Network Structure and the Efficiency Gains from Mergers: Evidence from U.S. Freight Railroads *

Yanyou Chen
University of Toronto

This Version: *May 31, 2025*

Abstract

The trade-off between market power and efficiency gains is central to antitrust analyses of mergers, but empirical evidence quantifying efficiencies remains limited. Using transaction-level data from U.S. freight railroads (1985–2005), this article quantifies merger-induced cost efficiencies, driven mainly by eliminating inter-railroad interchange costs and reoptimization of routing and resource allocation within integrated networks. I develop a spatial equilibrium model with oligopolistic competition to assess equilibrium merger effects. Counterfactual analyses show mergers reduce shipment costs by 12.9% and prices by 8.8%. Markups increase by 7.2%, driven primarily by non-merging firms reallocating resources away from regions where merged firms achieve large cost reductions.

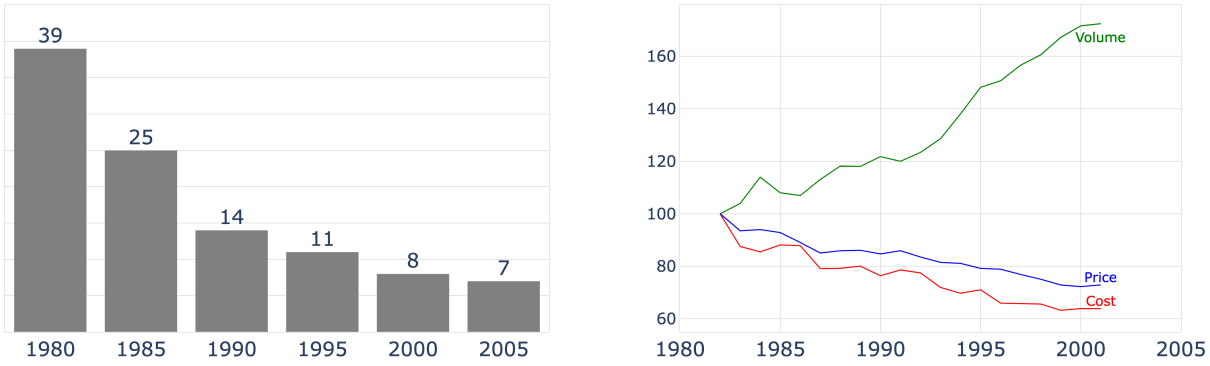
Keywords: Cost Efficiency, Oligopolistic Competition, Mergers, Market Power, Transportation

*I am deeply indebted to Allan Collard-Wexler, Daniel Yi Xu, James W. Roberts, Christopher D. Timmins, Bryan Bollinger, and Joy Tianjiao Tong for their valuable comments, continuous support, and encouragement throughout this research project. I am also grateful to the anonymous referees, whose insightful feedback significantly improved the quality of this work. I thank Victor Aguirregabiria, Heski Bar-Isaac, Yao Luo, Robert McMillan, Eduardo Souza-Rodrigues, and seminar participants at Minnesota, Wharton (Business Economics and Public Policy), Indiana Kelley, Tuck School of Business, Queen's University, Western University, Monash University, Peking University, University of Toronto, Stanford University, Microsoft Research, IO Canada, and EARIE for helpful discussions and constructive suggestions. I would also like to express my gratitude to participants in the Railway Management Certificate Program for generously sharing their valuable insights about the freight railroad industry. Funding support from the Social Sciences and Humanities Research Council Insight Development Grant is gratefully acknowledged. Any remaining errors are solely my responsibility. Disclaimer: The opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the U.S. Surface Transportation Board. All results have been carefully reviewed to ensure no confidential information is disclosed.

1 Introduction

The trade-off between market power and efficiency gains is central to antitrust analysis of mergers. While economists have recognized this tension at least since Williamson (1968), empirical evidence quantifying merger-induced efficiencies remains limited.[1] This gap is particularly pronounced in network-based industries, where mergers can improve operational efficiency by eliminating inter-firm coordination costs and optimizing resource allocation across integrated networks. My article addresses this gap by providing empirical evidence of efficiency gains resulting from mergers in the U.S. freight railroad industry.

The U.S. freight railroad sector provides a uniquely suitable setting for examining merger-induced efficiency gains due to its extensive historical consolidation and clear suggestive evidence of efficiency improvements. Between 1985 and 2005, substantial merger activity reduced the number of Class I railroads from 39 to 7, increasing the market share of the four largest firms from 66% to 94% (Panel (a) of Figure 1). Over this period, real shipment prices declined by approximately 20%, and shipment volumes doubled (Panel (b) of Figure 1). Given the absence of significant technological advancements during this timeframe, these observed price declines indicate efficiency gains attributable to mergers, making railroads an exemplary industry for this analysis.



(a) Number of Class I Railroads

(b) Shipment Price, Cost, and Volume

Figure 1: U.S. Freight Railroads

Notes: The data source used to construct the figures is the Class I Railroad Annual Report. Class I railroads are defined as those "having annual carrier operating revenues of $250 million or more in 1991 dollars." According to the Association of American Railroads, Class I railroads accounted for more than 95% of U.S. freight railroad industry revenues. Volume in panel (b) is measured in gross ton-miles carried, price is defined as freight revenue per ton-mile, and average cost is operating cost per ton-mile.

[1] As highlighted by Asker and Nocke (2021) in the *Handbook of Industrial Organization*, "Surprisingly little progress has been made in providing empirical characterizations of merger-specific efficiencies." Notable exceptions include Ashenfelter et al. (2015) for the U.S. beer industry and Clark and Samano (2022) for the electricity distribution industry.

Using comprehensive transaction-level data—including detailed shipment records, firm-level operational expenditures, and geographic information—I first provide reduced-form evidence that railroad mergers significantly reduce shipment prices. Difference-in-differences analysis reveals an average post-merger price decline of about 9%, with notably larger reductions (around 12%) occurring specifically on routes previously involving inter-railroad interchanges. These larger reductions suggest that eliminating interchange inefficiencies plays an important role in the observed merger gains. Additional analyses indicate that the efficiency gains predominantly stem from these scope-related operational improvements rather than increased shipment density alone.

To rigorously quantify these efficiency gains and explore their underlying mechanisms, I introduce a novel spatial equilibrium model of oligopolistic competition specifically tailored to network-based industries. This model captures two important dimensions of competitive interactions not extensively explored in previous literature: first, the interdependence among markets within each firm's transportation network; and second, strategic interactions among oligopolistic firms, as each firm simultaneously considers rivals' resource allocation, pricing, and routing decisions.

I utilize four data sources in my analysis. First, I use the Carload Waybill Sample, a transaction-level dataset providing information on the shipment price, the types of commodities being transported, the total invoiced weight, equipment used, the railroads involved, and the specifics of origin, destination, and interchange points for every load.[2] Second, the Class I Railroad Annual Report provides yearly data on firm characteristics and aggregate operational statistics such as annual maintenance spending. Third, the Commodity Flow Survey contains information about shipment volumes across various transport modes for distinct origin–destination markets. Fourth, I obtain geographic data from the Department of Transportation regarding all U.S. rail lines and their affiliated railroad firms. This enables me to retrospectively map out the rail network of each railroad firm for any specified year. I have acquired all four datasets from 1985 to 2005.

The reduced-form analyses indicate that efficiency gains primarily arise from the elimination of inter-railroad interchanges and subsequent network-wide route optimization, referred to in this article as scope-related efficiencies. Specifically, efficiency improvements are largest on routes transitioning from multi-carrier to single-carrier operations after mergers. In contrast, single-carrier routes without prior interchanges exhibit smaller efficiency improvements, and these improvements do not significantly differ based on the magnitude of

[2]In railroading, an "interchange" refers to the location or facility where two or more different railroad companies meet to exchange cars and freight. At these points, one railroad's cars can be transferred to another railroad's tracks for continuing their journey. I present a real example of an interchange in Section 2.1.

shipment volume increases. These findings indicate that increased shipment density alone is not the primary driver of efficiency gains observed during the railroad merger period from 1985 to 2005. These insights directly motivate the structural model, designed to quantify the equilibrium impacts of these operational adjustments.

I then develop an empirical model that captures railroad firms' annual strategic decisions, including pricing, routing, and resource allocation. In this framework, each railroad firm's network is represented as a graph, where nodes correspond to market locations and arcs represent physical track segments individually owned by the firm. The model explicitly incorporates two essential dimensions: first, the interdependence of markets within a firm's network—allowing resource allocations on specific track segments (arcs) to benefit all traffic passing through them—and second, strategic interactions among separate firms' networks. To reflect realistic industry practices, decision-making within each firm is decentralized: marketing teams set prices for local markets, while operational teams optimize routing and allocate maintenance resources across the firm's entire network. Customers' discrete choices over competing railroads or alternative transportation modes (e.g., trucking) determine market demand. In equilibrium, railroad firms simultaneously choose optimal pricing, routing, and resource allocation, taking into account competitors' strategies and market interdependencies, thus enabling precise quantification of merger-induced efficiencies.

I estimate demand for each origin–destination market using detailed waybill data. Regarding firms' decisions, I estimate three key cost parameters. First, the *impact of distance* parameter determines how shipment costs vary with travel distance. Second, the *interchange cost* parameter captures the additional costs incurred at inter-railroad interchanges. Third, the *resource allocation efficiency* parameter reflects how effectively maintenance and operational resources allocated to track segments (arcs) reduce shipment costs. In my model, a positive *resource allocation efficiency* parameter indicates that arcs receiving greater resource allocations have lower marginal shipment costs, capturing economies of scope from consolidating multiple routes onto common segments. Conversely, a zero value of this parameter would suggest no economies of scope, implying equal importance of all routes with no gains from consolidation.

These cost parameters are subsequently estimated using simulated method of moments. For the *impact of distance* parameter, I target the average shipping cost per mile in dollar terms. To estimate the *interchange cost* parameter, I use the average price difference between interchange and non-interchange routes. Lastly, in the ideal scenario for estimating the *resource allocation efficiency* parameter, I would observe price differences resulting from varying levels of resource allocation on different arcs. However, the actual resource allocation amount for each arc is unobservable. To address this, I leverage the positive

correlation between node connectedness and resource allocation. If a station is physically located at a more central position within the network, it will likely receive greater resource allocation. Therefore, in my analysis, I estimate the *resource allocation efficiency* parameter by analyzing the price differences across markets that result from varying degrees of node connectedness at their origins and destinations.

The shipment cost per carload-mile is estimated at \$1.20, and the cost per interchange is estimated at \$217, representing approximately 21% of the average shipment price. This finding underscores the significant expense associated with interchanges and highlights the potential for substantial cost efficiency gains by eliminating interchange costs following a merger. Finally, the estimated value of the *resource allocation efficiency* parameter is significantly different from 0, indicating the presence of economies of scope. Moreover, the estimated *resource allocation efficiency* parameter is less than 1, suggesting diminishing marginal returns in resource allocation for a specific arc.

Next, using the estimated parameters, I simulate the equilibrium outcomes before and after each merger among Class I railroads from 1985 to 2005. The results show that, on average, shipment cost falls by 12.9%, shipment price falls by 8.8%, and the markup increases by 7.2%. My counterfactual results closely align with the reduced-form evidence presented. Regarding the increased markup post-merger, my counterfactual analysis shows that reducing the number of firms in local markets is not the primary driving factor. Instead, the increased markup is primarily driven by non-merging firms reallocating resources away from regions where the merged firm achieves a significant cost reduction. As a consequence, this leads to additional growth in the merged firm's market share at the local market, resulting in greater markup increases in those regions. These results highlight the importance of considering equilibrium effects in merger analysis.

I investigate the key network characteristics that give rise to larger changes in shipping costs following mergers. I use two centrality measures to describe the nodes in a network: *degree centrality*, which captures the total number of links connected to the node within the network, and *betweenness centrality*, which evaluates the number of shortest paths that pass through the node. In both centrality measures, my findings show that markets with low centrality in the pre-merger network and high centrality in the post-merger network typically experience a greater reduction in shipping costs following a merger. Intuitively, when a node is positioned at the outer edges of the network before the merger and subsequently becomes more central in the post-merger network, it experiences a greater improvement in efficiency. Furthermore, the results show that an increase in betweenness centrality leads to a larger reduction in shipment costs and a greater increase in markups after the merger when compared to a similar increase in degree centrality.

Lastly, to shed light on future merger policies, I investigate which types of mergers are more likely to improve welfare. I focus on the degree of overlap and complementarity between the two networks involved in the merger. The degree of overlap is measured by the overall percentage of markets that both merging firms were present in prior to the merger. As for complementarity, it is calculated as the overall percentage of interchanges carried out between the merging firms. The findings indicate that a high degree of overlap leads to significant cost reductions and substantial increases in markups. A high degree of complementarity also leads to significant cost reductions but only a modest increase in markups.

Related Literature

This article contributes to four broad strands of literature: (i) merger analysis, particularly regarding merger-induced efficiency gains; (ii) transportation networks; (iii) the railroad industry; and (iv) network competition.

First, this article contributes to the literature on merger-related efficiency gains. A merger increases the combined firm's market power, enabling potential price increases, but it may also create operational efficiencies that reduce marginal costs, thereby allowing price reductions. Although extensive research has analyzed the price effects of mergers across industries—including airlines (e.g., Borenstein, 1990; Kim and Singal, 1993; Peters, 2006), ready-mix concrete (Collard-Wexler, 2014), and hospitals (Dafny, 2009; Dafny et al., 2019)—empirical studies explicitly quantifying merger-induced efficiencies remain scarce. Notable exceptions include Ashenfelter et al. (2015), Jeziorski (2014), and Clark and Samano (2022), along with ongoing analyses in the power plant industry by Demirer and Karaduman (2024). This article contributes to this limited body of evidence by directly quantifying efficiency gains from mergers in a network-based industry and demonstrating the critical role of equilibrium effects in shaping these merger outcomes.

Second, this article contributes to the rapid growing literature on transportation. Fajgelbaum and Schaal (2020) study the optimal infrastructure investment decisions of a social planner. Meanwhile, Buchholz (2022) investigates search frictions and regulations in taxi markets, and Brancaccio et al. (2020, 2023) analyze world oceanic shipping, with an emphasis on the influence of geography on shipment costs and on correcting inefficiencies stemming from the industry's decentralized nature. In general, the existing literature on transport markets typically takes one of two approaches, either looking at the decisions of a social planner or assuming the presence of numerous atomistic players, such as individual taxi drivers or container ships. While these approaches are effective for studying their respective sectors, in industries with a high degree of market concentration, such as railroads, it is necessary

to further consider oligopolistic competition. Building on the groundwork established by these studies, my article incorporates oligopolistic competition among transport firms. It allows each firm to make optimized decisions across the network while taking into account interactions with its rivals. This expands the scope of the literature in a relevant way and facilitates the analysis of mergers.

Third, my article contributes to the literature on the freight railroad industry—see Grimm and Winston (2000) and Gallamore and Meyer (2014) for comprehensive reviews. Existing studies typically focus on aggregate cost or price indices or evaluate merger effects by examining individual markets independently.[3] However, virtually no prior research has employed a structural framework that explicitly incorporates network-wide interdependencies and strategic firm interactions to quantify equilibrium merger effects. This article fills this gap by structurally modeling these aspects, explicitly accounting for market power and equilibrium effects—an essential consideration for appropriately analyzing the Williamson trade-off. This equilibrium-based perspective represents a key distinction relative to complementary work by Degiovanni and Yang (2023), who extend the framework introduced here to explore different aspects of railroad efficiency gains, but place comparatively less emphasis on equilibrium market power considerations. Furthermore, as demonstrated through the reduced-form analysis presented in this article, scope-related operational efficiencies constitute the primary source of merger-induced gains observed between 1985 and 2000. In contrast, the absence of additional mega-mergers between 2000 and 2010 likely shifted the industry's subsequent focus toward optimizing single-line routes through increased traffic density, which forms the central emphasis of the analysis in Degiovanni and Yang (2023) for periods from 2010 onward.

Finally, this article relates to the literature on network competition. For example, Ciliberto, Cook, and Williams (2019) analyze how consolidation affects airline network connectivity using various centrality measures, while Aguirregabiria and Ho (2012) examine factors influencing the adoption of hub-and-spoke networks in the U.S. airline industry. Unlike airline networks, where routes typically involve at most one intermediate stop, railroad shipments often traverse multiple nodes within a single firm's network. My approach incorporates this routing complexity, making resource allocation and routing decisions directly dependent on geographic location and the underlying network structure.

Regarding monopolistic distribution networks, Holmes (2011) studies Wal-Mart's choice

[3] Friebel et al. (2014) document network reductions and labor downsizing following deregulation and mergers. Prater et al. (2012) examine the sufficiency of intramodal rail freight competition and its impact on rates. Event studies such as Kwoka and White (1997) on the Union Pacific–Southern Pacific merger and Pittman (2009) on the Santa Fe–Southern Pacific merger proposal provide case-specific insights into merger efficiencies.

of locations and infers the magnitude of density economies, while Houde, Newberry, and Seim (2023) study the investment decisions leading to the expansion of Amazon's distribution network by considering the benefits of economies of density and vertical integration into package sortation. I contribute to this discussion by accounting for interactions among networks owned by different firms, specifically focusing on how these interactions affect overall welfare changes following mergers in equilibrium.

In sum, this article provides empirical evidence on the efficiency gains arising from mergers in the railroad industry, identifying and quantifying the key mechanisms driving these gains. Methodologically, it introduces a spatial equilibrium framework that explicitly incorporates market interdependencies and strategic interactions among firms, important for analyzing mergers within oligopolistic network-based industries. This framework can be broadly applied beyond railroads, enabling analyses of efficiency and market-power trade-offs in other network contexts.

The remainder of the article is organized as follows. Section 2 describes the industry background and the main datasets used in the article, and Section 3 provides suggestive evidence of merger gains after railroad mergers. Section 4 constructs the structural model of firm pricing, routing, and network optimization decisions in a rail network. Section 5 presents the estimation results and assesses the validity of the model, while Section 6 presents the counterfactual experiments and results. Section 7 concludes.

2 Industry Background and Data

2.1 A Running Example: Train 9-698-21

The story of Train 9-698-21 serves as a running example to illustrate fundamental concepts used in this paper. Figure 2 depicts three railroad companies: the Santa Fe railway, the Burlington Northern railway, and the Union Pacific railway. This simplified network comprises four nodes—Los Angeles, Avard Oklahoma, Claremore Oklahoma, and Memphis.

The Santa Fe railway operates Train 9-698, which runs from Los Angeles to Memphis.[4] The train commences its journey from Los Angeles promptly, proceeds without any delay and reaches Avard Oklahoma. Upon arrival at Avard Oklahoma, the train must be interchanged between Santa Fe railway and Burlington Northern railway. Initially, the train encountered a five-hour wait to exchange railcars with a tardy train from Richmond. Subsequently, the crew discovered that Burlington Northern had not dispatched a locomotive to retrieve the train. Upon contacting Burlington Northern, the responsible party responded by stating that

[4]Original story adapted from *Trains* magazine "Twenty-four hours at Supai Summit".

"you are not my first priority, so you'll have to wait." After additional delay, a locomotive eventually arrived and retrieved the train, which was subsequently transported by Burlington Northern until it finally reached its destination in Memphis.



Figure 2: The Story of Train 9-698-21

The story of Train 9-698 underscores two potential sources of efficiency gains achievable through mergers. First, mergers eliminate interchange costs, reducing delays and coordination problems. Second, the elimination of interchanges enables merged firms to re-optimize routing decisions and consolidate traffic through previously joint-operated routes, thereby enhancing overall network efficiency.

Figure 2 further illustrates two distinct service structures in railroad operations that will be central to subsequent analyses. A *single-line service* refers to a shipment fully managed by a single railroad company from origin to destination. For instance, the Union Pacific Railroad exclusively handles shipments from Claremore, Oklahoma, to Memphis. In contrast, a *joint-line service* involves multiple railroad companies, each responsible for a different segment of the route. An example is the shipment from Los Angeles to Memphis, which is jointly managed by the Santa Fe Railroad (Los Angeles to Avard) and Burlington Northern Railroad (Avard to Memphis).

The network-based nature of railroad operations also leads to two broader implications. First, since railroad firms manage interconnected networks, re-optimizing routing decisions in one market necessarily impacts operational decisions across the entire network. As a result, evaluating individual origin–destination markets in isolation will fail to capture these crucial interdependencies. Second, equilibrium responses from non-merging competitors must be considered. Competing railroads strategically adjust their pricing and operational decisions in response to merger-induced network changes. My structural modeling framework incorporates these network-wide interdependencies and competitive equilibrium responses, aiming to provide a comprehensive framework for evaluating railroad mergers.

