Reinforce or Reallocate?
Railroad Infrastructure Investment Decisions Under Climate Risk

Derek Thiele*

This Version: March 26, 2026

PRELIMINARY DRAFT

Abstract

This paper studies how climate risk shapes freight railroad infrastructure investment decisions. Using shipment-level data on freight rail markets, I document that short-run disruptions from climate hazards lead primarily to operational adjustments rather than immediate price or quantity responses. Over longer horizons, persistent increases in hazard exposure are associated with declining rail flows. I develop a network model of the freight rail industry to capture the underlying mechanisms and quantify the cost of climate disruptions. The findings suggest that climate hazards meaningfully raise effective shipping costs on affected routes. The framework is used to evaluate the welfare implications of alternative infrastructure adaptation policies under current and projected climate risk.

JEL Codes: R40, R41, R42, H54, Q54, L92

Keywords: freight rail, climate risk, infrastructure investment, network resilience, natural disasters, transportation economics

*Department of Economics, University of Toronto. d.thiele@mail.utoronto.ca. I am grateful to my advisors Stephan Heblich and Jeffrey Sun for their guidance and support. The Confidential Carload Waybill Sample was provided by the Surface Transportation Board under a data use agreement. All errors are my own. I declare no conflicts of interest. This paper is a preliminary draft circulated for discussion purposes. Several components remain in progress, including the full structural estimation of model parameters, counterfactual policy analysis, and the empirical results section. The model framework, identification strategy, and preliminary estimates presented here are subject to revision. Comments are welcome.

1 Introduction

Efficient, safe, and reliable transportation networks are essential for the movement of commodities across space. However, these networks are susceptible to disruptions caused by extreme weather events, making the movement of goods more costly. As climate change increases the frequency and severity of such events, a natural question becomes how transportation infrastructure can be adapted to improve network resiliency. Unlike people and firms, transportation infrastructure is difficult to relocate, but the network structure of these systems means that targeted investments in the right places could have substantial returns. Climate risk is also spatially heterogeneous — some parts of the network face greater exposure to hazards than others — meaning that the returns to adaptation will vary considerably across the network. I explore both of these issues in the context of the U.S. freight rail network, asking two related questions. First, what are the economic costs associated with climate hazard risks to rail infrastructure? Second, how should infrastructure investment be allocated across the network to most effectively improve resiliency to climate shocks? I focus on two possibilities: reinforcing links against climate hazards, and reallocating infrastructure investment to hedge against climate risk.

While recent work examines infrastructure allocation under future climate risks such as coastal flooding (Balboni, 2024; Hsiao, 2025), I study a distinct setting in which climate risk takes the form of stochastic, transitory cost shocks whose spatial correlation and network propagation are central to both the costs of disruption and the returns to infrastructure investment. Inland climate hazards are harder to anticipate, temporary in nature, and propagate through the network in ways that depend critically on its topology. I frame climate change's impact as effectively increasing the risk of stochastic transitory cost shocks to segments of the network. These shocks are spatially correlated and their frequency and severity will grow over time. The U.S. freight rail network is a particularly compelling

setting for this question. Rail infrastructure is uniquely exposed to climate hazards like flooding and extreme heat. Unlike road networks, freight rail is privately owned and operated by a small number of firms, making infrastructure modification inherently different — decisions to reinforce or reallocate must be weighed against private incentives that may not align with social welfare. Together, these features make the topic of network infrastructure adaptation both economically important and largely unexplored.

This paper makes three contributions. First, using proprietary data on rail shipments, I examine how prices and quantities respond to extreme weather events, providing evidence on how railroads react ex-post to hazard shocks. I further provide suggestive evidence on how underlying climate risk manifests as persistent variation across origin-destination markets. Second, I develop a spatial model of the U.S. freight rail network that captures the features central to the interaction between climate risk and transportation costs — including market power, congestion, and optimal routing behavior. Third, I use this framework to evaluate the returns to alternative infrastructure investment policies, showing how policies that jointly account for climate risk and network characteristics (capacity reallocation and/or direct reinforcement) can improve social welfare relative to the status quo.

The main findings of this paper are twofold. First, I find limited evidence that freight prices and quantities adjust in the short run following climate-induced network disruptions. This is consistent with railroads absorbing shocks through operational adjustments such as rerouting and service delays, and with the prevalence of long-term contracts that insulate prices from transitory cost shocks. Second, in the long run, however, quantities do respond to persistent changes in hazard exposure: origin-destination (OD) pairs experiencing increases in climate risk over time see declining flows, with no corresponding recovery when risk falls. I use this empirical pattern to estimate the structural cost of disruption, finding that climate hazards raise the effective shipping cost on affected routes

by approximately 61% on average. I apply my estimated model to evaluate different infrastructure adaptation policies, comparing their welfare implications under current and projected future climate risk.

One strand of literature this paper contributes to is spatial climate economics. Climate change and its associated risks have uneven impacts across space, giving rise to interesting interactions with the spatial distribution of economic activity that a growing body of work has begun to examine (Balboni & Shapiro, 2025). A subset of this literature focuses specifically on how climate-related disruptions and hazard risks shape economic activity. Castro-Vincenzi et al. (2024) studies how flooding and other types of climate risk cause Indian firms to structure their supply chains, finding evidence of diversification. In a similar vein, Balboni et al. (2024) find evidence of firm adaptation to flood risks in Pakistan — in response to flooding events, firms relocate, shift towards less risky suppliers, and use less flood-prone roads. Most recently, Massoni (2025) studies how production networks adapt in response to extreme weather events in the context of international maritime shipping, finding that cyclones lead to short but intense disruptions through which firms adapt by choosing alternative routes. While these papers highlight how climate-related disruptions shape economic activity, their focus is on firm adaptation rather than adaptation in the transportation sector itself. My paper complements this literature by shifting the lens to transportation infrastructure — asking not how firms respond to climate risks, but how the networks they rely on should be invested in and configured to mitigate those risks.

Most similar to my own work is Balboni (2024), who studies how transportation infrastructure should be optimally allocated in Vietnam in light of future sea level rise, finding that climate-aware allocation rules shift resources away from coastal areas and improve welfare. My paper is related in spirit but differs in several important ways. I focus on inland climate hazards — particularly riverine flooding and extreme heat — which are harder to anticipate and whose impacts are typically temporary and network-

dependent, rather than the permanent, spatially predictable threat of coastal inundation. Furthermore, my setting — the U.S. freight rail network — is privately owned and operated, introducing market structure considerations absent in Balboni (2024)'s framework. Finally, and most importantly, by framing climate risk as stochastic cost shocks to network segments, my model captures re-routing, congestion, and the spatial propagation of disruptions through the network in ways that a simpler approach does not.

Another strand of literature this paper contributes to is spatial transportation economics, which documents how transportation networks shape the spatial distribution of economic activity and the spatially heterogeneous returns to infrastructure improvements. Donaldson and Hornbeck (2016) provide an early example, evaluating how the historic U.S. rail network facilitated economic growth using a market access sufficient statistic approach. More recently, methodological advances have enabled analytical evaluation of optimal routing behavior and infrastructure allocation within quantitative spatial models. Fagjelbaum and Schaal (2020) develop a framework for determining optimal infrastructure allocations with a globally optimal solution obtainable via linear programming. Most relevant to my own work is Allen and Arkolakis (2022), which establishes a tractable framework for optimal routing with appealing analytical properties that I build directly upon. My paper contributes to this literature by introducing stochastic climate shocks into this class of models — segments of the network are subject to random cost shocks whose distribution depends on underlying climate risk. This extension has nontrivial theoretical implications for how traffic is routed across the network, how disruptions propagate spatially, and ultimately how infrastructure should be allocated in the presence of risk. To my knowledge, this is the first paper to explore these implications within the Allen and Arkolakis (2022) framework.

This paper also contributes to a literature on the economics of modern railroads. Degiovanni and Yang (2023) explore the role of capacity constraints and infrastructure in-

vestments in shaping transportation costs, paying careful attention to the market structure of the US rail industry. Chen (2024) explores how efficiency gains would arise from railroad mergers, driven by economies of scope and the elimination of interchange costs. Fuchs and Wong (2025) study multimodal transportation networks and intermodal substitution between road and rail within a quantitative spatial setting that extends Allen and Arkolakis (2022). My paper takes insights from each of these works while contributing to this relatively small literature by embedding climate risk directly into a model of railroad infrastructure investment.

Most closely related to my empirical setting is Du and Wilson (2026), who study how weather affects US railway safety and performance. They find that temperature extremes significantly increase accident rates, that operators respond by reducing speeds and throughput, and that these local adaptation responses propagate through the network as congestion externalities. These findings inform several of the key assumptions underlying my own framework: that climate hazards increase the cost of using affected segments and that disruptions spill over to other parts of the network. While their focus is on safety and short-run operational responses, my paper is complementary in its focus on the longer-run question of how infrastructure should be invested in and configured in light of these risks.

The rest of this paper proceeds as follows. Section 2 provides institutional background on the U.S. freight rail industry and the challenges posed by climate risks. Section 3 outlines the main data sources used for motivating facts and model estimation. Section 4 presents motivating facts that inform my modeling choices. Section 5 develops the model and explores relevant comparative statics. Section 6 outlines how I estimate and parameterize my model. Section 7 applies the model to evaluate counterfactual investment policies. I conclude in Section 8.

2 Institutional Background

The U.S. freight rail network spans over 140,000 miles and accounts for roughly one-third of all domestic freight ton-miles (Association of American Railroads, 2024). The industry's modern form was shaped decisively by the Staggers Rail Act of 1980, which deregulated freight rail pricing and is widely credited with rescuing an industry battered by decades of regulated stagnation: in the two decades that followed, inflation-adjusted rates fell by nearly 40 percent and volumes surged (Eakin et al., 2009). Freight rail is also increasingly important from an environmental standpoint — rail moves freight at roughly four times the fuel efficiency of trucking, positioning it as a key component of a lower-carbon freight system.

The commodity composition of U.S. freight rail reflects the network's comparative advantage in moving heavy, bulk goods over long distances. The largest commodity groups by tonnage include coal, chemicals, grain and agricultural products, and nonmetallic minerals, while intermodal traffic (shipments that utilize both rail and trucking) is the largest by revenue (Association of American Railroads, 2024). These commodities differ substantially in their intermodal substitutability. Intermodal traffic competes directly with long-haul trucking and is relatively price-sensitive; shippers of bulk commodities like coal, grain, and chemicals are often "captive" to rail, with limited alternatives.

Since the Staggers Act, consolidation has reduced the number of Class I railroads from over forty to six. As illustrated in 1, regional duopolies have emerged: BNSF and Union Pacific in the West, CSX and Norfolk Southern in the East, and CN and CPKC serving cross-border corridors. Each railroad owns and maintains its own track. While interchange agreements allow freight to cross railroad boundaries, doing so requires physical car handoffs between carriers and introduces additional costs and delays. The STB serves as a regulatory backstop for captive shippers, with authority to adjudicate rate disputes

when a railroad exercises market dominance.

