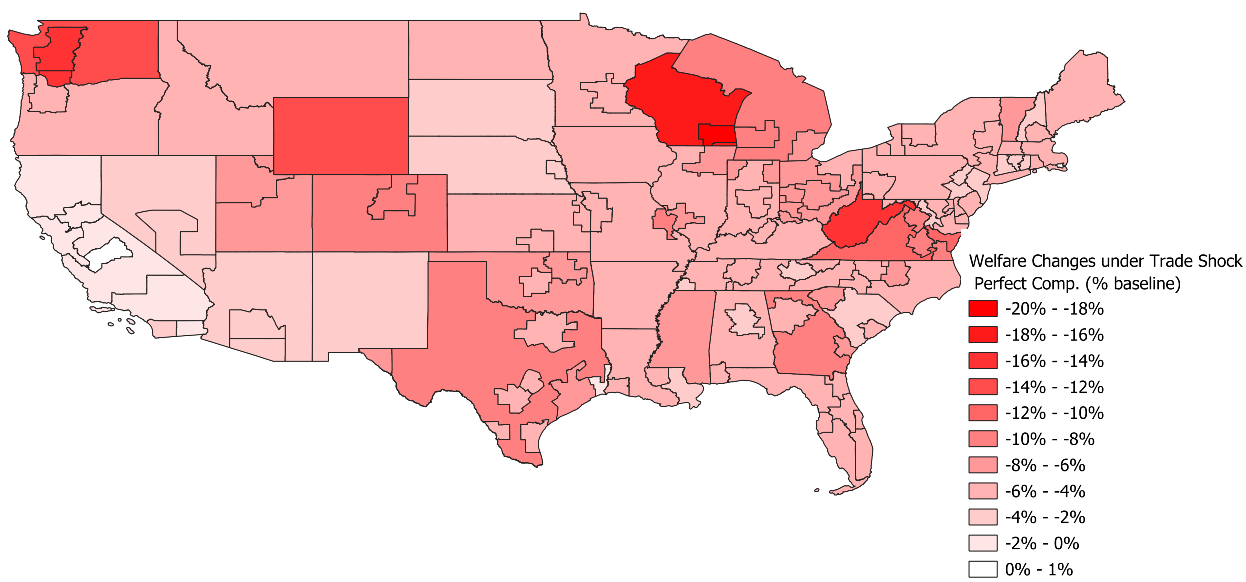

Notes: This figure presents percent changes in real welfare as a result of an international trade shock meant to emulate the spike in shipping prices witnessed during the Covid-19 pandemic. Panel A reports welfare changes under the status-quo, while Panel B presents welfare changes under the perfect freight market competition. Darker colors signify greater changes– red denotes a decline in value, while blue denotes a gain.

