

Regulation and the Cost of Moving Goods

Bentley Coffey* Patrick A. McLaughlin† Pietro F. Peretto‡

March 20, 2026

Abstract

We study how federal regulations affect the freight transportation industry in the United States over the last 50 years. In our specification, firms produce ton-miles of transportation from 3 input factors: labor, fuel, and capital. Building on the structural model of regulatory accumulation and investment in productivity developed by Coffey, McLaughlin, and Peretto (2020), firms can invest in improving each individual factor-specific productivity. Using RegData-based counts of regulatory restrictions by industry paired with common public data sources, we estimate the coevolution of labor, fuel, and capital productivity with the accumulation of regulatory restrictions. Regulatory restrictions are associated with lower labor productivity in all four modes, lower fuel productivity in air and rail, and lower capital productivity in air, rail, and water; trucking also shows negative but less precisely estimated effects on fuel and capital. Using our estimates, we perform a different counterfactual simulation for each mode where it experiences a 5 percent increase in regulatory restrictions in 2018. The counterfactuals indicate that such an increase raises unit costs and prices by roughly 0.8–2.3 percent and reduces quantities shipped in the targeted mode by about 1.4–4.1 percent. Because of the growth process, these losses persist and compound over time. Our results provide industry-level microfoundations for the macroeconomic evidence that regulatory accumulation depresses investment, productivity, and output for a crucial economic activity of transporting goods.

*Coffey: Department of Economics, University of South Carolina, USA (bentleygcoffey@gmail.com).

†McLaughlin: Hoover Institution at Stanford University, USA (pm2@stanford.edu).

‡Peretto: Department of Economics, Duke University, USA (peretto@econ.duke.edu).

1 Introduction

Freight transportation links modern production centers into supply chains. Nearly every good consumed in the United States spends part of its life on a truck, train, plane, or ship. Even small changes in the cost of moving freight propagate through supply chains, ultimately affecting aggregate output and welfare. Regulatory policy is an important determinant of shipping costs. U.S. transportation industries have experienced both explicit deregulation — most notably the Motor Carrier Act of 1980 and the Staggers Rail Act — and a slower but persistent accumulation of safety and environmental requirements. These regulatory constraints act as an effective tax on moving goods and can even overwhelm efficiency improvements in transportation technology. This paper examines how regulatory accumulation has affected the cost structure, innovation incentives, and modal allocation of U.S. freight transportation.

A growing literature documents that regulation, when measured as a stock that accumulates over time, has sizable adverse effects on aggregate performance. Dawson and Seater [2013] construct a time-series index of federal regulation and find that regulation added since 1949 has reduced annual U.S. output growth by roughly one percentage point on average, primarily through slower total factor productivity growth. Much of the more recent work uses versions of RegData, which transform legal text into quantitative measures of regulatory restrictions by industry and over time. Al-Ubaydli and McLaughlin [2017] introduce the original RegData database and McLaughlin and Sherouse [2019] introduce machine learning approaches to extend the panel data, which now supports many empirical applications. Coffey et al. [2020] embed these data in a dynamic growth model in which regulatory accumulation distorts firms’ incentives to invest in productivity; they estimate that federal regulatory accumulation since 1980 reduced the level of U.S. GDP in 2012 by roughly one quarter relative to a counterfactual with the 1980 stock of regulation frozen. Subsequent work by Coffey and McLaughlin [2021] and McLaughlin and Wong [2024] shows that systematic reductions in regulation — via regulatory budgeting in British Columbia or reductions in state regulatory stocks (respectively) — are associated with measurably higher growth.

Regulation is also not neutral across the income distribution. Chambers et al. [2019] document that higher regulatory burdens are associated with higher state-level poverty rates, and Chambers and O’Reilly [2022] show that regulation is associated with higher income inequality, consistent with the view that compliance costs are regressive and tilt labor demand toward compliance-intensive activities. In network industries such as rail, the level and

composition of regulation matter for safety as well as efficiency. Ellig and McLaughlin [2016] use RegData-based indices of economic regulation by the Interstate Commerce Commission and Surface Transportation Board and safety regulation by the Federal Railroad Administration to show that partial economic deregulation under the Staggers Act is associated with improved railroad safety, while the marginal contribution of subsequent expansions in safety regulation is comparatively modest once railroads' capital expenditures were less constrained by the economic regulations of the Interstate Commerce Commission.

Transportation economics provides complementary evidence on the consequences of regulation and deregulation in freight markets. Winston [1981] and related work estimate disaggregate freight demand models in which shippers choose modes based on prices, transit times, and service characteristics. Winston et al. [1990] show that surface freight deregulation generated large reductions in average rates, sizable productivity gains, and substantial shipper benefits as competition and efficiency were restored to both rail and trucking. Our analysis builds on this sectoral literature but connects it explicitly to the regulatory accumulation agenda: rather than studying one-time deregulatory events, we examine how the evolving stock of regulation affects factor productivities and the allocation of traffic across modes.

A complementary theoretical rationale for focusing on transportation comes from Coffey et al. [2025], which develops a multi-industry endogenous-growth model in which transportation is not a downstream industry but rather an in-house area for indispensable innovation that is integral to the production process. In that framework, intermediate goods producers must both make and ship their goods; that shipping is subject to per-unit transportation costs which do not automatically fall with innovation in the factory. As a result, transportation acts as a bottleneck in the innovation-growth nexus: when per-unit shipping costs are high (because of regulation, infrastructure constraints, or fuel price pressures), productivity improvements in manufacturing fail to scale into proportionate sales gains and that long-run growth can actually stall. Conversely, reductions in transportation costs amplify the returns to process innovation by enabling firms to reach larger markets at lower marginal cost. This perspective underscores that understanding economic performance requires analyzing production and distribution jointly rather than treating transportation as an auxiliary sector. It also suggests that government-induced distortions in freight transportation, such as regulatory accumulation, may have especially large economic consequences relative to distortions in more insular industries.

Using RegData U.S. 6.0 restrictions mapped to NAICS codes 481–484 (Makridis and

McLaughlin [2026]), we quantify the stock and evolution of federal regulatory restrictions facing air, rail, truck, and water freight carriers. We merge these measures with BEA data on value added, capital stocks, depreciation, wages, and employment. We use BTS data on prices (average revenue per ton-mile) and quantities (ton-miles) as well as our somewhat novel source for the number of competitors: a web-scraped count of firms with Class I railroads and establishment counts from County Business Patterns for the other freight modes. Other data sources, such as the Transportation Energy Data Book and FRED, complete our data set. We then embed these data in a structural model with two key components. On the supply side, we assume a nested CES technology in which firms combine labor with a composite of nonhuman services produced from fuel and capital, which allows us to infer input productivities and substitution parameters from observed cost shares and input prices. On the demand side, we adopt the discrete/continuous choice framework of Dube et al. [2025], which is equivalent to a BLP-esque logit RUM demand system over modes with an explicit representative-consumer interpretation. This structure yields mode-specific markups for each firm and demand elasticities that we then use to simulate the effects of regulatory shocks on prices, quantities, and modal substitution.

The empirical strategy proceeds in four steps. First, we estimate demand for freight transportation across modes, using mode-level prices and average transit times (constructed from major north-south routes) to recover own-price elasticities and implied markups using unit costs as an instrument for price. Second, given the estimated markups and observed revenues, we recover input cost shares and construct measures of labor, fuel, and capital productivity by mode and year. Third, we relate the evolution of these factor productivities to regulatory accumulation, controlling for market scale, input prices, and the real interest rate. Finally, we use the estimated parameters to simulate counterfactual regulatory policies— a 5 percent increase in regulatory restrictions for a particular mode in 2018—and trace the implications for costs, prices, quantities, and modal substitution. In doing so, we provide sector-level microfoundations for the macro evidence that regulatory accumulation slows productivity growth and worsens distributional outcomes.

2 Data

From the Bureau of Economic Analysis (BEA), we have collected the available annual data on each mode (m) as an industry: air, rail, water, and truck as NAICS 481, 482, 483, and 484

(respectively).¹ The collected variables include value-added to GDP (Y_{mt}), investment (I_{mt}), the current cost of the capital stock $[PK]_{mt}$ for each mode as an industry, an index measure of the capital stock (K_{mt})², depreciation of the capital stock (D_{mt}), employment measured as the labor hours of FTEs (L_{mt}), and the wage bill as total compensation of employees (W_{mt}). We also collect a few other variables from the BEA: nominal GDP per capita (y_t), total final goods produced (Y_{Gt}) measured as value-added to nominal GDP (from NAICS 11 through 31-33, except for 21), and total final goods consumed (G_t).

Most of the transformations that we perform on BEA data are simple and need no explanation (eg dividing total compensation of employees by the number of FTEs to get an annual wage and then again by 2000 to get an hourly wage instead of an annual wage); however, our transformations involving the capital variables need some explanation. To compute the depreciation rate, we follow Bai et al. [2022] who were in contact with BEA to discover the way that they generate the depreciation variable from microeconomic estimates of depreciation in each industry; the resulting estimated depreciation rate (δ_{mt}) is given by:

$$\delta_{mt-1} = \frac{D_{mt}}{K_{mt-1} + \frac{1}{2}I_{mt-1}} \quad (1)$$

Dividing the current cost of the capital stock by our rescaled index of the capital stock naturally yields the asset price of capital:

$$P_{mt} = \frac{[PK]_{mt}}{K_{mt}} \quad (2)$$

¹The other industries within the 2 digit NAICS code of 48 (ie Transit and Ground Passenger Transportation, Pipeline Transportation, Other Transportation and Support Activities [such as Sightseeing], Warehousing and Storage) pick up enough of the similar [yet not quite comparable] economic activities so that this remains a fairly clean comparison between the modes for freight transportation. Prior to 1997 adoption of NAICS codes, the SIC codes map fairly well into the NAICS (although truck transportation has to be scaled down via BEA's method to keep the series from being discontinuous). However, the NAICS code 481, covers the air transportation of both freight and passengers. Lumping passenger and freight transportation by air is understandable, especially when we consider that passenger airlines routinely devote unoccupied space in their holds to high value freight. We deal with this problem by dividing the quantities for air transportation by the fraction of ton-miles attributable to passengers, multiplying the observed passenger-miles from BTS data by the average weight of American adults in the periodic NHANES waves from 1971 to 2023 (with interpolation between years), plus an assumed 40 pounds of luggage (including carry-ons) per passenger. With these measurement issues, the data on air transportation of freight should be considered quite noisy and we have carefully investigated our results to make sure that everything is not being driven by the air data.