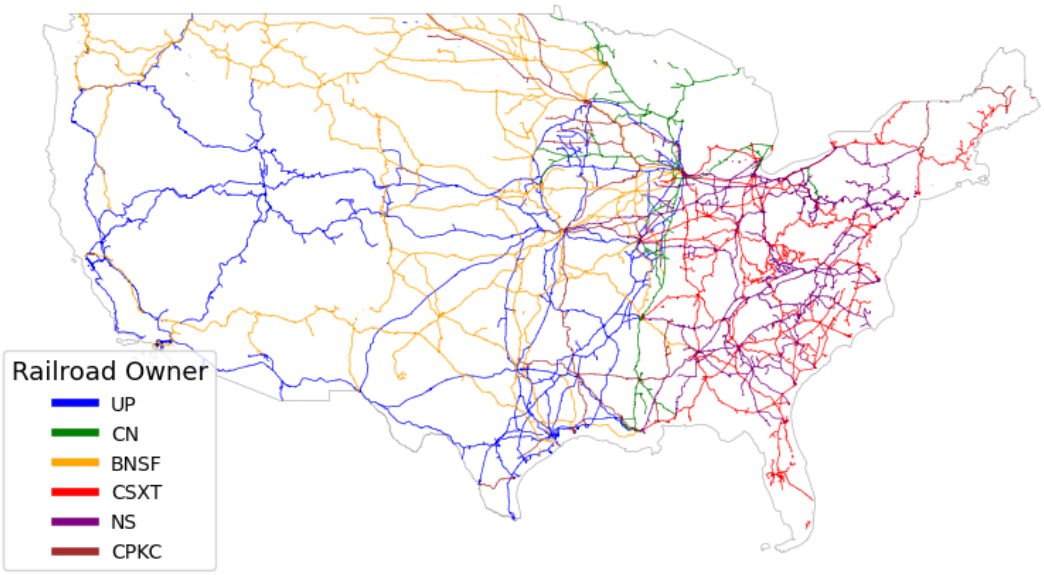


Figure 1: Class I U.S. Freight Rail Network

Railroads have also aggressively adopted precision-scheduled railroading (PSR) since the mid-2010s — a set of operating practices aimed at maximizing asset utilization and cutting costs. PSR has delivered significant efficiency gains but has done so largely by shedding excess capacity: low-density lines have been abandoned or sold off, locomotive and car fleets trimmed, and network slack reduced. The result is a leaner network that is more vulnerable to disruptions. Pricing is governed largely by long-term bilateral contracts between carriers and shippers, locking in rates and volumes and reflecting relationship-specific investments on both sides.

Freight rail infrastructure is long-lived, spatially fixed, and difficult to modify outside existing rights-of-way. To minimize construction costs, historical routes were often laid along riverbeds and through floodplains — terrain that today coincides with elevated climate hazard exposure. Rail infrastructure is susceptible to a range of climate hazards: flooding inundates track and damages ballast and signals; extreme heat causes rail buck-

ling; landslides block corridors and damage structures. These hazards are spatially heterogeneous, and the combination of historical routing decisions and uneven climate risk exposure implies substantial variation in vulnerability across the network.

The long asset lives involved mean that investment decisions made today will have impacts well into a future with uncertain and likely worsening climate hazards. Railroads can respond to climate risk in two broad ways. *Ex-post*, they can repair damage and restore service after a hazard event, potentially rerouting traffic in the interim — a reactive approach that dominates current practice. *Ex-ante*, they can *reinforce* existing infrastructure against climate hazards through track elevation, improved drainage, flood walls, or heat-resistant rail. Alternatively, they can *reallocate* capacity across the network by adding parallel track or passing sidings on lower-risk corridors, increasing throughput and expanding routing options, while concentrating less investment in corridors with greater hazard exposure. This paper focuses on these latter two strategies, asking how ex-ante investment policy should be designed to optimally improve network resilience to climate risk.

3 Data

3.1 Carload Waybill Sample

The primary source of railroad data is the Confidential Carload Waybill Sample (CWS), a 2.7% stratified sample of all carload waybills for U.S. rail traffic, covering any major carrier with more than 4,500 carloads annually. The unit of observation is a waybill — a record of a single carload shipment — containing information on origin, destination, interchanges, trip distance, shipment quantity, revenue, commodity type, and the railroads involved. I use the full unmasked version of the CWS for years 1990 to 2020, comprising

nearly 20 million observations. Origins and destinations are reported both at the Bureau of Economic Analysis (BEA) area level and at the station level via SPLC codes; I match SPLC codes to approximate geographic locations using data scraped from Railinc.

From the CWS, I construct a dataset of railroad market shares for each carrier in each origin-destination (OD) market, defined at the BEA level. To approximate total market size — including both rail and truck shipments — I draw on the 2017 Commodity Flows Survey (CFS), which provides estimates of freight flows by mode between BEA areas, allowing me to recover rail's share of each market and scale quantities accordingly.

3.2 Rail Network

I obtain a shapefile of the U.S. freight rail network from the North American Transportation Atlas Database (NTAD), which provides a detailed characterization of rail infrastructure across the continental United States. For each of the six Class I railroads, I construct a simplified network from this shapefile, retaining only the edges belonging to each carrier's system. The simplification proceeds in two steps. First, I collapse degree-2 nodes — nodes connecting exactly two segments — merging adjacent edges and summing their lengths, which reduces the node count while preserving the network's topological structure. Second, I apply a density-based spatial clustering algorithm (DBSCAN) to group nodes within 15 km of each other into a single node located at the cluster centroid, connected to adjacent clusters by the shortest available link. This step collapses complex infrastructure in major rail hubs and urban centers into single nodes, preserving the macro inter-city structure of the network while discarding intra-city detail that is not relevant for long-haul freight. For each edge in the simplified network, I compute its length and average track count, aggregating over the original edges from which it was constructed.

I use this simplified network to infer routes for individual waybills. The CWS does

not directly report the route taken; instead, each waybill records the sequence of states traversed and total trip distance. For a given origin-destination-railroad triple, I define a route as a unique combination of states traversed and total distance. I then compute the k shortest loopless paths on the subnetwork restricted to the reported states, and assign the waybill to the route(s) whose total distance most closely matches the reported distance, keeping all routes within a tolerance band and assigning equal probability across matched routes. Match rates vary across carriers, ranging from 82% (BNSF) to 93% (NS). The inferred routes are used to construct segment-level flow measures for the simplified network.

3.3 Climate Hazard Data

The primary source of climate hazard data is Ghoreishi et al. (2025), a new national dataset that spatially assigns historical natural hazard events to individual segments of the U.S. rail network for the period 2000–2024. The dataset covers flash floods, riverine floods, extreme heat events, and landslides, and records hazard type, timing, duration, severity, and associated injuries and fatalities for each event. I use the full distribution of events to characterize the spatial pattern of climate risk across the network, estimating segment-level hazard probabilities.

4 Motivating Facts

Fact 1: Hazard exposure is heterogeneous across space

Figure 2 displays the spatial distribution of hazards across the three event types. It illustrates that there is significant heterogeneity across space. Flooding is geographically dis-

perse, with meaningful exposure across nearly the entire continental network. Extreme heat exhibits a south-central concentration, with the highest densities along the Mississippi corridor. Landslides, while the least frequent in aggregate, are heavily concentrated in the Pacific Northwest and Appalachian corridors. The fact that climate risk varies across the network implies there is scope for spatial adaptation through infrastructure allocation choices. The returns to investment will depend on the level of underlying risk, such that two segments of equal importance in the network can have differing returns.

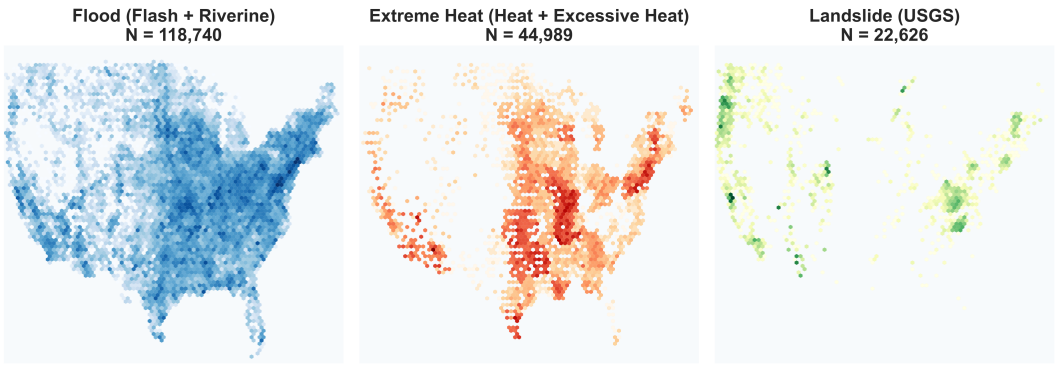


Figure 2: Spatial Distribution of Hazard Events

Fact 2: Hazard exposure is spatially correlated

Figure 3 plots the probability that a given segment on the rail network experiences exposure to a climate hazard at the same time. I find that the probability of co-occurrence decays with distance, consistent with spatial correlation. The implication is that disruptions tend to cluster geographically, limiting the scope for rerouting and suggesting that the full distribution of climate outcomes, not just mean exposure, matters for understanding the costs of climate risk to the network.

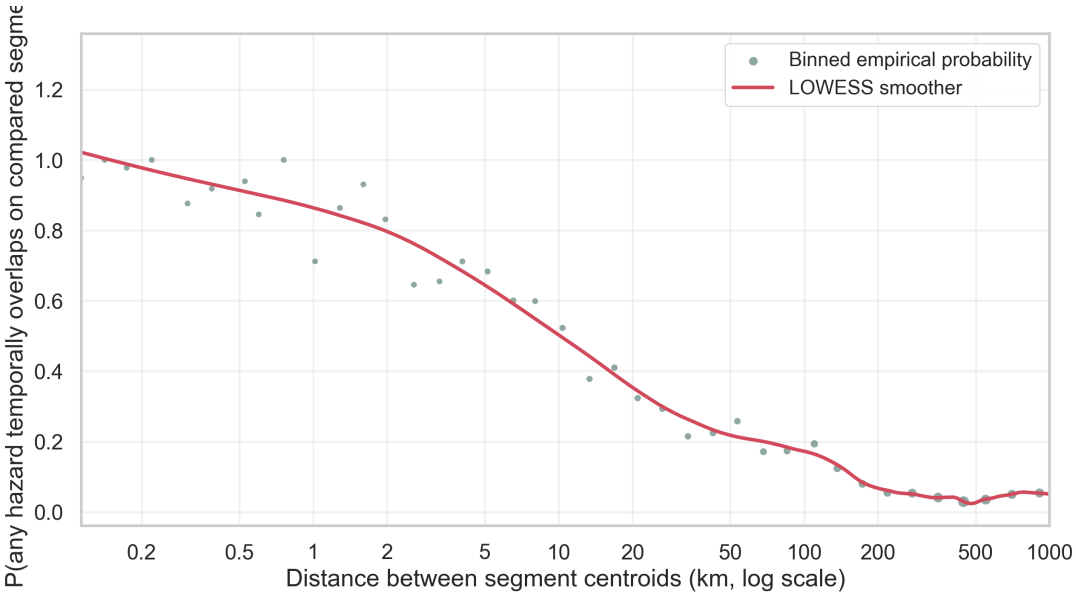


Figure 3: Hazard Co-occurrence Probability vs. Distance

Fact 3: Important rail corridors are exposed to hazard risk

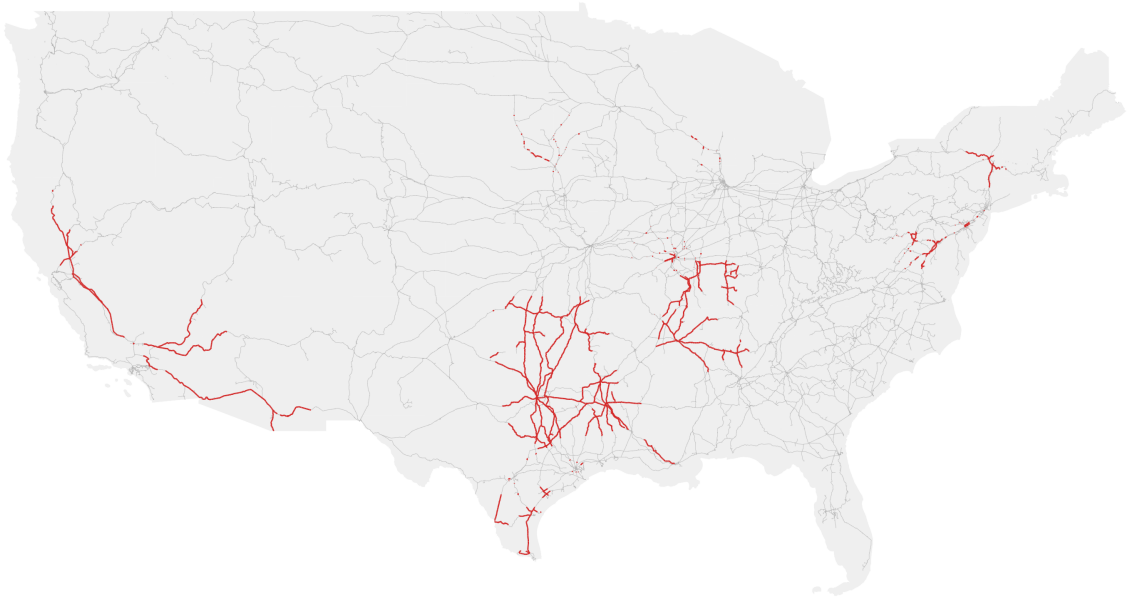


Figure 4: High Flow Corridors At Risk

Figure 4 highlights rail segments that are simultaneously high-flow and high-hazard, defined as edges at or above the 75th percentile on both flow centrality and cumulative hazard event exposure over 2015–2024. These doubly-exposed segments are not rare: they appear prominently along the Pacific Coast, through the Gulf Coast and lower Mississippi corridor, and in parts of the Northeast. The geographic concentration of these segments in economically critical corridors suggests that hazard risk and freight importance are not independent in the cross-section. Furthermore, these segments represent the corridors where the returns to adaptation investment are plausibly largest, and where the consequences of inaction may be highest.