Figure 11: Geographic Import Incidence of an International Trade Shock

(a) **Status-Quo Freight Pricing**

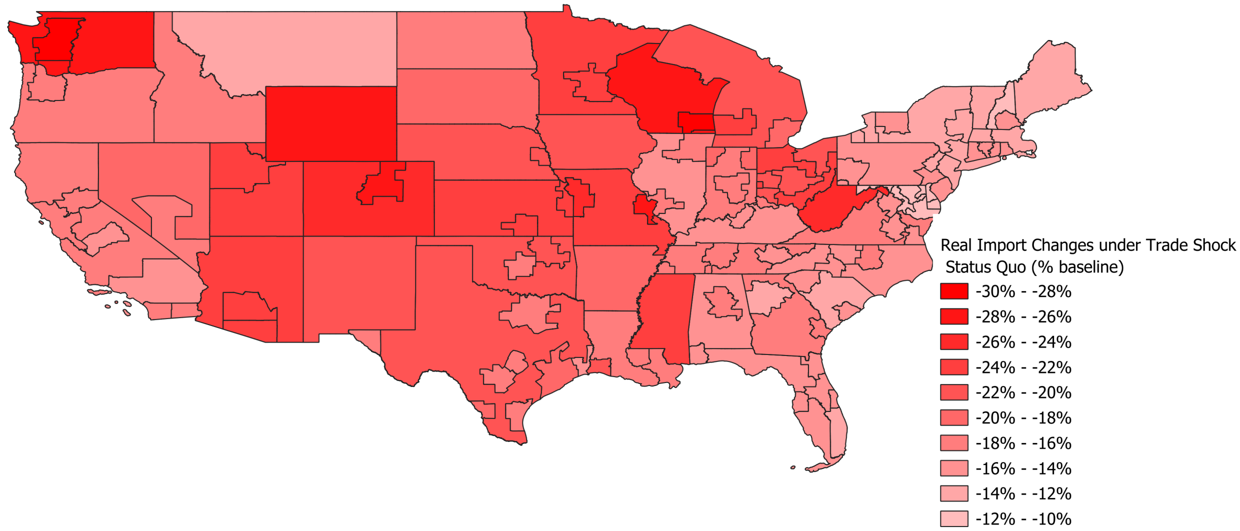

(b) **Perfect Freight Market Competition**

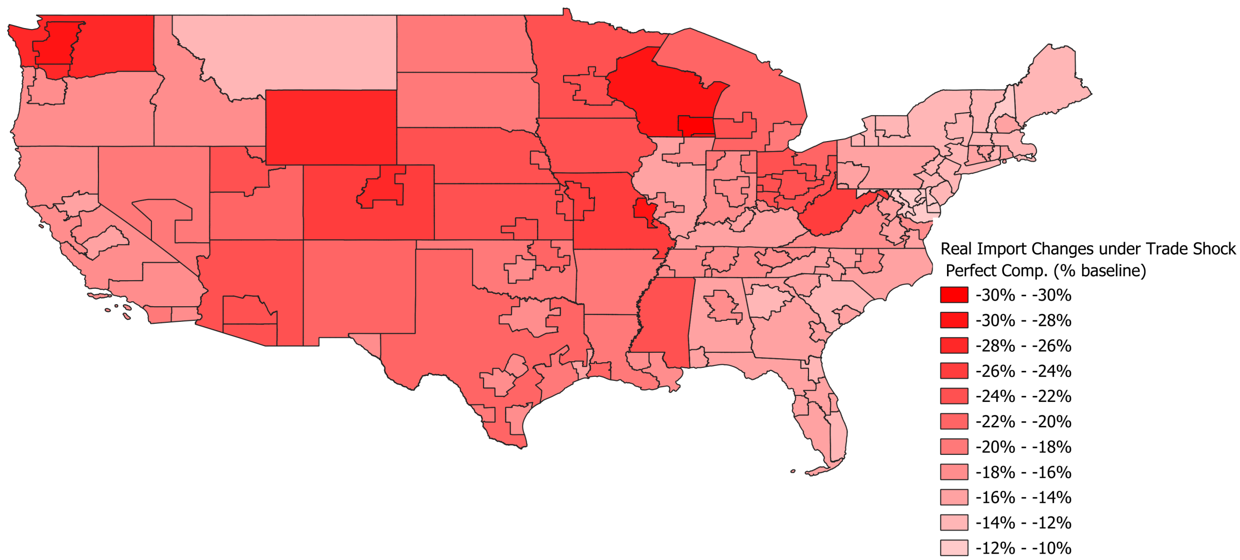

Notes: This figure presents percent changes in real imports as a result of an international trade shock meant to emulate the spike in shipping prices witnessed during the Covid-19 pandemic. Panel A reports import changes under the status-quo, while Panel B presents import changes under the perfect freight market competition. Darker colors signify greater changes– red denotes a decline in value, while blue denotes a gain.