²We normalize BEA's capital stock quantity index so that its value in 2017 equals estimates of the number of self-propelled vehicles in each industry. This rescaling is for convenience of interpretation and is entirely innocuous for our results.

Yet, our analysis relies on the rental price of capital (r_{mt}) instead of the asset price. The conversion should be fairly straightforward in theory. A simple arbitrage argument should equate the rental price of capital to the cost of buying the capital, using it for a single period so that it depreciates, and then sell it the next period (discounting it to present value) by the interest rate:

$$r_{mt} = P_{mt} - P_{mt+1}(1 - \delta_{mt-1})/(1 + i_t) \quad (3)$$

In practice, owners of capital at time t do not know its price at time $t+1$; hence, we use the expected price next period as predicted by the asset price in the current period. We accomplish this by regressing the first difference of the log of capital asset prices on an intercept and 2 dummies, one indicating that the year is at least as recent as 1981 and another indicating that the year is at least as recent as 2020. This specification appears necessary and appropriate after examining the time path of asset prices, which increases linearly in logs (indicating exponential growth) but with clear changes in slope at these thresholds. We suspect that the change in trend at those thresholds is due to a change in BEA's data collection methodology but cannot rule out structural changes to discounting with 1981 and 2020 corresponding to our observed maximum and minimum interest rates (respectively). Figure 1 plots the time paths of our computed rental prices of capital.

Ever since it became standard in macroeconomics to model innovation as the result of costly investment by monopolistic competitive firms seeking to reduce their costs a la Romer [1990], firms' market power (ie number of firms) is an important variable for explaining the amount of investment-driven innovations that we observe [at the industry-level]. As in Coffey et al. [2020], our primary source measuring market power is the industry-level number of establishments in the US, which we take from the County Business Patterns data collected by the US Census Bureau. However, that data source is missing a measure for rail transport. Hence, we have scraped the web for the life-cycle of every rail line that has ever been classified as Class I, which are the major rail lines (according to a revenue threshold that has risen over the course of the last century) subjected to the maximal amount regulations in the books. We then compute the number of independently owned and operated rail lines each year as our measure of rail transport firms.

Figure 2 displays two alternative reconstructions of the Class I railroad count, one manual and one algorithmic, which align closely over the recent decades used in our analysis. Our count of Class I firms will overstate market power because it excludes the Class II and III

carriers (i.e., shortline and regional railroads). Nonetheless, we believe that our count of Class I firms is the best available measure.³ We should note that the BTS data on the rail industry, our source for ton-miles and prices, is also exclusively constructed from Class I railroad data.

From the Bureau of Transportation Statistics (BTS) by mode and year, we have collected a measure of price: average revenue per ton (p_{mt}). As can be seen in Figure 3, the levels across the industries are separated by an order of magnitude with air transport being an order of magnitude above truck transport, which is then an order of magnitude above rail transport and water transport. All prices are growing at roughly the same rate since the mid-2000s. Prior to that, all prices were growing along that same trend except for rail transport. From the early 1980s to the mid-2000s, the price (i.e., the nominal price) of rail transport actually fell. From BTS we also collected the quantity (Q_{mt}) of freight measured in ton-miles as graphed in Figure 4. Air transport of freight rose from the lowest initial level at the fastest growth rate until a shock in fuel prices in the mid-2000s pushed its trend downward until the end of that shock in the mid teens when it rebounded to its initial trend (with a notable upward blip during the pandemic when so many consumers increased their online orders of goods from sellers like Amazon). Truck transport starts at the highest level and trends upward (albeit at a slower growth rate than air transport) except for a response to the fuel price shock that mimicked the temporary drop experienced by air transport. Rail transport follows truck transport until the fuel price shock, where it remains roughly flat (instead of falling like the other modes) before its trend begins falling in the mid-2000s. Water transport was roughly flat until the early nineties and has been declining ever since.

Although BTS appears to sometimes import some data on fuel consumption (eg from EIA), a more complete source for annual fuel consumption (F_{mt}) by mode is the periodic publication of the Transportation Energy Data Book by researchers at Oak Ridge National Laboratory; unfortunately, the last published edition of that publication ends with 2019 data and the latest edition has been delayed due to budget cuts. The movement of the relevant fuel prices (eg diesel, jet fuel, and residual fuel oil) mimics the well-known movement of

³Class II and III railroads are smaller railroads, by revenue and trackage network (distance and area); these railroads largely consist of lines that were less profitable and spun off from Class I railroads once enabled by the deregulatory Staggers Act of 1980. There is no consistent time series for the number of Class II and III railroads; even the American Short Line and Regional Railroad Association (ASLRRA), which is the trade association that represents the interests of Class II and III railroads, does not maintain a database of the number of firms per year for the last 5 decades. Fortunately, both industry experts and ASLRRA's available years of cross-sectional data point to the consolidation pattern observed for Class I rail having been matched in Class II and III rail where a few larger firms (such as Patriot Rail and Genesee & Wyoming) have been steadily purchasing many smaller rail lines in the last few decades.

gasoline prices during this time period. Fuel prices fell slightly from the early 1980s to the early 1990s before resuming their upward trend with a notable shock (exceeding the trend) from the mid-2000s to the mid-2010s.

From FRED, we collected annual data on the average yields for Moody’s seasoned BAA rated corporate bonds as the risk-adjusted nominal interest rate relevant for firm investment, as well as fuel price (f_{mt}) by mode.⁴ From RegData 6.0, we collected the number of regulatory restrictions (R_{mt}) applicable each year to each mode (m) as an industry. Figure 5 graphs our data on regulatory restrictions. Although regulatory restrictions tend to grow fairly steady over time for each industry, they are separated by an order of magnitude in their levels (i.e., rail and water have 10 times more than trucking while air has 10 times more than rail and water). There are also two notable exceptions in trend that aid with identification: trucking enjoys a prolonged period of deregulation from the late 1970s to the mid 1990s and all modes experienced a regulatory freeze during the first Trump administration. Identification is also aided by the aforementioned exogenous shock raising fuel prices trend for a decade.

Although we are primarily focused on supply-side shifts due to costs from regulations, we are also interested in the extent to which regulatory costs can induce modal substitutions. Hence, we need to adequately control for confounding trends in shifting demand. This requires us to collect a reasonable measure of transit time because quicker transit time has been the primary driver for the growing market share of air and truck transportation relative to rail and water transportation. Even though those slower two modes are far cheaper, richer consumers are increasingly willing to pay more to get their final goods sooner via air or truck.⁵ For each mode (m), we construct a proxy for transit time (T_m) using approximate route-level transit times compiled from publicly available web sources for three major north-south routes (NYC-Miami, LA-Seattle, and Houston-Chicago). Those transit times appear in Table 1.

⁴Rail and truck use diesel also exclusively but air freight relies on jet fuel, which is closer to kerosene, and water uses a mix of distillate fuels and a much less refined residual fuel oil. Naturally, all of these fuel prices are so strongly correlated to the price of a barrel of crude oil that missing observations can be readily inferred via a simple forecast (or backcast).

⁵Rail and water transportation are inherently cheaper due to advantages conveyed by the laws of physics; the low friction of steel wheels on steel rails (as well as the drafting effect minimizing air resistance) means that trains can be remarkably fuel efficient and the high density of water largely nullifies the cost of the transporting higher tonnage on ships (fuel usage increases with the volume displaced below the water line but hydrodynamic drag can be greatly reduced via our long evolved hull designs).

3 Empirical Design

We begin by specifying the static production technology that transforms input factors into ton-miles of output. Next we characterize the firm's dynamic problem of choosing levels of investment that yield gains in input factor productivity. Then we conclude this section by specifying the demand for transportation, which determines the firms' markup factors and substitution across modes.