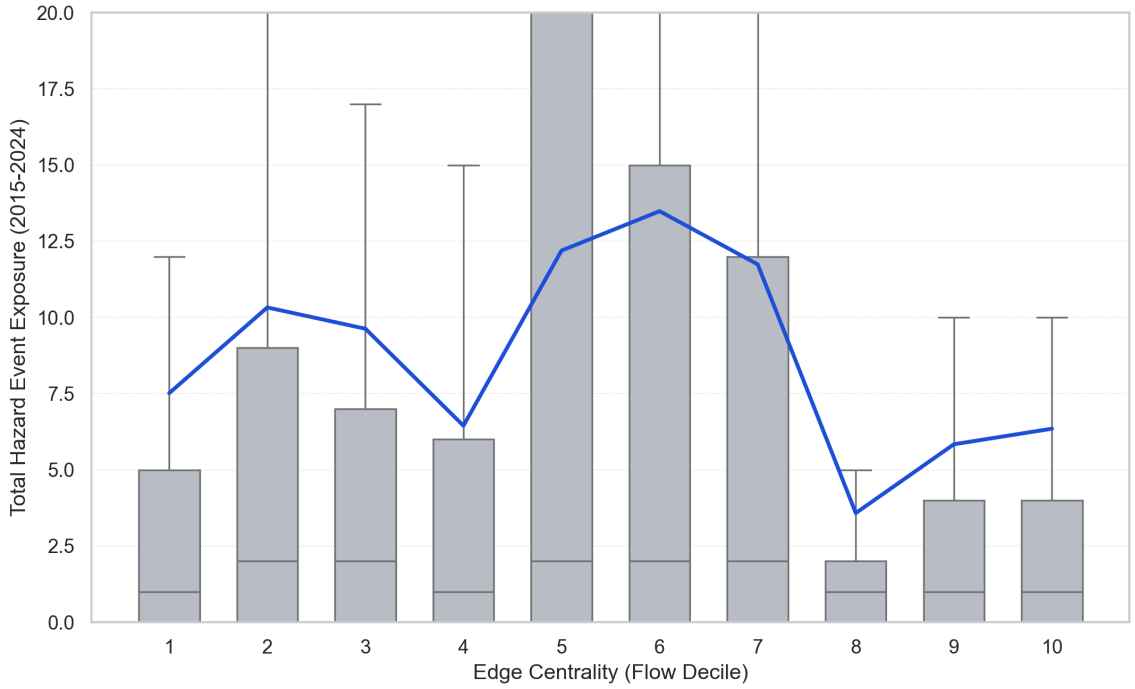


Figure 5: Variation in Hazard Risk Across Edge Centrality

Figure 5 provides a distributional view of the same relationship. Plotting total hazard exposure against edge centrality decile, mean exposure (blue line) rises through the middle deciles before declining at the upper end of the flow distribution. This relationship suggests that the very highest-flow edges are not the most hazard-exposed on average, but also that a substantial share of moderately central edges carry significant risk. Read

alongside Figure 4, this pattern suggests that the overlap between high flow and high hazard is concentrated in a specific subset of corridors rather than diffuse across the network.

Fact 4: Hazard risk has changed over time and unevenly across space

Figure 6 plots the distribution of the five-year rolling average of hazard exposure days per year across all track miles. The median is near zero throughout, but the mean rises from around 1.5 days in 2005 to a plateau of roughly 2.5–3.0 days during 2012–2015, dips through 2018–2020, and then climbs sharply through 2024. The sustained rise at the end of the sample — visible even after five-year smoothing — suggests that recent years represent a genuine increase in network-wide hazard exposure rather than the effect of isolated episodes, consistent with documented trends in the intensification of extreme weather events (IPCC, 2021). Under continued climate change, such exposure is projected to intensify further, with higher frequencies and severity of flooding, extreme heat, and related hazards across much of the continental United States (IPCC, 2021).

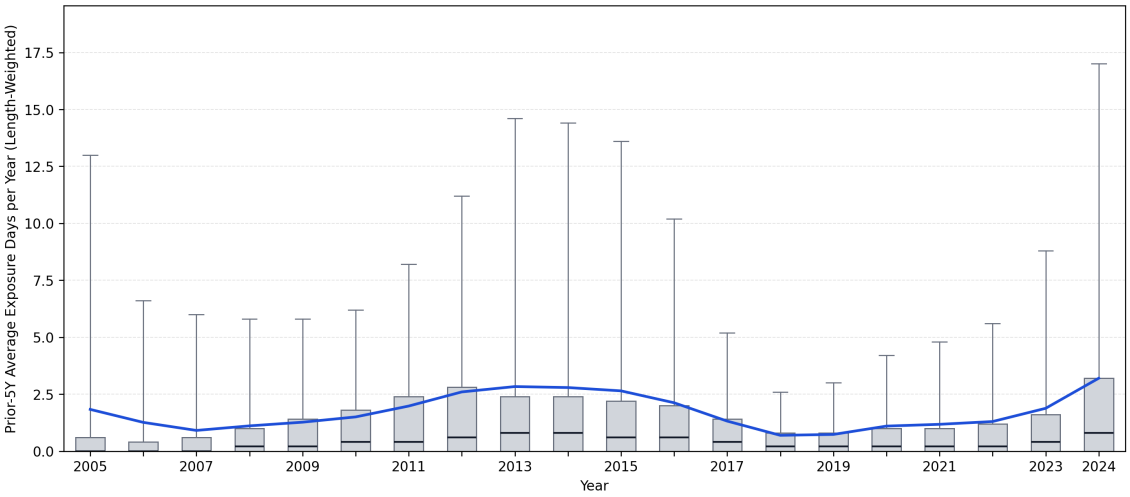


Figure 6: Distribution of Hazard Exposure Over Time

Figure 7 plots the spatial distribution of changes in hazard exposure between 2015 and

2024. Areas shaded dark red experienced more hazard days on average; areas shaded dark blue experienced fewer. The figure reveals spatial heterogeneity masked by the aggregate trends in Figure 6. The Southwestern U.S. saw large increases in hazard days over this period, while parts of the Midwest saw declines.

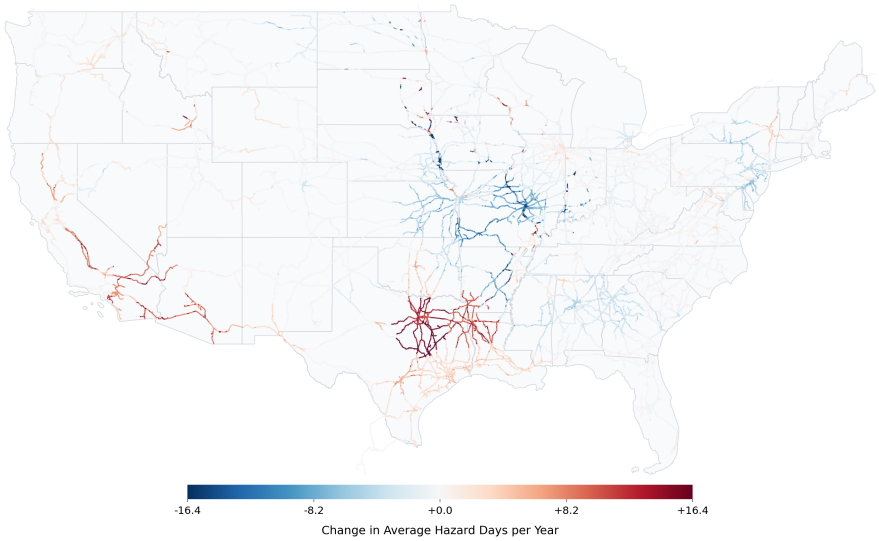


Figure 7: Change in Hazard Exposure Across Space

Fact 5: Disruptions do not cause short-run price responses and lead to limited quantity adjustments

To assess short-run responses to individual disruption events, I estimate an event study around hazard episodes from Ghoreishi et al. (2025). The unit of observation is a BEA-level origin-destination pair observed weekly within a ± 8-week window around each event. The treatment variable measures the share of potential flows on each OD pair routed through network segments affected by the hazard. I examine two outcomes – quantities shipped and revenue per ton-mile – across multiple specifications that vary both the criteria for admitting events and the sample of waybills included, with the spot-market sample isolating shipments most likely to exhibit detectable short-run responses

given their flexible pricing and scheduling relative to contracted flows.[1]

Event study estimates around hazard episodes reveal no detectable short-run response in shipping prices. Across all specifications, revenue per ton-mile remains statistically and economically flat throughout the event window, including in the spot-market sample where short-run price adjustment would be most likely to appear. This suggests that climate-related disruptions do not translate into observable price movements in rail freight markets.

Quantity responses are similarly limited. In the full sample, estimated effects on shipped tonnage are generally close to zero. In the spot-market sample, some specifications exhibit modest declines in the weeks immediately following disruption onset, but these effects are small and not stable across alternative event definitions. Overall, the evidence indicates that short-run railroad responses to climate hazards do not operate through meaningful price adjustment and involve, at most, limited quantity contraction.

This pattern is consistent with institutional features of rail freight markets, in particular the prevalence of long-term contracts that limit short-run price adjustment and dampen reallocation of contracted flows. The somewhat stronger responses observed in the spot-market sample are also consistent with adjustment occurring primarily on the margin of more flexible shipments. Taken together, the results suggest that railroads absorb climate-driven disruptions largely through operational adjustments, such as rerouting and scheduling changes, rather than through market price signals or substantial reductions in aggregate throughput.

[1] In Appendix F, I provide more details about the methodology.

Fact 6: In the long run, quantities respond to persistent changes in risk exposure

Table 2 in Section 6 shows that OD-pair quantities decline with increases in expected hazard exposure, consistent with railroads pricing the cost of risk into long-run routing and capacity decisions. Future work will examine price responses and incorporate more comprehensive evidence on how quantities and prices vary with properties of the risk distribution.

5 Model

5.1 Overview

I model equilibrium pricing and competition among railroads within each OD market, taking total freight demand at the OD level as given. This framework takes the geographic distribution of economic activity as fixed, abstracting from the possibility that firms may relocate production or sourcing in anticipation of future climate risks. Within each OD market, railroads compete as oligopolists, offering differentiated shipping services with trucking serving as the outside option. Demand for each railroad's service depends on its price relative to competitors and the outside option, with shippers choosing among alternatives according to a multinomial logit/CES structure.