Figure 12: Geographic Export Incidence of an International Trade Shock

(a) **Status-Quo Freight Pricing**

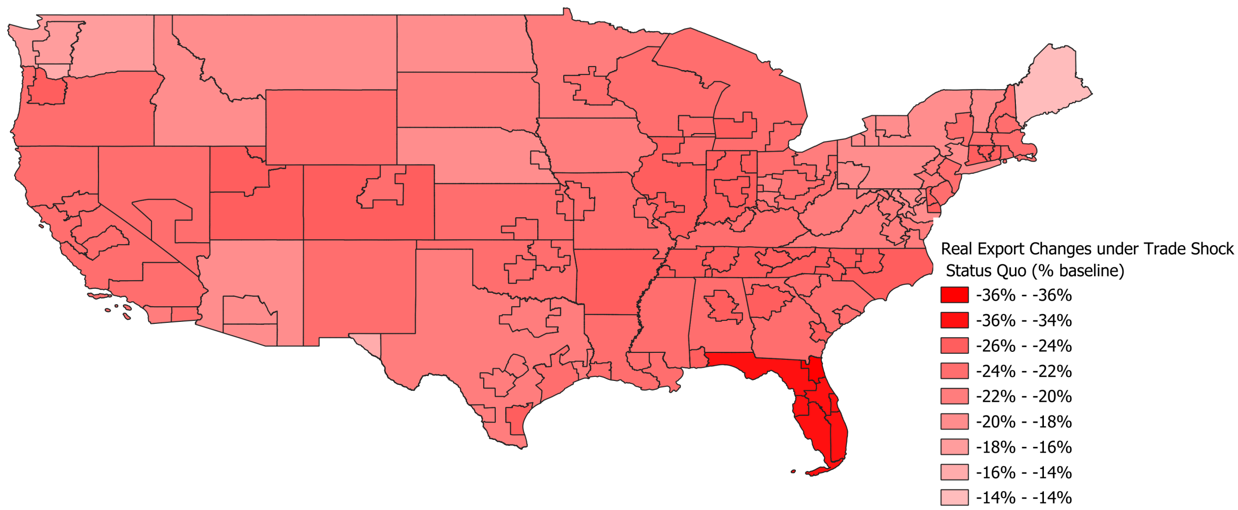

(b) **Perfect Freight Market Competition**

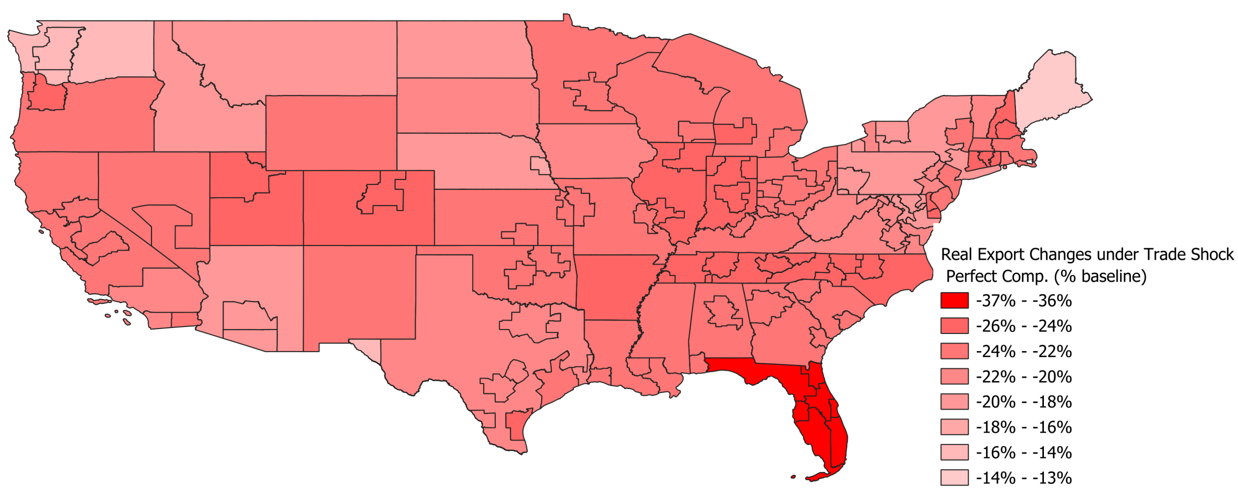

Notes: This figure presents percent changes in real exports as a result of an international trade shock meant to emulate the spike in shipping prices witnessed during the Covid-19 pandemic. Panel A reports export changes under the status-quo, while Panel B presents export changes under the perfect freight market competition. Darker colors signify greater changes– red denotes a decline in value, while blue denotes a gain.

A Mathematical Appendix

A.1 Simplifying Trade Shares

I begin with the trade shares from the nested logit structure:

$$\begin{aligned}\pi_{ir(e,x,m)|j} =& \pi_{r|iexmj}\pi_{m|iexmj}\pi_{ex|ij}\pi_{i|j} \\ =& \exp\Big(- \ln A_i - \theta\big[\ln c_i + (\mu_{irj} + \kappa_r + u_{irj})/(\varphi\rho\zeta) + (\alpha_{iem} + \gamma_{xmj})/(\rho\zeta) + (\eta_{ie} + \nu_{xj})/\zeta\big] \\ & - (1-\varphi)J_{iexmj} - (1-\rho)V_{iexj} - (1-\zeta)I_{ij} - Q_j\Big).\end{aligned}$$

Unfortunately, this formulation is not a tenable regression specification. The inclusive values used to identify the correlation coefficients – J_{iexmj}, V_{iexj}, and I_{ij} – are not directly observed in the data; they could be calculated, though at substantial effort. Fixed effects do not offer a viable solution, as the origin-by-destination-by-mode fixed effect required to capture J_{iexmj} would absorb all of the non-simulated variation present in the data. However, some simple manipulation yields a much more tractable specification:

$$\begin{aligned}\ln \pi_{ir(e,x,m)|j} =& - \ln A_i - \theta\big[\ln c_i + (\mu_{irj} + \kappa_r + u_{irj})/(\varphi\rho\zeta) + (\alpha_{iem} + \gamma_{xmj})/(\rho\zeta) + (\eta_{ie} + \nu_{xj})/\zeta\big] \\ & + \theta\big[\eta_{ie} + \nu_{xj}\big] - \theta\big[\eta_{ie} + \nu_{xj}\big] + \rho\zeta V_{iexj} - \rho\zeta V_{iexj} \\ & - (1-\varphi)J_{iexmj} - (1-\rho)V_{iexj} - (1-\zeta)I_{ij} - Q_j \\ =& - \ln A_i - \theta\big[\ln c_i + (\mu_{irj} + \kappa_r + u_{irj})/(\varphi\rho\zeta) + (\alpha_{iem} + \gamma_{xmj})/(\rho\zeta) + \eta_{ie} + \nu_{xj}\big] \\ & + (1-\zeta)\big[- \theta(\eta_{ie} + \nu_{xj})/\zeta + \rho V_{iexj} - I_{ij}\big] - (1-\varphi)J_{iexmj} - (1-\rho\zeta)V_{iexj} - Q_j \\ =& - \ln A_i - \theta\big[\ln c_i + (\mu_{irj} + \kappa_r + u_{irj})/(\varphi\rho\zeta) + (\alpha_{iem} + \gamma_{xmj})/(\rho\zeta) + \eta_{ie} + \nu_{xj}\big] \\ & + (1-\zeta)\ln \pi_{ex|ij} - (1-\varphi)J_{iexmj} - (1-\rho\zeta)V_{iexj} - Q_j \\ =& - \ln A_i - \theta\big[\ln c_i + (\mu_{irj} + \kappa_r + u_{irj})/(\varphi\rho\zeta) + (\alpha_{iem} + \gamma_{xmj})/(\rho\zeta) + \eta_{ie} + \nu_{xj}\big] \\ & + \theta\big[\alpha_{iem} + \gamma_{xmj}\big] - \theta\big[\alpha_{iem} + \gamma_{xmj}\big] + \varphi\rho\zeta J_{iexmj} - \varphi\rho\zeta J_{iexmj} \\ & + (1-\zeta)\ln \pi_{ex|ij} - (1-\varphi)J_{iexmj} - (1-\rho\zeta)V_{iexj} - Q_j \\ =& - \ln A_i - \theta\big[\ln c_i + (\mu_{irj} + \kappa_r + u_{irj})/(\varphi\rho\zeta) + \alpha_{iem} + \gamma_{xmj} + \eta_{ie} + \nu_{xj}\big] \\ & + (1-\rho\zeta)\big[- \theta(\alpha_{iem} + \gamma_{xmj})/(\rho\zeta) + \varphi J_{iexmj} - V_{iexj}\big] \\ & + (1-\zeta)\ln \pi_{ex|ij} - (1-\varphi\rho\zeta)J_{iexmj} - Q_j\end{aligned}$$