Figure 3 illustrates the complex network owned by seven major Class I railroads, each

independently managing extensive physical track networks.[5]

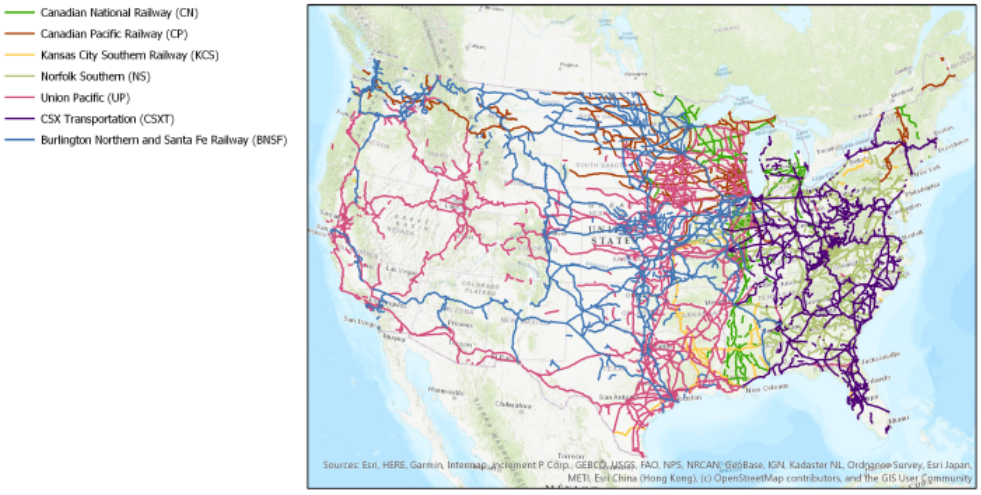


Figure 3: U.S. Class I Railroads

Additional historical context on railroad mergers is provided in Appendix A, and a detailed comparison of freight transportation by railroad and other transportation modes is presented in Appendix C.

2.2 Data

I utilize four datasets in my analysis. First, the Carload Waybill Sample, which contains detailed shipment information, including commodities carried, total billed weight, equipment used, participating railroads, and origin, destination, and interchange locations for each load. The Carload Waybill Sample accounts for approximately 2% of total freight railroad shipment transactions and includes detailed price information in its confidential version. It is designed to represent all U.S. rail traffic and is a stratified sample of waybills for carload shipments that move on U.S. railroads, collected annually by the Surface Transportation Board (STB). The stratification ensures that the sample adequately represents the diversity of rail traffic, including different types of commodities, shipment sizes, and routes.[6] I have roughly 12 million total observations for the past twenty years. Second, the Class I Railroad Annual Report R-1 dataset. This dataset contains yearly data on firm characteristics and aggregate

[5]BNSF and Union Pacific primarily compete in the western region, CSXT and Norfolk Southern compete in the east, Canadian railroads CP and CN connect freight between Canada and the U.S., and Kansas City Southern (KCS) facilitates trade between Mexico and the U.S.

[6]Waybill data have broad applications in, among other things, rate cases, exemption decisions, and analyses supporting regulations. Waybill data are also used by transportation practitioners, consultants, and law firms in preparing verified statements to be submitted in formal proceedings before the Board or other public agencies. For more details, see STB's official descriptions here.

operational statistics such as annual maintenance spending. Third, the Commodity Flow Survey. This dataset contains information on shipment volumes for different transportation modes, which I account for as share of outside options in my demand model.[7] Fourth, geographic information from the Department of Transportation. I utilize this dataset to obtain information on all U.S. rail lines and their associated railroad companies, allowing me to trace back through time and reconstruct the rail network of each railroad firm between 1985 and 2005. All four datasets cover the time span from 1985 to 2005.

Table 1: Summary Statistics of Variables

	Mean	Std. Dev.	25th Percentile	Median	75th Percentile
Price per Railcar ($)	1,034	1,399	384	703	1,266
Shipment Weight (Tons per Railcar)	54	46	16	26	102
Travel Distance (Miles)	1,045	773	404	854	1,647
Number of Waybills (Carrier-Origin-Destination-Date)				12,113,581	

Source: The Surface Transportation Board, Carload Waybill Sample

First, I provide summary statistics for the key variables from the waybill data. Table 1 shows the results. The shipment price per carload varies from $384 to $1,266 at the 25th and 75th percentiles, with a mean price of $1,034. The shipment weight per carload ranges from 16 tons to 102 tons at the 25th and 75th percentile, with an average weight of 54 tons. The average travel distance per shipment is 1,045 miles, with the 25th percentile at 404 miles and the 75th percentile at 1,647 miles. This indicates that railroad shipments are mainly long-distance. The median shipment price per ton-mile is 2.65 cents, a value that is similar to the industry-reported price per ton-mile.[8]

Next, I provide insight into the competitive landscape of the industry. In accordance with the Surface Transportation Board, a market is defined as an origin–destination pair where the origins and destinations correspond to the Bureau of Economic Analysis (BEA) economic areas. The contiguous United States comprises a total of 170 BEA areas.

Table 2 shows the number of markets and the percentage of interchange lines over the years. The number of Class I railroads decreased significantly from 39 to 7 due to the wave of mergers that occurred between 1985 and 2005. However, the first column of Table 2 shows that the number of active origin–destination markets remained relatively stable during the

[7]The U.S. Census Bureau conducts the Commodity Flow Survey every five years. I have access to data from the years 1993, 1997, 2002, 2007, and 2012.

[8]According to the AAR, the mean price per ton-mile in the United States was 2.32 cents in 2001.

Table 2: Summary Statistics of Market Competition

Year	Number of Markets	Percentage of Interchange Lines	Number of Competitors in an o–d Market			Number of Waybills
			mean	25th percentile	75th percentile	
1985	12,088	41%	3	1	3	262,703
1990	11,835	35%	2	1	3	323,570
1995	11,632	26%	2	1	3	453,802
2000	11,732	14%	2	1	2	544,738
2005	11,611	11%	2	1	2	611,033

Source: The Surface Transportation Board, Carload Waybill Sample

examined period, decreasing only slightly from 11,835 to 11,611 from 1990 to 2005. Thus, the impact of mergers on the extensive margin seems to be insignificant. The percentage of interchange lines decreased from 41% to 11%, showing that firms eliminated a considerable number of them after the mergers. The average number of competitors in each origin–destination market decreased slightly from 3 to 2 between 1985 and 2005, indicating that firms primarily engage in oligopolistic competition in most local markets. Finally, the total number of waybills in the waybill sample increased from about 263,000 in 1985 to 611,000 in 2005, reflecting that the total volume of railroad shipment more than doubled during this period.[9] Comparing railroads with other transportation modes, Figure C.1 illustrates the total ton-miles of freight transported by each transportation mode between 1980 and 2011. Although the shipment volume also rose for other transportation modes, like trucking, the proportion of railroad shipment increased during the examined period.

3 Suggestive Evidence for Merger Gains

To provide initial evidence of efficiency gains from railroad mergers, I first examine the evolution of shipment prices around merger events.

I begin by presenting the estimation results of merger-related price effects using the following specification:

$$\log P_{j,odt} = \gamma_j + \gamma_{od} + \gamma_t + \delta_1 D_{j,odt} + X'_{j,odt}\beta + \epsilon_{j,odt}. \tag{1}$$

[9] A year-by-year analysis reveals very similar findings.

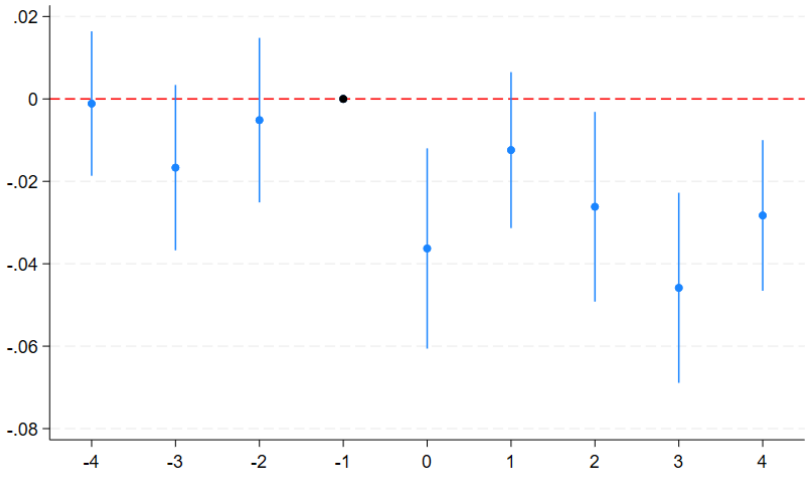
In this specification, the dependent variable $P_{j,odt}$ denotes the shipment price charged by railroad firm j from origin o to destination d at time t. The explanatory variables include origin–destination fixed effects (γ_{od}), firm-specific fixed effects (γ_j), and year fixed effects (γ_t). Importantly, the merger indicator $D_{j,odt}$ equals one if firm j experienced a merger at or before time t, and zero otherwise.[10] Additionally, the vector $X_{j,odt}$ controls for shipment attributes, such as railcar ownership and commodity type. The term $\epsilon_{j,odt}$ represents the unobserved error component.

The primary estimated effect from equation (1) is a statistically significant 9.5% reduction in shipment prices (Panel (a) of Figure 4), consistent with prior analyses of freight railroad mergers. For instance, the Surface Transportation Board (STB) documented price decreases of approximately 11% for coal shipments and 6% for other commodities following the Union Pacific–Southern Pacific merger.[11]

	(1) Log Price
Post Merger	-0.095***
	(0.014)
N	12,108,248
Firm FE	Y
Year FE	Y
O–D Route FE	Y

Standard errors in parentheses. Clustered at O–D route level. *** $p < 0.001$

(a) OLS Results



(b) Dynamic Effects Estimated using Sun and Abraham (2021)

Figure 4: Estimated Effect of Railroad Merger on Log Shipment Price

Because merger events occur at different points in time for different railroad firms across various geographic regions, the setting naturally constitutes a staggered difference-in-differences design. Such staggered timing introduces potential biases due to heterogeneous treatment effects (De Chaisemartin and d'Haultfoeuille, 2020). To address these potential biases, I im-

[10] For railroads involved in multiple mergers (multiple treatments), $D_{j,odt}$ switches to 1 at the first merger event (first treatment).

[11] In comparison, these price reductions in the U.S. railroad industry are notably larger than those observed in other industries. For instance, merger-induced price reductions in the brewing industry averaged approximately 2% (Ashenfelter et al., 2015). This underscores the particular importance of merger-induced cost efficiencies within the railroad industry.

plement the event-study approach proposed by Sun and Abraham (2021), which explicitly accounts for the staggered nature of treatment timing and heterogeneity in treatment effects.

Panel (b) of Figure 4 displays the dynamic effects on shipment prices from four years before to four years after mergers. Several important patterns emerge. First, there is no statistically significant pre-trend in prices in the four years preceding mergers, supporting the assumption that mergers were not anticipated through price adjustments. Immediately following mergers, we observe a notable price reduction, averaging around 4%. However, the magnitude and consistency of price reductions strengthen substantially one to two years after mergers occur. This delayed impact likely reflects the time required for railroad firms to successfully eliminate interchange inefficiencies, re-optimize routing, and consolidate operations across newly integrated networks. The operational challenges encountered in early years following mergers, such as those documented for the Union Pacific—Southern Pacific merger, illustrate the complexities of realizing these efficiency gains quickly.[12]

To ensure the robustness of these findings, Appendix D further examines price reductions across various commodities. The results consistently show significant and substantial price decreases, particularly for commodities predominantly shipped by rail, such as coal, chemicals, and construction materials, reinforcing the generalizability and reliability of the observed merger effects.

Given consistent evidence of price reductions following railroad mergers, the next step is to identify the primary sources underlying these efficiency gains. Clarifying these sources will inform which margins of adjustment are most critical to explicitly incorporate into my structural modeling approach.

Sources of Efficiency Gains

To identify these efficiency sources, I begin by drawing insights from industry experts. Interviews with railroad executives,[13] detailed industry reports,[14] and government investigations into past railroad mergers[15] consistently highlight two primary types of efficiency gains: scope-related and scale-related efficiencies.

Scope-related efficiencies primarily emerge from eliminating costly inter-railroad interchanges, enabling railroads to consolidate shipments and re-optimize traffic routing across

[12]See the 1997 Wall Street Journal article discussing the operational difficulties Union Pacific faced one year after its merger with Southern Pacific: Problems in Wake of Merger Continue for Union Pacific.

[13]See interviews of the Chief Operating Officer of CSXT in Ex Parte No. 711 (Sub-No.1) Reciprocal Switching, Opening Comments, CSX Transportation Inc. Also, books authored by former CEOs of Southern Pacific (Krebs, 2018) and Canadian National Railway (Harrison, 2005).

[14]Industry reports by the Association of American Railroads (AAR) discuss railroad scale economies.

[15]For example, STB Finance Docket No. 32760 documents efficiency sources from the Union Pacific and Southern Pacific merger.

newly integrated networks. Routes experiencing these efficiencies—referred to as *scope-routes*—originally required coordination between two separate railroads before the merger but transitioned to single-line operations afterward. This structural shift substantially reduces interchange delays, simplifies operational procedures, and facilitates network-wide optimization through improved traffic classification, more efficient train scheduling, and smoother operational transitions.

In contrast, scale-related efficiencies arise from increased shipment volumes on routes already operated as single-line services prior to the merger. I term these routes *single-carrier routes*. Scale-related gains reflect lower marginal costs resulting from higher shipment density, enabling railroads to leverage operational improvements such as increased utilization of unit trains and greater double-stacking of container operations.

To empirically distinguish between these two sources of efficiencies, I explicitly define the two types of routes as follows. *Scope-routes* are those that transition from multi-carrier services before the merger to single-carrier services afterward. By comparison, *single-carrier routes* remain under the operation of a single railroad both pre- and post-merger, experiencing varying degrees of shipment volume increases.

Table 3 presents the empirical analysis distinguishing these route types. The results (Panel A of Table 3) indicate that the most substantial price reductions occur on scope-routes, averaging approximately 12%, reflecting significant operational efficiencies from eliminating interchange points and optimizing routes. Single-carrier routes exhibit comparatively smaller price reductions (around 6%). Notably, single-carrier routes experiencing high shipment volume increases (top 10 percentile) do not exhibit statistically significantly larger price reductions than those with lower volume increases (approximately 7% versus 5%, respectively, with the difference not statistically significant). These findings suggest that the primary source of observed efficiency gains following mergers arises from scope-related operational improvements rather than solely from route-level density increases.

Furthermore, an examination of the underlying mechanisms reveals significant shipment volume growth on scope-routes following mergers, while single-carrier routes, on average, experience negligible or slightly negative changes in volume. Specifically, Panel B of Table 3 demonstrates that both total carloads and total billed weight significantly increase for scope-routes after mergers, whereas single-carrier routes, on average, see a small and statistically insignificant decline in shipment weights. These findings support the interpretation that post-merger traffic is reconsolidated, shifting volumes from pre-merger single-carrier routes toward scope-routes, indicating that scope-related operational improvements play a more crucial role in cost reductions than density-driven efficiencies alone.

These findings indicate that efficiency gains following mergers primarily stem from scope-

related operational improvements rather than merely from increased traffic density on single-carrier routes. Specifically, significant volume growth occurs predominantly on scope-routes, reflecting a deliberate strategy by merging railroads to reconsolidate and optimize shipments across their expanded networks. In contrast, single-carrier routes—regardless of their shipment volume increases—do not achieve comparable price reductions, emphasizing that operational efficiencies arising from the elimination of interchange and optimized routing are central drivers of merger-related gains during the 1990–2005 period.

Table 3

Reduced-Form Analysis of Merger Effects

	(1)	(2)
Panel A: Price Effect of Mergers	log(Price)	log(Price)
Scope-Routes	−0.1206***	−0.1209***
	(0.0183)	(0.0184)
Single-Carrier Routes	−0.0588***	
	(0.0137)	
Single-Carrier Routes (High Volume Increase)		−0.0685***
		(0.0224)
Single-Carrier Routes (Low Volume Increase)		−0.0534***
		(0.0154)
Observations	12,108,248	12,108,248
R^2	0.747	0.747
	(1)	(2)
Panel B: Volume Changes after the Merger	log(Carloads)	log(Weight)
Scope-Routes	0.0726***	0.0507***
	(0.0168)	(0.0164)
Single-Carrier Routes	0.0079	-0.0187
	(0.0163)	(0.0192)
Observations	1,103,565	1,103,547
R^2	0.396	0.447

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Control variables include period fixed effects, firm fixed effects, route fixed effects, railcar ownership (privately owned by shippers or owned by railroads), and commodity types. Standard errors are clustered at the route level (origin—destination pair).

Guided by these empirical insights, the structural model developed in the next section

explicitly focuses on capturing efficiencies related to three key operational improvements: the elimination of interchange costs, network consolidation, and shipment flow re-optimization. The model particularly emphasizes the transition of shipments on scope-routes—formerly handled by separate merging railroads requiring interchange—into streamlined single-carrier services post-merger. While recognizing potential endogeneity concerns (for example, the decision about where to reroute or consolidate traffic could be endogenous to mergers), the structural approach is designed to directly address this endogeneity by explicitly modeling firm-level routing and consolidation decisions.

Complementary evidence from recent literature (Degiovanni and Yang, 2023) highlights the increasing importance of scale-related efficiencies from approximately 2010 onward. One plausible explanation for this shift is that scope-related efficiencies had largely been exhausted by earlier mergers, especially given the absence of major railroad mergers between 1999 and 2010, prompting the industry to focus more intensively on optimizing single-line routes through increased traffic density.

To further validate these quantitative findings, I conducted extensive qualitative fieldwork over one month, involving direct observations and interviews with key industry practitioners. Fieldwork included downstream customers (steel manufacturers and the UPS consolidation center in Chicago), classification yards and inter-modal terminals (e.g., BNSF facilities), regional railroads heavily involved in interchange (Lake State Railway Company), and technology centers specializing in route optimization and network management (The Transportation Technology Center in Pueblo, Colorado). These ethnographic insights corroborate the empirical analysis, reinforcing the conclusion that operational efficiencies associated with the elimination of interchange and optimized route consolidation were critical to achieving merger gains in the 1990–2005 period. Detailed field notes, observational records, and photographic documentation are presented in Appendix B.

4 Model

The empirical analysis indicates substantial price reductions following railroad mergers, particularly on interconnecting routes. To understand the underlying mechanisms driving these observed price changes, I introduce a novel spatial equilibrium model incorporating network interdependencies, oligopolistic competition, and decentralized internal decision-making. Railroad marketers set local market prices, while operational teams strategically determine routing and resource allocation decisions across the firm's network.

In what follows, I first specify consumer demand. Then, I define firm networks, profit and cost functions, and the constraints firms face. I clarify the decentralized decision-making

structure within each firm and finally present the equilibrium concept.

4.1 Demand

For each origin–destination market, customers choose a transportation service, denoted as s, to move their goods. Following a standard discrete-choice logit framework, customer i's utility from choosing service s in market $o \to d$ at time t is:

$$u_{is,odt} = \alpha \cdot p_{s,odt} + \beta \cdot \log TTM_{s,odt} + \gamma_{s(o)} + \gamma_{s(d)} + \gamma_m + \gamma_t + \xi_{s,odt} + \varepsilon_{is,odt}. \tag{2}$$

In this equation, $p_{s,odt}$ denotes the shipping price, and $TTM_{s,odt}$ is the total track mileage owned by the service providers in the origin and destination regions. Greater track mileage within the origin and destination regions facilitates easier transfers of goods, thereby reducing loading and unloading times. The model includes fixed effects for origin and destination railroads ($\gamma_{s(o)}$, $\gamma_{s(d)}$), market-specific characteristics (γ_m, including origin–destination distance), and time effects (γ_t). The term $\xi_{s,odt}$ represents unobserved service quality, while $\varepsilon_{is,odt}$ denotes customer-specific preference shocks, assumed to follow an extreme-value type I distribution. Customers also have an outside transportation option (non-rail), represented as $j_o = j_d = 0$.

Given these utilities, the market share $h_{s,odt}$ and corresponding demand $Q_{s,odt}$ for each service are computed as:

$$h_{s,odt} = \frac{\exp(\alpha p_{s,odt} + \beta \log TTM_{s,odt} + \gamma_{s(o)} + \gamma_{s(d)} + \gamma_m + \gamma_t + \xi_{s,odt})}{1 + \sum_{s'} \exp(\alpha p_{s',odt} + \beta \log TTM_{s',odt} + \gamma_{s'(o)} + \gamma_{s'(d)} + \gamma_m + \gamma_t + \xi_{s',odt})}, \tag{3}$$

$$Q_{s,odt} = M_{odt} \cdot h_{s,odt}, \tag{4}$$

where M_{odt} denotes the market size for the o–d market at time t.

Due to data limitations—specifically, the absence of trip duration information—I abstract away from modeling demand explicitly as a function of shipping duration. Further discussion of this assumption and its implications are provided in Appendix E.1.