A railroad's ability to compete in a given market depends on its shipping costs, which are determined by the underlying network structure. Each railroad owns and operates its own network, and shipping a good from origin to destination requires traversing a sequence of connected edges. Edge costs depend on physical characteristics — distance,

grade, and track capacity — and are subject to congestion: as more freight is routed along a segment, the marginal cost of using it rises Because of congestion, it is optimal for railroads to distribute flows across multiple routes connecting the same origin and destination rather than concentrate everything on a single least-cost path. Evidence of multiple routes being taken is visible in the waybill data, as documented by Degiovanni and Yang (2023).

Crucially, investment decisions are made at the edge level, since a single segment of track contributes to many routes across many different OD markets. This network structure is what makes infrastructure investment decisions fundamentally spatial in nature: improving or hardening a given segment has ripple effects on costs and flows throughout the network.

The model incorporates stochastic climate risk through disruptions at the segment level. Each segment faces a probability of being disrupted by a climate hazard event, drawn from a distribution that is spatially correlated — reflecting the fact that geographically proximate segments tend to face common flood or heat exposure. When a segment is disrupted, the cost of routing freight through it is scaled up by a disruption cost multiplier, reflecting the additional time and operational burden of traversing impaired infrastructure.

Pricing decisions are made ex-ante, based on expected costs over the full distribution of climate states, while routing decisions are made ex-post, after the realization of which segments are disrupted. This timing assumption reflects the institutional reality of the freight rail industry: the vast majority of shipments are governed by long-term contracts, so prices and quantities are set in advance and do not respond to short-run disruptions. Railroads instead adapt ex-post by adjusting how they route shipments — avoiding disrupted segments where possible, or continuing to use them at higher cost, reflecting the delays and reduced throughput associated with damaged or impaired infrastructure.

5.2 Setup

Markets are indexed by origin i and destination n and are of exogenous size M_{in}. Each market is serviced by a set of railroads $\mathcal{J}_{in}$, who offer a shipping service along their network. For now, I only consider single-railroad shipments and abstract away from interchanges.

Railroad $R \in \mathcal{J}_{in}$ charges a per-unit price of $p_{in,R}$ for their shipping service in market in. The quantity demanded is denoted by $Q_{in,R} = M_{in} \cdot d_{in,R}$, where $d_{in,R}$ is the market share given by:

$$d_{in,R} = \frac{\Lambda_R p_{in,R}^{-\alpha}}{1 + \sum_{R' \in \mathcal{J}_{in}} \Lambda_{R'} p_{in,R'}^{-\alpha}}$$

This can be micro founded with discrete choice and logit shocks, under which the outside option is trucking and Λ_R can be interpreted as a railroad-specific demand shifter. The elasticity of demand is given by $\varepsilon_{in,R} = \alpha(1 - d_{in,R})$.

Railroad R operates a network $\mathcal{G}_R = (\mathcal{N}_R, \mathcal{E}_R)$, where $\mathcal{N}_R$ is a set of nodes and $\mathcal{E}_R$ a set of edges. A route r between i and n is a sequence of connected nodes, where the kth node is denoted $e_{in,R,r}(k)$, with $e_{in,R,r}(0) = i$ and $e_{in,R,r}(K_{in,r}) = n$. Let $\mathcal{R}_{in,R}$ denote the set of available routes for railroad R in market in.

The world is subject to stochastic climate states $s \in \mathcal{S}$. In state s, the per-unit cost of shipping via route r is:

$$\tau_{in,R,r}(s) = m^{D_{in,R,r}(s)} \left(\prod_{k=0}^{K_{in,R,r}} \zeta_{e_{in,R,r}(k), e_{in,R,r}(k+1)} \right) Q_{in,R,r}^{\beta} \equiv c_{in,R,r}(s) \cdot Q_{in,R,r}^{\beta} \qquad (1)$$

where $\zeta_{jl} > 0$ is the cost of traversing edge (j, l), $Q_{in,R,r}$ is the quantity routed via r, $\beta > 0$

governs the degree of congestion, and $D_{in,R,r}(s) \in \{0,1\}$ indicates whether route r is disrupted in state s. The parameter $m > 1$ is a disruption cost multiplier: when a route is disrupted, its effective cost is scaled up, reflecting delays and reduced throughput along affected segments.

Two simplifying assumptions are embedded in this formulation. First, congestion enters at the route level rather than at the segment or rail yard level. In reality, a heavily-trafficked segment may become congested due to flows originating from many different OD pairs simultaneously, generating spillover costs across markets. Modeling congestion at the segment level would be more realistic, but would substantially complicate the cost aggregation problem; the route-level formulation preserves tractability while still capturing the key trade-off between cheaper and more congested routes. Second, the disruption cost multiplier m is assumed to be homogeneous across routes and events. In practice, the cost of a disruption likely varies with the severity and duration of the hazard, the characteristics of the affected segment, and the availability of nearby alternatives.

Given a total quantity $Q_{in,R}$ to ship in market in, the railroad minimizes costs by optimally distributing flows across routes. Cost minimization yields an aggregate shipping cost of:

$$\tau_{in,R}(Q_{in,R})(s) = \left(\sum_{r \in \mathcal{R}_{in,R}} c_{in,R,r}(s)^{-\frac{1}{\beta}} \right)^{-\beta} Q_{in,R}^{\beta} \equiv c_{in,R}(s) \cdot Q_{in,R}^{\beta} \tag{2}$$

which has the familiar CES aggregator form, with routes acting as imperfect substitutes. As individual route costs rise — due to disruption or congestion — flows are reallocated toward cheaper alternatives, though at diminishing returns governed by β.

Pricing decisions are made ex-ante, before the climate state is realized, while routing decisions are made ex-post. Railroad R therefore chooses $p_{in,R}$ to maximize expected

profits:

$$\max_{p_{in,R}} \left(p_{in,R} - \mathbb{E}[c_{in,R}(s)]\right) \cdot Q_{in,R}^{1+\beta} \tag{3}$$

The first-order condition yields a standard markup over expected marginal cost, with price given by:

$$p_{in,R} = \frac{\varepsilon_{in,R}}{\varepsilon_{in,R} - 1}(1+\beta)\mathbb{E}[c_{in,R}(s)]Q_{in,R}^{\beta}$$

Firms take into account how equilibrium quantities affect their marginal cost, and take as given the decisions of other firms, which enters through $\varepsilon_{in,R}$. Equilibrium prices and quantities across all firms are jointly determined as a fixed point.

Note that while one might expect firms to charge a risk premium reflecting exposure to disruption, the long-term contracting environment makes this unlikely in practice: prices are fixed for extended periods over which disruptions are transitory, so $\mathbb{E}[c_{in,R}(s)]$ is better interpreted as a markup over average costs across the contract horizon rather than a price that explicitly compensates for climate risk.

5.3 Network Costs and Climate Risk

To obtain tractable expressions for shipping costs and disruption risk, I follow Allen and Arkolakis (2022) and assume that $\mathcal{R}_{in,R}$ contains all possible routes between i and n, including those that revisit segments. While this may seem extreme, it is without loss of generality in the sense that optimal flows will assign negligible weight to inefficient routes; the assumption is made purely for analytical convenience and delivers closed-form expressions for costs that are straightforward to compute.

Under this assumption, the aggregate shipping cost and disruption risk probability admit the following representations. Define the matrix $\mathbf{A}_R = [\zeta_{jl,R}^{-1/\beta}]_{jl}$ of edge cost inverses, and let $\mathbf{B}_R = (\mathbf{I} - \mathbf{A}_R)^{-1}$ be the corresponding Leontief inverse. In state s, let $\mathbf{D}_R(s) = [D_{jl,R}(s)]_{jl}$ be the matrix of segment disruption indicators, and define $\mathbf{B}_R(s) = (\mathbf{I} - \mathbf{D}_R(s) \odot \mathbf{A}_R)^{-1}$, where $\odot$ denotes the Hadamard product. Denoting by $b_{in,R}$ the (i,n) element of $\mathbf{B}_R$ and $b_{in,R}(s)$ the corresponding element of $\mathbf{B}_R(s)$, the aggregate shipping cost is:

$$\tau_{in,R}(s) = \left[m^{-\frac{1}{\beta}} b_{in,R} + \left(1 - m^{-\frac{1}{\beta}}\right) b_{in,R}(s)\right]^{-\beta} Q_{in,R}^{\beta} \tag{4}$$

The cost expression has a natural interpretation: $b_{in,R}$ captures the baseline (undisrupted) connectivity between i and n, while $b_{in,R}(s)$ captures connectivity after disruptions have removed or impaired certain edges. As disruptions reduce $b_{in,R}(s)$ relative to $b_{in,R}$, the effective cost of shipping rises.

The model also yields tractable expressions for two measures of disruption exposure. The share of goods shipped along disrupted routes is:

$$\pi_{in,R}^{disrupted}(s) = \frac{m^{-\frac{1}{\beta}} \left(b_{in,R} - b_{in,R}(s)\right)}{m^{-\frac{1}{\beta}} b_{in,R} + \left(1 - m^{-\frac{1}{\beta}}\right) b_{in,R}(s)} \tag{5}$$

and the share of total shipments that are rerouted relative to the undisrupted optimum is:

$$\pi_{in,R}^{rerouted}(s) = \left(b_{in,R} - b_{in,R}(s)\right) \left(\frac{1}{b_{in,R}} - \frac{m^{-\frac{1}{\beta}}}{m^{-\frac{1}{\beta}} b_{in,R} + \left(1 - m^{-\frac{1}{\beta}}\right) b_{in,R}(s)}\right) \tag{6}$$

Both measures are increasing in the gap $b_{in,R} - b_{in,R}(s)$, which reflects the extent to which disruptions degrade network connectivity. When no segments are disrupted, $b_{in,R}(s) = b_{in,R}$ and both measures equal zero. Expected values $\mathbb{E}[\tau_{in,R}(s)]$, $\mathbb{E}[\pi_{in,R}^{D}(s)]$, and $\mathbb{E}[\pi_{in,R}^{R}(s)]$ are computed via simulation draws over the distribution of s.

Edge costs are parameterized as:

$$\zeta_e = \frac{\text{dist}_e^{\delta_1} \cdot \text{slope}_e^{\delta_2}}{\text{tracks}_e^{\gamma_1} \cdot \text{siding}_e^{\gamma_2}} \tag{7}$$

where dist_e is segment length in thousands of miles, slope_e is average grade, tracks_e is the number of parallel tracks, and siding_e is an indicator for the presence of passing sidings. Distance and grade increase costs, reflecting longer travel times and greater operational difficulty, while additional tracks and sidings reduce costs by expanding capacity and allowing trains to pass.

Disruption risk is modeled as follows. Segments are assigned to groups g based on their geographic location, with each group corresponding to a one-degree grid cell. The disruption indicator for segment e in state s is:

$$D_e(s) = \mathbf{1}\left\{Z_{g(e)}(s) \geq 0\right\} \tag{8}$$

where $\mathbf{Z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is a vector of latent risk factors. The mean parameter μ_g governs the baseline disruption probability for group g, with $p_g = \Phi(-\mu_g)$ giving the implied disruption probability. Spatial correlation is captured through $\boldsymbol{\Sigma}$, with $\Sigma_{gg} = 1$, $\Sigma_{gg'} = \rho$ if g and g' are spatially adjacent, and $\Sigma_{gg'} = 0$ otherwise. This parsimonious structure allows for correlated disruptions across nearby segments — reflecting the tendency of flood and heat events to affect contiguous areas — while remaining tractable for estimation and simulation.

5.4 Comparative Statics

In this section, I derive a set of comparative statics that illustrate how the core parameters of the model shape shipping costs and disruption exposure, and highlight the interactions

between network structure, congestion, and climate risk.