$$
\begin{aligned}
&= -\ln A_i - \theta\big[\ln c_i + (\mu_{irj} + \kappa_r + u_{irj})/(\varphi\rho\zeta) + \alpha_{iem} + \gamma_{xmj} + \eta_{ie} + \nu_{xj}\big] \\
&\qquad + (1-\rho\zeta)\ln \pi_{m|iexj} + (1-\zeta)\ln \pi_{ex|ij} - (1-\varphi\rho\zeta)J_{iexmj} - Q_j \\
&= -\ln A_i - \theta\big[\ln c_i + (\mu_{irj} + \kappa_r + u_{irj})/(\varphi\rho\zeta) + \alpha_{iem} + \gamma_{xmj} + \eta_{ie} + \nu_{xj}\big] \\
&\qquad + \theta\big[\mu_{irj} + \kappa_r + u_{irj}\big] - \theta\big[\mu_{irj} + \kappa_r + u_{irj}\big] \\
&\qquad + (1-\rho\zeta)\ln \pi_{m|iexj} + (1-\zeta)\ln \pi_{ex|ij} - (1-\varphi\rho\zeta)J_{iexmj} - Q_j \\
&= -\ln A_i - \theta\big[\ln c_i + \mu_{irj} + \kappa_r + u_{irj} + \alpha_{iem} + \gamma_{xmj} + \eta_{ie} + \nu_{xj}\big] \\
&\qquad + (1-\varphi\rho\zeta)\big[-\theta(\mu_{irj} + \kappa_r + u_{irj})/(\varphi\rho\zeta) - J_{iexmj}\big] \\
&\qquad + (1-\rho\zeta)\ln \pi_{m|iexj} + (1-\zeta)\ln \pi_{ex|ij} - Q_j \\
&= -\ln A_i - \theta\big[\ln c_i + \mu_{irj} + \kappa_r + u_{irj} + \alpha_{iem} + \gamma_{xmj} + \eta_{ie} + \nu_{xj}\big] \\
&\qquad + (1-\varphi\rho\zeta)\ln \pi_{r|iexmj} + (1-\rho\zeta)\ln \pi_{m|iexj} + (1-\zeta)\ln \pi_{ex|ij} - Q_j.
\end{aligned}
$$

A.2 Price Index

I now provide a detailed derivation of Equation (14). Let $\boldsymbol{p}_j(\omega)^\top = \begin{bmatrix} p_{irj_1}(\omega) & p_{irj_2}(\omega) & \dots \end{bmatrix}$ be the vector of prices for good ω in market j. Define similarly the vectors $\boldsymbol{\mu}_j$, $\boldsymbol{\tau}_j$, $\boldsymbol{\epsilon}_j(\omega)$, and $\boldsymbol{c}_j$. Let $\boldsymbol{z} \in \mathbb{R}^N_{++}$, where N is the number of unique trade routes serving j. Define the function:

$$
\begin{aligned}
G_j(z) &= \Pr\left(\boldsymbol{p}_j(\omega) \le z\right) \\
&= \Pr\left(\frac{\boldsymbol{c}_j \circ \boldsymbol{\tau}_j \circ \boldsymbol{\mu}_j}{\boldsymbol{\epsilon}_j(\omega)} \le \boldsymbol{z}\right) \\
&= \Pr\left(\frac{\boldsymbol{c}_j \circ \boldsymbol{\tau}_j \circ \boldsymbol{\mu}_j}{\boldsymbol{z}} \le \boldsymbol{\epsilon}_j(\omega)\right) \\
&= \Pr\left(\ln \boldsymbol{c}_j + \ln \boldsymbol{\tau}_j + \ln \boldsymbol{\mu}_j - \ln \boldsymbol{z} \le \ln \boldsymbol{\epsilon}_j(\omega)\right) \\
&= 1 - \Pr\left(\ln \boldsymbol{c}_j + \ln \boldsymbol{\tau}_j + \ln \boldsymbol{\mu}_j - \ln \boldsymbol{z} \ge \ln \boldsymbol{\epsilon}_j(\omega)\right) \\
&= 1 - F_j\left(\ln \boldsymbol{c}_j + \ln \boldsymbol{\tau}_j + \ln \boldsymbol{\mu}_j - \ln \boldsymbol{z}\right) \\
&= 1 - \exp\left(-\sum_{i\in\mathcal{S}}\left(\sum_{e,x\in\mathcal{S}^d}\left(\sum_{m\in\mathcal{M}}\left(\sum_{r\in\mathcal{R}^m_{ex}} \exp\left(\frac{-\ln A_i - \theta(\ln c_i + \ln \tau_{irj} + \ln \mu_{irj} - \ln z)}{\varphi\rho\zeta}\right)\right)^{\varphi}\right)^{\rho}\right)^{\zeta}\right) \\
&= 1 - \exp\left(-z^\theta \sum_{i\in\mathcal{S}}\left(\sum_{e,x\in\mathcal{S}^d}\left(\sum_{m\in\mathcal{M}}\left(\sum_{r\in\mathcal{R}^m_{ex}} \exp\left(\frac{-\ln A_i - \theta(\ln c_i + \ln \tau_{irj} + \ln \mu_{irj})}{\varphi\rho\zeta}\right)\right)^{\varphi}\right)^{\rho}\right)^{\zeta}\right) \\
&= 1 - \exp\left(-z^\theta Q_j\right).
\end{aligned}
$$

From the CES utility assumption, it follows,

$$\begin{aligned}
p_j =& \left(\int_\Omega p_j(\omega)^{1-\sigma} d\omega\right)^{\frac{1}{1-\sigma}} \\
p_j^{1-\sigma} =& \int_0^1 z^{1-\sigma} dG_j(z) \\
p_j^{1-\sigma} =& \int_{\mathbb{R}_{++}} z^{\theta-\sigma} \exp\left(-z^\theta Q_j\right) \theta Q_j dz \\
p_j^{1-\sigma} =& \Gamma\left(\frac{\theta+1-\sigma}{\theta}\right) Q_j^{\frac{\sigma-1}{\theta}} \\
p_j =& \Gamma\left(\frac{\theta+1-\sigma}{\theta}\right)^{\frac{1}{1-\sigma}} Q_j^{-\frac{1}{\theta}}.
\end{aligned}$$

A.3 Derivatives of trade shares

I begin with the derivative of the trade share along a route r with respect to its own markup:

$$\frac{\partial \pi_{ir|j}}{\partial \mu_{irj}} = -\left(\frac{\theta}{\varphi\rho\zeta}\right) \pi_{ir|j} \left(1 - (1-\varphi)\pi_{r|iexmj} - \varphi(1-\rho)\pi_{rm|iexj} - \varphi\rho(1-\zeta)\pi_{rexm|ij} - \varphi\rho\zeta\pi_{ir|j}\right).$$