4.2 The Firm's Problem

I analyze railroad firms' strategic pricing, routing, and resource allocation decisions on an annual basis, focusing on equilibrium interactions within their networks. At the beginning of each year, firm j operates a fixed network $\mathcal{N}_j = (\mathcal{Z}_j, \mathcal{A}_j)$, consisting of nodes $\mathcal{Z}_j$, which represent the regions served by firm j, and arcs $\mathcal{A}_j$, which represent rail segments connecting these regions. Firms simultaneously set optimal prices $\{p_{s,od}\}_{o,d \in \mathcal{Z}_j}$ for each origin–destination

market they serve, determine routing decisions $\{\mathcal{R}_{j,od}\}_{o,d\in\mathcal{Z}_j}$, and allocate resources (such as annual maintenance spending) $\{I_{j,ab}\}_{(a,b)\in\mathcal{A}_j}$ across their networks to minimize total operational costs. A routing decision $\mathcal{R}_{j,od}$ for each market from node o to node d corresponds to a sequence of connected arcs. For example, $\mathcal{R}_{j,od} = \{(o, m_1), (m_1, m_2), (m_2, d)\}$ indicates the chosen shipment route from origin o through intermediate nodes m_1 and m_2 to destination d.

Defining the Network. To construct a network $\mathcal{N}_j$ for a given railroad company j, a non-directed graph $(\mathcal{Z}_j, \mathcal{A}_j)$ is defined with nodes denoted by $\mathcal{Z}_j$ and arcs denoted by $\mathcal{A}_j$. The nodes in this network represent the centroids of the BEA regions, while the arcs correspond to the rail lines connecting each BEA economic area. The rail network of the Burlington Northern railway (BN) in 1994 is depicted in the left panel of Figure 5. A BEA area a is included in $\mathcal{Z}_j$ if BN has tracks in that region. Similarly, if two adjacent BEA areas are connected by tracks owned by BN, then $(a, b) \in \mathcal{A}_j$, as illustrated in the right panel of Figure 5. I follow the same process in constructing the network of each railroad firm for every year between 1985 and 2005, based on detailed geographic information for each rail line and information about the ancestry of rail lines obtained from the Federal Transit Administration.[16]

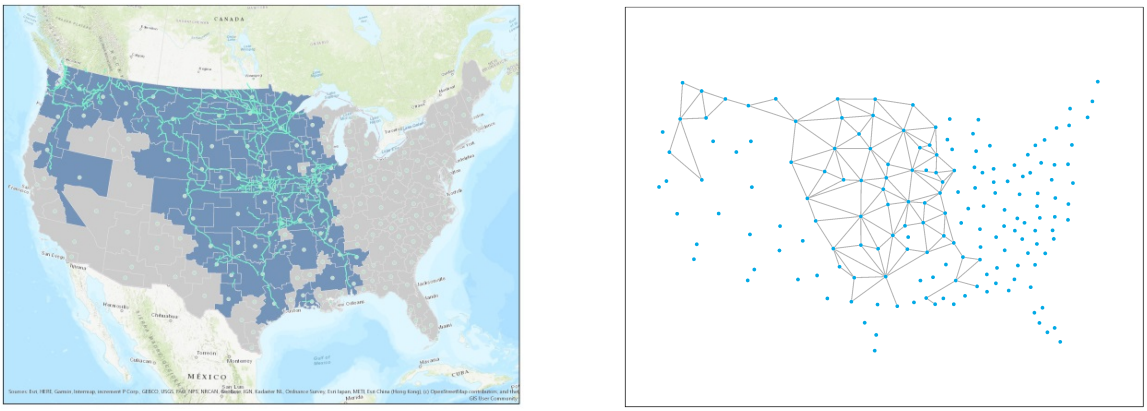


Figure 5: BN Rail Network

Profit Function. The profit function of firm j is defined as:

$$\begin{aligned}\pi_j(\boldsymbol{p_j}, \boldsymbol{\mathcal{R}_j}, \boldsymbol{I_j}) &= \sum_{s:s\in S(j)} p_{s,od} \cdot Q_{s,od}(p_{s,od}, \boldsymbol{p_{-s,od}}) - C(\boldsymbol{Q_j}, \boldsymbol{\mathcal{R}_j}, \boldsymbol{I_j}) \\ &= \sum_{s:s\in S(j)} (p_{s,od} - C_{s,od}(\boldsymbol{\mathcal{R}_j}, \boldsymbol{I_j})) \cdot Q_{s,od}(p_{s,od}, \boldsymbol{p_{-s,od}}). \qquad (5)\end{aligned}$$

[16]Historical information about railroad networks is obtained by comparing shapefiles of U.S. rail lines with merger history information from the Oak Ridge National Laboratory.

The vector $\boldsymbol{p_j}$, $\boldsymbol{\mathcal{R}_j}$, and $\boldsymbol{I_j}$ represent the decisions of firm j for prices, routing, and resource allocation, respectively. I use bold symbols to represent vectors in my model. $S(j)$ represents the set of services that the railroad firm j offers.[17] $C_{s,od}$ is the marginal cost of transportation.

I assume that the originating railroad firm in a joint-line service sets the price, abstracting from double marginalization. I provide detailed justification for this assumption in Section 4.3, when deriving equilibrium pricing decisions.

Cost Function. The marginal cost of transportation $C_{s,od}$ is specified as:

$$C_{s,od} = \begin{cases} \sum_{(a,b)\in\mathcal{R}_{j,od}} c_{j,ab}(\boldsymbol{I}_j), & \text{if } j_o = j_d, \\ \sum_{j'\in J(s)} \sum_{(a,b)\in\mathcal{R}_{j',od_{j'}(s)}} c_{j',ab}(\boldsymbol{I}_{j'}) + \#_{interchanges} \cdot \eta, & \text{if } j_o \neq j_d. \end{cases} \tag{6}$$

In the case of a single-line service ($j_o = j_d = j$), marginal cost is computed by summing the arc-level costs along the entire route provided by firm j. For a joint-line service involving multiple railroads ($j_o \neq j_d$), marginal cost additionally includes interchange costs. Here, $J(s)$ denotes the set of firms providing service s, $\#_{interchanges}$ represents the number of interchanges, and η is the cost per interchange. Most shipments involve at most one interchange; in the data, approximately 80% of joint-line trips have exactly one interchange, and around 95% of all trips have zero or one interchange.

In this model, interchanges specifically refer to connections between different railroad companies. Movements through multiple stations within the same railroad firm are not considered interchanges but are instead captured by the summation of arc-level transportation costs. Thus, if a shipment travels through internal *hubs* within the same railroad, it benefits from potentially lower arc-level transportation costs, reflecting operational efficiencies associated with these hubs.

The arc-level transportation cost, $c_{j,ab}$, depends on the arc-specific geographic characteristics, $\delta_{j,ab}$, and the investment resources allocated by firm j to arc (a, b). The efficiency parameter γ reflects the effectiveness of resources and has a positive value. Therefore, if firm j allocates more investment resources to arc (a, b), the arc-level transportation cost $c_{j,ab}$ will be smaller.

[17] For all $s \in S(j)$, $j_o(s) = j$, $j_d(s) = j$, or $j_i(s) = j$, where o is the origin, d is the destination, and i is the interchange station. There can be multiple interchanges for a shipment. $S(j)$ is directly obtained from the data.

Specifically, the arc-level transportation cost is defined as:

$$c_{j,ab} = \begin{cases} \frac{\delta \cdot Dist_{j,ab}}{I_{j,ab}^{\gamma}}, & \text{if } I_{j,ab} > 0 \\ \infty, & \text{if } I_{j,ab} = 0 \end{cases}. \tag{7}$$

When $I_{j,ab} = 0$, transportation along arc (a, b) is considered infeasible, resulting in an infinite transportation cost. Here, $Dist_{j,ab}$ denotes the distance between nodes a and b for firm j, and the parameter δ represents the shipping cost per mile.

Motivated by the reduced-form evidence presented in Section 3, the key modeling mechanism here explicitly captures *scope-related efficiencies* at the route level. Specifically, when railroad firms consolidate shipments onto common network segments—routes that simultaneously serve multiple origin–destination markets within the same firm—the resulting increase in resource concentration (such as enhanced infrastructure maintenance, better scheduling, and improved route-level coordination) benefits all traffic traveling through these segments. Following Fajgelbaum and Schaal (2020), this mechanism is captured by the scope efficiency parameter (γ). The allocation decisions $I_{j,ab}$ broadly represent operational enhancements, empirically approximated by firm-level maintenance expenditures. In Appendix F, I provide comparative statics to illustrate explicitly how changes in the scope parameter γ affect routing and resource allocation decisions.

Additionally, my model incorporates economies of scale arising from fixed costs. Each railroad company incurs fixed operational costs, including administrative and general overheads; thus, higher shipment volumes naturally generate economies of scale at the firm-wide aggregate level.[18]

Constraints. The optimization problem of each firm j is subject to two constraints.

First, the resource allocation constraint:

$$\sum_{(a,b)\in\mathcal{A}_j} I_{j,ab} \leq K_j, \tag{8}$$

where K_j represents the total annual maintenance spending.

Due to data limitations, I observe maintenance spending only at the firm level (K_j), while the segment-level allocation ($I_{j,ab}$) is not directly observed. The model therefore en-

[18] In my current analysis, I do not model congestion or economies of scale at the arc-level explicitly, due to computational constraints. Incorporating these arc-level nonlinearities would transform the optimal routing decision into a nonlinear optimization problem, significantly increasing computational complexity when repeatedly solving the equilibrium during estimation. Appendix E.2 discusses alternative modeling strategies for these considerations and evaluates their potential implications for equilibrium outcomes.

dogenously predicts how maintenance resources are allocated across segments. I provide detailed validation of these resource allocation predictions, including evidence from track abandonment patterns and observed maintenance practices, in Section 5.3.

The second constraint of firm j's optimization problem ensures balanced flow at each node. Specifically, for each service s, market o–d, and node $z \in \mathcal{Z}_j$:

$$\sum_{a \in \mathcal{Z}_j(z)} Q_{s,od} \cdot \mathbb{1}\{(a,z) \in \mathcal{R}_{j,o_j(s),d_j(s)}\} + D_{j,z} = \sum_{b \in \mathcal{Z}_j(z)} Q_{s,od} \cdot \mathbb{1}\{(z,b) \in \mathcal{R}_{j,o_j(s),d_j(s)}\}. \quad (9)$$

Here, $\mathcal{Z}_j(z)$ denotes the set of neighboring nodes of node z, and

$$D_{j,z} = \begin{cases} Q_{s,od} & \text{if } z = o \\ -Q_{s,od} & \text{if } z = d \\ 0 & \text{otherwise} \end{cases}$$

is the net demand at node z.

Intuitively, this constraint ensures that at every node within railroad j's network, the inflow of goods plus net demand equals the outflow, maintaining consistent traffic flows.

Note that this model abstracts away from the movement of empty railcars (backhauling) and specifies the balanced-flow constraint at the origin–destination level rather than at a more aggregated level. Further discussions of these assumptions, their implications, and potential extensions are detailed in Appendix E.3.

4.3 Equilibrium

Decisions within each railroad firm are decentralized into two stages. In the *marketing stage*, marketers set prices in each local market, taking routing and resource allocation decisions as given. In the subsequent *operational stage*, operational teams determine routing and resource allocation decisions across the entire network, given the quantities implied by marketing decisions.

These simultaneous pricing and network-wide operational decisions capture the interdependencies inherent in railroad operations: pricing decisions in one local market influence network-wide operational costs, while operational decisions in one region affect marginal costs in other regions.[19]

[19]To evaluate the assumption of decentralized decision-making, I derive and solve a monopoly firm's Lagrangian for centralized decisions. In Appendix E.4, I then compare equilibrium outcomes under centralized and decentralized structures.

A Nash equilibrium thus consists of pricing, routing, and resource allocation decisions that simultaneously satisfy the marketers' optimal pricing first-order conditions (taking network decisions as given) and operational teams' optimality conditions for routing and resource allocation (given marketing decisions). I formally define the Nash equilibrium after specifying optimal decisions for pricing, routing, and resource allocation.

Optimal Pricing Decisions. Firm j sets the price for each single-line service it fully provides, as well as for each joint-line service for which it is the originating railroad. The optimal pricing decision for service $s \in S(j)$ satisfies the standard first-order condition:

$$\frac{\partial \pi_j}{\partial p_{s,od}} = Q_{s,od} + (p_{s,od} - C_{s,od}) \cdot \frac{\partial Q_{s,od}}{\partial p_{s,od}} = 0. \tag{10}$$

In principle, it might be possible for a firm to offer both a single-line and a joint-line service in the same o–d market. However, given the inherent operational advantages of single-line services—such as fewer coordination issues and greater efficiency—this scenario rarely occurs in the data. Therefore, equation (10) omits summation over multiple services provided by the same firm in a given o–d market, because this scenario is empirically rare and does not significantly affect model solutions.

In the case of a joint-line service, only the originating railroad firm sets the price. For example, consider a service s involving one interchange and two railroad companies, where the total marginal cost is $C_{s,od} = C_{1,oi} + C_{2,id} + \eta$. Railroad firm 1 (the originating firm) bases its pricing decision $p_{s,od}$ on its own cost segment $C_{1,oi}$ (from the origin to the interchange station i), the subsequent shipment cost $C_{2,id}$ incurred by firm 2 (from interchange i to destination d), and the interchange cost η.

The assumption that the originating railroad firm sets the price in joint-line services aligns with standard industry practice and documented regulatory evidence. Interviews with railroad managers confirm that typically only one railroad firm determines prices, thus eliminating double marginalization (see Appendix B). Moreover, empirical research conducted by the Department of Justice (Alexandrov et al., 2018) provides additional evidence supporting the absence of double marginalization in pricing decisions for joint-line railroad services.

Including double marginalization would alter local pricing incentives, typically leading to higher pre-merger prices. Therefore, my model's predictions on price reductions from mergers represent conservative estimates (lower bounds), as removing double marginalization would further enhance merger-induced price reductions.

Optimal Resource Allocation Decisions. Optimal routing and resource allocation de-

cisions are jointly determined by operational teams within each railroad firm. In the operational stage, given quantities implied by marketers' pricing decisions, these teams minimize total operational costs by simultaneously choosing optimal resource allocation and routing across the network.

The operational team minimizes the shipment cost for quantities $\boldsymbol{Q}_j$:

$$\mathcal{L}_j(\boldsymbol{\mathcal{R}_j}, \boldsymbol{I_j}) = \sum_{s:s\in S(j)} C_{s,od}(R_{j,o_j(s)d_j(s)}, \boldsymbol{I_j}) \cdot Q_{s,od} + \lambda_j \left(\sum_{(a,b)\in\mathcal{A}_j} I_{j,ab} - K_j \right) + \sum_{s:s\in S(j), z\in\mathcal{Z}_j} \lambda_{j,s,z} \tag{11}$$
$$\left(D_{j,z} + \sum_{a\in\mathcal{Z}_j(z)} Q_{s,od} \cdot \mathbb{1}\{(a,z) \in \mathcal{R}_{j,o_j(s),d_j(s)}\} - \sum_{b\in\mathcal{Z}_j(z)} Q_{s,od} \cdot \mathbb{1}\{(z,b) \in \mathcal{R}_{j,o_j(s),d_j(s)}\} \right).$$

The vectors $\boldsymbol{\mathcal{R}_j}$, and $\boldsymbol{I_j}$ represent the decisions of firm j for routing and resource allocation, respectively. $C_{s,od}$ is the marginal cost of transportation and λ's are Lagrangian multipliers.

To obtain the optimal resource allocation decisions, we begin by taking the derivatives of the Lagrange function with respect to $I_{j,ab}$ for every arc $(a,b) \in \mathcal{A}_j$, yielding:

$$\frac{\partial \mathcal{L}_j}{\partial I_{j,ab}} = -\gamma \cdot \delta_{j,ab} \cdot \frac{\sum_{s:s\in S(j)} Q_{s,od} \cdot \mathbb{1}\{(a,b) \in \mathcal{R}_{j,o_j(s),d_j(s)}\}}{I_{j,ab}^{\gamma+1}} + \lambda_j. \tag{12}$$

The numerator of this expression calculates the total quantity that travels through arc (a,b) by summing over all the origin and destination markets for which firm j operates. The optimal allocation decision for each arc is then determined through the Kuhn-Tucker condition, which requires that for all $(a,b) \in \mathcal{A}_j$:

$$\frac{\partial \mathcal{L}_j}{\partial I_{j,ab}} \leq 0, \quad I_{j,ab} \geq 0, \quad \text{and } I_{j,ab} \frac{\partial \mathcal{L}_j}{\partial I_{j,ab}} = 0. \tag{13}$$

Here, I provide further intuition on the optimal allocation decisions. In equilibrium, for any non-zero $I^*_{j,ab}$ and $I^*_{j,a'b'}$, they satisfy

$$\begin{aligned} \frac{I^*_{j,ab}}{I^*_{j,a'b'}} &= \left[\frac{\delta_{j,ab} \cdot \sum_{s:s\in S(j)} Q^*_{s,od} \cdot \mathbb{1}\{(a,b) \in \mathcal{R}^*_{j,o_j(s),d_j(s)}\}}{\delta_{j,a'b'} \cdot \sum_{s:s\in S(j)} Q^*_{s,od} \cdot \mathbb{1}\{(a',b') \in \mathcal{R}^*_{j,o_j(s),d_j(s)}\}} \right]^{\frac{1}{\gamma+1}} \\ &= \left[\frac{\delta_{j,ab} \cdot q_{j,ab}(\boldsymbol{Q}^*_j, \boldsymbol{\mathcal{R}}^*_j)}{\delta_{j,a'b'} \cdot q_{j,a'b'}(\boldsymbol{Q}^*_j, \boldsymbol{\mathcal{R}}^*_j)} \right]^{\frac{1}{\gamma+1}} \end{aligned} \tag{14}$$

where $\delta_{j,ab}$ is the arc-specific geographic characteristics. $q_{j,ab}(Q, \mathcal{R})$ is the total shipment

amount that runs through arc (a, b), which depends on the total quantity that needs to be shipped from each origin–destination market, $\boldsymbol{Q}_j$, and the routing decisions, $\boldsymbol{\mathcal{R}}_j$.

Equation (14) shows that the optimal allocation of maintenance resources is proportional to the usage of the track. The greater the quantity being shipped through a particular track, the more maintenance resources will be allocated to it. The exact ratio of allocation is influenced by the value of the *resource allocation efficiency* parameter, γ, which affects the marginal returns of resource allocation on a specific arc.

Optimal Routing Decisions. The optimal routing decision is equivalent to finding the least expensive route to travel from origin o to destination d for each o–d market. This can be formulated as a linear-programming problem:

$$\min_{\mathcal{R}_{j,od}} \sum_{(a,b)\in\mathcal{R}_{j,od}} \frac{\delta Dist_{j,ab}}{I_{j,ab}^{\gamma}} \tag{15}$$

such that $\forall m' \in \mathcal{Z}_j$,

$$\mathbb{1}\{m' = o\} - \mathbb{1}\{m' = d\} + \sum_{a\in\mathcal{Z}_j(m')} \mathbb{1}\{(a, m') \in \mathcal{R}_{j,od}\} = \sum_{b\in\mathcal{Z}_j(m')} \mathbb{1}\{(m', b) \in \mathcal{R}_{j,od}\}.$$

This constraint is binding in equilibrium, and the linear programming (LP) algorithm used to determine the optimal routing decision enforces this constraint. LP algorithms, such as the Simplex method or Interior-Point methods, do not rely on setting derivatives to zero. Instead, they operate within the feasible region defined by the constraints, iteratively moving towards the optimal solution using linear algebra techniques. Consequently, the balanced-flow constraint appears in the LP but does not explicitly appear in the FOCs for the allocation decisions as in equation (12).

The Equilibrium. A Nash equilibrium in this decentralized setting involves strategic interactions among railroad firms, each making optimal pricing, routing, and resource allocation decisions, simultaneously satisfying the conditions outlined below.

First, equilibrium prices $p_{s,od}^*$ are determined by the marketing teams of each railroad firm. Specifically, for each railroad firm j and each service $s \in S(j)$, these equilibrium prices satisfy the first-order conditions derived from profit maximization, as characterized by equation (10). In choosing optimal prices, each firm accounts explicitly for how its pricing decisions affect equilibrium market outcomes and considers its own equilibrium routing and resource allocation decisions, $\mathcal{R}_{j,od}^*$ and $I_{j,ab}^*$, while taking as given the equilibrium pricing,

routing, and investment allocation decisions, $(p^*_{-s,od}, \mathcal{R}^*_{-j,od}, I^*_{-j,ab})$, of all competing firms.

Second, given these equilibrium prices $p^*_{s,od}$, the operational teams within each railroad firm simultaneously determine optimal resource allocation decisions $I^*_{j,ab}$. These decisions satisfy the Kuhn–Tucker optimality conditions described by equation (13), taking as given the equilibrium routing decisions $\mathcal{R}^*_{j,od}$ and shipment quantities implied by the equilibrium pricing decisions.

Third, at the same operational stage, and given equilibrium prices and resource allocations, the operational teams also choose optimal routing decisions $\mathcal{R}^*_{j,od}$ for each firm j. These routing decisions minimize the firm's total shipment costs by solving the linear programming problem characterized by equation (15), conditional on equilibrium investment allocations $I^*_{j,ab}$ and shipment quantities resulting from equilibrium pricing decisions $p^*_{s,od}$.