Proposition 1 *The share of goods rerouted in state s $\pi_{in,R}^{rerouted}(s)$ is increasing in m, while the share of goods disrupted, $\pi_{in,R}^{disrupted}(s)$ is decreasing in m.*

Proposition 1 reflects the idea that railroads will use disrupted routes less when the costs of using them are larger. Re-routing itself is costly, as it involves allocating more flows to costlier routes and leads to congestion spillovers. However, when disruption costs are sufficiently high, the benefit of avoiding impaired infrastructure outweighs these re-routing costs, and railroads shift flows accordingly. This gives rise to a trade-off at the heart of the model: a higher m implies less exposure to disrupted infrastructure, but greater congestion on the alternative routes that absorb the diverted flows.

Proposition 2 *$\tau_{in,R}(s)$ is increasing in disruption cost shifter m, but at a declining rate that converges to 0 as $m \to \infty$.*

As disrupted routes become more costly to use, average shipping costs rise. However, this effect is diminishing: as m increases, railroads reroute a greater share of flows away from disrupted segments, partially offsetting the direct cost increase. In the limit as $m \to \infty$, disrupted routes become effectively impassable, but because flows can be fully redirected through the remaining network, average shipping costs remain finite. Proposition 2 thus captures a natural dampening mechanism: the more severe the disruption penalty, the more aggressively railroads avoid it.

Proposition 3 *Let $\eta_{in,R}(s)$ denote railroad R's share of flows between i and n impacted by climate hazards in state s. Then the proportional increase in per-unit shipping cost relative to the undisputed baseline is given by:*

$$\frac{\tau_{in,R}(s)}{\tau_{in,R}} = \left[1 - (1 - m^{-\frac{1}{\beta}})\eta_{in,R}(s)\right]^{-\beta}$$

Proposition 3 provides a clean decomposition of how disruptions translate into higher shipping costs. The proportional cost increase depends on two objects: the share of flows impacted by the hazard, $\eta_{in,R}(s)$, and the disruption cost multiplier m. When no flows are affected, $\eta_{in,R}(s) = 0$ and costs are unchanged. As the share of impacted flows rises, the bracketed term falls below one, and costs rise at a rate governed by the congestion parameter β. The role of m is intuitive: a larger disruption penalty shrinks $m^{-1/\beta}$ toward zero, which amplifies the sensitivity of costs to $\eta_{in,R}(s)$. Conversely, as $m \to 1$, disruptions become costless. Notably, the proportional cost increase is also increasing in β: as congestion becomes more severe, re-routing flows around disrupted segments becomes costlier, since alternative routes are already heavily trafficked. Congestion thus acts as an amplification mechanism, magnifying the cost consequences of any given level of disruption exposure.

5.5 Climate Distribution and Model Outcomes

A property of particular interest is the role that the spatial correlation of climate hazards plays in shaping shipping costs. While the assumed risk distribution does not yield clean analytical results, it is straightforward to establish that expected shipping costs are increasing in hazard probabilities. The role of spatial correlation ρ is less immediate, and I use my estimated and parameterized model to evaluate it via simulation.

I first ask whether ρ affects the *level* of expected shipping costs. I find that, on average, ρ has no meaningful effect on expected costs (see Appendix Figure 10). This result is intuitive: spatial correlation governs the co-movement of hazards across cells, but leaves the marginal disruption probability of each segment unchanged. From the perspective of

any individual route, the expected number of disrupted segments is invariant to ρ, and so the expected cost of shipping is as well.

What ρ does affect, however, is the *distribution* of realized costs across states of the world. Higher spatial correlation concentrates disruptions geographically — rather than isolated, independent failures scattered across the network, correlated hazards produce spatially coherent events that simultaneously impair large contiguous portions of infrastructure. In such states, the network's capacity to absorb and reroute around disruptions is severely diminished: the alternative routes that would ordinarily provide redundancy are themselves compromised. The result is that high-ρ environments feature occasional catastrophic states in which costs spike across many origin-destination pairs, interspersed with many periods of near-normal operations. I find that the average variance of shipping costs is increasing in ρ (see Appendix Figure 11): when disruptions are correlated, the good states are roughly as good as under independence, but the bad states are substantially worse.

I also find that spatial correlation reduces the *cross-sectional dispersion* of expected costs across origin-destination pairs (see Appendix Figure 12). As ρ increases, cost shocks become more uniform across the network — OD pairs that would otherwise face idiosyncratic exposure instead face a common aggregate shock. From the perspective of railroads, this is particularly costly: under low correlation, a railroad operating across many OD pairs benefits from natural diversification, since disruptions affecting some routes are unlikely to coincide with disruptions elsewhere. Higher spatial correlation erodes this diversification benefit. In bad states of the world, railroads find their costs elevated across a broad range of markets simultaneously.

As shown by Proposition 3, disruption extent $\eta_{in,R}(s)$ acts as a sufficient statistic for the ex-post cost impact of hazards in state s. The ex-ante expected cost ratio $\frac{\mathbb{E}[\tau_{in,R}(s)]}{\tau_{in,R}}$, however, depends not just on $\mathbb{E}[\eta_{in,R}(s)]$ but on the entire distribution of $\eta_{in,R}(s)$, which is

determined by the underlying climate hazard distribution. Because $g(\eta) \equiv (1-(1-m^{-\frac{1}{\beta}})\eta)^{-\beta}$ is convex, the expected cost ratio is increasing under any mean-preserving spread of $\eta_{in,R}(s)$. Greater variance in disruption extent, driven by more frequent tail events, raises expected shipping costs even holding mean exposure fixed.

Taken together, these results underscore that the climate hazard distribution is the primitive object of interest for evaluating freight costs under climate risk. Changes in spatial correlation, tail thickness, or geographic concentration can raise expected costs and erode network resilience even when average hazard probabilities are unchanged. As climate change alters not only the frequency but the spatial structure of extreme events, these distributional features become increasingly relevant for understanding the long-run cost exposure of freight networks.

6 Estimation and Parameterization

6.1 Parameterization

For now, I focus on estimating the disruption cost parameter m and the climate hazard distribution. I parameterize the remaining parameters as follows. I set the distance cost parameter $\delta_1 = 1$, so that edge costs scale proportionately with length, and abstract away from the role of slope by setting $\delta_2 = 0$. For the track capacity parameter, I set $\gamma_1 = 0.3$, consistent with Degiovanni and Yang (2023), who find that an additional set of parallel tracks reduces costs by roughly 20% on average. I similarly abstract away from passing sidings by setting $\gamma_2 = 0$. The simplified rail network described in Section 3 provides the underlying edge-level values for distance and track count. BEA areas are matched to nodes on this network via nearest-neighbor matching of geographic centroids.

I set the congestion elasticity $\beta = 0.1$. While in reality, rail congestion is non-linear and depends on the level of capacity utilization, this approximation is reasonable in the context of freight rail, where capacity utilization is typically well below the threshold at which congestion effects become severe. It is also worth noting that railroads exhibit strong economies of density, which may partially offset congestion effects in practice, making a small value of β a reasonable baseline. Finally, a small value of β ensures that the model's predictions remain close to those of a least-cost routing framework, which is the natural limiting case as $\beta \to 0$. In Appendix D, I suggest an approach to identifying β through the waybill data and inferred routes. I find estimates that do not differ much from my parameterization.

I set the demand elasticity parameter $\alpha = 4$. Direct estimates of rail freight demand elasticities are limited and vary considerably across commodity, region, and time period. Wilson et al. (1988) estimate a structural own-price elasticity of approximately -1.5 for grain shipments in the post-deregulation period, and Chen (2024) reports an estimate of -1.6, though both likely understate elasticity in markets with stronger intermodal competition. Given that the effective elasticity in the model is $\varepsilon_{in,R} = \alpha(1 - d_{in,R})$, which is strictly less than α and varies across markets with the degree of captive demand, $\alpha = 4$ implies market-level elasticities broadly consistent with the upper end of the range reported in the literature. Formal estimation of α is left for a future draft.

Given the chosen value for α, I back out estimates for railroad-specific demand shifters Λ_R by choosing the values which produce market shares most consistent with observed information on prices and quantity shares from the CWS.

Table 1: Hazard Distribution Estimates

Period	N_{groups}	$\hat{\rho}$	$\overline{\mu}$	$sd(\mu)$	$\overline{p}$
2000-2010	708	0.7827	-2.5817	1.6339	0.0297
2011-2020	708	0.6990	-2.0067	1.0711	0.0508

6.2 Hazard Distribution Estimation

I estimate the hazard distribution separately for two periods: 2000–2010 and 2011–2020, allowing for the possibility that climate risk has shifted over time. Risk groups are defined as one-degree grid cells, with each segment assigned to the group corresponding to the cell in which its centroid falls.

For each group g, I estimate the hazard probability $\hat{p}_g$ as the average annual frequency of climate hazard events (flooding and extreme heat) occurring within the cell, weighted by episode duration in days and normalized by the number of days in a year. The group-level mean of the latent risk variable follows directly as $\hat{\mu}_g = \Phi^{-1}(\hat{p}_g)$.

I estimate the spatial correlation parameter ρ via maximum likelihood, exploiting information on the frequency of overlapping hazard episodes across geographically adjacent cells. Intuitively, ρ is identified by how often neighboring grid cells experience simultaneous hazard events relative to what would be expected under independence: a higher ρ implies greater co-movement in hazard status across adjacent cells. Details of the likelihood function are provided in Appendix C.

Table 1 summarizes the hazard distribution estimation results. The average cell-level hazard probability rises from approximately 0.03 in the 2000–2010 period to 0.05 in the 2011–2020 period, corresponding to an average increase of 7.7 expected hazard days per year. The estimated spatial correlation parameter $\hat{\rho}$ is large in both periods, reflecting the tendency for hazard events to affect contiguous portions of the network simultane-

ously. Appendix Figures 13 and 14 display the spatial distribution of estimated hazard probabilities for each period, revealing substantial geographic heterogeneity in baseline exposure.

Figure 8 maps the change in estimated hazard probability between the two periods. The increase in risk is widespread: the vast majority of cells experienced a rise in hazard probability, consistent with the growing frequency and severity of extreme weather events over this period. However, the magnitude of this increase varies considerably across space, with some regions – including the Midwest – experiencing substantially larger shifts than others.

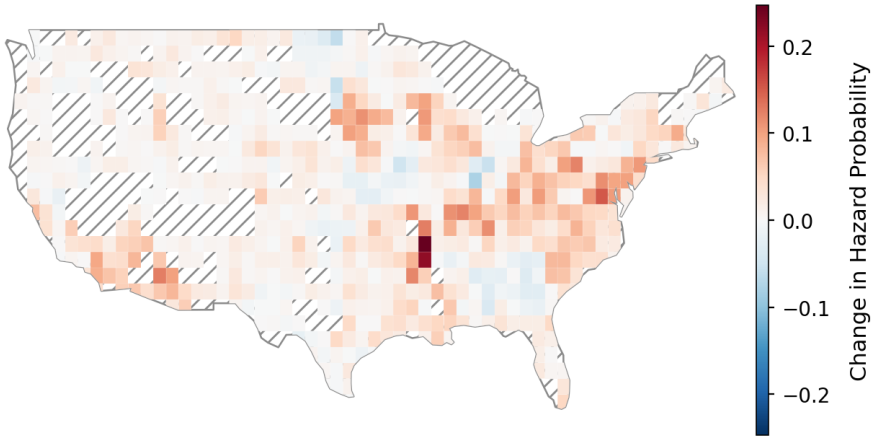


Figure 8: Change in Hazard Risk, 2000–2010 to 2011–2020

6.3 Disruption Cost Estimation

I estimate the disruption cost shifter m via simulated method of moments, targeting the empirical relationship between changes in disruption exposure and changes in observed rail flows. Let $\eta_{in,R} = \mathbb{E}[\eta_{in,R}(s)]$ denote the model-implied disruption exposure for OD pair in served by railroad R — the share of total flows on the network that either traverse disrupted routes or are rerouted around them, computed in expectation. Crucially, $\eta_{in,R}$ depends only on the underlying distribution of hazard risk and the network struc-

ture, and is invariant to the value of m. I run the following long-difference regression on waybill quantities between 2009 and 2019:

$$\Delta \log Q_{in,R} = \lambda \, \Delta \eta_{in,R} + u_{in,R} \tag{9}$$

where $\Delta \eta_{in,R}$ is the change in disruption exposure driven solely by the shift in the underlying hazard distribution between the two periods, holding all else fixed. The coefficient λ captures the average sensitivity of flows to changes in network-level disruption exposure. I then find the value of m that causes the model-implied analog of λ to match the empirical estimate.