Now, I consider the derivatives of trade shares with respect to markups along competing routes. Let $m(r), x(r)$ denote the mode, port of entry/exit for a particular route r. If $r \neq r'$ but $m(r) = m(r')$ and $x(r) = x(r')$, then

$$\frac{\partial \pi_{ir|j}}{\partial \mu_{ir'j}} = \left(\frac{\theta}{\varphi\rho\zeta}\right) \pi_{ir|j} \left((1-\varphi)\pi_{r'|iexmj} + \varphi(1-\rho)\pi_{r'm|iexj} + \varphi\rho(1-\zeta)\pi_{r'exm|ij} + \varphi\rho\zeta\pi_{ir'|j}\right).$$

If $r \neq r'$ and $m(r) \neq m(r')$ but $x(r) = x(r')$, then

$$\frac{\partial \pi_{ir|j}}{\partial \mu_{ir'j}} = \left(\frac{\theta}{\varphi\rho\zeta}\right) \pi_{ir|j} \left(\varphi(1-\rho)\pi_{r'm|iexj} + \varphi\rho(1-\zeta)\pi_{r'exm|ij} + \varphi\rho\zeta\pi_{ir'|j}\right).$$

If $r \neq r'$ and $m(r) \neq m(r')$ and $x(r) \neq x(r')$, then

$$\frac{\partial \pi_{ir|j}}{\partial \mu_{ir'j}} = \left(\frac{\theta}{\varphi\rho\zeta}\right) \pi_{ir|j} \left(\varphi\rho(1-\zeta)\pi_{r'exm|ij} + \varphi\rho\zeta\pi_{ir'|j}\right).$$

If $i \neq i'$, then

$$\frac{\partial \pi_{ir|j}}{\partial \mu_{i'r'j}} = \theta \pi_{ir|j} \pi_{i'r'|j}.$$

Finally, it should be noted that I assume all destinations enjoy completely independent freight markets, such that, for all $j \neq j'$

$$\frac{\partial \pi_{ir|j}}{\partial \mu_{i'r'|j'}} = 0.$$

A.4 Proof of Convergence of the EM Algorithm

It will prove convenient to define the following matrices:

$$\phi^\top = \begin{bmatrix} \lambda_i & \beta_m & 1-\zeta & 1-\rho\zeta & 1-\varphi\rho\zeta \end{bmatrix}$$
$$X_{irj} = \begin{bmatrix} \delta_{irj} & Miles_{irj} \times \mathbb{1}(m) & \ln \pi_{ex|ij} & \ln \pi_{m|iexj} & \ln \pi_{r|iexmj} \end{bmatrix}$$

where $\mathbb{1}$ denotes the indicator function. With these matrices, I re-define my primary estimating equation:

$$\ln \boldsymbol{\pi} = \boldsymbol{X}\phi + \boldsymbol{u}. \tag{22}$$

where $\boldsymbol{\pi}^\top = \begin{bmatrix} \pi_{ir_1|j} & \pi_{ir_2|j} & \dots \end{bmatrix}$, $\boldsymbol{X}^\top = \begin{bmatrix} X_{ir_1j} & X_{ir_2j} & \dots \end{bmatrix}$, and $\boldsymbol{u}^\top = \begin{bmatrix} u_{ir_1j} & u_{ir_2j} & \dots \end{bmatrix}$ denote the matrix representation of each variable. Obviously, Equation (22) presents the complete-information estimating equation. Unfortunately, elements of both $\boldsymbol{X}$ and $\ln \boldsymbol{\pi}$ are unobserved.[47] I therefore rely on the incomplete-information formulation:

$$\ln \boldsymbol{\pi}^{(p)} = \boldsymbol{X}^{(p)}\phi + \boldsymbol{u}, \tag{23}$$