3.1 Static Production Technology

We specify production as a nested CES because of its well-known flexibility and generality. We begin by specifying the outer nest, which is a CES production technology for firm j producing mode m ton-miles (Q_{jm}) of transportation services from human labor (L) and nonhuman services (S):

$$Q_{jm} = \left[(A_{Ljm}L_{jm})^{\theta_{LSm}} + (S_{jm})^{\theta_{LSm}} \right]^{\frac{1}{\theta_{LSm}}} \quad (4)$$

Where A_{Ljm} is labor productivity. For brevity, we will suppress some subscripts. The unit cost function for producing 1 unit of Q is given by:

$$c_Q = \left[\left(\frac{w}{A_L} \right)^{\frac{\theta_{LS}}{\theta_{LS}-1}} + c_S^{\frac{\theta_{LS}}{\theta_{LS}-1}} \right]^{\frac{\theta_{LS}-1}{\theta_{LS}}} \quad (5)$$

Where c_S is the unit cost of nonhuman services. The conditional factor demands are:

$$L = \left[\frac{\left(\frac{w}{A_L} \right)^{\frac{\theta_{LS}}{\theta_{LS}-1}}}{\left(\frac{w}{A_L} \right)^{\frac{\theta_{LS}}{\theta_{LS}-1}} + c_S^{\frac{\theta_{LS}}{\theta_{LS}-1}}} \right] \frac{c_Q Q}{w} \quad S = \left[\frac{c_S^{\frac{\theta_{LS}}{\theta_{LS}-1}}}{\left(\frac{w}{A_L} \right)^{\frac{\theta_{LS}}{\theta_{LS}-1}} + c_S^{\frac{\theta_{LS}}{\theta_{LS}-1}}} \right] \frac{c_Q Q}{c_S} \quad (6)$$

The bracketed terms in these conditional factor demands are the unit cost share attributable to labor and nonhuman services, respectively. We then specify the inner nest, which produces nonhuman services from a CES production technology that combines fuel (F) and capital (K):

$$S = \left[(A_F F)^{\theta_{FK}} + (A_K K)^{\theta_{FK}} \right]^{\frac{1}{\theta_{FK}}} \quad (7)$$

Where A_F and A_K are fuel and capital productivities, respectively. The unit cost function for S is then just:

$$c_S = \left(\left(\frac{f}{A_F} \right)^{\frac{\theta_{FK}}{\theta_{FK}-1}} + \left(\frac{r}{A_K} \right)^{\frac{\theta_{FK}}{\theta_{FK}-1}} \right)^{\frac{\theta_{FK}-1}{\theta_{FK}}} \quad (8)$$

Where f is the fuel price and r is the rental price of capital. The conditional factor demands are given by:

$$F = \left[\frac{\left(\frac{f}{A_F} \right)^{\frac{\theta_{FK}}{\theta_{FK}-1}}}{\left(\frac{f}{A_F} \right)^{\frac{\theta_{FK}}{\theta_{FK}-1}} + \left(\frac{r}{A_K} \right)^{\frac{\theta_{FK}}{\theta_{FK}-1}}} \right] \frac{c_S S}{f} \quad K = \left[\frac{\left(\frac{r}{A_K} \right)^{\frac{\theta_{FK}}{\theta_{FK}-1}}}{\left(\frac{f}{A_F} \right)^{\frac{\theta_{FK}}{\theta_{FK}-1}} + \left(\frac{r}{A_K} \right)^{\frac{\theta_{FK}}{\theta_{FK}-1}}} \right] \frac{c_S S}{r} \quad (9)$$

Substituting the expression for c_S into c_Q yields an expression for the unit cost of Q that depends on the prices and productivities of all 3 inputs:

$$c_Q(w, r, f; A) = \left[\left(\frac{w}{A_L} \right)^{\frac{\theta_{LS}}{\theta_{LS}-1}} + \left(\left(\frac{f}{A_F} \right)^{\frac{\theta_{FK}}{\theta_{FK}-1}} + \left(\frac{r}{A_K} \right)^{\frac{\theta_{FK}}{\theta_{FK}-1}} \right)^{\frac{\theta_{LS}}{\theta_{FK}} \left[\frac{\theta_{FK}-1}{\theta_{LS}-1} \right]} \right]^{\frac{\theta_{LS}-1}{\theta_{LS}}} \quad (10)$$

Where A is a vector of the 3 productivities. We do not directly observe the productivities $\{A_L, A_F, A_K\}$, the substitutability parameters $\{\theta_{LS}, \theta_{FK}\}$, or the quantity of nonhuman services $\{S\}$. Although we could estimate the substitutability parameters from our data in principle, from which the productivities (and quantity of nonhuman services) are easily computed, the functional forms identifying those substitutability parameters introduce a relatively large amount of instability in the unobserved stochastic term (ie noise) to get reliable results from a relatively limited number of data points. Hence, we proceed under the assumption that those input factors are perfect complements (ie $\theta_{LS} = \theta_{FK} = -\infty$). Although this may at first appear extremely rigid, we note that it only restricts factors to not be substitutable in the short run. In the long run, a firm can respond to a spike in an input factor by investing in innovation in that factor's productivity so that less of it is needed in the input mix. Moreover, Leontief production reduces the productivities to simple

and intuitive measures: the average product of each input factor.⁶

3.2 Dynamic Problem of Investing in Input Factor Productivity

Having specified the firm's static technology, we can turn to its dynamic problem of investing in improvements to productivity (as well as its endogenous price setting). We can specify the instantaneous profit of firm j offering transportation services of ton-miles via mode m :

$$\pi_j = [p_j - c_{Qj}(w, r, f; A)] Q_j - [I_{Lj} + I_{Fj} + I_{Kj}] \quad (11)$$

The firm can invest in R&D to improve the productivity of its factors of production but regulatory shocks can dampen/enhance the efficacy of that investment:

$$\frac{\dot{A}_{Lj}}{A_{Lj}} = \xi_L(R) \frac{I_{Lj}}{w} \quad \frac{\dot{A}_{Fj}}{A_{Fj}} = \xi_F(R) \frac{I_{Fj}}{w} \quad \frac{\dot{A}_{Kj}}{A_{Kj}} = \xi_K(R) \frac{I_{Kj}}{w} \quad (12)$$

The Current Value Hamiltonian (CVH) for firm j is given by:

$$CVH = [p_j - c_{Qj}(w, r, f; A)] Q_j - [I_{Lj} + I_{Fj} + I_{Kj}] + \nu_{Lj} \dot{A}_{Lj} + \nu_{Fj} \dot{A}_{Fj} + \nu_{Kj} \dot{A}_{Kj} \quad (13)$$

where the state variables are the productivities and the corresponding costate variables are the shadow value of the marginal unit of that respective productivity. The price of transportation services and the investments are the control variables. The first order condition for the price yields the familiar markup over marginal cost (and that markup factor depends on the [own] price elasticity of demand, η):

$$p_j = \left(\frac{\eta_j}{\eta_j - 1} \right) c_{Qj} \quad (14)$$

Linearity of the Hamiltonian in investment then results in a familiar bang-bang structure.

⁶In defense of our assumption of perfect complements, we should note that BEA employs the same Leontief assumptions in constructing their periodic Input-Output tables describing the flow of intermediate goods between industries. We also believe that industry experts would not be alarmed at this specification, especially once they consider that firms can invest in the efficiency of each input. In the general case, each productivity is equal to the input factor's average product scaled by an expenditure ratio raised to the power of the relevant substitutability parameter; when viewed from the perspective of a skeptical economist constructing the empirical design, the role of that expenditure ratio scaling the average product is to effectively control for the substitution between factors in response to relative price changes. Under our Leontief assumptions, any change in the average product of an input factor over time is due to investment and/or regulation (and not a within-period response to new information on input factor prices).

Following standard arguments, we have 3 possible cases for each channel of investment: the marginal benefit of investment is greater than its marginal cost, equal to its marginal cost, or less than its marginal cost. We can rule out the first case because that would imply that a firm would try to invest an infinite amount of resources into productivity and that cannot be an equilibrium. In the case that the marginal benefit is less than the marginal cost, we would get no investment in productivity. In the case that we may get some reasonable amount of positive investment, then the marginal benefit of a productivity investment equals its marginal cost:

$$\frac{\nu_{Lj}\xi_L(R)A_{Lj}}{w} = 1 \quad \frac{\nu_{Fj}\xi_F(R)A_{Fj}}{w} = 1 \quad \frac{\nu_{Kj}\xi_K(R)A_{Kj}}{w} = 1 \quad (15)$$

Our maximization of the CVH yields arbitrage equations equating the interest rate (net the risk of a death shock to the firm) to the dividend-price ratio plus capital gain/loss of an investment in productivity, which takes the form of the following 3 differential equations governing the costate variables:

$$i + \delta_j = \frac{\frac{\partial \pi_j}{\partial A_{Lj}}}{\nu_{Lj}} + \frac{\dot{\nu}_{Lj}}{\nu_{Lj}} \quad i + \delta_j = \frac{\frac{\partial \pi_j}{\partial A_{Fj}}}{\nu_{Fj}} + \frac{\dot{\nu}_{Fj}}{\nu_{Fj}} \quad i + \delta_j = \frac{\frac{\partial \pi_j}{\partial A_{Kj}}}{\nu_{Kj}} + \frac{\dot{\nu}_{Kj}}{\nu_{Kj}} \quad (16)$$

Finally, each of these differential equations has a terminal condition that the discounted stream of the product of the state variable and costate variable converges to 0 in the long run. Taking logs and time derivatives of our equations describing the shadow price of productivity investments under the case of some investment, we get:

$$\frac{\dot{\nu}_{Lj}}{\nu_{Lj}} = \frac{\dot{A}_{Lj}}{A_{Lj}} + \frac{\dot{w}}{w} \quad \frac{\dot{\nu}_{Fj}}{\nu_{Fj}} = \frac{\dot{A}_{Fj}}{A_{Fj}} + \frac{\dot{w}}{w} \quad \frac{\dot{\nu}_{Kj}}{\nu_{Kj}} = \frac{\dot{A}_{Kj}}{A_{Kj}} + \frac{\dot{w}}{w} \quad (17)$$

Substituting these 3 equations, as well as the equations for the costate variables and the technology for increasing productivity, into the preceding arbitrage equations gives us:

$$\begin{aligned} i + \delta_j &= -\frac{\partial c_{Qj}}{\partial A_{Lj}} Q_j \frac{\xi_L(R)A_{Lj}}{w} + \xi_L(R) \frac{I_{Lj}}{w} + \frac{\dot{w}}{w} \\ i + \delta_j &= -\frac{\partial c_{Qj}}{\partial A_{Fj}} Q_j \frac{\xi_F(R)A_{Fj}}{w} + \xi_F(R) \frac{I_{Fj}}{w} + \frac{\dot{w}}{w} \\ i + \delta_j &= -\frac{\partial c_{Qj}}{\partial A_{Kj}} Q_j \frac{\xi_K(R)A_{Kj}}{w} + \xi_K(R) \frac{I_{Kj}}{w} + \frac{\dot{w}}{w} \end{aligned} \quad (18)$$