To allow for asymmetry in the response, I also estimate an augmented specification that distinguishes OD pairs for which disruption exposure increased from those for which it decreased:

$$\Delta \log Q_{in.R} = \lambda^{+} \left(\Delta \eta_{in.R} \times \mathbf{1}\{\Delta \eta_{in,R} \geq 0\} \right) + \lambda^{-} \left(\Delta \eta_{in,R} \times \mathbf{1}\{\Delta \eta_{in,R} < 0\} \right) + u_{in,R} \tag{10}$$

Table 2: Quantity Response to Changes in Risk

	(1)	(2)	(3)
$\hat{\lambda}$	-0.25*		
	(0.14)		
$\hat{\lambda}^{+}$		-0.47**	-0.51**
		(0.20)	(0.22)
$\hat{\lambda}^{-}$			0.13
			(0.29)
N (OD pairs)	7,521	7,521	7,521
R^2	0.0005	0.0008	0.0009

Table 2 reports the results of the long-difference regressions. Specification (1) pools increases and decreases in risk exposure Specifications (2) and (3) allow for asymmetry by separating OD pairs that experienced an increase in risk exposure from those that ex-

perienced a decrease. The estimated response to a risk increase is substantially larger in magnitude: a 10 percentage point increase in risk exposure is associated with a 5.1% decline in flows, while the coefficient on a risk decrease is statistically insignificant. The asymmetry is consistent with agents forming expectations about long-run climate risk that are not easily revised downward: even when a given OD pair experiences a reduction in measured hazard exposure over the period, shippers and railroads may rationally discount this as short-term fluctuation.

Table 2 reports the results of the long-difference regressions. Specification (1) pools increases and decreases in risk exposure, yielding an estimated $\hat{\lambda} = -0.25$. Specifications (2) and (3) allow for asymmetry by separating OD pairs that experienced an increase in risk exposure from those that experienced a decrease. The estimated response to a risk increase is substantially larger in magnitude: Specification (3) yields $\hat{\lambda}^{+} = -0.51$, implying that a 10 percentage point increase in risk exposure is associated with a 5.1% decline in flows, while the coefficient on a risk decrease is small statistically insignificant. The asymmetry is consistent with agents forming expectations about long-run climate risk that are not easily revised downward: even when a given OD pair experiences a reduction in measured hazard exposure over the period, shippers and railroads may rationally discount this as short-term fluctuation.

I select Specification (3) as the preferred specification and target $\hat{\lambda}^{+} = -0.51$ in the SMM procedure. I find the value of m that matches the model-implied quantity response to changes in risk to the empirical one by running an analogous regression on model-implied data and searching over values of m until the simulated coefficient matches the empirical target. This search has a unique solution, since the magnitude of the model-implied quantity response is monotonically increasing in m. The procedure yields a preliminary estimate of $\hat{m} = 1.617$, implying that the effective cost of using a disrupted route is approximately 61.7% higher than an undisrupted one.

This result suggests that climate hazards on average create substantial short-run increases in usage costs for affected routes, though not so large as to make disrupted routes impossible to use. This is consistent with the nature of most reported hazard events: not all cause major infrastructure damage, and events vary considerably in the extent to which they impair a given segment. In many cases, railroads respond by reducing speeds and accepting delays in the movement of goods along disrupted routes, which increases costs but does not preclude their use entirely.

The estimate of $\hat{m}$ is likely sensitive to the choice of β and α. A larger congestion elasticity β amplifies the cost impact of marginal changes in risk exposure — implying that a smaller value of m would be needed to match the same empirical moment. Similarly, a larger demand elasticity α implies that shippers are more responsive to changes in effective shipping costs, generating a larger quantity response for a given value of m, and again biasing the estimate upward. A natural next step is therefore to conduct sensitivity analyses over these parameters, and ultimately to estimate β and α credibly rather than externally calibrating them.

More broadly, the current estimation procedure targets a single empirical moment, while the model and available data together imply a richer set of identifying restrictions that could support more robust estimation. Future drafts will incorporate additional moments, such as cross-sectional variation in the quantity response across OD pairs with different network positions, or heterogeneity in the response across event types and severities.

A further consideration is the time horizon over which hazard probabilities are defined. The current approach computes annual disruption probabilities, but the model's implications for $\hat{m}$ are not invariant to this choice: defining risk at a finer frequency (e.g. monthly) would yield lower baseline probabilities, but would require a correspondingly larger value of m to generate the same quantity response.

7 Counterfactuals

7.1 Climate Projections

A natural extension of this analysis is to replace the current hazard distribution with projections from future climate scenarios. The frequency of extreme weather events is expected to increase with global mean temperature, and to do so nonlinearly: rarer events face the largest proportional increases, so that tail risks that appear manageable today could grow substantially under even moderate warming (IPCC, 2021). These changes are also spatially uneven, meaning some parts of the network will face much larger shifts in exposure than others. Incorporating scenario-based hazard projections into the counterfactual framework is a planned component of a future draft.

7.2 Predicting Network Flows

Since a first-order approximation of the returns to infrastructure investment is proportional to segment-level flows, understanding how those flows vary with risk is of direct economic interest. Appendix E describes how segment-level flows are computed for any climate state s. Expected flows can be computed by simply taking the expectation over many draws from the climate distribution.

Figure 9 displays the change in model-predicted expected flow importance percentile for each segment, comparing a riskless baseline to the present-day risk environment, for BNSF and UP. Darker green segments gain relative importance with the introduction of risk; darker red segments lose it. Segments that gain relative importance tend to lie in lower-risk areas, reflecting traffic reallocation away from more exposed parts of the network. Changes in segment-level flows reflect two forces: OD pairs exposed to high

hazard ship less as market quantities fall, reducing flows on segments that serve them; and shippers substitute away from routes passing through high-risk segments, further reallocating flows toward lower-risk alternatives. In absolute terms, flows fall nearly everywhere across the network, as risk raises effective transport costs.[2] However, risk produces a meaningful reallocation of relative flow importance, suggesting that climate risk is a non-trivial determinant of the returns to infrastructure investment.

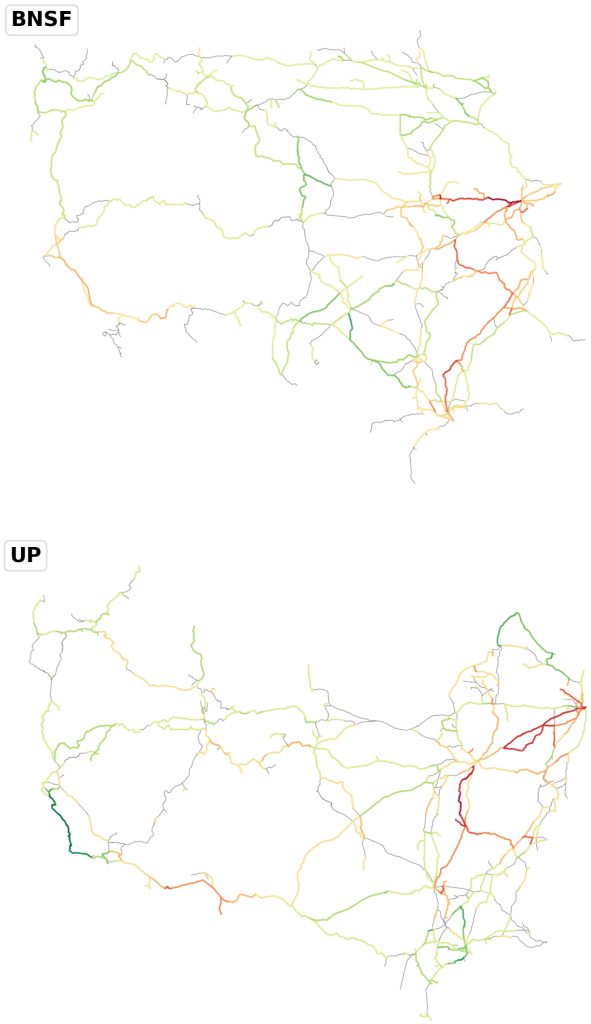


Figure 9: Effect of Risk on Network Flows Ranking, by Railroad

[2]See Appendix Figure 15

7.3 Infrastructure Adaptation Policy

I formalize adaptation decisions at the segment level along two dimensions. For each segment e, a railroad may choose the quantity of parallel track $\text{track}_e \in \{1, 2, \ldots, \bar{T}_e\}$ and whether to reinforce the segment against climate hazards, $A_e \in \{0, 1\}$. These two instruments correspond to the two modes of adaptation emphasized throughout the paper. Increasing track_e expands capacity and reduces the segment cost ζ_e, improving the segment's ability to absorb flows. Setting $A_e = 1$ eliminates disruption risk on the segment entirely, so that $D_e(s) = 0$ for all $s \in \mathcal{S}$, effectively hardening the link against climate hazards regardless of the realized state of the world. I plan to evaluate two classes of counterfactual policy. A *reallocation-centered policy* will identify segments for which track capacity should be expanded or reduced, redirecting investment toward parts of the network where the returns to capacity are highest given current and projected climate risk. A *reinforcement-centered policy* will identify segments for which hardening against climate hazards (setting $A_e = 1$) would generate sufficient reductions in expected shipping costs to justify the associated investment cost.

For each class of policy, I plan to evaluate welfare implications under alternative climate risk projections, comparing outcomes against the status quo baseline. The benefit of any adaptation policy is measured by its reduction in expected long-run transportation costs, which in the model maps to consumer surplus via the demand system and to producer surplus via firm profits. The cost is the associated construction and maintenance expenditure required to implement the policy. A policy is welfare-improving if the present discounted value of transportation cost savings exceeds the upfront and ongoing investment costs. By varying the scope and targeting of each policy, I aim to characterize which parts of the network offer the highest returns to adaptation investment, and how those returns vary as climate risk evolves over time.

8 Conclusion

This paper studies how climate risk affects transportation costs and provides a framework for evaluating how railroad infrastructure investment should adapt in response. Using shipment-level data and a network model of freight rail routing, I document that short-run disruptions from climate hazards lead primarily to operational adjustments rather than immediate price or quantity responses. Over longer horizons, however, persistent increases in hazard exposure are associated with declining rail flows. Calibrating the model to these patterns yields a preliminary estimate that climate disruptions raise the effective cost of using affected routes by roughly 61 percent on average.

These results are preliminary and several components remain in progress. Future drafts will refine the empirical analysis, estimate additional model parameters, and use the framework to evaluate counterfactual infrastructure policies, particularly the relative benefits of reinforcing vulnerable segments versus reallocating network capacity under evolving climate risk.

References

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

Association of American Railroads. (2024). *Freight rail facts & figures.* Retrieved from https://www.aar.org/data-center/

Balboni, C. (2024). In harm's way? infrastructure investments and the persistence of coastal cities. *American Economic Review*, *115*(1), 77-116.

Balboni, C., Boehm, J., & Waseem, M. (2024). Firm adaptation in production networks:

evidence from extreme weather events in pakistan. *PEDL Working Paper*.

Balboni, C., & Shapiro, J. (2025). Spatial environmental economics. *NBER Working Paper*.