Given a set of parameters, the steps to compute the equilibrium are as follows:

Algorithm 1 Iterative Algorithm for Prices and Investment Decisions

Require: Initial tolerances ϵ_p, ϵ_I, initial prices $p^0_{s,od}$, initial investments $I^0_{j,ab}$

Initialize iteration counters: $k \leftarrow 0$, $m \leftarrow 0$
repeat
    Set $k \leftarrow k + 1$
    Given $p^{k-1}_{s,od}$, initialize investment decisions $I^{k,0}_{j,ab} \leftarrow I^{k-1}_{j,ab}$
    repeat
        Set $m \leftarrow m + 1$
        Compute optimal routing decisions using linear programming given $p^{k-1}_{s,od}, I^{k,m-1}_{j,ab}$
        Derive new investment allocation decisions $I^{k,m}_{j,ab}$ from Kuhn-Tucker conditions
    until $\|I^{k,m}_{j,ab} - I^{k,m-1}_{j,ab}\| < \epsilon_I$
    Set $I^k_{j,ab} \leftarrow I^{k,m}_{j,ab}$
    Given $I^k_{j,ab}$, derive new prices $p^k_{s,od}$ from first-order conditions
until $\|p^k_{s,od} - p^{k-1}_{s,od}\| < \epsilon_p$
return Optimal prices $p^k_{s,od}$, optimal investments $I^k_{j,ab}$, and corresponding routing decisions

When computing equilibrium, I impose an empirical constraint ensuring that firms maintain positive equilibrium market shares in local markets where their market shares are positive in the data. This empirical restriction prevents equilibria involving unrealistic sorting of firms into exclusive regional monopolies—such as one firm operating exclusively in the east and another exclusively in the west—thus ensuring the model's predictions align closely with observed industry patterns.

5 Estimation

I begin by presenting the argument for identification, followed by an explanation of how I estimate the parameters for both demand and cost. I estimate demand and cost parameters separately. The estimation of demand parameters is conducted using the two-stage least squares method, while the cost parameters are estimated through the simulated method of moments.

5.1 Demand Estimation

The demand model outlined in equation (4) shows that the difference between the logarithm of the observed market share for each service and the logarithm of the share for the outside good can be calculated as follows:

$$\log(h_{s,odt}) - \log(h_{0,odt}) = \alpha \cdot p_{s,odt} + \beta_1 \cdot \log TTM_{s,odt} + \zeta_{s,odt}. \tag{16}$$

In this equation, $h_{s,odt}$ represents the market share of a specific service s that serves the market route from origin o to destination d at a specific time t. The variable $h_{0,odt}$ represents the market share of the outside alternative. In my regression, the share of the outside good is constructed by calculating the shipment shares of other transportation modes, such as trucking and water shipping, using data from the Commodity Flow Survey.

The estimates of α and β_1 can be obtained from a linear instrumental variables regression of differences in log market shares on prices and service characteristics. The error term in equation (16) represents the unobserved market-specific demand shocks. Since I assume that railroads observe and account for this deviation, it will influence the market-specific markup and bias the estimate of price sensitivity. To address the issue of endogeneity, a series of steps are taken. Firstly, fixed effects for both origin and destination railroads, market fixed effects, and time fixed effects are incorporated. Furthermore, instrumental variables (IVs) are employed to mitigate the endogeneity problem.

The first set of instrumental variables used is commonly referred to as the "BLP Instruments." Much of the previous research treats the endogeneity problem by assuming the characteristics space is exogenous or predetermined.[20] Therefore, characteristics of other services will be correlated with price since the markup of each service will depend on the distance from the nearest neighbour. Here I use the average track miles and the distance to other railroads in the same o–d market as instruments.

For my second set of IVs, I utilize mergers between railroad firms as proxies to capture

[20] See Berry (1994), Berry, Levinsohn, and Pakes (1995), and Bresnahan et al. (1997).

changes in market power within local markets. Considering that railroad companies operate on a national scale, the decision to merge between two railroad companies is unlikely to be driven by local demand shocks in a single origin–destination market. Thus, I argue that the merger decision is orthogonal to local demand shocks. To examine the impact of mergers on market concentration, I analyze the changes in concentration following railroad mergers. To quantify these changes, I draw upon the methodologies presented in the works of Garmaise and Moskowitz (2006) and Dafny, Duggan, and Ramanarayanan (2012) to construct a simulated change in the Herfindahl-Hirschman Index ($sim\Delta HHI$). The $sim\Delta HHI$ represents the projected change in the HHI that would have occurred after the merger if no other factors were to change. By utilizing the projected change instead of the actual change in the HHI, I aim to isolate the post-merger market share adjustments that may be influenced by local market conditions. The simulated change in the HHI is derived as a result of this approach.

$$\begin{aligned} sim\Delta HHI_{odt} &= (TargetShare_{odt-1} + AcquirorShare_{odt-1})^2 \\ &\quad - (TargetShare^2_{odt-1} + AcquirorShare^2_{odt-1}) \\ &= 2 \times TargetShare_{odt-1} \times AcquirorShare_{odt-1}. \end{aligned}$$

Here, $TargetShare_{odt-1}$ represents the market share held by the target firm in the o–d market one year prior to the merger. Similarly, $AcquirorShare_{odt-1}$ indicates the market share of the acquiring firm in the o–d market before the merger took place.

Estimation Results. Table 4 presents the results. All standard errors are clustered at the market level (origin–destination). Columns (1)–(3) display the ordinary least squares results. The coefficient on price and the resulting own-price elasticities are relatively low, with absolute values less than 1. The logit demand structure does not impose a constant elasticity; therefore, the estimates imply a different elasticity for each service–market–year combination. Some statistics of the own-price elasticity distribution are shown at the bottom of each column. Two sets of instrument variables were explored to deal with the endogeneity problem. Column (4) uses BLP instruments described above as IVs in the same regression. Column (5) uses simulated changes in HHI as IVs. Finally, column (6) uses both sets of IVs. Columns (4)–(6) also include market, firm, and time fixed effects.

Several conclusions can be drawn from the results in Table 4. First, once IVs are used, the coefficient on price and the implied own-price elasticity both increase in absolute value. This is predicted by theory and holds in a wide variety of studies such as Nevo (2000) and Nevo (2001). Second, the "BLP instruments" seem to generate results almost identical to those produced by using the simulated change in HHI. The estimated price elasticity of

demand, as shown in column (6), indicates an estimated price coefficient of approximately -0.72, resulting in an average price elasticity of demand of around -1.60. This estimated average price elasticity is comparable to estimates reported in the transportation literature. For instance, Wilson et al. (1988) estimated an own-price elasticity of demand for U.S. freight rail of -1.46. Beuthe et al. (2001) estimated a demand elasticity of -1.77 for total tonnes.

Table 4: Results of Demand Estimation

	OLS			IV		
Variables	(1)	(2)	(3)	(4)	(5)	(6)
Price	0.240***	−0.280***	−0.281***	−0.708***	−0.681***	−0.720***
	(0.013)	(0.012)	(0.012)	(0.059)	(0.059)	(0.059)
Log Track Miles	0.485***	0.364***	0.334***	0.360***	0.358***	0.360***
	(0.016)	(0.015)	(0.016)	(0.016)	(0.016)	(0.016)
Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
o–d Market Fixed Effect	–	Yes	Yes	Yes	Yes	Yes
Firm Fixed Effect	–	–	Yes	Yes	Yes	Yes
Instruments						
BLP instruments	–	–	–	Yes	–	Yes
Predicted ΔHHI	–	–	–	–	Yes	Yes
First-stage F-statistic	–	–	–	11.18	10.05	11.17
Own price elasticity						
Mean	0.53	–0.62	–0.62	–1.57	–1.51	–1.60
Standard errors	0.28	0.33	0.33	0.83	0.80	0.85
Median	0.51	−0.60	−0.60	–1.51	−1.45	−1.53

Note: Demand estimates are based on 30,058 market–service–year observations in 1993, 1997, 2002, 2007. Figures in parentheses are standard errors. *** p<0.01, ** p<0.05, * p<0.1

5.2 Estimation of Cost Parameters

There are three key cost parameters to be estimated. The first, the *impact of distance* parameter δ, determines how shipment costs vary with travel distance. The second, the *interchange cost* parameter η, relates to costs incurred per interchange. Finally, the *resource allocation efficiency* parameter γ evaluates how effectively resource allocations reduce shipment costs.

Estimation of the cost parameters $\theta = \{\delta, \eta, \gamma\}$ is implemented according to the following procedures. First, given a parameter value, the industry equilibrium is solved, resulting in optimal decisions for pricing, routing, and investment allocation, represented as $(\boldsymbol{p}(\theta), \boldsymbol{\mathcal{R}}(\theta), \boldsymbol{I}(\theta))$. After solving the equilibrium, I obtain the simulated moments as $\Gamma^S(\theta)$. Next, the estimate $\hat{\theta}$ is obtained through the simulated method of moments minimizes the weighted distance between the data moments and the simulated moments using the following objective function:

$$L(\theta) = \min_{\theta}[\Gamma^d - \Gamma^S(\theta)]' W [\Gamma^d - \Gamma^S(\theta)],$$

where W represents a positive-definite matrix.[21]

I target four data moments Γ^d. The first moment is the average price per mile for shipments. The second moment is the average price difference between a route that has interconnections and one that does not. The third and fourth moments pertain to how the degree and betweenness centrality of a node impact shipment prices.

The degree centrality of a particular node i represents the total number of other nodes in the network $\mathcal{N}_j$ of firm j that have a direct connection to node i. Mathematically, the degree centrality of node i in network $\mathcal{N}_j$ can be expressed as:

$$d_i(\mathcal{N}_j) = \sum_{k \in \mathcal{A}_j} \mathbb{1}\{a_{ik} = 1\} \tag{17}$$

On the other hand, betweenness centrality captures how frequently a node appears on the shortest path connecting an origin to a destination. Its calculation is defined as follows:

$$b_i(\mathcal{N}_j) = \sum_{o,d \in \mathcal{Z}_j} \frac{\mathbb{1}\{i \in l(o,d)\}}{(Z_j - 1)(Z_j - 2)}. \tag{18}$$

Here, Z_j denotes the total number of nodes in network $\mathcal{N}_j$, and $l(o,d)$ represents the shortest path from origin o to destination d. Importantly, $l(o,d)$ refers specifically to the shortest geographic path, rather than the least-cost path. Therefore, betweenness centrality captures relative geographic positions, rather than the direct operational costs that depend on endogenous resource allocation decisions.

More specifically, the third and fourth data moments are obtained through regression analysis (represented by m_1 and m_2 in equation 19). In this analysis, the shipment price per mile for each origin-destination market is regressed on the degree centrality and betweenness

[21] I use the inverse of the variance–covariance matrix as the weighting matrix and estimate it using a bootstrap procedure.

centrality of the origin station:

$$P_{s,odt} = \gamma_{s(o)} + \gamma_{s(d)} + \gamma_m + \gamma_t + a_1 \cdot interchanges + m_1 \cdot d_o(\mathcal{N}_{j(s)}) + m_2 \cdot b_o(\mathcal{N}_{j(s)}) + \epsilon_{s,odt} \tag{19}$$

In this equation, $P_{s,odt}$ represents the shipment price for a specific origin–destination market at a given time. I account for the fixed effects related to railroads, origin-destination markets, and time periods, respectively. The variable *interchanges* captures the total number of interchanges. $d_o(\mathcal{N}_{j(s)})$ and $b_o(\mathcal{N}_{j(s)})$ represent the degree centrality and betweenness centrality of the origin station, respectively.

The rationale behind the identification argument is as follows: The first data moment focuses on quantifying the impact of travel distance on average shipping prices. As the value of δ increases, so does the average shipping cost per mile in dollar terms. Therefore, the first moment helps determine the value of the *impact of distance* parameter δ. The second moment compares the average price difference between interconnecting and non-interconnecting routes. When the interchange cost η increases, the price difference between these routes also increases. Therefore, conditional on the values of δ and γ, the second moment identifies the *interchange cost* parameter, η. Lastly, the value of γ represents the effectiveness of resource allocation and significantly influences the last two moments related to network measures. Intuitively, when $\gamma = 0$, resources have no impact on shipping costs, and only the travel distance between the origin and destination matters. As a result, there is no benefit in consolidating traffic, and the network structure does not affect shipping expenses, conditional on travel distance. Consequently, the values of m_1 and m_2 in equation (19) will be insignificantly different from 0. On the other hand, when $\gamma \neq 0$, the network structure plays a role in traffic consolidation, thereby affecting the effects of degree and betweenness centrality on prices. Therefore, the values of m_1 and m_2 are influenced by the value of the *resource allocation efficiency* parameter, γ.

Column (2) of Table 5 shows the data moments. The average shipping price per loaded car per mile is \$0.65, comparable to the number published by the Association of American Railroads. The average price difference between an interconnecting and a non-interconnecting route in the data is \$264.43 per loaded car. As a benchmark, the average shipment price in the data is \$1034 per loaded car. Thus, the average price for joint-line service is about 26% higher than of average shipment price. The estimated impact of betweenness centrality on price is −\$0.30 in the data. Comparing a station at the 95th percentile of betweenness centrality to one at the 5th percentile, the price difference is approximately 14%.[22] Similarly,

[22]Betweenness centrality ranges from 0 to 1. In the data, the 95th percentile of betweenness centrality is 0.32, while the 5th percentile is 0.04. Consequently, the difference can be calculated as $-0.30 \times (0.32 - 0.04) =$

Table 5: Comparison of Data and Simulated Moments

	(1) Identification	(2) Data Moments	(3) Simulated Moments
Average shipping price (per loaded car per mile)	pin down δ	\$0.65	\$0.65
Average difference of price between interconnecting route and non-interconnecting route (per loaded car per mile)	pin down η	\$0.26	\$0.24
Moments related to network measures			
m_1 (coefficient of degree centrality)	pin down γ, δ	−\$0.0014	−\$0.0016
		(0.0008)	(0.0000)
m_2 (coefficient of betweenness centrality)	pin down γ, δ	−\$0.2984	−\$0.3017
		(0.0094)	(0.0083)

comparing a station at the 95th percentile to one at the 5th percentile in terms of degree centrality results in approximately a 2% difference in price.

Estimation Results. Table 6 reports the point estimates and their 95% confidence intervals. Following Andrews, Gentzkow, and Shapiro (2017), I use finite differencing to calculate standard errors for the estimated parameters.[23]

The estimated cost per efficient mile of shipment amounts to \$1.20 per carload. The interchange cost is estimated to be \$217, representing approximately 21% of the average shipment price. This finding underscores the substantial expense associated with interchange and highlights the potential cost-efficiency gains from mergers. Finally, the estimated γ significantly differs from 0, indicating that markets exhibit strong interdependence. Additionally, the fact that γ is less than 1 suggests diminishing marginal returns in investment allocation across segments.

−0.09. Relative to the average shipping price of 0.65, this change amounts to approximately a 14% difference.

[23]Here is how I chose the step size for finite difference estimates of the standard errors: I started by setting the step size h to a small percentage of each parameter's value (approximately 10^{-3}) and performed a convergence study. I gradually reduced h and checked for convergence in the derivative estimates. The estimates became noisy at around 10^{-6}, indicating rounding errors dominated. After testing various step sizes, I settled on a step size of 10^{-5} times each parameter's value.

Table 6: Estimation Results for Cost Parameters

	Point Estimate	95% Confidence Interval
δ	1.20	$[1.10, 1.29]$
η	217	$[155, 279]$
γ	0.17	$[0.14, 0.20]$

5.3 Model Validation

Column (3) of Table 5 compares the simulated moments to the data moments. In general, the simulated moments closely match the data, highlighting the model's effectiveness in fitting observed market patterns.

To further evaluate my model, I conduct an out-of-sample assessment of its predictive performance, specifically examining the accuracy in predicting price changes following mergers. In the data, the average observed price reduction after a merger is approximately 9.79%. My model's predicted average price reduction is 10.51%, closely aligning with the observed changes. Additionally, the model demonstrates strong predictive accuracy, correctly identifying approximately 63% of markets with the largest price reductions (top 10% of markets by magnitude), despite relying on only three cost parameters. This result underscores the model's capability to effectively capture the underlying economic patterns in the data.

Moreover, while actual maintenance allocation at the segment level remains unobserved, I validate the model's resource allocation predictions using observable indicators. First, I examine track abandonment patterns ($I_{j,ab} = 0$) and find the model accurately predicts post-merger abandonment locations. Second, the optimal allocation condition derived in equation (14) predicts maintenance allocation proportional to track usage, a relationship empirically confirmed by regulatory evidence from the Surface Transportation Board (STB) and Federal Railroad Administration (FRA).[24] These validation exercises collectively illustrate the robustness and predictive reliability of my model, particularly its ability to capture key resource allocation decisions within railroad firms.

[24]See the FRA's technical monograph on estimating maintenance costs for mixed high-speed passenger and freight rail corridors, available here.

6 Counterfactual Experiments

In my counterfactual analyses, I explore three key questions. First, I calculate the average gains from mergers between Class I railroads and analyze the sources of cost reductions and markup increases post-merger. Second, I examine the merger gains within a single merger, identifying which markets benefit and which suffer, and exploring the reasons behind this. Third, I compare the outcomes across various mergers, and assess the impact of different network structures of the merging parties on the extent of merger gains.

6.1 Efficiency Gains from Mergers

I calculate average merger gains by comparing equilibrium outcomes before and after each merger. Specifically, for each of the 12 mergers between Class I railroads occurring from 1985 to 2005, I simulate equilibrium outcomes pre- and post-merger and compute the resulting changes in shipment prices, costs, and markups. Figure G.1 illustrates the participating railroads and depicts the networks of the two merging parties in each case. Average merger gains are then calculated by weighting outcomes according to post-merger shipment quantities, averaged across all origin–destination markets and mergers.[25]

The results show that, on average, mergers reduce shipment costs by 12.9%, shipment prices by 8.8%, and increase markups by 7.2%. The simulated price reductions in the baseline model closely align with observed data. While mergers result in increased profitability and higher markups for the merged firms, the large efficiency gains generated also benefit consumers, who experience an average price reduction of 8.8%.

Unpacking the Black Box

The real strength of the proposed model is that it allows me to study the equilibrium effects of mergers and conduct counterfactual experiments to examine the sources of cost reductions and markup increases following a merger.

To unpack the "black box" of how the network structure is related to the level of merger gains, I conduct the following analysis. Let MG be a measure of merger gains and NC a measure of network centrality. The structural parameters are represented by θ. In the baseline experiment, I consider two network structures, NC_0 and NC_1 (pre- and post-merger),

[25] Multiple equilibria in merger simulations: The estimation of cost parameters relies only on necessary conditions of Nash equilibrium; hence, multiple equilibria would not affect the properties of the estimated parameters. However, multiple equilibria could negatively impact merger simulations. Although I cannot analytically prove uniqueness, I conduct numerical searches by varying initial values when computing post-merger equilibria, finding no instances of multiple equilibria.

calculating the merger gains as

$$MG = f(NC_1, \theta_0) - f(NC_0, \theta_0). \tag{20}$$

Here, MG represents the merger gains obtained when the network structure of the merging firms changes from NC_0 to NC_1. In my analysis, I use three measures for MG: price, cost, and markup. The function $f(\cdot, \theta)$ corresponds to the equilibrium model described in section 4. θ_0 represents the parameters estimated in section 5, which I refer to as the baseline parameters.

To examine the sources of merger gains, I conduct three counterfactual analyses by introducing different variations in θ. In the first counterfactual, denoted as $\theta_1 \equiv (\delta_0, 0, 0)$, I eliminate interchange costs and economies of scope. Here, the investigation focuses on merger gains given θ_1, calculated as:

$$MG = f(NC_1, \theta_1) - f(NC_0, \theta_1). \tag{21}$$

The intuition here is to envision a scenario with no interchange costs and no economies of scope initially. In this scenario, the benefit of the merger would come solely from shorter travel distances. We would then explore the extent of the merger gains after combining the two railroad networks.

In the second counterfactual, denoted by $\theta_2 \equiv (\delta_0, \eta_0, 0)$, I allow for interchange costs. In the third counterfactual, denoted by $\theta_3 \equiv (\delta_0, 0, \gamma_0)$, I allow for economies of scope. Together, these counterfactual analyses facilitate a comprehensive exploration of how variations in structural parameters interact with the network structure to influence merger gains.

Table 7 shows how average merger gains change in the different counterfactuals.[26] Column (1) shows the baseline results using the estimated parameters $(\delta_0, \eta_0, \gamma_0)$.

Column (2) of Table 7 presents the results from the first counterfactual, in which I account only for travel distance and set both interchange cost and economies of scope parameters to zero $(\delta_0, 0, 0)$. The results show that when considering only travel distance as a factor, cost reduces by a small amount of 1.9% on average post merger. In local markets, the merged firm gains additional market power as both the acquiring and acquired firms operate in those markets, resulting in a reduction in the number of firms following the merger. However, the results show that merger-generated concentration in local markets only increases the markup slightly, by 0.7%. This outcome is not surprising given that, out of the twelve mergers examined, jointly owned markets typically represent only 3% to 6% (with a maximum of

[26] All numbers reported in Table 7 are weighted averages, weighting by post-merger quantity for each local o–d market.