We perform some algebraic manipulation on the partial derivatives of unit cost with respect to productivity to produce:

$$\begin{aligned}\frac{\partial c_{Qj}}{\partial A_{Lj}} &= -\frac{1}{A_L Q_j} \left[\frac{wL_j}{c_{Qj} Q_j} \right] c_{Qj} Q_j \\ \frac{\partial c_{Qj}}{\partial A_{Fj}} &= -\frac{1}{A_F Q_j} \left[\frac{fF_j}{c_{Qj} Q_j} \right] c_{Qj} Q_j \\ \frac{\partial c_{Qj}}{\partial A_{Kj}} &= -\frac{1}{A_K Q_j} \left[\frac{rK_j}{c_{Qj} Q_j} \right] c_{Qj} Q_j\end{aligned}\tag{19}$$

Where the terms in brackets are the factor cost shares, which can be readily replaced with the more complicated form of cost shares involving production parameters, prices, and productivities that we derived above. Substituting these into our equation characterizing some positive investment and substituting in the firm's pricing strategy and imposing symmetry across firms yields:

$$\begin{aligned}i + \delta_j &= \left(\frac{\xi_L(R)}{\eta_j/(\eta_j-1)} \right) \frac{Y}{wN} \left[\frac{wL}{c_{Qj} Q_j} \right] + \xi_L(R) \frac{I_L}{wN} + \frac{\dot{w}}{w} \\ i + \delta_j &= \left(\frac{\xi_F(R)}{\eta_j/(\eta_j-1)} \right) \frac{Y}{wN} \left[\frac{fF}{c_{Qj} Q_j} \right] + \xi_F(R) \frac{I_F}{wN} + \frac{\dot{w}}{w} \\ i + \delta_j &= \left(\frac{\xi_K(R)}{\eta_j/(\eta_j-1)} \right) \frac{Y}{wN} \left[\frac{rK}{c_{Qj} Q_j} \right] + \xi_K(R) \frac{I_K}{wN} + \frac{\dot{w}}{w}\end{aligned}\tag{20}$$

Where Y is the industry's value added to GDP (which, in theory, should equal revenue given that the technology for transportation services has been specified here as not using any intermediate goods). Because the measure of the nominal interest rate in our data has been adjusted for risk, we treat this sum on the left-hand side as our measure. Note that the wage rate appears in all 3 equations in a role of transforming nominal terms into real terms. In a general equilibrium setting, using the wage as the numeraire is most convenient for this class of models. We can then do some algebra to get the following equations that determine the real investment per firm as a function of the scale of the firm's market (capturing the scale of the benefit from investment) and the real interest rate (which captures that the cost of investment is borne before the fruits of investment):

$$\begin{aligned}\frac{I_L}{wN} &= \max \left\{ 0, \left(\frac{1}{\eta_j/(\eta_j-1)} \right) \frac{Y}{wN} \left[\frac{wL}{c_{Qj} Q_j} \right] - \left[\frac{(i+\delta)-\dot{w}/w}{\xi_L(R_m)} \right] \right\} \\ \frac{I_F}{wN} &= \max \left\{ 0, \left(\frac{1}{\eta_j/(\eta_j-1)} \right) \frac{Y}{wN} \left[\frac{fF}{c_{Qj} Q_j} \right] - \left[\frac{(i+\delta)-\dot{w}/w}{\xi_F(R_m)} \right] \right\} \\ \frac{I_K}{wN} &= \max \left\{ 0, \left(\frac{1}{\eta_j/(\eta_j-1)} \right) \frac{Y}{wN} \left[\frac{rK}{c_{Qj} Q_j} \right] - \left[\frac{(i+\delta)-\dot{w}/w}{\xi_K(R_m)} \right] \right\}\end{aligned}\tag{21}$$

This formulation with the max function captures the case of 0 investment. In practice, we do not observe 0 investment in the data and thus do not bother with propagating the max function into the rest of the empirical design. These 3 equations for investment can be rearranged so that the growth of productivity is an outcome variable that depends directly on investment, which then depends on the benefits of investment (a function of the scale of the market, the firm's market power, and the share of unit cost) and the [opportunity] cost of investment (ie the risk-adjusted real interest rate):

$$\begin{aligned}
\ln \frac{A_{Lmt+1}}{A_{Lmt}} &= \frac{\xi_L(R_m)I_{Lm}}{w_m N_m} = \left(\frac{\xi_L(R_{mt})}{\frac{\eta_j}{\eta_j-1}} \right) \frac{Y_{mt}}{w_{mt} N_{mt}} \left[\frac{w_{mt} L_{mt}}{c_{Qmt} Q_{mt}} \right] - \left(i_t + \delta_m - \frac{\dot{w}_{mt}}{w_{mt}} \right) \\
\ln \frac{A_{Fmt+1}}{A_{Fmt}} &= \frac{\xi_F(R_{mt})I_{Fm}}{w_{mt} N_{mt}} = \left(\frac{\xi_F(R_{mt})}{\frac{\eta_j}{\eta_j-1}} \right) \frac{Y_{mt}}{w_{mt} N_{mt}} \left[\frac{f_{mt} F_{mt}}{c_{Qmt} Q_{mt}} \right] - \left(i_t + \delta_m - \frac{\dot{w}_{mt}}{w_{mt}} \right) \\
\ln \frac{A_{Kmt+1}}{A_{Kmt}} &= \frac{\xi_K(R_{mt})I_{Km}}{w_{mt} N_{mt}} = \left(\frac{\xi_K(R_{mt})}{\frac{\eta_j}{\eta_j-1}} \right) \frac{Y_{mt}}{w_{mt} N_{mt}} \left[\frac{r_{mt} K_{mt}}{c_{Qmt} Q_{mt}} \right] - \left(i_t + \delta_m - \frac{\dot{w}_{mt}}{w_{mt}} \right)
\end{aligned} \tag{22}$$

We can specify that the function of regulations is linear and then these pieces provide a nearly operable regression. To transform it into a bivariate regression so that we can focus exclusively on the role of regulation, we can sum the factor productivity growth across the 3 input factors of production and then divide by real investment per firm to obtain a single regression equation. Alternatively, we could subtract the risk-adjusted real interest rate from the growth rate of each factor productivity and then divide through by real costs of the factor per firm to obtain 3 separate regression equations:

$$\begin{aligned}
\left[\frac{\ln A_{Lt+1} - \ln A_{Lt} + \left(i_t + \delta_m - \frac{\dot{w}_{mt}}{w_{mt}} \right)}{(w_{mt} L_{mt}) / (N_{mt} w_{mt})} \right] &= \xi_{0L} + \xi_{1L} [R_{mt}] + \varepsilon_{Lmt} \\
\left[\frac{\ln A_{Ft+1} - \ln A_{Ft} + \left(i_t + \delta_m - \frac{\dot{w}_{mt}}{w_{mt}} \right)}{(f_{mt} F_{mt}) / (N_{mt} w_{mt})} \right] &= \xi_{0F} + \xi_{1F} [R_{mt}] + \varepsilon_{Fmt} \\
\left[\frac{\ln A_{Kt+1} - \ln A_{Kt} + \left(i_t + \delta_m - \frac{\dot{w}_{mt}}{w_{mt}} \right)}{(r_{mt} K_{mt}) / (N_{mt} w_{mt})} \right] &= \xi_{0K} + \xi_{1K} [R_{mt}] + \varepsilon_{Kmt}
\end{aligned} \tag{23}$$

3.3 Demand, Markup, and Modal Substitution

All that is left to finish the model is to specify demand so that we can both compute the markup factor and estimate the loss of business (some of which is results from modal shift) due to regulations. We employ the ARUM model of Dube et al. [2025], which is both consistent with a logit RUM and also cleanly aggregates to a representative consumer's

demand function; it's built on a static [indirect] utility function with additively separable terms including an idiosyncratic draw of taste for the mode of transportation:

$$U_{imt} = g(T_m, y_t; \delta) - \sigma \ln p_m + v_{imt} \quad (24)$$

T captures the relatively slow transit time of water and rail transportation in comparison to truck and air, GDP per capita (y) appears to capture wealth effects increasing the importance of transit time, and $g()$ is some flexible functional form. Assuming that the idiosyncratic draw, v_{imt} , is an iid draw from a Gumbel distribution, we get the common result of McFadden's Multinomial Logit market shares:

$$S_m = \frac{e^{g(T_m, y_t; \delta) - \sigma \ln p_m}}{\sum_{m=1}^M e^{g(T_m, y_t; \delta) - \sigma \ln p_m}} \quad (25)$$