Castro-Vincenzi, J., Khanna, G., Morales, N., & Pandalai-Nayar, N. (2024). Weathering the storm: Supply chains and climate risk. *NBER Working Paper*.

Chen, Y. (2024). Network structure and the efficiency gains from mergers: Evidence from the u.s. freight railroads. *Working Paper*.

Degiovanni, P., & Yang, R. (2023). Economies of scale and scope in railroading. *Working Paper*.

Donaldson, D., & Hornbeck, R. (2016). Railroads and american economic growth: A "market access" approach. *Quarterly Journal of Economics*, *131*(2), 799-858.

Du, X., & Wilson, A. J. (2026). Weather and U.S. railways: Risk, adaptation, and congestion. *Journal of Public Economics*, *255*, 105572. doi: 10.1016/j.jpubeco.2026.105572

Eakin, K., Meitzen, M., Bozzo, T., & Schoech, P. (2009). *A study of competition in the U.S. freight railroad industry and analysis of proposals that might enhance competition.* Retrieved from https://www.stb.gov/wp-content/uploads/files/docs/competitionStudy/Executive%20Summary.pdf

Fagjelbaum, P., & Schaal, E. (2020). Optimal transport networks in spatial equilibrium. *Econometrica*, *88*(4), 1411-1452.

Fuchs, S., & Wong, W. F. (2025). Multimodal transport networks. *Working Paper*.

Ghoreishi, B., Staes, B. M., Liu, C., & Wang, H. (2025). Mapping climate extremes and infrastructure vulnerability: A new database for U.S. railroads. *Scientific Data*. Retrieved from https://www.nature.com/articles/s41597-025-06205-z doi: 10.1038/s41597-025-06205-z

Hsiao, A. (2025). Sea level rise and urban adaptation in jakarta. *Working Paper*.

IPCC. (2021). Weather and climate extreme events in a changing climate. In *Climate change 2021: The physical science basis. contribution of working group I to the sixth assessment report of the intergovernmental panel on climate change* (pp. 1513–1766). Cambridge,

United Kingdom and New York, NY, USA: Cambridge University Press. doi: 10 .1017/9781009157896.013

Massoni, H. (2025). Climate trade costs: Extreme weather, transportation, and supply chains. *Working Paper*.

Wilson, W. W., Wilson, W. W., & Koo, W. W. (1988). Modal competition in grain transport. *Journal of Transport Economics and Policy*, *22*(3), 319–337.