12%) of the total markets owned by the two merging firms.

Table 7: Average Merger Gains

Percentage Change in:	Baseline (1) Distance + Interchange Cost + Economies of Scope $(\delta_0, \eta_0, \gamma_0)$	Unpacking the Black Box		
		(2) Distance $(\delta_0, 0, 0)$	(3) Distance + Interchange Cost $(\delta_0, \eta_0, 0)$	(4) Distance + Economies of Scope $(\delta_0, 0, \gamma_0)$
Price	−8.8%	−1.4%	−2.8%	−3.8%
Cost	−12.9%	−1.9%	−3.4%	−7.2%
Markup	7.2%	0.7%	0.7%	6.9%

Column (3) of Table 7 shows the results of the second counterfactual, in which I now allow for interchange cost but keep economies of scope turned off. Because merger eliminates interchange cost, the post-merger cost reduction is now larger, at 3.4% compared to the results for the first counterfactual. Meanwhile, there is no difference in increase of markup, which is still 0.7%. This means that reinstating interchange cost will increase efficiency gains after merger but will have no impact on firms' market power. Nevertheless, incorporating interchange cost alone is insufficient to explain all of the cost reduction in the results; it only explains 26% of the total cost reduction found in the baseline results.[27]

Column (4) of Table 7 shows the results of the third counterfactual, in which I turn off interchange cost but allow for economies of scope. The operational decisions for each local markets have become interdependent since any traffic passing through arc (a, b) can now make use of the allocated resources at that specific arc. After the merger, the merged firm can consolidate traffic and better utilize resources. Given economies of scope, we observe a higher cost reduction at 7.2% post merger. This shows that economies of scope are the major factor in driving down shipment cost after merger. Moreover, the inclusion of economies of scope leads to a significantly higher markup increase. Specifically, after the merger, the markup increase by 6.9%, in contrast to the mere 0.7% increase observed in the first two counterfactual scenarios.

To further investigate what drives up the markup post-merger, I first examine how markup could be affected in the model. The markup for each firm j at market o–d is

[27]The baseline result shows a cost reduction of 12.9%. In the scenario where only interchange costs are considered, the cost reduction is 3.4%, amounting to 26% (calculated as 3.4/12.9) of the total cost reduction.

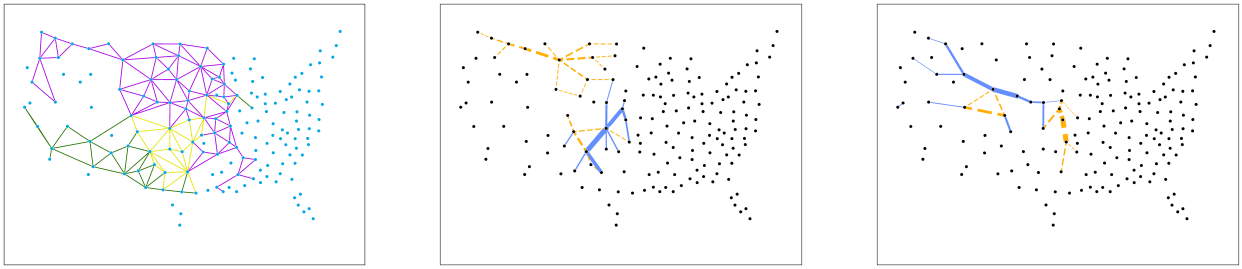
calculated from the first-order condition of the firm's optimization problem as

$$p_{j,od} - c_{j,od} = -\frac{1}{\alpha(1 - h_{j,od})} \tag{22}$$

where α is the price coefficient, and $h_{j,od}$ is the market share of firm j at market od. Equation (22) shows that the markup increases only when the local market share $h_{j,od}$ rises. One possible explanation is that the merger reduces the number of firms in local markets, thereby increasing a firm's market share in those markets. However, the first two counterfactuals show that merger-generated concentration in local markets only increases the markup slightly, by 0.7%.

I then investigate how the pricing and investment allocation decisions change for both merging and non-merging firms. The merger eliminates interchange costs and strengthens the connections of the merged firm, resulting in significant cost reductions, particularly in the interconnected areas of the two merging firms. I find that non-merging competitors reallocate their resources away from such regions in equilibrium. Assuming all other factors remain constant, when the merged firm's costs decrease in a given market, its market share in that market will increase. Consequently, the competitors' allocated resources to that market will have reduced utilization. This leads to a decrease in the competitors' marginal return on capital in that market, prompting them to reallocate their resources to other markets. When resources are moved away, the shipment cost of non-merging competitors will further increase in those markets. This results in further increases in the local market share of the merged firm. Hence, there will be a large increase in markup for the merged firm in the areas where the merged firm exhibits a large cost reduction post merger.

For example, Figure 6 shows that after the merger between Burlington Northern Railway (BN) and Santa Fe Railway (SF), the merged firm allocated more resources to improve the connectivity between the integrated sections of their network. In the west/eastbound direction, more resources were allocated to the arcs between Amarillo, TX and Omaha, NE. In the north/southbound direction, more was allocated to the arcs between South Dakota and Kansas. Panel (b) of Figure 6 shows the changes in allocation of resources for BNSF post merger. Meanwhile, in equilibrium, BNSF's primary competitor, Union Pacific Railway (UP), reallocates resources away from the South Dakota and Kansas areas and redirects them towards the Northwest corridor. Panel (c) of Figure 6 illustrates the resource allocation changes for UP subsequent to the BN and ATSF merger. As a result of all these changes, the market share of BNSF in local areas near South Dakota and Kansas will further increase.



(a) BNSF Network (b) Changes of BNSF (c) Changes of UP

Figure 6: Changes in Allocation of Resources After ATSF–BN Merger

Notes: Panel (a) shows the combined network of the two merging firms. The purple areas in the northwest represent the network solely owned by BN, while the green areas in the south represent the network solely owned by SP. The yellow areas indicate the overlapping region of the two networks. Panel (b) shows the changes in resource allocation for BNSF's network after the merger, while Panel (c) shows the changes in resource allocation for UP's network. In Panels (b) and (c), the solid blue line represents increased allocation after merger, while the dashed yellow line represents decreased allocation. The line thickness represents the magnitude of change. Changes in allocation are calculated by comparing the equilibrium allocation of resources post merger with that pre merger.

6.2 Within A Merger: Which Markets Benefit and Which Suffer?

Next, I investigate within a merger, how the extent of merger gain differs across different markets, taking into account the location of these markets in the network before and after the merger.

I derive the centrality measures for each node in the pre-merger network and in the post-merger network.[28] For instance, before the merger of the Burlington Northern railway and the Santa Fe railway, Kansas City was situated on the edge of Santa Fe's network, but afterwards, it found itself at the core of the combined network of Santa Fe and Burlington Northern. I will measure the degree centrality of Kansas City in both the pre-merger and post-merger networks, and then calculate the change in its network centrality resulting from the merger.

To summarize the findings, I perform a regression analysis where I examine the relationship between the simulated log-price changes and the changes in centrality measures after the merger.

$$\Delta \log P_{s,od} = \alpha_0 + \alpha_1 \cdot \mathbb{1}_{interchange} + \alpha_2 \cdot \Delta NC_o + \epsilon_{s,od} \tag{23}$$

[28] I use two centrality measures in my analysis. Firstly, the *degree centrality* measures the total number of links a node has in a given network, as defined in equation (17). Secondly, the *betweenness centrality* measures the number of paths travelling through each node, as defined in equation (18).

Here, $\Delta \log P_{s,od}$ represents the simulated price changes before and after the merger. The term $\mathbb{1}_{\text{interchange}}$ is an indicator function that determines whether service s is a joint-line service provided by the two merging firms. ΔNC_o measures the changes in network centrality of the node before and after the merger. Intuitively, ΔNC_0 measures the change in centrality regarding the location of each station within a firm's network. In my analysis, ΔNC is either the changes in degree centrality ΔDC or changes in betweenness centrality ΔBC.

The baseline results show that if post-merger degree centrality increases by one, shipment cost will further decrease by 0.53%. To better interpret the results, I calculate the difference of the merger gains between nodes in the 95th and 5th percentiles of changes in degree centrality. The results show that a node at the 95th percentile of changes in degree centrality has an extra 1.59 percent cost reduction and an extra 0.3 percent increase in markup post merger, compared to a node at the 5th percentile of changes in degree centrality. By comparison, a node at the 95th percentile of changes in betweenness centrality has an extra 5.17 percent cost reduction and an extra 1.17 percent increase in markup post merger compared to a node at the 5th percentile of changes in betweenness centrality.

The results suggest that when a node is positioned at the outer edges of the network before the merger and subsequently becomes more central in the post-merger network, it experiences a greater improvement in efficiency. Additionally, the baseline findings indicate that, in comparison to degree centrality, nodes that demonstrate a greater increase in betweenness centrality benefit from a larger reduction in shipment cost and an increase in markup after the merger. For example, consider the case of Kansas City, which was initially located on the periphery of Santa Fe's network but served as a significant interchange point between Burlington Northern Railway and Santa Fe Railway. The merger amplifies the importance of Kansas City in the rail network, particularly when measured by the change in betweenness centrality. As a result, the Kansas City station enjoys a significantly larger efficiency gain following the merger of Burlington Northern Railway and Santa Fe Railway.

Unpacking the Black Box

Next, I investigate why the degree and betweenness centrality measures yield different results. To do so, I first conduct the same counterfactual experiments as described in equations (20) and (21). After obtaining the equilibrium outcomes from each of these counterfactual experiments, I re-run the regression as outlined in equation (23) and obtain the estimated results. To make the results easier to interpret, I report the difference of the merger gains between nodes at the 95th and 5th percentiles of changes in a centrality measure. For example, column (1) of Table 8 shows that in the baseline model, a node at the 95th percentile of changes in degree centrality has an extra 1.59 percent cost reduction post merger compared

to a node at the 5th percentile of changes in degree centrality.

Table 8 shows the results of these calculations.[29] Panel I of Table 8 shows how changes in centrality affect cost reduction post merger. The first counterfactual turns off interchange cost and economies of scope. Column (2) of Table 8 shows the results of this counterfactual: When only distance matters, increases in both degree and betweenness centrality result in higher cost reduction post merger. Furthermore, the degree and betweenness centrality measures produce similar results in this scenario. In the second counterfactual I reincorporate interchange cost. Column (3) shows no significant change regarding the effect of changes in centrality after doing so.[30]

Table 8: Merger Gains and Centralities

	Baseline	Unpacking the Black Box		
	(1) Distance + Interchange Cost + Economies of Scope $(\delta_0, \eta_0, \gamma_0)$	(2) Distance $(\delta_0, 0, 0)$	(3) Distance + Interchange Cost $(\delta_0, \eta_0, 0)$	(4) Distance + Economies of Scope $(\delta_0, 0, \gamma_0)$
Panel I: Cost				
Δ Degree Centrality	-1.59%	-3.00%	-2.91%	-2.34%
Δ Betweenness Centrality	-5.17%	-2.67%	-2.86%	-4.73%
Panel II: Price				
Δ Degree Centrality	-2.28%	-3.36%	-3.24%	-2.85%
Δ Betweenness Centrality	-4.19%	-2.50%	-2.67%	-3.81%
Panel III: Markup				
Δ Degree Centrality	0.30%	0.03%	0.03%	0.33%
Δ Betweenness Centrality	1.17%	0.05%	0.05%	1.20%

The third counterfactual turns off interchange cost but allows for economies of scope. Column (4) shows that under this counterfactual, changes in betweenness centrality have a greater effect on cost reduction compared to the first two counterfactuals, while changes in degree centrality have a smaller effect. Betweenness centrality measures the total number of

[29] Table G.1 in Appendix G shows the full regression results of merger gains on changes in network centrality.

[30] Regression results in Table G.1 show that the coefficient of indicator of interchange becomes much larger in the second counterfactual, meaning that the effect of interchange cost is mainly absorbed by the fixed effect of interlines, leaving the effect of changes in centrality unaffected.

paths that travel through each node. Therefore, with economies of scope present, firms will concentrate resources on nodes with high betweenness centrality, thus maximizing utilization of resources. As a result, nodes with higher changes in betweenness centrality will exhibit a greater cost reduction post merger. On the other hand, nodes with higher changes in degree centrality are likely to be located on the periphery of the pre-merger network. These nodes will have better routing options and hence shorter travel distances post merger. However, resources will not necessarily be reallocated to these nodes post merger.

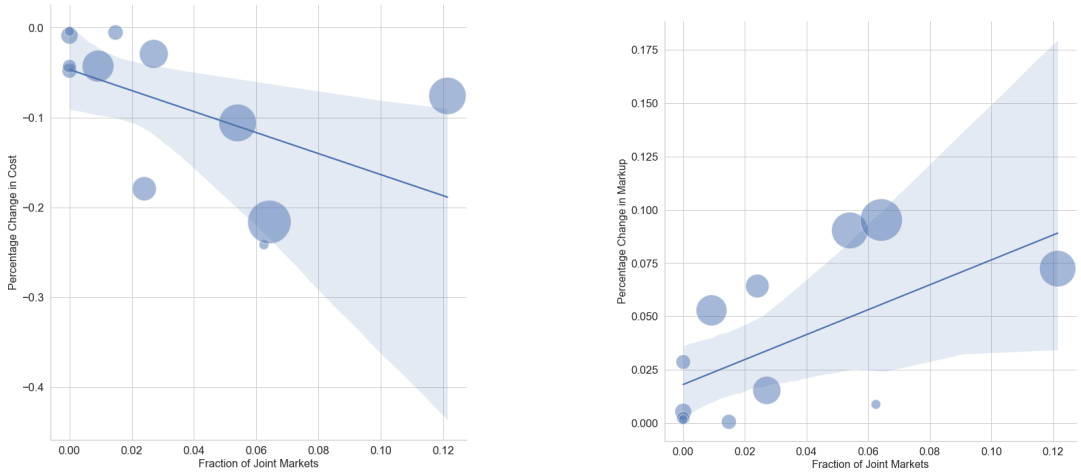
Panel III of Table 8 shows how changes in centrality affect changes in markup post merger. The results show that none of the changes in either degree or betweenness centrality have a strong impact on changes in markup in the absence of economies of scope. This is consistent with our earlier finding that merger-induced concentration in local markets results in only a very small increase in markup. Column (4) shows the results of the third counterfactual, turning off interchange cost while economies of scope remain present. We can see that changes in centrality have a much greater effect on changes in markup in this counterfactual. This is because nodes with higher betweenness centrality benefit more from reallocation of resources and hence greater cost reduction post merger when economies of scope are present. As previously explained, non-merging competitors tend to move resources away from regions where the merged firm experiences a large efficiency gain. Therefore, nodes with greater increases in betweenness centrality post merger will have greater increases in local market share and hence in markup.

6.3 Across Mergers: What Kind of Mergers Are Better?

Last, I examine the impact of different network structures of the merging parties on the extent of merger gains, comparing the outcomes across various mergers. I focus on the extent of overlap and complementarity between the two networks involved in the merger. I measure the degree of overlap by calculating the overall percentage of markets operated by both merging firms prior to the merger. As for complementarity, I count the number of joint-line services (interlines) offered by the merging firms and calculate the proportion of interlines.

Figure 7 shows the relation between merger gains and degrees of overlap. The size of each circle in the figure corresponds to the total number of origin–destination markets involved in a specific merger. Panel (a) of the figure illustrates that higher degrees of overlap lead to more substantial cost reductions after the merger. In a merging network with a higher degree of overlap, the merged firm benefits more from economies of scope by eliminating redundant lines and consolidating traffic and resources. Panel (b) demonstrates that a higher degree of

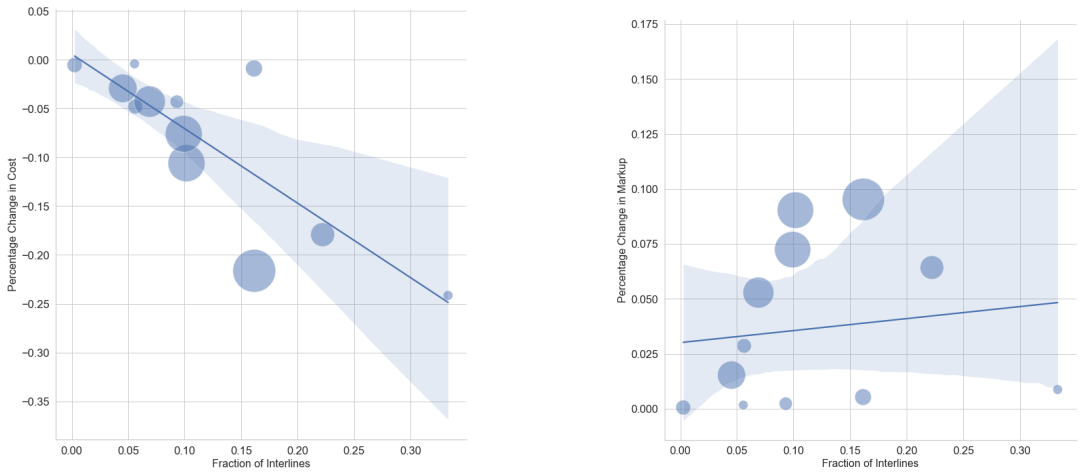
overlap also results in a larger increase in markup.



(a) Cost Changes and Degree of Overlap **(b)** Markup Changes and Degree of Overlap

Figure 7: Degree of Overlap and Average Merger Effects

Figure 8 shows the relationship between merger gains and degrees of complementarity. Panel (a) shows that higher degrees of complementarity lead to more significant cost reductions after the merger. When a merging network exhibits a higher degree of complementarity, the merged firm benefits from a greater reduction in interchange costs.



(a) Cost Changes and Degree of Complementarity (b) Markup Changes and Degree of Complementarity

Figure 8: Degree of Complementarity and Average Merger Effects

Comparing Panel (b) in Figure 8 with Panel (b) in Figure 7, we can observe that although a higher degree of complementarity also contributes to a larger increase in markup, the impact is comparatively less significant when compared to a higher degree of overlap.

Given the limited sample of twelve mergers, establishing causality is challenging. The results suggest that higher complementarity in merging networks often results in notable cost reductions and a modest increase in markup. Conversely, greater network overlap might lead to larger cost savings and a more substantial markup increase.

7 Conclusion

This article quantifies cost efficiencies resulting from mergers among U.S. freight railroads. These efficiency gains primarily stem from eliminating costly inter-railroad interchanges and realizing economies of scope by operating larger, integrated networks. To capture the equilibrium effects of such mergers, I introduce a spatial equilibrium model of oligopolistic competition that explicitly considers two key aspects: market interdependencies within a network, and strategic interactions among competing firms operating separate networks.

Applying this framework to the U.S. railroad industry demonstrates that mergers significantly reduce shipment costs and result in notable increases in markups. Counterfactual analyses further reveal that the observed markup increases are not driven mainly by reductions in the number of competing firms in local markets. Instead, these markups arise predominantly from non-merging competitors reallocating resources away from areas where merged firms achieve substantial equilibrium cost savings. Despite increased markups, these mergers ultimately pass through efficiency benefits to customers, resulting in lower shipment prices.

Additionally, the findings provide policy-relevant insights for evaluating future mergers in network-based industries. Specifically, the analysis highlights markets most likely to experience significant benefits or harms from mergers and identifies characteristics of mergers associated with greater efficiency gains. Markets represented by nodes situated at network peripheries before mergers that become more centrally located afterward tend to realize significant cost reductions and larger increases in markups, particularly when evaluated using betweenness centrality measures. Across various mergers, both a higher degree of overlap and greater complementarity between merging firms' networks are associated with substantial cost reductions. However, mergers involving extensive network overlap are also linked to notably larger markup increases.

References

Aguirregabiria, V. and C.-Y. Ho (2012). A dynamic oligopoly game of the us airline industry: Estimation and policy experiments. *Journal of Econometrics 168*(1), 156–173.

Alexandrov, A., R. W. Pittman, and O. Ukhaneva (2018). Pricing of complements in the us freight railroads: Cournot versus coase. *US Department of Justice, Antitrust Division, Economic Analysis Group Discussion Paper EAG*.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics 132*(4), 1553–1592.

Ashenfelter, O. C., D. S. Hosken, and M. C. Weinberg (2015). Efficiencies brewed: Pricing and consolidation in the us beer industry. *The RAND Journal of Economics 46*(2), 328–361.

Asker, J. and V. Nocke (2021). Collusion, mergers, and related antitrust issues. In *Handbook of Industrial Organization*, Volume 5, pp. 177–279. Elsevier.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica*, 841–890.

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

Beuthe, M., B. Jourquin, J.-F. Geerts, and C. K. à Ndjang'Ha (2001). Freight transportation demand elasticities: A geographic multimodal transportation network analysis. *Transportation Research Part E: Logistics and Transportation Review 37*(4), 253–266.

Borenstein, S. (1990). Airline mergers, airport dominance, and market power. *American Economic Review 80*(2), 400–404.