Following Dube et al. [2025], this then results in a simple demand structure for transportation services from mode m at time t :

$$Q_m(p, Y_G) = \frac{\psi_G [\mu G + (1 - \mu) Y_G]}{p_m} S_m \quad (26)$$

Where Y_G is the value of all goods produced and G is the expenditure on goods consumed (we estimate a mixture of goods produced and goods consumed, with weight μ , since both get transported) and we have suppressed the time subscript for brevity. The parameter ψ_G times the weighted average of Y_G and G is just the total expenditure on transportation (a constant ψ_G fraction would be consistent with a Cobb-Douglas structure); the δ parameters govern how willingness to pay differs with transit time (and may evolve over the years as we get wealthier) and the remaining parameter (σ) sizes the magnitude of the dispersion of the unobserved idiosyncratic determinants of willingness to pay relative to the scale of the price (in logs). These parameters can all be estimated in a regression of log revenue on the log remnants of the right hand side (with unit cost serving as an IV for price) or even a more sophisticated variant on logit such as the apparatus from Berry et al. [1995] (which has its own IVs) and an additional regression step to recover (ψ_G, μ) . It is then a fairly straightforward exercise in comparative statics (using the following equations) to compute the effect of an increase in marginal cost on the price that a carrier charges, the quantity of ton-miles sold, and the amount of business that substitutes away to a different mode:

$$\begin{aligned}
\frac{\partial Q_m}{\partial p_m} &= \frac{\psi_G Y_G}{p_m} \left[\frac{\partial S_m}{\partial p_m} - \frac{S_m}{p_m} \right] < 0 & \frac{\partial Q_m}{\partial p_{-m}} &= \frac{\psi_G Y_G}{p_m} \left[\frac{\partial S_m}{\partial p_{-m}} \right] > 0 \\
\frac{\partial S_m}{\partial p_m} &= -\frac{\sigma}{p_m} S_m (1 - S_m) < 0 & \frac{\partial S_m}{\partial p_{-m}} &= \frac{\sigma}{p_{-m}} S_m S_{-m} > 0
\end{aligned} \tag{27}$$

$$\frac{dp_m}{dc_m} = \frac{\eta_m}{\eta_m - 1 + c_m \left[\frac{\partial \eta_m / \partial p_m}{\eta_m - 1} \right]} \tag{28}$$

$$\eta_m = -\frac{\partial Q_m}{\partial p_m} \left[\frac{p_m}{Q_m} \right] = 1 + \sigma (1 - S_m) \tag{29}$$

$$\frac{\partial \eta_m}{\partial p_m} = -\sigma \frac{\partial S_m}{\partial p_m} = \frac{\sigma^2}{p_m} S_m (1 - S_m) > 0 \tag{30}$$

4 Results

The empirical estimates reveal several interesting systematic relationships between regulatory accumulation, factor productivities, prices, and shipper behavior. Across all freight modes, increases in regulatory restrictions are associated with declines in at least nearly all input productivity measures, though the magnitude differs by mode. All four modes (air, rail, truck, and water) experience significant reductions in labor productivity, consistent with regulatory requirements diverting labor effort toward compliance-related tasks such as reporting and inspections. Three of the four (air, rail, and trucking) exhibit significant declines in fuel productivity, while all four show significant reductions in capital productivity. These productivity effects are not merely transitory; because productivity growth compounds over time, even modest annual reductions induced by regulatory accumulation can substantially alter the long-run trajectory of transportation efficiency.

Given the estimated demand elasticities, reductions in factor productivities translate into higher marginal and unit costs. Naturally, this results in carriers charging higher prices. Even with somewhat imperfect substitutability between modes, we would expect to see shippers respond to these price changes by reallocating shipments across modes. Because shippers substitute toward alternatives whose relative costs have changed less due to uneven regulation across modes, shocks to one mode might propagate to others through the choices of shippers operating multimodal supply chains.

To quantify these effects on unit costs, prices, market shares, intramodal quantities, and

intermodal substitution, we conduct a counterfactual policy experiment in which regulatory restrictions for a single mode increase by five percent in 2018. The resulting first-year changes in unit costs, prices, market shares, and quantities are summarized in the counterfactual policy table below. The experiment highlights three patterns. First, regulatory shocks translate rapidly into increased operating costs, with rail facing the largest first-year unit-cost increase (2.3%), followed by air (1.3%), water (1.1%), and truck (0.8%). Second, targeted-mode quantities fall by 1.4% to 4.1%, with the largest decline in rail and the smallest in trucking. Third, although we expected to see competing modes expand as shippers reallocate traffic, the results actually showed net quantity losses in the competing modes, indicating that total freight throughput more than offsets any substitution effect. Taken together, these results imply that regulatory accumulation exerts its effects through both productivity and allocation channels, and that the first-year impacts observed in the counterfactual represent a lower bound on the long-run cost of additional regulatory burdens.

4.1 Demand-side Results

We begin by examining our estimation of demand, which is important because the implied elasticity is then an input into our simulated policy experiments where it governs the amount of substitution between modes due to increased costs due to increases in regulatory restrictions. We note that our demand specification is a simple variant from the BLP family of Berry, Levinsohn, and Pakes (1995); hence, we first invert the mode shares (S_{mt}) at each point in time to uncover the implied value of that BLP common parameter for each mode at each point in time (Δ_{mt}). Then we regress that common parameter on a carefully constructed set of transit time and income controls as well as the predicted log price from a first stage of a 2SLS procedure (the IV that we use for log price is the log of unit cost, although we could alternatively use the prices of the input factors) to attain the crucial substitution parameter (σ). The results of that first stage regression appear in Table 2.

Almost all parameter signs are as anticipated and the goodness of fit measures are basically the same (yielding qualitatively similar predicted prices). Given the surprising sign on the rental price of capital, we prefer the specification based on unit cost. Our results from the second stage are summarized in Table 3.

Our more parsimonious specification relies on a covariate designed to capture the value of each mode's (time-invariant) transit time in 2000, the deviation of log GDP per capita from its value in 2000, and an interaction between those two variables. To construct the

value of the transit time in 2000, we begin with a reasonable initial guess for the price elasticity and then compute the sum of the common parameter and that price term so that we can explore how it changes with [a monotonic transformation of] transit time.⁷ We have used an initial guess of -3 because that is the smallest round number (in absolute value) corresponding to a monotonic ordering of transit time (ie controlling for price, buyers should be willing to pay more for shorter transit times); elasticities that are much closer to 0 imply much higher markup factors for price over marginal cost. The resulting parameter estimates for our parsimonious specification are each statistically significant and have the signs that theory would suggest. The key result is the estimated coefficient for log price (σ) and that estimate is fairly robust with respect to local perturbations in both our specification and the aforementioned initial guess of σ .

Table 3 also gives results for a more conventional specification of modal fixed effects and polynomial expansions of log GDP per capita (also relative to 2000, just to preserve comparability). This less parsimonious specification increases the goodness of fit from $R^2 = 0.94$ to $R^2 = 0.99$ and most parameters are still statistically significant with signs that are consistent with theory. However, there are some aberrations that suggest this more conventional specification suffers from overfitting. Neither of the quadratic nor cubic polynomial terms are statistically significant and the cubic term actually has a negative point estimate (which runs contrary to a theorist's prior expectations). More concerning is that the modal fixed effects, which are a more flexible way of capturing each mode's relative desirability, are not monotonic in transit time. Where this specification really differs is in the estimated parameter on log price, which implies a markup of over 2.35 times marginal cost. We should be clear that, in order to get this large of a markup (or larger), we had to either use a sprawling specification that is likely to overfit the data and/or seed our parsimonious specification with an initial guess that organized the modes in an order that violates monotonicity in transit time. Ultimately, this led us to select our more parsimonious specification and its suggested markup of around 1.35; we note that this markup is more consistent with the 1.5 markup factor that we computed from a recent research report on a survey of the trucking industry.⁸

⁷In particular, we perform the Hyman method for Hermite interpolation splines on the 4 different transit times because it preserves the underlying monotonicity, positivity, and convexity of the shape in the data. Figure 9 shows this specification.

⁸When we examine the ratio of price to our computed unit cost for truck transport, we find that a markup factor of 1.35 and 1.5 to be between the 25th and 40th percentile across the years (respectively). In theory, we would actually expect rail to have a higher markup factor than trucking because individual firms have a higher market share (i.e. the mode's share of transportation divided by the number of competitors is larger for rail than for trucking) and we computed a markup factor closer to 2 for a similar research report surveying the rail industry. However, we do not know enough about the underlying data in those reports (eg what is

The final element of demand to estimate is how the aggregate quantity of freight changes with the average price of transport (weighted by the modal shares). We should note, in addition to assuming exogenous prices on inputs, that our partial equilibrium assumption binds here. Essentially, we are assuming that the expenditure on transportation is tied to the expenditure on goods and that expenditure is fixed. The results of our regression analysis appear in Table 4. With a slightly better goodness of fit, we prefer the specification that includes both the value of goods produced and consumed because the difference is exports and imports, which naturally both involve transportation. These parameter estimates imply that transportation expenditure averages around 11.5% of the value of goods produced (i.e., $\psi = 0.115$) and that the mixing weight on the value of goods consumption is 89.5% (i.e., $\mu = 0.895 = 0.103/0.115$).

4.2 Effects of Regulations on the Average Productivity of all Input Factors via Investment

We can now turn to the factor productivities as our outcome variables.

Our first approach follows Coffey, McLaughlin, and Peretto (2020) in using investment as an intermediary between regulatory restrictions and factor productivities. According to that theoretical model, real investment per firm increases linearly with real GDP per firm.

Figure 6-8 show the timepaths of labor, fuel, and capital productivity by mode. The main feature that stands out across all three of these figures is the growth in all three productivities for rail transport that begins in roughly 1980 and flattens out around 2008. We note that those years coincide with two of the historically largest (in terms of regulatory consequence) rail transport-related acts of Congress, the Staggers Act of 1980 (which was deregulatory) and the Rail Safety Improvement Act of 2008 (which was regulatory). The Staggers Act removed significant portions of economic regulations of the rail industry. Ellig and McLaughlin [2016] show that the freedom to allocate capital where needed, instead

included in each category) to compute a figure that can be fairly compared with our computations from our own data. There are a variety of reasons to be concerned with how we compute markup factors from our data and most are related to how we measure capital. Rail unit costs might be overstated because of the difference between ownership versus rental and 2 of the 4 modes (i.e., rail and water) were shedding some of their capital stock during our period of study (some by liquidation and some by deferred maintenance of depreciating assets). Ultimately, BEA only provides an index for the quantity of capital and we have scaled that up by the number of prime mover vehicles in each industry in the base year. Our efforts are most likely to be closest to the truth in trucking where there is relatively little capital beyond the trucks, yet our efforts are much more problematic with valuable capital in other industries (e.g., miles of tracks in a rail network). Air freight is the most problematic mode because BEA makes no distinction between freight and passenger transport due to the economies of scope across those markets.

of where mandated—a direct result of the partial economic deregulation from the Staggers Act—explains nearly all of the drastic improvements in safety witnessed in the industry over the following thirty years. We speculate that the same freedom to reallocate capital could explain rail’s post-1980 productivity growth seen in these figures as well as the slowdown in productivity growth post-2008, when more regulatory constraints to capital allocation were introduced with the Rail Safety Improvement Act.