A Additional Figures and Tables

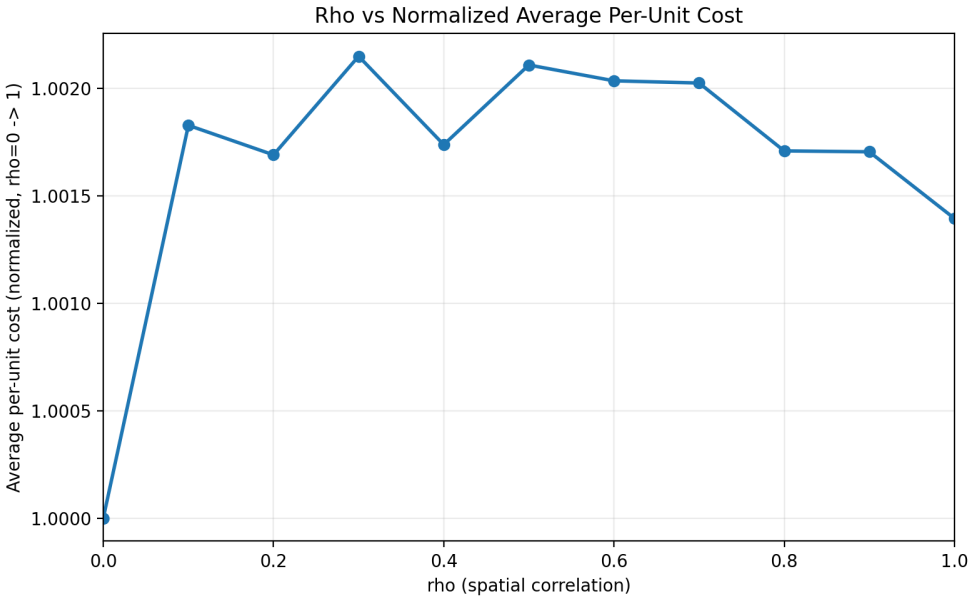


Figure 10: ρ vs. Normalized Average Per-Unit Cost

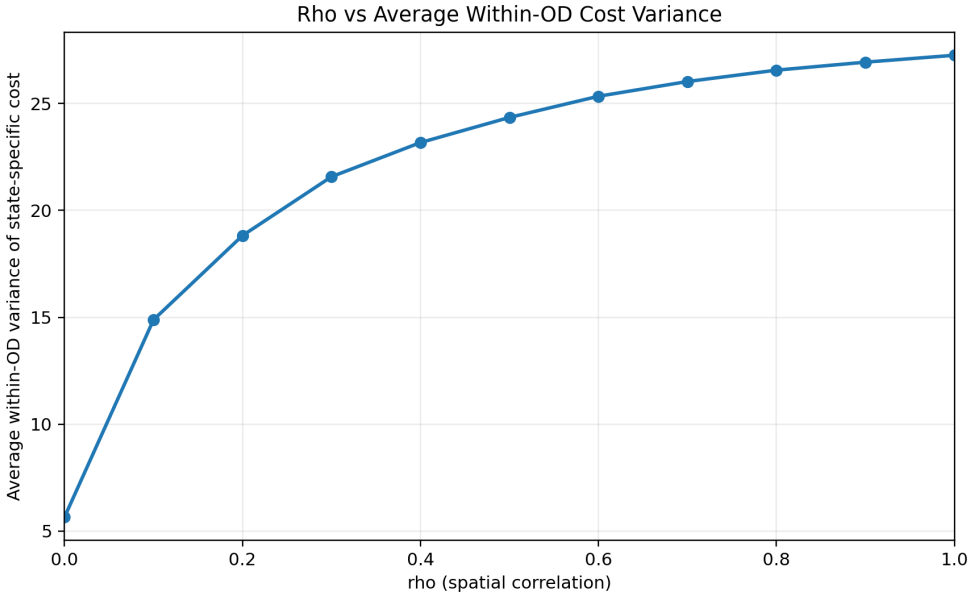


Figure 11: ρ vs. Average Within-OD Cost Variance

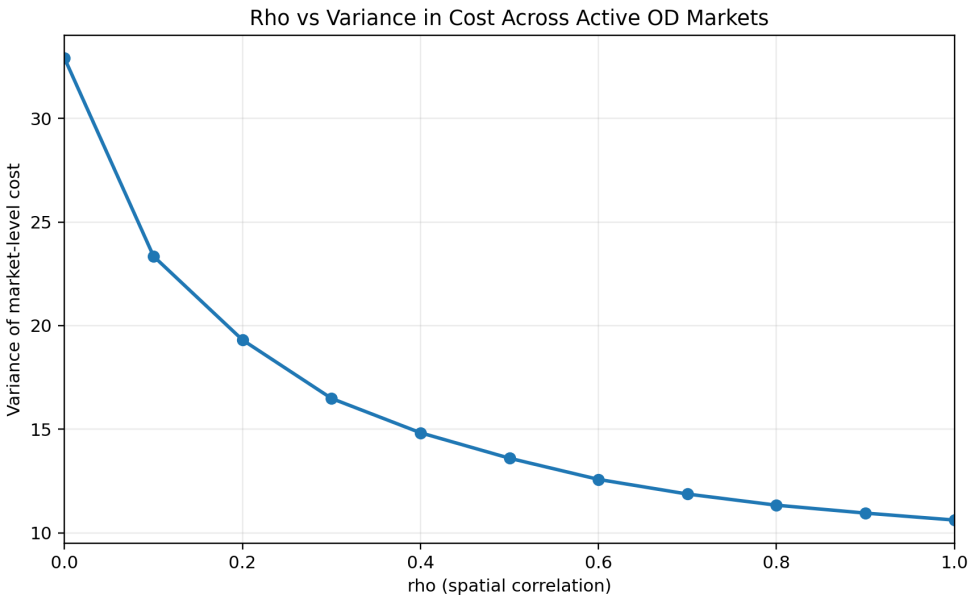


Figure 12: ρ vs. Expected Cost Dispersion

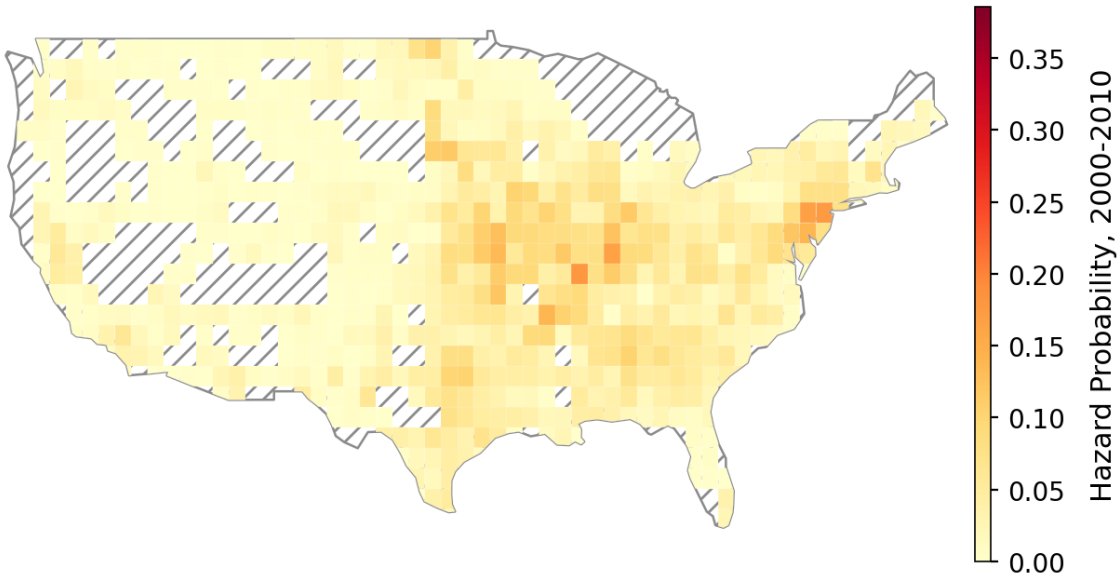


Figure 13: Spatial Distribution of Hazard Risk, 2000-2010

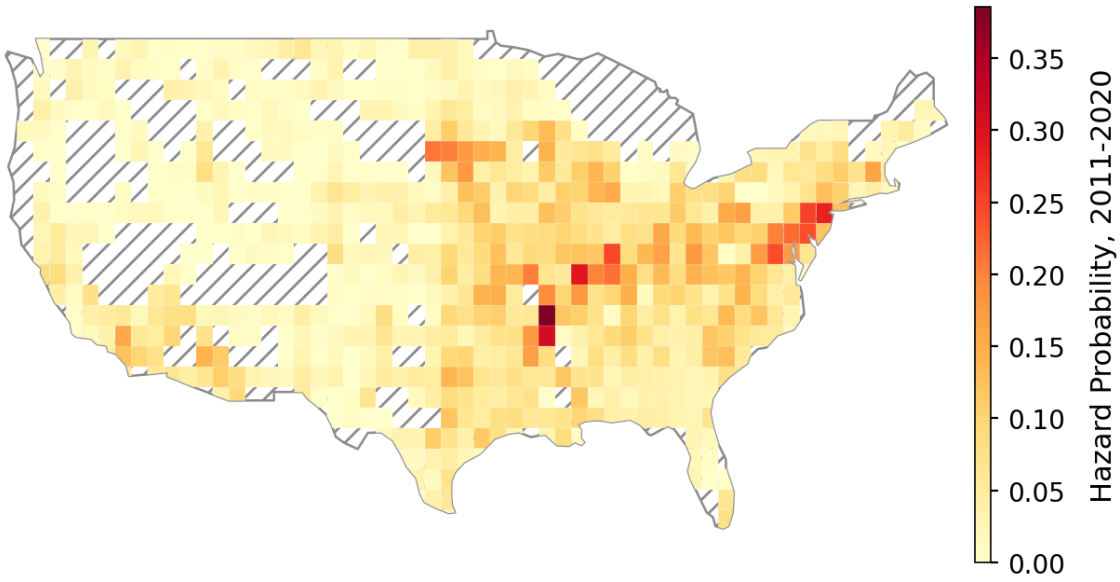


Figure 14: Spatial Distribution of Hazard Risk, 2011-2020

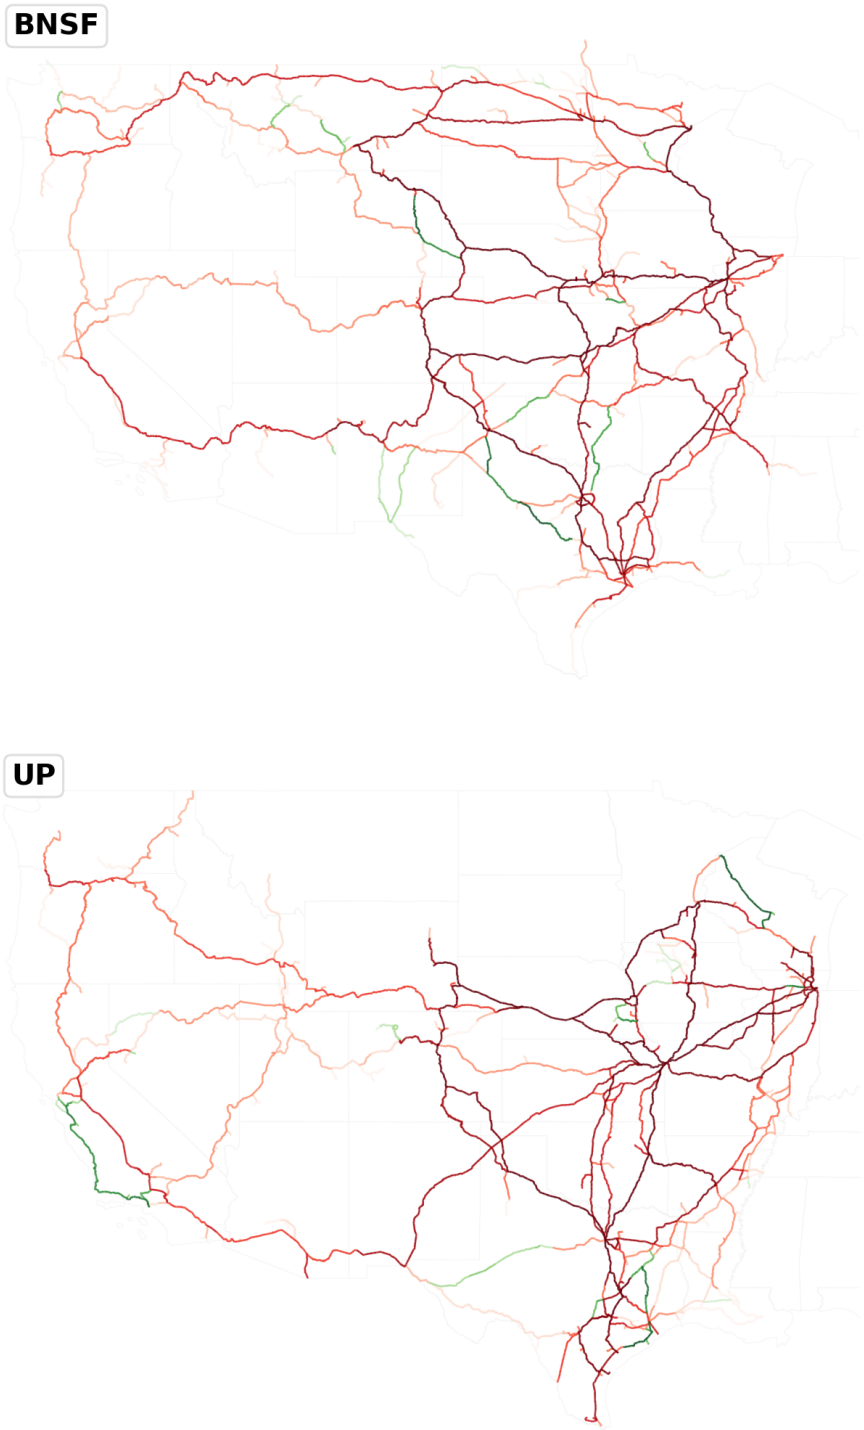


Figure 15: Effect of Risk on Network Flows, by Railroad

B Proofs of Propositions

To be added.

C Maximum Likelihood Estimation of ρ

To be added.

D Identification of Congestion Parameter β

Abstracting away from risk, the model implies the following relationship between route shares and relative costs:

$$\frac{s_{in,R,r}}{s_{in,R,r'}} = \left(\frac{c_{in,R,r}}{c_{in,R,r'}} \right)^{-\frac{1}{\beta}} \tag{11}$$

For waybills with an inferred route, I construct model-implied cost estimates $\hat{c}_{in,R,r}$ and estimate β via the log-linearized regression:

$$\log \left(\frac{s_{in,R,r}}{s_{in,R,r'}} \right) = -\frac{1}{\beta} \log \left(\frac{\hat{c}_{in,R,r}}{\hat{c}_{in,R,r'}} \right) + u_{in,R,r} \tag{12}$$

The error term reflects measurement error in the route inference procedure, as well as ignoring hazard risk, which is plausibly orthogonal to model-implied costs, yielding consistent identification of β.

Table 3: Inferred-route cost dispersion estimates, 2018–2020

Railroad	Coef. on $\log$(cost ratio)	SE	β	95% CI for β	N
BNSF	-5.8253	0.1958	0.1717	[0.1611, 0.1838]	20,770
CSXT	-6.0566	0.1715	0.1651	[0.1564, 0.1748]	13,072
NS	-7.6229	0.1677	0.1312	[0.1258, 0.1371]	16,134
UP	-7.0942	0.2058	0.1410	[0.1334, 0.1495]	25,026
Pooled	-6.5254	0.1086	0.1532	[0.1484, 0.1584]	75,037

Notes: Unit of observation is origin-destination-railroad route cell. Shares are computed within origin-destination-railroad. Route cost is the average across inferred candidate routes of the multiplicative iceberg cost with segment multiplier $1 + (\text{dist miles}/1000)/\text{track}^{0.3}$. The denominator route is the minimum-reported-distance route within each origin-destination-railroad. $\beta = -1/\hat{\gamma}$, where $\hat{\gamma}$ is the coefficient on $\log$(cost ratio).

Table 3 reports estimates of β from the route share regression. The pooled estimate is $\hat{\beta} = 0.153$, with consistent estimates across carriers ranging from 0.131 to 0.172. The baseline value of $\beta = 0.1$ lies modestly below this range, so the parameterization is conservative relative to the data.

E Segment-level Flows

It can be shown that the share of trips between i and n that use segment jl for railroad R is given by:

$$x_{in,jl,R} = \varsigma_{jl,R}^{-\frac{1}{\beta}} \frac{m^{-\frac{1}{\beta}} b_{ij,R} b_{ln,R} + (1 - m^{-\frac{1}{\beta}})(1 - D_{jl,R}(s)) b_{ij,R}(s) b_{ln,R}(s)}{m^{-\frac{1}{\beta}} b_{in,R} + (1 - m^{-\frac{1}{\beta}}) b_{in,R}(s)}$$

Aggregating over all origin-destination pairs, and weighing by quantity, segment flows is equal to:

$$X_{jl}(s) = \sum_{in} Q_{in,R} \cdot x_{in,jl,R}(s)$$

F Short-Run Responses to Hazard Disruptions

To estimate short-run responses of freight rail flows to climate-related disruptions, I employ a stacked event study design centered on discrete hazard episodes from Ghoreishi et al. (2025). The unit of observation is a BEA-level origin-destination pair od observed in week t within a symmetric ± 8 week window around an event occurring at time a.

F.1 Empirical specification

For each event a, I estimate:

$$y_{odt}^{a} = f\left(\sum_{\substack{h=-8 \\ h\neq -1}}^{8} \{\lambda_h \mathbf{1}\{t = a+h\} + \delta_h \cdot D_{od}^{a} \cdot \mathbf{1}\{t = a+h\}\} + \alpha_{od}^{a} + \gamma_{w(t)} + \varepsilon_{odt}^{a}\right), \tag{13}$$

where h indexes event time and the omitted period is $h = -1$. The coefficients of interest, δ_h, capture differential responses by exposure intensity relative to the week immediately preceding the disruption.

The function $f(\cdot)$ corresponds to the estimation method: I use PPML for quantity outcomes so that zero flows are retained, and log-linear OLS for price outcomes on the positive-flow sample.

The specification includes event-by-OD fixed effects α_{od}^{a}, which absorb all time-invariant differences across OD pairs within each event window, and week-of-year fixed effects

$\gamma_{w(t)}$ to control for seasonality. The λ_h terms capture the average event-time path common to all OD pairs within an episode, while the δ_h coefficients isolate how this path varies with exposure.

F.2 Treatment construction

The treatment variable D^a_{od} measures the share of pre-event flows on OD pair od that are routed through network segments affected by event a. Its construction proceeds in two steps, mapping from station-level flows and routes to BEA-level OD markets.

First, for each station-level OD pair p and route $r \in \mathcal{R}_p$, define:

$$D^a_{pr} = \begin{cases} 1 & \text{if route } r \text{ intersects any segment affected by event } a, \\ 0 & \text{otherwise.} \end{cases}$$

Let π^a_{pr} denote the share of pre-event flows on station pair p assigned to route r, and let ρ^a_{odp} denote the share of pre-event flows on BEA OD pair od that are attributed to station pair p. The OD-level treatment intensity is then:

$$D^a_{od} = \sum_{p \in \mathcal{S}_{od}} \rho^a_{odp} \sum_{r \in \mathcal{R}_p} \pi^a_{pr} \cdot D^a_{pr}. \tag{14}$$

This measure lies in $[0, 1]$ and captures the fraction of baseline flows that rely on routes exposed to the disruption. Variation in D^a_{od} arises from heterogeneity in route structure across OD pairs and their interaction with the spatial footprint of the hazard.

F.3 Interpretation

The coefficients δ_h measure how outcomes evolve over event time for more-exposed OD pairs relative to less-exposed OD pairs within the same disruption episode. Because D^a_{od} is continuous, the estimates should be interpreted as semi-elasticities with respect to exposure intensity rather than average treatment effects of a binary intervention.

F.4 Event definition and specification grid

Hazard episodes must be admitted based on thresholds for severity, duration, and minimum exposure. I therefore estimate the model across a grid of admissible event definitions. These specifications should be interpreted as sensitivity exercises over plausible definitions of economically meaningful disruptions, rather than conventional robustness checks, since changing the thresholds alters both the number and composition of events in the sample.

In practice, more restrictive definitions yield fewer events and therefore rely on limited identifying variation, making those specifications closer to case-study evidence.

F.5 Results

Across the full specification grid, I find no evidence of short-run price responses. Estimates of revenue per ton-mile are close to zero at all horizons and are statistically insignificant across all event definitions, including in the spot-market sample where prices are most flexible.

Quantity responses are also limited. In the full sample, estimates are generally near

zero throughout the event window. In the spot-market sample, some specifications exhibit modest negative effects in the weeks immediately following disruption onset, but these effects are small and not stable across alternative event definitions.

Taken together, these results indicate that short-run responses to climate-related disruptions do not operate through observable price adjustments and involve, at most, limited reductions in aggregate shipped quantities. Instead, adjustment appears to occur primarily along operational margins not directly captured by prices or OD-level shipment volumes.