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

Brancaccio, G., M. Kalouptsidi, T. Papageorgiou, and N. Rosaia (2023). Search frictions and efficiency in decentralized transport markets. *The Quarterly Journal of Economics 138*(4), 2451–2503.

Bresnahan, T. F., S. Stern, and M. Trajtenberg (1997). Market segmentation and the sources of rents from innovation: Personal computers in the late 1980s. *The RAND Journal of Economics*, S17–S44.

Buchholz, N. (2022). Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry. *The Review of Economic Studies 89*(2), 556–591.

Ciliberto, F., E. E. Cook, and J. W. Williams (2019). Network structure and consolidation in the us airline industry, 1990–2015. *Review of Industrial Organization 54*(1), 3–36.

Clark, R. and M. Samano (2022). Incentivized mergers and cost efficiency: Evidence from the electricity distribution industry. *The Journal of Industrial Economics 70*(4), 791–837.

Collard-Wexler, A. (2014). Mergers and sunk costs: An application to the ready-mix concrete industry. *American Economic Journal: Microeconomics 6*(4), 407–447.

Dafny, L. (2009). Estimation and identification of merger effects: An application to hospital mergers. *The Journal of Law and Economics 52*.

Dafny, L., M. Duggan, and S. Ramanarayanan (2012). Paying a premium on your premium? consolidation in the us health insurance industry. *American Economic Review 102*(2), 1161–1185.

Dafny, L., K. Ho, and R. S. Lee (2019). The price effects of cross-market mergers: Theory and evidence from the hospital industry. *The RAND Journal of Economics 50*(2), 286–325.

De Chaisemartin, C. and X. d'Haultfoeuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–2996.

Degiovanni, P. and R. Y. Yang (2023). Economies of scale and scope in railroading. Technical report, Working Paper.

Demirer, M. and Ö. Karaduman (2024). Do mergers and acquisitions improve efficiency? evidence from power plants. Technical report, National Bureau of Economic Research.

Fajgelbaum, P. D. and E. Schaal (2020). Optimal transport networks in spatial equilibrium. *Econometrica 88*(4), 1411–1452.

Friebel, G., G. McCullough, and L. P. Angulo (2014). Patterns of restructuring the us class 1 railroads from 1984 to 2004. *Journal of Transport Economics and Policy 48*(1), 115–135.

Gallamore, R. E. and J. R. Meyer (2014). *American Railroads*. Harvard University Press.

Garmaise, M. J. and T. J. Moskowitz (2006). Bank mergers and crime: The real and social effects of credit market competition. *The Journal of Finance 61*(2), 495–538.

Grimm, C. and C. Winston (2000). Competition in the deregulated railroad industry: Sources, effects, and policy issues. *Deregulation of Network Industries: What's Next*, 41–71.

Harrison, E. H. (2005). *How We Work and Why*. Canadian National Railway Company.

Holmes, T. J. (2011). The diffusion of wal-mart and economies of density. *Econometrica 79*(1), 253–302.

Houde, J.-F., P. Newberry, and K. Seim (2023). Nexus tax laws and economies of density in e-commerce: A study of amazon's fulfillment center network. *Econometrica 91*(1), 147–190.

Jeziorski, P. (2014). Estimation of cost efficiencies from mergers: Application to us radio. *The RAND Journal of Economics 45*(4), 816–846.

Kim, E. H. and V. Singal (1993). Mergers and market power: Evidence from the airline industry. *American Economic Review*, 549–569.

Krebs, R. D. (2018). *Riding the Rails: Inside the Business of America's Railroads*. Indiana University Press.

Kwoka, J. E. and L. J. White (1997). Manifest destiny? the union pacific–southern pacific merger.

Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, 395–421.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica 69*(2), 307–342.

Peters, C. (2006). Evaluating the performance of merger simulation: Evidence from the us airline industry. *The Journal of Law and Economics 49*(2), 627–649.

Pittman, R. (2009). Railway mergers and railway alliances: Competition issues and lessons for other network industries. *Competition and Regulation in Network Industries 10*(3), 259–278.

Prater, M. E., K. Casavant, E. Jessup, B. Blanton, P. Bahizi, D. Nibarger, and I. Weingram (2012). Rail competition changes since the staggers act. In *Journal of the Transportation Research Forum*, Volume 49.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.

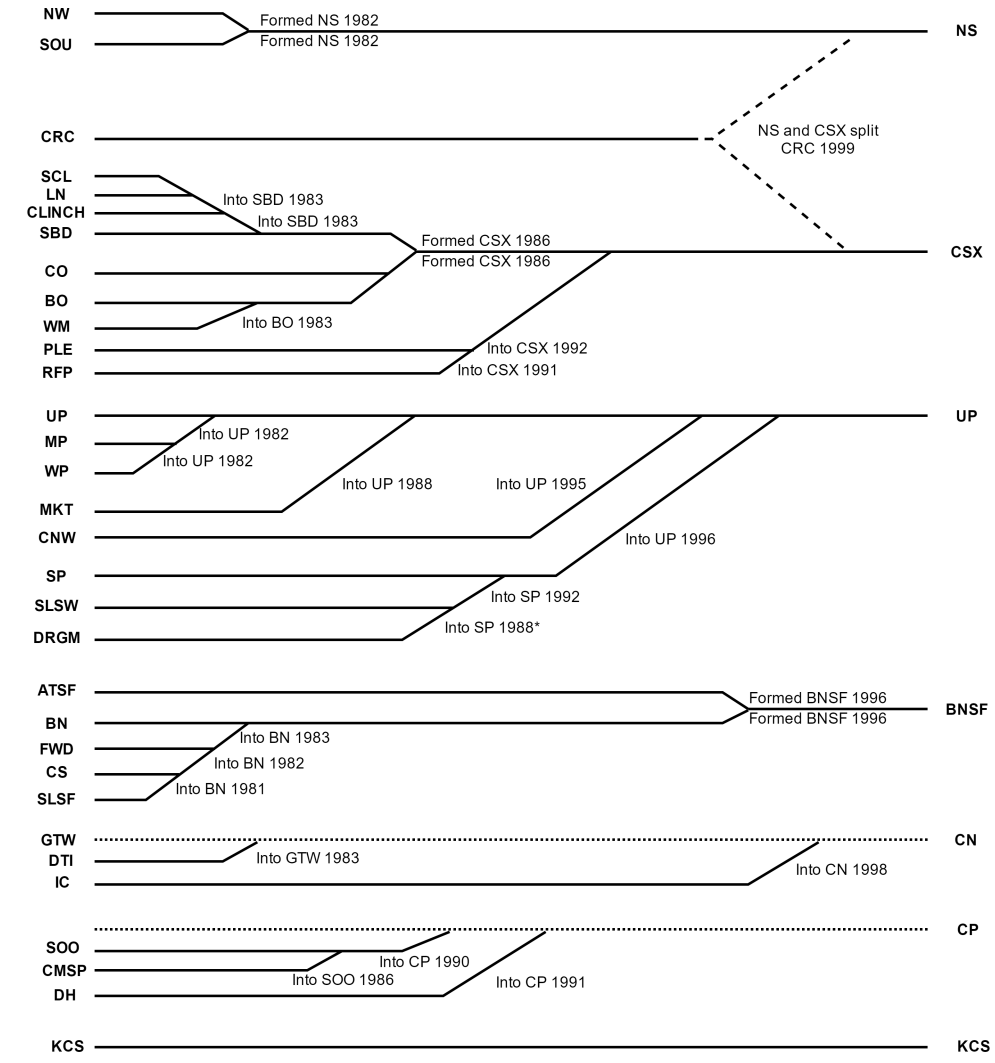
Williamson, O. E. (1968). Economies as an antitrust defense: The welfare tradeoffs. *American Economic Review 58*(1), 18–36.

Wilson, W. W., W. W. Wilson, and W. W. Koo (1988). Modal competition in grain transport. *Journal of Transport Economics and Policy*, 319–337.

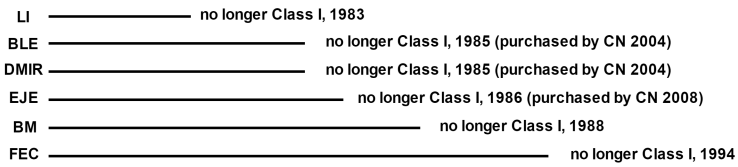
Online Appendices: Not For Publication

A History of American Railroads

Figure A.1: Merger History of Railroads



Exiting firms:



B Fieldwork and Ethnographic Observations

To cross-validate the empirical findings presented in this study, I conducted one month of fieldwork involving extensive interviews and observations across various segments of the railroad industry. This included visits with downstream customers such as steel manufacturers (EVRAZ North America, Inc.) and the UPS consolidation center in Chicago, operational observations at the BNSF classification yard and intermodal terminal, interactions with Lake State Railway Company (LSRC), a regional railroad actively engaged in interchange operations, and observations at a technology center specializing in route optimization and network management in Pueblo, Colorado. The qualitative insights consistently reinforced the quantitative results, highlighting the critical role of eliminating interchanges and optimizing routing and resource allocation as primary sources of merger-related efficiencies. Below, I document key insights from these interactions.

Figure B.1: Field Observation: Railroad Operations, Track Maintenance, and R&D



Notes: The author participating in a training session at the Transportation Technology Center in Pueblo, Colorado, focused on railroad operations, track maintenance procedures, and research and development practices.

Interview with Business Development Manager

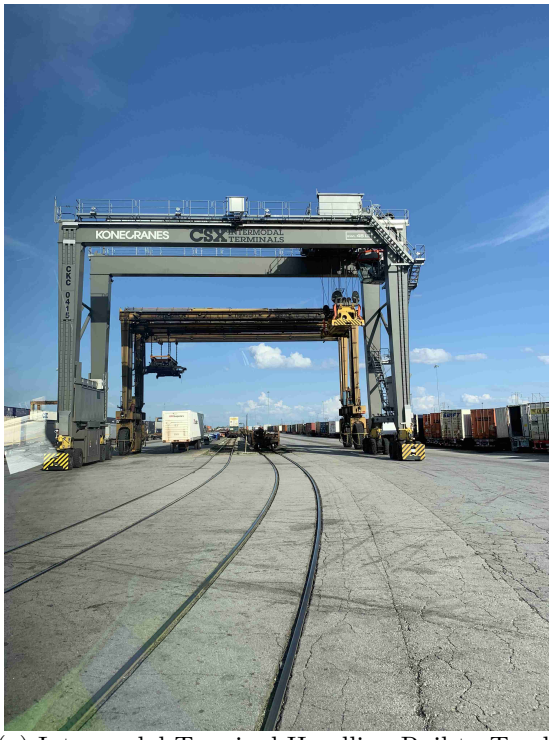
The interview illuminated two critical aspects of railroad economics: pricing decisions and interchange contracts.

First, the pricing department receives operational cost estimates for each origin–destination market from the costing department in the form of Excel files. Prices are then strategically

set to maximize profitability by applying an optimal margin constrained by market tolerance. Notably, when making pricing decisions, the pricing department does not take into account how changes in demand might subsequently affect costs and operational decisions. Instead, they rely solely on the cost figures provided by the costing department.

Second, regarding interchange contract negotiations, it usually favors the originating railroad, but specific circumstances can shift bargaining power. An example shared during the interview involved Canadian National requesting a shipment quote from a connecting railroad for transportation from Buffalo to New York. However, the connecting railroad's marketing representative insisted on quoting only from Chicago to New York to optimize their own revenue. The scenario emphasized that operational priorities and perceptions of value differ significantly across railroads, succinctly summarized by the manager's remark: "The hot stuff of one person is not the hot stuff of the other."

Figure B.2: Field Observations: Intermodal Operations and Consolidation Facilities



(a) Intermodal Terminal Handling Rail-to-Truck Transfers



(b) Customer Consolidation Center Operations

Interview with Train & Terminal Operations Manager

Interchange activities were highlighted as inherently costly due to frequent incentive misalignments and coordination problems. As a short-line railroad, LSRC regularly exchanges railcars with Class I railroads. A recurring operational inefficiency arises when partner companies, particularly during peak seasons when locomotives are in high demand, leave railcars just outside the designated interchange points, removing their own locomotives prematurely

to ensure timely availability elsewhere. Consequently, LSRC must allocate their locomotives to reposition these railcars, incurring additional operational burdens and inefficiencies.

A detailed example from industry literature reinforces the challenges of inter-railroad coordination. The article described Train 9-698-21, an express freight train primarily servicing UPS and J.B. Hunt, intended to meet stringent delivery timelines in Los Angeles. Operational difficulties arose due to misaligned incentives under a pre-merger haulage agreement between Santa Fe and Burlington Northern (BN). Although Santa Fe sold the service, BN operated the trains east of Avard without including these trains in its internal performance metrics, resulting in persistent unreliability. Eventually, customers such as J.B. Hunt were advised by Santa Fe to consider alternative arrangements until the railroads could streamline coordination efforts post-merger.

Additional qualitative support emerged from discussions with the Terminal Superintendent at Conrail, highlighting interchange-related inefficiencies stemming from divergent incentives and priorities. Daily disputes over timing, delays, and communications were frequent, exacerbated by fluctuating shipment volumes. These challenges underscored the complexities and additional operational costs inherently associated with inter-railroad coordination.

These ethnographic insights comprehensively validate the empirical findings regarding merger gains derived from the elimination of interchange costs and the reconsolidation of traffic through route reoptimization—processes that are unlikely to be successfully implemented prior to mergers due to significant operational coordination challenges and incentive misalignment among independent railroad firms.

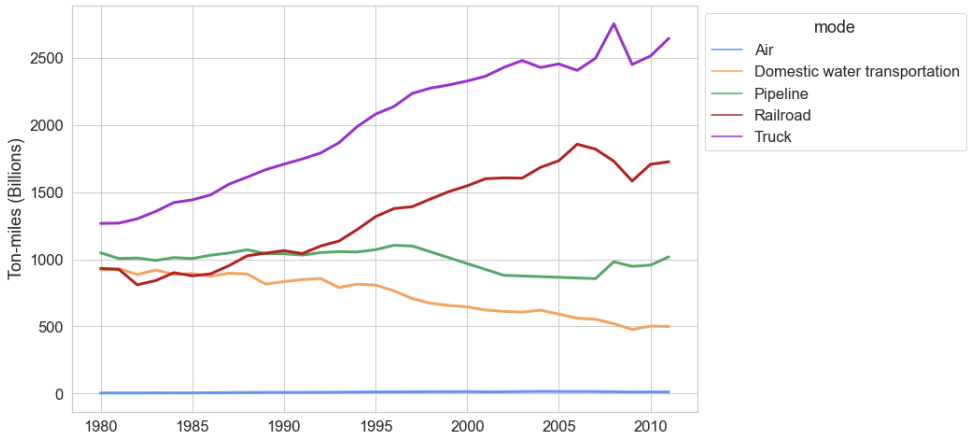
Figure B.3: Field Observations: Classification Yard Operations



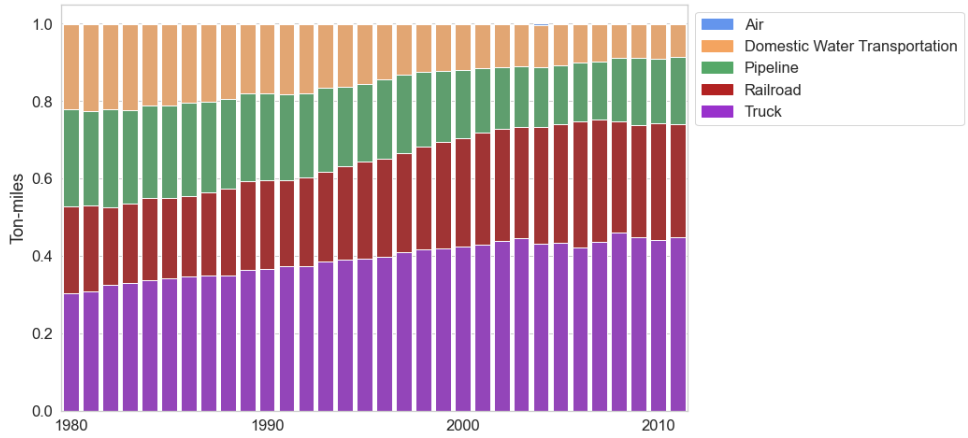
Notes: Author at the classification yard observing and learning procedures for railcar classification, interchange operations, and yard management.

C Railroad Freight vs. Other Transportation Modes

Figure C.1 illustrates the total freight transported (measured in ton-miles) by various transportation modes in the U.S. from 1980 to 2011. Initially, the railroad industry accounted for a relatively modest share, around 20%, even smaller than pipelines at the beginning of the 1980s. However, following the deregulation of 1980, the share of freight carried by railroads steadily increased, reaching approximately 33% prior to the financial crisis. According to the American Association of Railroads, railroads handle an even larger share—about 40%—when considering intercity ton-mile shipments exclusively, making it the dominant mode for long-distance freight transportation.



(a) Total U.S. Freight by Mode (Ton-miles), 1980–2011



(b) Mode Shares of Total U.S. Freight (Ton-miles), 1980–2011 (%)

Source: Bureau of Transportation Statistics

Figure C.1: Comparison of Railroad Freight and Other Transportation Modes in the U.S.

D Robustness Check for Reduced-form Analysis

To rule out alternative explanations—specifically, whether the observed price reductions might reflect competitive pressures from alternative transportation modes such as trucking—I conduct additional commodity-specific analyses. The rationale is that commodities vary significantly in their reliance on rail transportation relative to other modes. For example, the Commodity Flow Survey (CFS) of 2012 indicates that coal shipments are predominantly transported by rail (94.8%), with minimal trucking involvement (1.5%), while food products rely heavily on trucking (76.2%), with only 23.5% transported by rail. If trucking competition drove the observed price reductions, we would expect larger reductions for commodities more heavily transported by trucking, such as food, relative to commodities like coal.

Table D.1: Effect of Merger on Price Change (by Commodities)

	(1)	(2)
	Log Price	Log Price
	(Coal)	(Food or Kindred Products)
Indicator of Merger	-0.179^{***}	-0.052^{***}
	(0.028)	(0.014)
Ownership of Railcar	-0.096^{***}	-0.132^{***}
(Private)	(0.027)	(0.008)
Ownership of Railcar	-0.021	-0.144^{***}
(Trailer Train)	(0.071)	(0.016)
N	1,002,552	882,066
Firm Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes
o–d Route Fixed Effects	Yes	Yes

Standard errors in parentheses. Clustered at route level. *** $p < 0.001$

Table D.1 presents results that contradict the alternative explanation. Railroad mergers significantly reduce prices for both coal and food products, with the price reduction being considerably larger for coal. This finding is inconsistent with the hypothesis that trucking competition drives these price reductions.

Additionally, commodity-specific regressions using detailed Standard Transportation Commodity Classification (STCC) categories confirm consistent price reductions across various commodities. Notably, commodities heavily dependent on rail transport—such as coal, chemicals, and construction materials—show especially pronounced and statistically significant reductions. Detailed regression results for each commodity type are comprehensively reported in Table D.2. These findings strongly support the conclusion that the observed price

decreases primarily reflect merger-related efficiencies within the railroad industry, rather than external competitive pressures from other transportation modes.

In summary, the robustness checks presented in this appendix reinforce the main results regarding significant merger-driven efficiency gains in the railroad industry. The consistent price reductions across diverse commodity categories provide compelling evidence that these effects are driven by internal operational improvements following mergers, rather than external competitive dynamics.

Table D.2: Effect of Merger on Price Change (by Commodities)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Field Crops	Metallic Ores	Coal	Nonmetallic Minerals	Food or Kindred Products	Apparel or Textile Products	Lumber or Wood Products	Furniture or Fixtures	Pulp, Paper	Newspapers
Indicator of Merger	−0.009	0.048	−0.179***	−0.036	−0.052***	0.049	0.016	−0.023	−0.013	−0.057
	(0.014)	(0.062)	(0.028)	(0.029)	(0.014)	(0.047)	(0.013)	(0.036)	(0.015)	(0.042)
Observations	466,222	93,316	1,002,552	371,035	882,066	46,409	487,275	34,095	483,952	15,933
Number of marketID	6,982	1,086	1,360	3,697	10,766	1,210	8,145	1,694	8,441	780
Adjusted R-squared	0.178	0.144	0.251	0.234	0.266	0.767	0.274	0.829	0.299	0.667
Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
o–d Route FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
	Chemicals	Petroleum	Plastics Products	Clay, Concrete, Stone Products	Primary Metal Products	Fabricated Metal Exc.	Machinery	Electrical Machinery	Transportation Equipment	Miscellaneous Products
Indicator of Merger	−0.114***	−0.115***	−0.056**	−0.045***	−0.112***	−0.104***	−0.035	−0.146***	−0.022	0.022
	(0.031)	(0.024)	(0.028)	(0.015)	(0.027)	(0.035)	(0.037)	(0.038)	(0.035)	(0.035)
Observations	634,684	158,774	62,197	323,910	354,322	24,385	23,343	70,889	1,097,641	21,963
Number of marketID	10,462	4,175	2,087	6,670	6,639	2,009	2,072	1,996	7,177	1,163
Adjusted R-squared	0.157	0.210	0.730	0.262	0.163	0.558	0.413	0.569	0.321	0.765
Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
o–d Route FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

	(21)	(22)	(23)	(24)	(25)	(26)	(27)	(28)
	Waste or Scrap Materials	Miscellaneous Freight Shipment	Containers	Mail	Shipper Association	Miscellaneous Mixed	Small Packaged	49
Indicator of Merger	−0.000	−0.103*	0.049	0.036*	−0.179***	−0.170***	0.019	−0.072***
	(0.022)	(0.055)	(0.052)	(0.019)	(0.035)	(0.034)	(0.041)	(0.018)
Observations	341,973	60,443	660,061	43,965	48,523	3,434,108	62,487	762,745
Number of marketID	7,843	2,775	2,747	903	1,245	5,821	633	9,562
Adjusted R-squared	0.195	0.412	0.426	0.611	0.468	0.569	0.670	0.224
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
o–d Route FE	Y	Y	Y	Y	Y	Y	Y	Y
o–d Route Cluter	Y	Y	Y	Y	Y	Y	Y	Y

E Model Limitations and Extensions

E.1 Discussion of Demand Modeling Assumptions

In my demand specification, I assume that customer choices depend on shipping prices and track mileage (which proxies for the ease of transferring goods to railroads), but do not explicitly depend on trip duration. In practice, interchanges between railroad companies not only incur monetary costs but may also extend the shipment duration, potentially impacting customer decisions.