Our first direct test of the relationship between regulation and productivity via the investment channel is presented in Table 5. Because the model links real investment per firm to overall productivity growth, we begin by averaging the growth rates of labor, fuel, and capital productivity into a single summary measure. This is restrictive, but it provides a transparent first-pass test of whether regulatory accumulation weakens the productivity payoff to investment before allowing the effects to differ by factor in the next subsection. We regress the ratio of the growth of average productivity to real investment per firm on regulatory restrictions. For each mode, the dependent variable is standardized within mode, and regulatory restrictions are rescaled in thousands for ease of interpretation. In the first column, we impose a common slope on regulatory restrictions across all modes while allowing mode-specific intercepts; in the second, we allow the slope on regulatory restrictions to differ by mode.

The pooled estimate in the first column has the anticipated negative sign and is precisely estimated: an additional 1,000 regulatory restrictions is associated with a 0.324 standard-deviation decline in the ratio of average productivity growth to real investment per firm. When the slope is allowed to vary by mode, the negative relationship remains statistically significant only for rail, where an additional 1,000 restrictions is associated with a 0.155 standard-deviation decline in the standardized outcome. The coefficients for truck and water are also negative but lack statistical significance, while the air coefficient is close to zero. Figure 10 is visually consistent with this pattern: the fitted relationships slope downward in all four modes, but the rail relationship is the clearest.

This approach has used a simple mean averaging of productivity growth across the input factors. In doing so, it prevents regulations from having a different impact on different factors. To relax that undesirable assumption, we must dispense with using investment as an intermediary between regulation and productivity growth, despite the intuitive and theoretical appeal. Unfortunately, we do not observe separate investments into different productivities. Thus we proceed with our specification design that dispenses with investment, which stacks the input factors and can be estimated as a pooled regression or separately by

mode.

4.3 Effects of Regulations on the Productivity of Specific Input Factors

By setting aside investment (because we cannot separate investments into the different productivities) and looking directly at the relationship between regulation and productivity, we obtain the sharper results presented in Table 6. The dependent variable for the regressions presented in Table 6 is the mode-specific ratio of the growth of input factor productivity (plus the real interest rate) to the real cost of the input factor per firm. The ratio is transformed with a standardized normalization (i.e., demeaned by the within-mode mean and then divided by the within-mode standard deviation) prior to analysis. We regress the growth in input factor productivity ratio on regulatory restrictions, which is also mode-specific. As in previous regressions, regulatory restrictions have been rescaled to be in thousands.

The results of these regressions strongly suggest that pooling industries is not appropriate for this part of the analysis, because the magnitudes differ meaningfully across modes. All point estimates have the anticipated negative sign. Regulation has a negative and statistically significant effect on labor productivity in all four modes. For fuel productivity, the estimated effect is negative and statistically significant in air and rail, negative and weakly significant in trucking, and negative but not statistically distinguishable from zero in water. For capital productivity, the estimated effect is negative and statistically significant in air, rail, and water, with a smaller negative effect in trucking that is significant at the 10 percent level. Thus, only one of the twelve regulatory coefficients—water fuel productivity—is not statistically significant at conventional levels. Figure 11 shows these relationships for each mode and productivity type, with dots showing observations for each mode-year and the line reflecting the estimated coefficient.

Having estimated all of the relevant parameters of our econometric design, we can now proceed with conducting counterfactual experiments on transportation regulations. We consider a counterfactual experiment of a 5 percent increase in the regulatory restrictions constraining all firms of a particular mode of transportation for the year 2018, which affects investment in productivity in 2018 and then costs starting in 2019. Table 7 presents the computed variables of interest for our counterfactual experiments. A 5 percent change is not at all uncommon; RegData statistics for these industries (as depicted in Figure 5) indicate 47 separate instances of year-over-year changes of 5 percent or more.

The counterfactual regulatory policy experiment provides a concrete way to interpret the magnitudes of the estimated effects. A five percent increase in regulatory restrictions applied to a single freight mode in 2018 raises the unit cost of that mode by between 0.8 and 2.3 percent in the first year, with the low and high coming from trucking and rail, respectively. Because pass-through is close to complete, these cost increases lead to similar percentage increases in freight prices. The resulting quantity responses are sizable: air and water shipments fall by about 2.3 and 2.0 percent, respectively, while rail and truck shipments decline by roughly 4.1 and 1.4 percent. These are economically meaningful adjustments for the first year alone in response to a single-year, 5 percent change in the regulatory stock.

The reallocation of freight toward other modes is small and generally negative. When any one mode is subjected to higher regulatory burdens, the aggregate quantity shipped by the remaining modes declines modestly, by only 0.01 to 2.4 percent, depending on which mode is targeted. This asymmetry implies that the regulatory shock reduces total freight throughput rather than merely reshuffling it across modes, and we speculate that this result reflects the intermodal realities of freight.

Moreover, the experiment captures only the first-year effects of the regulatory increase. Because regulatory accumulation also reduces the growth rate of factor productivities, the losses compound over time. The higher cost and lower quantity in the first year are followed by subsequent years in which productivity remains permanently lower than in the counterfactual, so that the level of output in the freight sector—and therefore the level of goods movement in the economy—is structurally depressed. Interpreted in this light, the counterfactual experiment illustrates how seemingly modest incremental increases in regulation can have meaningful and persistent consequences for both the scale and composition of freight activity. Because the regulatory shock lowers productivity growth, the first-year level loss does not disappear in later years; instead, the freight sector remains on a permanently lower productivity path unless the added restrictions are later removed.

These findings also have implications for antitrust policy in freight transportation. Our demand estimates imply that competition in freight markets is often intermodal rather than purely intramodal. Rail carriers compete not only with one another, but also with trucking and, in some corridors, with water transportation. That does not mean substitution across modes is instantaneous or one-for-one; indeed, our counterfactual exercises suggest that higher transportation costs mostly reduce overall freight activity in the short run. It does mean, however, that merger analysis should be cautious about defining relevant markets too narrowly within a single mode or inferring market power directly from mode-specific

concentration.

Moreover, our results suggest that antitrust analysis in freight transportation should give substantial weight to dynamic efficiencies. Across several specifications, higher regulatory burdens are associated with weaker productivity performance, and the counterfactual exercises show that cost increases translate into higher prices and lower quantities transported. Those findings underscore how important cost-reducing innovation is in this sector. One economic rationale for specialized freight firms is that scale and network density can support investments in logistics, routing, equipment, and other process innovations that many individual shippers could not efficiently undertake in-house. For that reason, merger review should weigh potential long-run efficiency gains and innovation incentives alongside any short-run concern about increased concentration (Creanza [2025]).

5 Conclusion

This paper develops and estimates a structural model of production and demand to quantify how regulatory accumulation affects the costs, productivity, and modal allocation of U.S. freight transportation. Using RegData-based measures of regulatory restrictions for air, rail, truck, and water freight carriers, matched to detailed data from BEA, BTS, the Transportation Energy Data Book, and FRED, we recover mode-specific markups, factor cost shares, and input productivities over time. The demand estimates imply own-price elasticities consistent with markups of roughly 1.25 on average, with somewhat higher markups for air and rail. These figures are broadly consistent with independent evidence on pricing behavior in transportation markets and with the markup magnitudes embedded in Coffey et al. [2020] aggregate model of regulatory accumulation and growth.

The supply-side results show that regulatory accumulation is not a neutral friction that simply scales up costs. Instead, it selectively depresses factor productivity in ways that differ across modes and inputs. Regulatory restrictions are associated with lower labor productivity in all four modes, lower fuel productivity in air and rail, and lower capital productivity in air, rail, and water; trucking also shows negative but less precisely ($p < 0.1$) estimated effects on fuel and capital. These findings are consistent with regulation diverting labor toward compliance tasks, reducing the effective utilization of capital, and constraining energy-saving innovations and operational practices. Because productivity growth compounds over time, the regulatory effects we estimate imply a persistent wedge in the cost structure of the freight

sector that widens as the regulatory stock grows. This is precisely the channel emphasized in the broader regulatory-accumulation literature, in which regulation reduces the return to investment in productivity-enhancing activities and thereby shifts the economy onto a lower long-run growth path.

The counterfactual policy experiments highlight the short-run and medium-run consequences of incremental regulatory changes. A 5 percent increase in regulatory restrictions applied to a single mode raises that mode’s unit costs and prices by roughly 0.8 to 2.3 percent and reduces its quantity of freight shipped by about 1.4 to 4.1 percent in the first year. The effects on other modes are small and generally negative, indicating that higher transportation costs reduce overall freight activity by enough to more than offset any induced substitution. Because regulation also reduces the growth rate of factor productivities, the gap between actual and counterfactual output paths widens over time: the first-year quantity losses reported in our table represent only a lower bound on the long-run cost of regulatory accumulation.

These findings connect the freight sector to a broader body of evidence that regulation has regressive and distributionally non-neutral effects. Higher freight costs ultimately feed into higher delivered prices for goods, with particular consequences for regions and households that are especially reliant on long-distance shipments. The negative productivity effects we document for rail and water are therefore likely to propagate to bulk commodities, agriculture, and heavy industry, sectors that are central to the economic fortunes of many lower-income regions. This perspective is highly consistent with the results in Chambers et al. [2019] and related work using FRASE-based measures of regulatory incidence.