where $\ln \boldsymbol{\pi}^{(p)} = \mathbb{E}\big(\ln \boldsymbol{\pi}|\phi^{(p)}\big)$ and $\boldsymbol{X}^{(p)} = \mathbb{E}\big(\boldsymbol{X}|\phi^{(p)}\big)$ denote the expected values of $\ln \boldsymbol{\pi}$ and $\boldsymbol{X}$ for some, fixed value of $\phi^{(p)}$. Estimating Equation (23) via OLS[48] yields the next entry in the EM sequence, $\phi^{(p+1)}$; denote the residuals from this estimation $\boldsymbol{u}^{(p+1)} = \ln \boldsymbol{\pi}^{(p)} - \boldsymbol{X}^{(p)}\phi^{(p+1)}$. From

[47] Specifically, δ_{irj} and $\pi_{r|iexmj}$ are unobserved in my data.

[48] Note that OLS estimation satisfies the criteria of the EM process, as OLS estimates are also maximum-likelihood estimates.

here, define the likelihood function:

$$Q\left(\phi^{(p+1)}|\phi^{(p)}\right) = \mathbb{E}\left[\ln f\left(\boldsymbol{X}|\phi^{(p+1)}\right)\Big|\phi^{(p)}\right]$$

where f is the probability of $\boldsymbol{X}$ given ϕ. In words, $Q\left(\phi^{(p+1)}|\phi^{(p)}\right)$ defines the log-likelihood of observing $\phi^{(p+1)}$ given $\phi^{(p)}$.

Per Wu (1983), a sufficient condition for convergence of the EM process is the existence of a forcing function $c > 0$ such that

$$Q\left(\phi^{(p+1)}|\phi^{(p)}\right) - Q\left(\phi^{(p)}|\phi^{(p)}\right) \geq c\left\|\phi^{(p+1)} - \phi^{(p)}\right\| \ \forall p. \tag{24}$$

Note that this inequality holds trivially when $\phi^{(p+1)} = \phi^{(p)}$. I now focus on the case $\phi^{(p+1)} \neq \phi^{(p)}$. Given the assumption that $u_{irj} \sim N(0, v)$, it follows that

$$Q\left(\phi^{(p+1)}|\phi^{(p)}\right) = -(n/2)\ln(2\pi) - (n/2)\ln v - \big(1/(2v)\big)\boldsymbol{u}^{(p+1)\top}\boldsymbol{u}^{(p+1)}, \tag{25}$$

where n is the number of rows in $\boldsymbol{X}$. Note also that, thanks to the OLS estimation,

$$\phi^{(p+1)} = \left(\boldsymbol{X}^{(p)\top}\boldsymbol{X}^{(p)}\right)^{-1}\left(\boldsymbol{X}^{(p)\top}\ln\boldsymbol{\pi}^{(p)}\right). \tag{26}$$

Combining equations (24) - (26) yields a characterization for c that satisfies Wu's criterion:

$$\frac{\left(\boldsymbol{u}^{(p+1)\top}\boldsymbol{u}^{(p+1)} - \boldsymbol{u}^{(p)\top}\boldsymbol{u}^{(p)}\right)}{2v\left\|\left(\boldsymbol{X}^{(p)\top}\boldsymbol{X}^{(p)}\right)^{-1}\left(\boldsymbol{X}^{(p)\top}\ln\boldsymbol{\pi}^{(p)}\right) - \phi^{(p)}\right\|} \geq c > 0. \tag{27}$$

The existence of c therefore hinges on three conditions: i) v must be non-zero and finite; ii) $\boldsymbol{X}^{(p)\top}\boldsymbol{X}^{(p)}$ must be invertible (in other words, $\boldsymbol{X}$ must have non-zero variance); and iii) $\boldsymbol{X}^{(p)}$ and $\ln\boldsymbol{\pi}^{(p)}$ must have non-zero, finite covariance. That is, the same conditions required for OLS estimation guarantee convergence of the EM process.