Ideally, one might consider a model where demand directly depends on trip duration. However, due to data constraints—specifically, the absence of recorded shipment durations—I adopt a simplified approach that models the cost associated with interchanges as monetary rather than temporal.

Despite this simplification, the implications for consumer welfare remain conceptually consistent with a duration-based model. Specifically, in my model, reduced interchange costs following a merger translate directly into lower shipment prices and higher consumer welfare. In contrast, in a duration-focused model, a reduction in interchange costs might not always yield lower prices directly, but it would typically reduce shipment durations, similarly enhancing consumer welfare. Hence, while the predicted price changes might differ across these modeling approaches, their implications for consumer welfare are largely aligned.

Thus, the current modeling choice provides a conservative estimate of consumer welfare benefits, acknowledging that incorporating duration explicitly might further enhance the estimated benefits arising from mergers.

E.2 Congestion and Other Scale Economies

In my model, the arc-level transportation cost is defined as:

$$c_{j,ab} = \begin{cases} \frac{\delta Dist_{j,ab}}{I_{j,ab}^{\gamma}} & \text{if} \quad I_{j,ab} > 0 \\ \infty & \text{if} \quad I_{j,ab} = 0 \end{cases},$$

When $I_{j,ab} = 0$, the transportation on the arc (a, b) is assumed to be impossible, resulting in an infinite transportation cost.

This specification is equivalent to setting $\tau = 0$ in equation (E.1):

$$c_{j,ab} = \frac{\delta Dist_{j,ab}}{I_{j,ab}^{\gamma}} \cdot Q_{j,ab}^{\tau}, \tag{E.1}$$

where $Q_{j,ab}$ is the quantity shipped through arc (a, b). If $\tau > 0$, we have congestion; if $\tau < 0$, we have economies of scale. As explained in Section 4.2, I set $\tau = 0$ to simplify computation, thereby keeping the optimal routing problem as a linear programming problem.

If incorporating economies of scale (or congestion) becomes necessary, I can use another method that does not impose the parametric form as in equation (E.1), thereby allowing me to continue framing the routing as linear programming problems. Interviews with practitioners suggest that they address congestion by solving optimization decisions layer by layer.

For instance, they prioritize the main route from St. Louis to Chicago, determining the optimal routing for this key origin–destination first. Subsequently, they plan the next most important route, such as from New Orleans to Chicago, taking into account the congestion already introduced by the St. Louis to Chicago route. This process continues sequentially until all routes are optimized. I can replicate this layer-by-layer optimization approach in my model to incorporate economies of scale and congestion.

Specifically, in my current model the optimal routing decision is equivalent to finding the least expensive route to travel from origin o to destination d for each o–d market. This can be formulated as a linear-programming problem:

$$\min_{\mathcal{R}_{j,od}} \sum_{(a,b)\in\mathcal{R}_{j,od}} \frac{\delta Dist_{j,ab}}{I^{\gamma}_{j,ab}}$$

such that $\forall m' \in \mathcal{Z}_j$,

$$\mathbb{1}\{m' = o\} - \mathbb{1}\{m' = d\} + \sum_{a\in\mathcal{Z}_j(m')} \mathbb{1}\{(a,m') \in \mathcal{R}_{j,od}\} \leq \sum_{b\in\mathcal{Z}_j(m')} \mathbb{1}\{(m',b) \in \mathcal{R}_{j,od}\}.$$

The optimal routing problem is then solved through LP algorithms, which operate within the feasible region defined by the constraints, iteratively moving towards the optimal solution using linear algebra techniques.

Directly incorporating economies of scale or congestion parametrically as in equation (E.1), causes the optimal routing problem to no longer be a linear programming problem, because the marginal cost now varies with shipment quantity $Q_{j,ab}$ and is no longer constant. However, this challenge can be addressed by employing a layer-by-layer optimization approach. Specifically, for a target shipment from origin o to destination d, the quantity travelled through arc (a,b) is $Q_{j,ab}$, while the quantity already traveling trough arc (a,b) from previous layers is $Q'_{j,ab}$. Thus, the arc-level cost can be rewritten as

$$\begin{aligned} c_{j,ab} &= \frac{\delta Dist_{j,ab}}{I^{\gamma}_{j,ab}} \cdot \left(\sum_{(o',d')\neq(o,d)} q_{j,ab}(o',d') \right)^{\tau} \\ &= \frac{\delta Dist_{j,ab}}{I^{\gamma}_{j,ab}} \cdot (Q'_{j,ab})^{\tau} \end{aligned} \tag{E.2}$$

In this setup, changes in the routing decision for a specific shipment from o to d affect the quantity traveling through arc (a,b) for that particular shipment but not the quantities previously determined by earlier layers. This method allows us to maintain the optimal routing as a linear programming problem because, conditional on the quantities from previous layers, the arc-level cost of the shipment remains constant.

The main drawback of this approach lies in the subjective decisions regarding the sequence in which routes are optimized in the layer-by-layer process. In practice, operators typically prioritize routes based on shipment volume and customer needs. If I were to optimize based on the sequence of shipment quantities, my results are unlikely to differ significantly from the current model. This is because, for the heaviest traffic flows that are optimized first,

the arc-level cost remains identical to that specified in my existing model. For subsequent layers optimized later, while there may be slight variations in routing options, their smaller quantities mean the overall impact on optimal decisions would be minimal.

E.3 Balanced-flow Constraint

In modeling railroad routing and resource allocation decisions, I abstract from the movement of empty railcars. Specifically, my model captures shipments from node a to node b but does not explicitly consider the return trips of empty railcars from b back to a. A potential consequence of this simplification is neglecting price discounts often observed in the industry for return shipments on routes with significant traffic imbalances—known as 'backhauling.' While origin–destination market fixed effects partially capture this phenomenon in the optimal pricing decision, aggregate traffic flow patterns may shift significantly in equilibrium under counterfactual scenarios, potentially altering regions with imbalanced flows. This equilibrium shift and its implications for backhauling are not explicitly modeled here.

Furthermore, the balanced-flow constraint is defined at the origin–destination level rather than at a more aggregated station or regional level. An aggregate-level modeling approach, allowing regions more flexibility in selecting commodity sources, would require developing a comprehensive general equilibrium model, including regional supply decisions. Given the current model's complexity and available data, I chose not to pursue this approach. Additionally, many of the customers, such as power plants, operate under long-term contracts, and manufacturing plants typically face high relocation costs and delays in adjusting their sourcing decisions.

Therefore, the results presented should be interpreted as the implied welfare and efficiency gains conditional on current locations and established sourcing arrangements. These estimates thus represent a conservative lower bound for the potential welfare implications if broader market adjustments were explicitly considered.

E.4 Centralized Decision Making

Lagrangian. Given the optimization problem, the lagrangian for firm j is

$$\mathcal{L}_j(\boldsymbol{p_j}, \boldsymbol{\mathcal{R}_j}, \boldsymbol{I_j}) = \sum_{s:s\in S(j)} \left(p_{s,od} - C_{s,od}(R_{j,o_j(s)d_j(s)}, \boldsymbol{I_j})\right) \cdot Q_{s,od} + \lambda_j \left(\sum_{(a,b)\in\mathcal{A}_j} I_{j,ab} - K_j \right) + \sum_{s:s\in S(j), z\in\mathcal{Z}_j} \lambda_{j,s,z} \tag{E.3}$$
$$\left(D_{j,z} + \sum_{a\in\mathcal{Z}_j(z)} Q_{s,od} \cdot \mathbb{1}\{(a,z) \in \mathcal{R}_{j,o_j(s),d_j(s)}\} - \sum_{b\in\mathcal{Z}_j(z)} Q_{s,od} \cdot \mathbb{1}\{(z,b) \in \mathcal{R}_{j,o_j(s),d_j(s)}\} \right).$$

The vectors $\boldsymbol{p_j}$, $\boldsymbol{\mathcal{R}_j}$, and $\boldsymbol{I_j}$ represent the decisions of firm j for prices, routing, and resource allocation, respectively. $C_{s,od}$ is the marginal cost of transportation and λ's are Lagrangian multipliers.

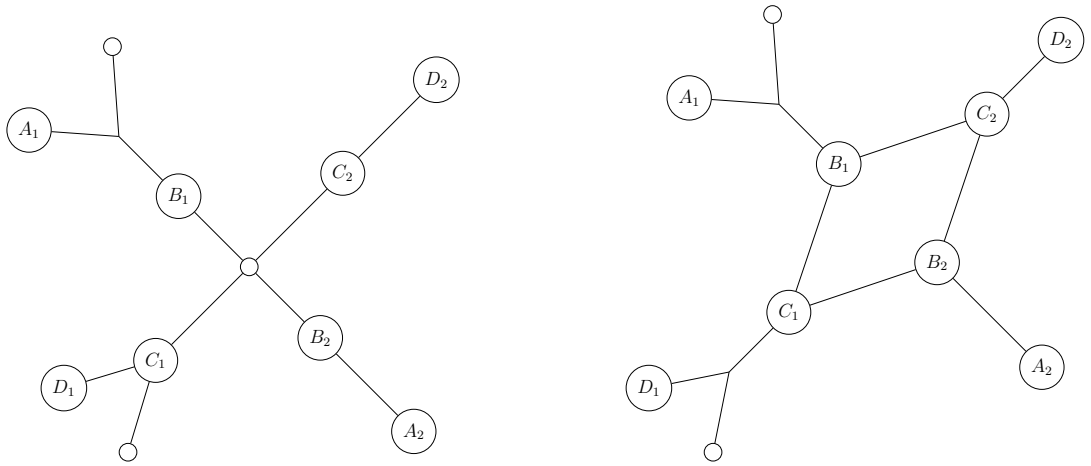
In the model, I assume that when railroad firms set prices in local markets, they do not consider how changes in quantity would affect routing and allocation decisions. This

assumption translates to setting the own-cost and cross-cost effects to zero in the first-order-condition with respect to price:

$$Q_{s,od} + p_{s,od} \cdot \frac{\partial Q_{s,od}}{\partial p_{s,od}} - \frac{\partial Q_{s,od}}{\partial p_{s,od}} \cdot \left[C_{s,od} + \underbrace{\frac{\partial C_{s,od}}{\partial Q_{s,od}} Q_{s,od}}_{\text{own-cost effect}} + \underbrace{\sum_{s' \in S(j), s' \neq s} \frac{\partial C_{s',o'd'}}{\partial Q_{s,od}} Q_{s,o'd'}}_{\text{cross-cost effects}} \right] = 0.$$

To investigate the impact of this assumption on equilibrium results, I derived solutions for a monopoly by incorporating the own- and cross-cost effects, which I refer to as the full model. In contrast, the benchmark model has the own- and cross-cost effects set to zero.

Below I compare the equilibrium results of the two models. There are three scenarios: no economy of scope ($\gamma = 0$), economy of scope ($\gamma \neq 0$) and the topology of the network is a tree, and economy of scope and non-tree network topology. The distinction between a tree and a non-tree network topology is illustrated in Figure E.1, where panel (a) depicts a tree topology, while panel (b) shows a non-tree topology in which B_1 — C_1 — B_2 — C_2 — B_1 forms a loop. The primary difference between the two is the number of routing options available. In a tree network topology, only one routing option exists for each o–d market.



(a) Tree (b) Non-Tree

Figure E.1: Examples of Tree and Non-Tree Topologies

Proposition 1 *When there is no economy of scope ($\gamma = 0$), the optimal decision for each o–d market is independent, and the benchmark model is equivalent to the full model.*

Proof: The marginal cost of service s from o to d is specified as

$$C_{s,od}(\boldsymbol{I}_j, \mathcal{R}_{j,od}) = \sum_{(a,b)\in\mathcal{R}_{j,od}} \frac{\delta_0 Dist_{j,ab}}{I_{j,ab}^{\gamma}}$$
$$= \sum_{(a,b)\in\mathcal{R}_{j,od}} \delta_0 Dist_{j,ab} \qquad \text{if } \gamma = 0. \tag{E.4}$$

When $\gamma = 0$, maintenance allocation $\boldsymbol{I}$ no longer affects cost, and the optimal routing is to choose the shortest o–d path for each service s. Hence, the routing decision for each service s is independent. From equation E.4, we can see that the choice of optimal routing depends only on the travel distance between each arc (a, b) and is irrelevant to the quantity being shipped. Therefore, $\frac{\partial C(\boldsymbol{Q},\boldsymbol{R},\boldsymbol{I})}{\partial Q_{s,od}} = \sum_{(a,b)\in\mathcal{R}^*_{j,od}} \delta_0 Dist_{j,ab} = C^{**}_{s,od}$. The second part of the equation holds because the optimal routing for firm j from o to d is the same in the two models when $\gamma = 0$. Therefore, both models' optimal strategies are equivalent when there is no economy of scope ($\gamma = 0$).

Next I discuss the case when $\gamma \neq 0$. When the topology of a network is a tree, there is only one routing option between each o–d market. Therefore, $R^*_{j,od}$ is irrelevant to the choice of prices and the allocation decision.

Lemma 1 *For any service s in market o–d, the own-cost effect $\frac{\partial C_{s,od}}{\partial Q_{s,od}} < 0$. The cross-cost effect for market m–n, $\frac{\partial C_{s,mn}}{\partial Q_{s,od}} > 0$ if $\mathcal{R}_{s,od} \cap \mathcal{R}_{s,mn} = \emptyset$. Otherwise the sign of the cross-cost effect is uncertain.*

Proof: Equation E.4 shows that $C_{s,od}$ is inversely proportional to maintenance allocation $I_{j,ab}, (a, b) \in \mathcal{R}_{j,od}$. Therefore, the own- or cross-cost effect narrows down to how firms allocate the fixed amount of resources $I_{j,ab}$, where $\sum_{(a,b)\in\mathcal{A}_j} I_{j,ab} \leq K_j$. Based on the Kuhn-Tucker conditions in 13, we know that for any non-zero $I_{j,ab}$ and $I_{j,a'b'}$,

$$\frac{I_{j,ab}}{I_{j,a'b'}} = \left[\frac{Dist_{j,ab} \cdot q_{j,ab}}{Dist_{j,a'b'} \cdot q_{j,a'b'}}\right]^{\frac{1}{1+\gamma}}. \tag{E.5}$$

where $q_{j,ab}$ is the total amount of shipment running through arc (a, b), $q_{j,ab} = \sum_{s:s\in S(j)} Q_{s,od} \cdot \mathbb{1}\{(a, b) \in \mathcal{R}_{j,o_j(s),d_j(s)}\}$. Therefore, equation E.5 shows that the amount of allocation to arc (a, b) is proportional to the traffic going through arc (a, b). Therefore, when $Q_{s,od}$ increases, $\forall(a, b) \in \mathcal{R}_{j,od}, I_{j,ab}$ increases and hence $C_{s,od}$ decreases. That is, the own-cost effect $\frac{\partial C_{s,od}}{\partial Q_{s,od}} < 0$. For market m to n, if m to n uses a totally different route than o to d (i.e. $\mathcal{R}_{j,od} \cap \mathcal{R}_{j,mn} = \emptyset$), then resources will be allocated away, yielding a higher shipment cost from m to n. Consequently, $\frac{\partial C_{s,mn}}{\partial Q_{s,od}} > 0$. However, if $\mathcal{R}_{j,od} \cap \mathcal{R}_{j,mn} \neq \emptyset$, m to n will partially benefit from the reallocation to arcs (a, b) where $(a, b) \in \mathcal{R}_{j,od} \cap \mathcal{R}_{j,mn}$. Hence, the net effect on $C_{s,mn}$ is ambiguous if $\mathcal{R}_{j,od} \cap \mathcal{R}_{j,mn} \neq \emptyset$.

Lemma 2 *When the demand of one single market dominates ($Q_{mn} \gg Q_{od}, \forall od \neq mn$), $I_{j,ab} \approx 0$ if $(a, b) \notin \mathcal{R}_{j,mn}$, and $I_{j,ab} \approx K_j \cdot \frac{Dist_{j,ab}^{\frac{1}{1+\gamma}}}{\sum_{(a',b')\in\mathcal{R}_{j,mn}} Dist_{j,a'b'}^{\frac{1}{1+\gamma}}}$ if $(a, b) \in \mathcal{R}_{j,mn}$.*

Proof: From equation E.5, we know that

$$\begin{aligned}\frac{I_{ab}}{I_{a'b'}} &= \left[\frac{Dist_{j,ab}\cdot q_{j,ab}}{Dist_{j,a'b'}\cdot q_{j,a'b'}}\right]^{\frac{1}{1+\gamma}} \\ &= \left[\frac{Dist_{j,ab}\cdot \sum_{s:s\in S(j)} Q_{s,od}\cdot \mathbb{1}\{(a,b)\in\mathcal{R}_{j,od}\}}{Dist_{j,a'b'}\cdot \sum_{s:s\in S(j)} Q_{s,od}\cdot \mathbb{1}\{(a',b')\in\mathcal{R}_{j,od}\}}\right]^{\frac{1}{1+\gamma}}. \end{aligned} \tag{E.6}$$

If $Q_{mn} \gg Q_{od}, \forall od \neq mn$, from equation E.6 we know that

$$\frac{I_{ab}}{I_{a'b'}} \approx 0$$

$\forall (a,b) \notin \mathcal{R}_{j,mn}$ and $(a',b') \in \mathcal{R}_{j,mn}$. Similarly, if $(a,b) \in \mathcal{R}_{j,mn}$ and $(a',b') \in \mathcal{R}_{j,mn}$, then $q_{ab} \approx q_{a'b'} \approx Q_{mn}$. Therefore,

$$\frac{I_{a'b'}}{I_{ab}} \approx \frac{Dist_{j,a'b'}^{\frac{1}{1+\gamma}}}{Dist_{j,ab}^{\frac{1}{1+\gamma}}}.$$

Because $\sum_{(a,b)} I_{ab} = K$, we have $I_{j,ab} \approx K \cdot \frac{Dist_{j,ab}^{\frac{1}{1+\gamma}}}{\sum_{(a',b')\in\mathcal{R}_{j,mn}} Dist_{j,a'b'}^{\frac{1}{1+\gamma}}}$ if $(a,b) \in \mathcal{R}_{j,mn}$ and $I_{j,ab} \approx 0$ if $(a,b) \notin \mathcal{R}_{j,mn}$.

Proposition 2 *When there is economy of scope ($\gamma \neq 0$) and the topology of the network is a tree, if the demand of a single market dominates, then the difference between the benchmark model and the full model is negligible. Otherwise, the direction of the bias of the benchmark model is ambiguous, and it depends on the level of demand in each origin–destination market.*

Proof: From Lemma 2, we know that when one market dominates, all the resources will be allocated to minimize the shipment cost for that dominant market in both models. Given that the network is a tree, the optimal routing decisions $\boldsymbol{R}_j^*$ are irrelevant to the choice of prices and allocation decisions. Therefore, the own- and cross-cost effect in the full model will be close to zero because the marginal change in quantities will have minimal effect on the routing and allocation decision. Hence, the optimal pricing decisions will essentially be identical in both models. Therefore, when a single market dominates, the difference between the benchmark model and the full model is negligible. From Lemma 1, we can see that the own-cost effect is negative while the cross-cost effect is mostly positive. Hence, the net effect and therefore the difference in pricing in the full and benchmark models are ambiguous.

Next, I discuss the case when $\gamma \neq 0$ and when the topology of a network is non-tree. In this scenario, the routing decision becomes non-trivial and has a significant impact on the allocation of resources. With a fixed set of routing decisions, we can determine the optimal allocation and pricing decisions. Since the set of origin-destination markets is finite, there is a limited number of possible routing options through those markets. Therefore, we can evaluate all possible routing decisions and identify the optimal solutions.