From a policy standpoint, our analysis supports a shift away from purely marginal, rule-by-rule evaluation toward a more systematic management of the accumulated stock of regulation. Regulatory budgeting, sunset provisions with meaningful review, and periodic regulatory clean-ups focused on major infrastructure sectors such as transportation could arrest or reverse the growth of low-value requirements that depress productivity without delivering commensurate safety or environmental benefits. The experience of British Columbia’s regulatory budget, documented by Coffey and McLaughlin [2021], and emerging evidence from U.S. states in McLaughlin and Wong [2024] suggest that such institutional reforms can raise growth without undermining core regulatory objectives. By clarifying the channels through which regulatory accumulation affects factor productivity and modal allocation in freight, this paper provides a microeconomic foundation for the view that systematic deregulation of low-value requirements can deliver meaningful gains in productivity, growth, and ultimately

the welfare of households that bear the costs of a more heavily regulated transportation system.

Figures




Figure 1: Time path of our computed rental rate of capital by mode.




Figure 2: Time path of the number of Class I railroad firms under two alternative reconstruction methods, one manual (solid garnet) and one algorithmic (dashed red-orange). The discrepancy between them reflects different judgments about whether a rail line is an independent firm.




Figure 3: Time path of revenue per ton-mile by transportation mode from the Bureau of Transportation Statistics.

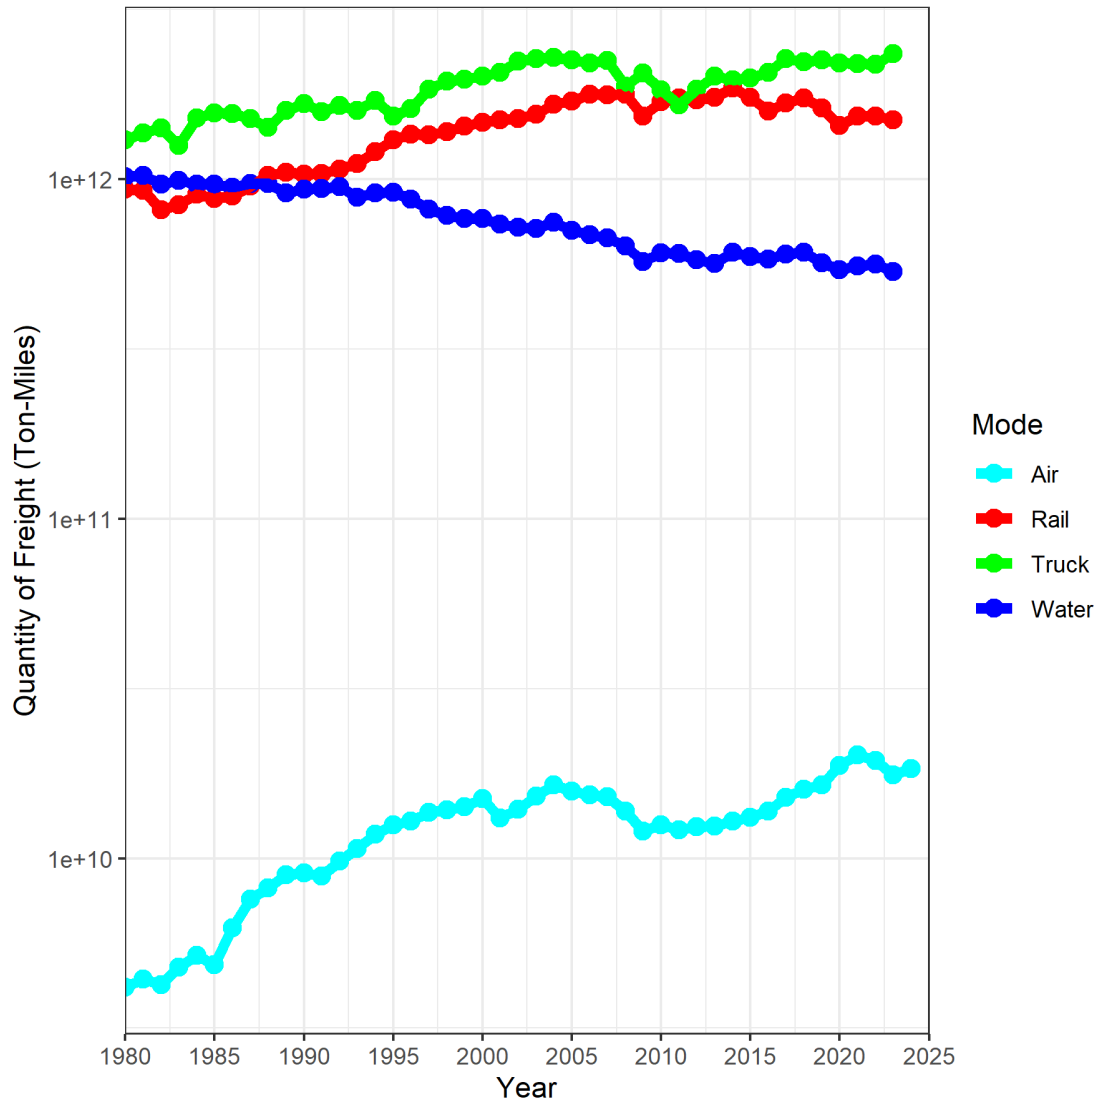


Figure 4: Time path of freight quantity (in ton-miles) by transportation mode from the Bureau of Transportation Statistics.




Figure 5: Time path of number of regulatory restrictions by transportation mode from RegData.




Figure 6: Time path of average product of labor by mode, computed by dividing BTS quantity of ton-miles by FTEs from BEA.




Figure 7: Time path of average product of fuel by mode, computed by dividing BTS quantity of ton-miles by fuel usage from Transportation Energy Data Book.




Figure 8: Time path of average product of capital by mode, computed by dividing BTS quantity of ton-miles by our rescaling of BEA's capital quantity index.

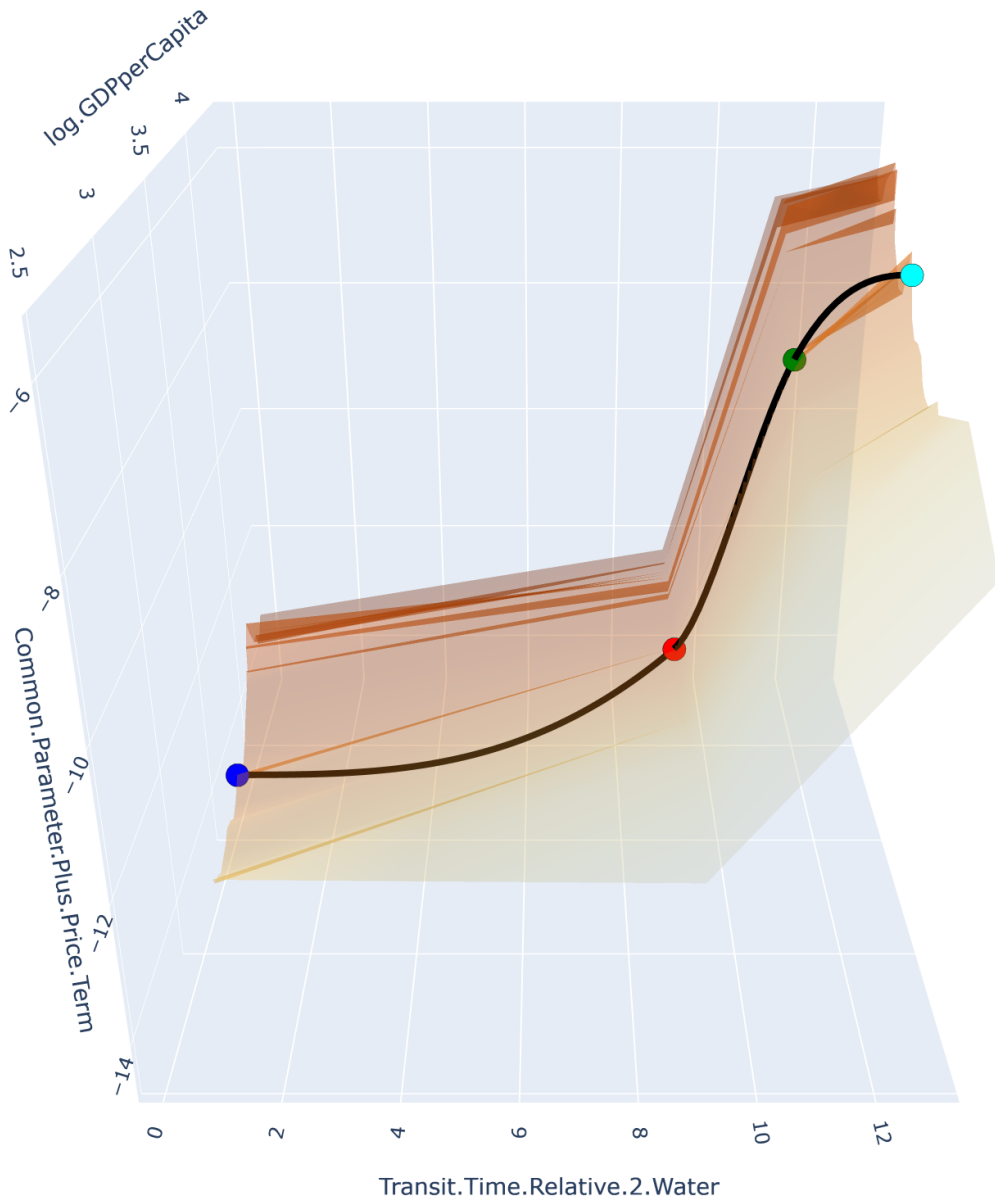


Figure 9: Parsimonious specification of the value of a mode as a function of transit time at different levels of log GDP per capita, assuming an initial guess on the log price coefficient of $\sigma = -3$. The large colored dots correspond to the values in the year 2000 and the black curve is a Hermite Interpolation Spline that smoothly connects these points.