Lemma 3 *When the demand of one single market dominates ($Q_{mn} \gg Q_{od}, \forall od \neq mn$), $I_{j,ab} \approx 0$ if $(a,b) \notin \mathcal{R}_{j,mn}$, and $I_{j,ab} \approx K \cdot \frac{Dist_{j,ab}^{\frac{1}{1+\gamma}}}{\sum_{(a',b')\in\mathcal{R}_{j,mn}} Dist_{j,a'b'}^{\frac{1}{1+\gamma}}}$ if $(a,b) \in \mathcal{R}_{j,mn}$. The routing of m to n is obtained by solving the shortest travel distance between m and n in both models.*

Proof: We first show that Lemma 2 still holds when the topology of the network is non-tree. Given a set of routing decisions $\boldsymbol{\mathcal{R}}_{(\cdot)}$, the optimal allocation decisions in both models will be solved through

$$I_{j,ab} = \left[\frac{\gamma}{\lambda_j} \cdot \delta_0 Dist_{j,ab} \cdot q_{j,ab}\right]^{\frac{1}{1+\gamma}}$$

and, following the proof in Lemma 2, we can easily show that when the demand of one single market dominates ($Q_{mn} \gg Q_{od}, \forall od \neq mn$), $I_{j,ab} \approx 0$ if $(a,b) \notin \mathcal{R}_{j,mn}$, and $I_{j,ab} \approx K_j \cdot \left(Dist_{j,ab}^{\frac{1}{1+\gamma}} / \sum_{(a',b')\in\mathcal{R}_{j,mn}} Dist_{j,a'b'}^{\frac{1}{1+\gamma}}\right)$ if $(a,b) \in \mathcal{R}_{j,mn}$. Given the solutions to the optimal allocation decisions, the shipment cost can be rewritten as

$$\begin{aligned} C_{s,od}(\boldsymbol{I}_j) &= \sum_{(a,b)\in\mathcal{R}^*_{j,od}} \frac{\delta_0 Dist_{j,ab}}{I^{\gamma}_{j,ab}} \\ &= \sum_{(a,b)\in\mathcal{R}^*_{j,od}} \delta_0 Dist_{j,ab}^{\frac{1}{1+\gamma}} \sum_{(a',b')\in\mathcal{R}_{j,mn}} Dist_{j,a'b'}^{\frac{\gamma}{1+\gamma}}. \end{aligned}$$

The last equality is obtained by subtracting the solutions of $I_{j,ab}$ into the equation. Without loss of generality, we can assume that $Dist_{ab} = Dist_{a',b'}, \forall (a',b')$. Therefore, we have

$$C_{s,od}(\boldsymbol{I}_j) = \delta_0 Dist_{j,ab}^{\frac{1}{1+\gamma}} \cdot (\#\mathcal{R}_{od})^2$$

where $\#\mathcal{R}_{od}$ stands for the total number of arcs gone through by routing $\mathcal{R}_{od}$. Then to minimize shipment cost, the optimal routing decision is to find the shortest travel distance between o to d, which is irrelevant to the allocation and pricing decision.

Proposition 3 *When there is economy of scope ($\gamma \neq 0$) and the topology of the network is non-tree, if the demand of a single market dominates, then the difference between the benchmark model and the full model is negligible. Otherwise, the direction of the bias of the benchmark model is ambiguous and depends on the level of demand in each origin–destination market.*

Proof: Similar to the proof in Proposition 2, based on Lemmas 1 and 3 the equilibrium results of the full and benchmark models will be approximately the same when one market dominates. Otherwise, the direction of the bias is ambiguous.

Propositions 2 and 3 demonstrate that when there is an economy of scope ($\gamma \neq 0$) and no single market dominates, the difference between the benchmark model and the full model varies depending on the demand level in each origin-destination market, and the direction of the benchmark model's bias is uncertain. Therefore, I conduct numerical simulations to

gain a better understanding of the magnitude of the difference in scenarios where no single market dominates.

In the numerical simulation, I assume a logit demand, the distance of all the arcs equals 1, and the parameter of the economy of scope $\gamma = 1$. Therefore, the shipment cost for each o–d market is

$$c_{s,od}(\boldsymbol{I}_j) = \sum_{(a,b)\in\mathcal{R}^*_{j,od}} \frac{1}{I_{j,ab}}.$$

I also assume that the total capital equals the number of edges in the numerical example. The constraint on allocation is that $\sum I_{j,ab} = K$. I derive the analytical solutions of the optimal strategies for the benchmark model and the full model for the tree and non-tree network topology cases. Based on the analytical solutions, I calculate the optimal strategies for the benchmark model and the full model in various situations. The main parameter I change in the simulation is the total mass of demand M_{od} for each o–d market.

Table E.1: Equilibrium Prices and Costs under Tree

Market	**Prices**			**Costs**		
	Benchmark Model	Full Model	Difference	Benchmark Model	Full Model	Difference
Panel I: $M_{od} = 1$ **for all four markets**						
$A_1 \to A_2$	5.56	5.56	1E-05	3.22	3.22	9E-06
$B_1 \to B_2$	3.61	3.61	-3E-06	0.71	0.71	-4E-06
$C_1 \to C_2$	3.61	3.61	-3E-06	0.71	0.71	-4E-06
$D_1 \to D_2$	5.56	5.56	1E-05	3.22	3.22	9E-06
Panel II: $M_{A_1A_2} = 100, M_{od} = 1$ **for all the other markets**						
$A_1 \to A_2$	4.17	4.17	3E-07	1.50	1.50	-4E-08
$B_1 \to B_2$	3.46	3.46	-9E-07	0.50	0.50	-1E-08
$C_1 \to C_2$	NA	NA	NA	NA	NA	NA
$D_1 \to D_2$	NA	NA	NA	NA	NA	NA

Numerical Simulations: Topology of the Network is a Tree

In this numerical exercise, I assume that the topology of the network is the same as illustrated in panel (a) of Figure E.1. Assume that there are four origin–destination markets, $A_1 \to A_2$, $B_1 \to B_2$, $C_1 \to C_2$, and $D_1 \to D_2$.

Table E.1 compares the equilibrium price and cost results between the two models. In Panel I of Table E.1 when all four markets have an equal total mass of demand ($M_{od} = 1, \forall od$), the prices of $A_1 \to A_2$ and $D_1 \to D_2$ are higher than the prices of $B_1 \to B_2$ and $C_1 \to C_2$. This is because the former two markets have longer travel distances and more resources are allocated to (B_1, B_2) and (C_1, C_2), yielding lower cost for the latter two markets. In Panel II when the demand of $A_1 \to A_2$ dominates ($M_{A_1A_2} = 100$), we know that all resources are allocated to (A_1, B_1), (B_1, B_2), and (B_2, A_2). Hence, the cost of $C_1 \to C_2$ and $D_1 \to D_2$ is close to infinity and marked as NA. Equivalently, markets from $C_1 \to C_2$ and $D_1 \to D_2$ will not be served. Table E.1 shows that the difference in the equilibrium prices and costs between the benchmark model and the full model is very small.

Table E.2: Equilibrium Prices and Costs under Non-Tree

Market	Prices			Costs		
	Benchmark Model	Full Model	Difference	Benchmark Model	Full Model	Difference
Panel I: $M_{od} = 1$ **for all three markets**						
$A_1 \to C_1$	4.09	4.09	0.00	1.38	1.38	0.00
$A_2 \to C_1$	3.91	3.91	0.00	1.14	1.14	0.00
$C_1 \to C_2$	3.91	3.90	0.01	1.14	1.14	0.00
Panel II: $M_{A_1C_1} = 10, M_{od} = 1$ **for all the other markets**						
$A_1 \to C_1$	3.58	3.58	0.00	0.68	0.68	0.00
$A_2 \to C_1$	5.52	5.51	0.01	3.17	3.16	0.01
$C_1 \to C_2$	4.21	4.21	0.00	1.55	1.55	0.00
Panel II: $M_{A_2C_1} = 10, M_{od} = 1$ **for all the other markets**						
$A_1 \to C_1$	5.52	5.51	0.01	3.17	3.17	0.01
$A_2 \to C_1$	3.58	3.58	0.00	0.68	0.68	0.00
$C_1 \to C_2$	4.21	4.19	0.02	1.55	1.55	0.00

Numerical Simulations: Topology of the Network is Non-Tree

Here I assume that the topology of the network is the same as illustrated in Panel (b) of Figure E.1. Assume that there are three origin–destination markets, $A_1 \to C_1$, $A_2 \to C_1$, and $C_1 \to C_2$. Table E.2 compares the results of equilibrium prices and costs between the two models. In Panel I, all three markets have an equal total mass of demand ($M_{od} = 1, \forall od$). In Panel II, the market from A_1 to C_1 dominates, and in Panel III, the market from A_2 to C_1 dominates. We can see that in all scenarios, the difference between the equilibrium prices and costs between the benchmark and the full model is very small.

In conclusion, the results of the numerical simulations show that the difference in equilibrium outcomes between the full and benchmark models is minimal, regardless of whether the network topology is a tree or non-tree.

F Comparative Statistics

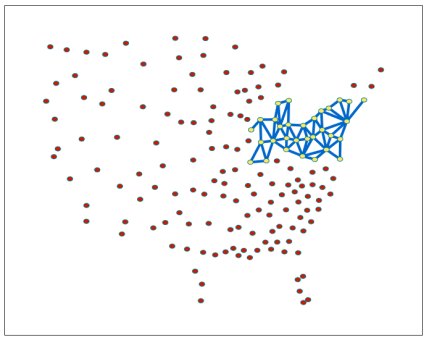
Figure F.1 illustrates how routing and resource allocation decisions of a hypothetical railroad firm vary with different values of the scope parameter, γ. Panels (a), (c), and (e) depict the resource allocation decisions, with line thickness representing the quantity of resources allocated. Panels (b), (d), and (f) illustrate the optimal routing decisions from origin 3 to destination 7 as the value of γ changes.

The comparative statics are presented for three scenarios: $\gamma = 0$, 0.5, and 0.85. When $\gamma = 0$, there are no economies of scope—resource allocation does not influence routing decisions, rendering markets independent. Consequently, resources are evenly distributed across all arcs owned by Conrail (Panel a), and the optimal route from node 3 to 7 simply follows the shortest path ($3 \to 10 \to 7$, Panel b).

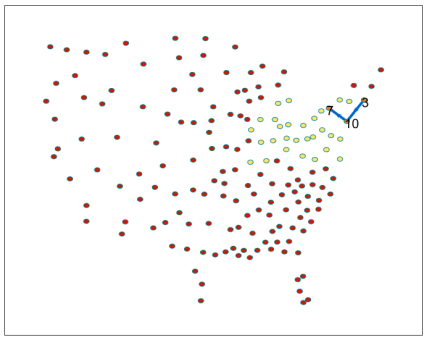
As γ becomes positive, economies of scope emerge, making it cost-effective to consolidate traffic and focus resources on high-volume routes. Comparing Panel (c) with Panel (a), and Panel (e) with Panel (c), we observe increasing consolidation of resources along major routes as γ rises.

Correspondingly, routing decisions adapt to leverage the improved cost efficiency provided by resource-intensive "major" routes (analogous to choosing highways over country roads). The routing shifts from the direct shortest path in Panel (b) to more consolidated paths: $3 \to 10 \to 8 \to 7$ in Panel (d), and further extending to $3 \to 10 \to 9 \to 8 \to 7$ in Panel (f).

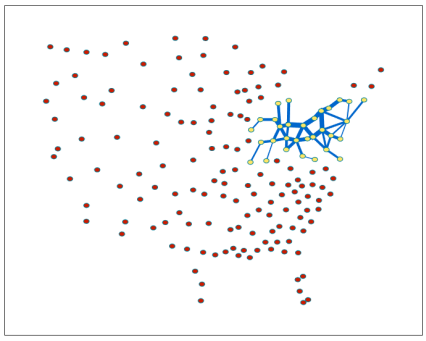
In summary, these comparative statistics demonstrate how higher values of γ enhance the effectiveness of resource allocation, thereby significantly influencing optimal routing and resource distribution strategies.



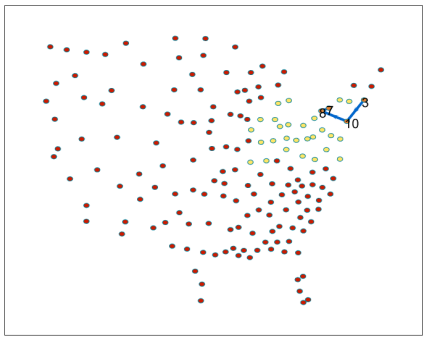
(a) Allocation of Resources, $\gamma = 0$



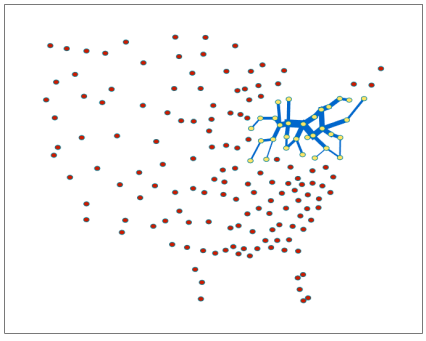
(b) Routing of $3 \to 7$, $\gamma = 0$



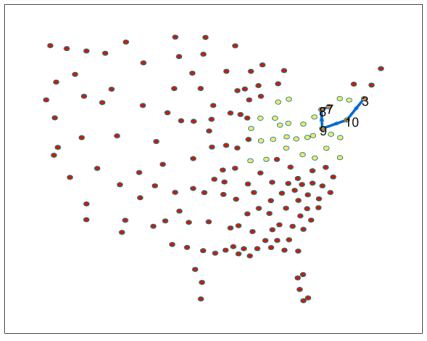
(c) Allocation of Resources, $\gamma = 0.5$



(d) Routing of $3 \to 7$, $\gamma = 0.5$



(e) Allocation of Resources, $\gamma = 0.85$



(f) Routing of $3 \to 7$, $\gamma = 0.85$

Figure F.1: How Value of γ Affects Allocation and Routing Decisions

G Additional Counterfactual Results

Table G.1 shows the full regression results of merger gains on changes in network centrality.

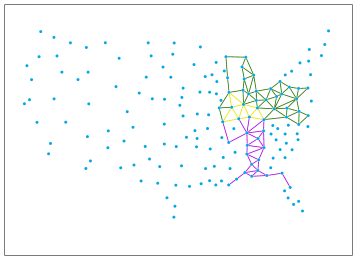
Table G.1: Merger Gains and Centralities (Regression Results)

	Baseline	Unpacking the Black Box		
	(1) Distance + Interchange Cost + Economies of Scope	(2) Distance	(3) Distance + Interchange Cost	(4) Distance + Economies of Scope
		Panel I: $\Delta \log(\text{Price})$		
Δ Degree Centrality	-0.0076^{***}	-0.0112^{***}	-0.0108^{***}	-0.0095^{***}
	(0.0010)	(0.0005)	(0.0006)	(0.0009)
Δ Betweenness Centrality	-0.000154^{***}	-0.000092^{***}	-0.000098^{***}	-0.000140^{***}
	(0.000010)	(0.000007)	(0.000008)	(0.000009)
Indicator of Interchange	-0.3946^{***}	-0.0572^{***}	-0.2816^{***}	-0.0481^{***}
	(0.0033)	(0.0016)	(0.0019)	(0.0030)
		Panel II: $\Delta \log(\text{Cost})$		
Δ Degree Centrality	-0.0053^{***}	-0.0100^{***}	-0.0097^{***}	-0.0078^{***}
	(0.0011)	(0.0005)	(0.0006)	(0.0010)
Δ Betweenness Centrality	-0.000190^{***}	-0.000098^{***}	-0.000105^{***}	-0.000174^{***}
	(0.000011)	(0.000008)	(0.000009)	(0.000010)
Indicator of Interchange	-0.4354^{***}	-0.0621^{***}	-0.2990^{***}	-0.0581^{***}
	(0.003385)	(0.001725)	(0.001997)	(0.003046)
		Panel III: Δ log(Markup)		
Δ Degree Centrality	0.0010^{***}	0.0001^{***}	0.0001^{***}	0.0011^{***}
	(0.0003)	(0.0000)	(0.0000)	(0.0002)
Δ Betweenness Centrality	0.000043^{***}	0.000002^{***}	0.000002^{***}	0.000044^{***}
	(0.000003)	(0.000000)	(0.000000)	(0.000003)
Indicator of Interchange	0.0125^{***}	0.0002^{***}	0.0003^{***}	0.0100^{***}
	(0.0009)	(0.0001)	(0.0001)	(0.0008)

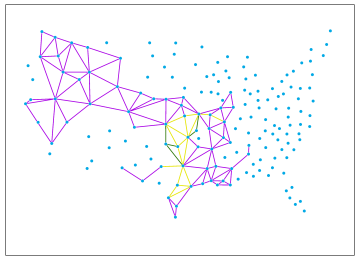
Notes: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure G.1: Networks for Each Merger Case

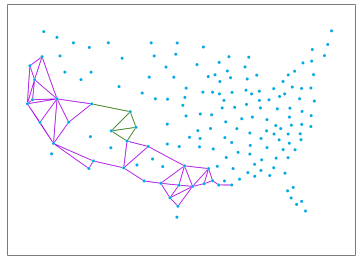
Note: Figure G.1 shows the networks of the two merging parties in each merger case between Class I railroads from 1985 to 2005. There were 12 mergers in total. Within each merger (firm1 + firm2), the network of firm1 is marked in green, that of firm2 is marked in purple, and the overlapping part is marked in yellow. For example, in panel (a) the network solely owned by COBO before the merger is marked in green, the network solely owned by SBD is marked in purple, and the overlapping region is marked in yellow.



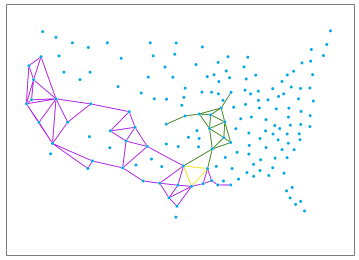
(a) 1986, COBO + SBD



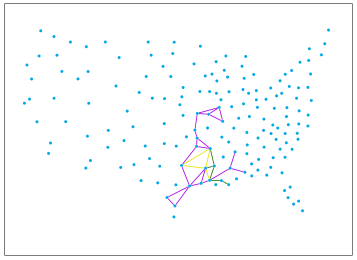
(b) 1988, MKT + UP



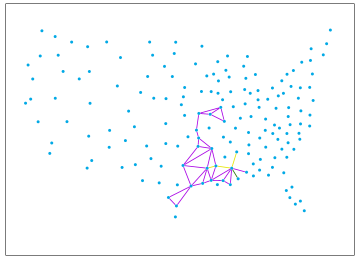
(c) 1988, DRGW + SP



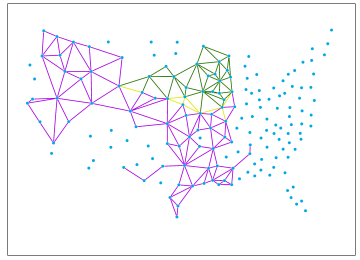
(d) 1992, SSW + SP



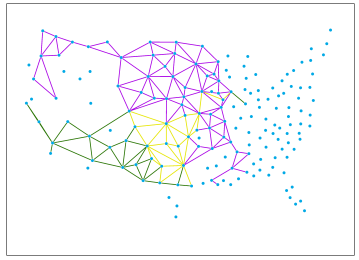
(e) 1992, LA + KCS



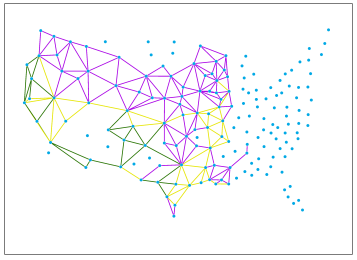
(f) 1993, MSRC + KCS



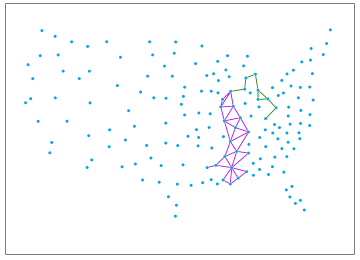
(g) 1995, CNW + UP



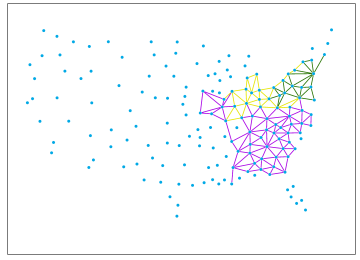
(h) 1996, ATSF + BN



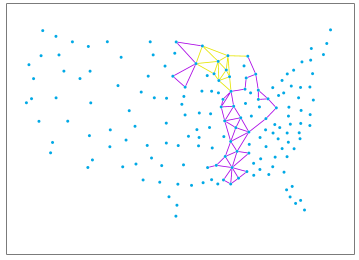
(i) 1996, SP + UP



(j) 1998, CN + IC



(k) 1999, CR + NS



(l) 2004, WC + CN