Figure 10: Productivity Growth (averaged across 3 input factors) over the Real Investment per Firm plotted versus levels of regulatory restrictions. For each mode, the outcome variable has been transformed with a standardized normalization (ie demeaned by the within-mode mean and then divided by the within-mode standard deviation).




Figure 11: Input Factor Productivity Growth over the Real Factor Cost per Firm plotted versus levels of regulatory restrictions, separated out for each of the 3 factors and 4 modes. For each mode, the outcome variable has been transformed with a standardized normalization (ie demeaned by the within-mode mean and then divided by the within-mode standard deviation).

Tables

Mode	NYC-MIA	HOU-CHI	LA-SEA	Mode Average
Air	1	1	1	1
Truck	3	3	3	3
Rail	4	6	5	5
Water	7	21	14	14
Route Average	3.75	7.75	5.75	5.75

Table 1: For each mode (m), we construct a proxy for Transit Time (T_m) using approximate route-level transit times compiled from publicly available web sources for three major north-south routes (NYC-Miami, LA-Seattle, and Houston-Chicago). Transit times are measured in days.

	$\ln(\text{Price})$	$\ln(\text{Price})$
<i>Air</i>	1.27* (0.72)	-1.34*** (0.18)
<i>Rail</i>	-2.53*** (0.62)	-2.07** (0.92)
<i>Truck</i>	-1.87*** (0.42)	-0.45** (0.20)
<i>Water</i>	-2.73*** (0.65)	-1.84*** (0.24)
$\ln(\text{Wage})$	0.55*** (0.08)	
$\ln(\text{FuelPrice})$	0.21*** (0.04)	
$\ln(\text{CapitalRentalPrice})$	-0.24*** (0.05)	
$\ln(\text{UnitCost})$		1.10*** (0.17)
<i>Rail</i> \times $\ln(\text{UnitCost})$		-0.67** (0.32)
<i>Truck</i> \times $\ln(\text{UnitCost})$		-0.44** (0.19)
<i>Water</i> \times $\ln(\text{UnitCost})$		-0.54*** (0.179)
R^2	0.997	0.996
N	133	125

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Results for Regressing price on supply-side determinants of price as the first stage of the 2SLS.

	$\Delta_{mt} S_{mt} = \frac{e^{\Delta_{mt}}}{\sum_n e^{\Delta_{nt}}}$	$\Delta_{mt} S_{mt} = \frac{e^{\Delta_{mt}}}{\sum_n e^{\Delta_{nt}}}$
Intercept	-6.34*** (0.04)	-0.52*** (0.20)
$[\ln(GDPperCapita_t) - \ln(GDPperCapita_{2000})]$	0.40*** (0.08)	2.24*** (0.30)
[Value of Transit Time in 2000]		0.90*** (0.03)
[Value of Transit Time in 2000] $\times [\ln(GDPperCapita_t) - \ln(GDPperCapita_{2000})]$		0.13*** (0.03)
$[\ln(GDPperCapita_t) - \ln(GDPperCapita_{2000})]^2$	0.02 (0.14)	
$[\ln(GDPperCapita_t) - \ln(GDPperCapita_{2000})]^3$	-0.18 (0.18)	
Air	-6.34*** (0.04)	
Rail	-3.89*** (0.30)	
Truck	-2.64*** (0.19)	
Water	-4.79*** (0.34)	
Fitted $\ln(Price)$	-0.73*** (0.08)	-2.88*** (0.07)
R^2	0.997	0.944
N	160	160

* p < 0.1, ** p < 0.05, *** p < 0.01

Table 3: Results for Regressing BLP Common Parameter on predicted log price (the second stage in a 2SLS) and other determinants of modal choice.

	$\sum_m p_{mt}q_{mt}$	$\sum_m p_{mt}q_{mt}$	$\sum_m p_{mt}q_{mt}$
Intercept	-68.459*** (11.672)	-2.528*** (8.391)	
<i>ValueOfGoodsConsumed</i>		0.114*** (0.003)	
<i>ValueOfGoodsProduced</i>	0.158*** (0.004)		0.115*** (0.003)
<i>(ValueOfGoodsConsumed)</i> <i>- (ValueOfGoodsProduced)</i>			0.103*** (0.013)
R^2	0.978	0.984	0.998
N	31	31	31

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Results of regressing expenditure on freight transportation (summed across all modes) on the value of goods consumed and the value of goods produced (all variables in billions of dollars).

	Ratio of [Growth in Input Factor Productivity (Averaged Across all 3 Input Factors)] to [Real Investment per Firm]	
<i>Air</i>	1.568* (0.842)	0.209 (0.909)
<i>Rail</i>	0.303 (0.225)	1.445*** (0.422)
<i>Truck</i>	0.058 (0.162)	0.387 (0.489)
<i>Water</i>	0.343 (0.241)	1.182 (0.951)
<i>Regulatory Restrictions</i> (All Modes)	-0.324*** (0.017)	
<i>Regulatory Restrictions</i> $\times 1\{Air\}$		-0.004 (0.018)
<i>Regulatory Restrictions</i> $\times 1\{Rail\}$		-0.155*** (0.042)
<i>Regulatory Restrictions</i> $\times 1\{Truck\}$		-0.216 (0.259)
<i>Regulatory Restrictions</i> $\times 1\{Water\}$		-0.112 (0.089)
R^2	0.023	0.097
N	156	156

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Results from regressing the ratio of the growth of average productivity to real investment per firm on regulatory restrictions. For each mode, the outcome variable has been transformed with a standardized normalization (i.e., demeaned by the within-mode mean and then divided by the within-mode standard deviation); regulatory restrictions have been rescaled to be in thousands.

	Ratio of [Growth in Input Factor Productivity (Plus Real Interest Rate)] to [Real Costs of Factor Input per Firm]			
	(Air)	(Rail)	(Truck)	(Water)
<i>Labor</i>	1.974** (0.890)	1.63*** (0.014)	0.820* (0.419)	2.347*** (0.879)
<i>Fuel</i>	2.403*** (0.890)	1.324*** (0.360)	0.912* (0.490)	0.706 (0.949)
<i>Capital</i>	1.887** (0.890)	1.163*** (0.323)	0.755* (0.419)	2.009** (0.879)
<i>Regulatory Restrictions</i> $\times 1\{Labor\}$	-0.041** (0.018)	-0.163*** (0.029)	-0.423** (0.203)	-0.217*** (0.080)
<i>Regulatory Restrictions</i> $\times 1\{Fuel\}$	-0.050*** (0.018)	-0.142*** (0.036)	-0.509* (0.260)	-0.067 (0.088)
<i>Regulatory Restrictions</i> $\times 1\{Capital\}$	-0.039** (0.018)	-0.115*** (0.029)	-0.390* (0.203)	-0.185** (0.080)
R^2	0.134	0.342	0.091	0.101
N	117	125	125	125

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Results from regressions of the ratio of the growth of input factor productivity (plus the real interest rate) to the real cost of the input factor per firm on regulatory restrictions. For each mode, the outcome variable has been transformed with a standardized normalization (ie demeaned by the within-mode mean and then divided by the within-mode standard deviation); regulatory restrictions have been rescaled to be in thousands.

Counterfactual Regulatory Policy Experiment	Targeted Mode	1st Year Change in Unit Cost for Targeted Mode	1st Year Change in Price for Targeted Mode	1st Year Change in Targeted Mode's Market Share	1st Year Change in Targeted Mode's Quantity Sold	1st Year Change in Other Modes' Quantity Sold
↑ R by 5%	Air	↑ c by 1.3%	↑ p by 1.3%	↓S by 0.01pp	↓ Q_m by 2.3%	↓ Q_{-m} by 0.01%
↑ R by 5%	Rail	↑ c by 2.3%	↑ p by 2.3%	↓S by 1.52pp	↓ Q_m by 4.1%	↓ Q_{-m} by 2.41%
↑ R by 5%	Truck	↑ c by 0.8%	↑ p by 0.8%	↓S by 0.54pp	↓ Q_m by 1.4%	↓ Q_{-m} by 1.09%
↑ R by 5%	Water	↑ c by 1.1%	↑ p by 1.1%	↓S by 0.34pp	↓ Q_m by 2.0%	↓ Q_{-m} by 0.40%

pp: percentage points

Table 7: Counterfactual regulatory policy experiment following a 5 percent increase in a mode's regulatory stock in 2018. Reported values are first-year effects in 2019 when no other mode's regulatory stock changes.

References

- O. Al-Ubaydli and P. A. McLaughlin. Regdata: A numerical database on industry-specific regulations for all united states industries and federal regulations, 1997–2012. *Regulation & Governance*, 11(1):109–123, 2017.
- H. Bai, E. X. N. Li, C. Xue, and L. Zhang. Asymmetric investment rates. *NBER Working Paper No. 29957*, 2022.
- S. Berry, J. Levinsohn, and A. Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.
- D. Chambers and C. O’Reilly. Regulation and income inequality in the united states. *European Journal of Political Economy*, 72, 2022.
- D. Chambers, P. A. McLaughlin, and L. Stanley. Regulation and poverty: An empirical examination of the relationship between the incidence of federal regulation and the occurrence of poverty across the us states. *Public Choice*, 180:131–144, 2019.
- B. Coffey and P. A. McLaughlin. Regulation and economic growth: Evidence from british columbia’s experiment in regulatory budgeting. *Mercatus Working Paper*, 2021.
- B. Coffey, P. A. McLaughlin, and P. Peretto. The cumulative cost of regulations. *Review of Economic Dynamics*, 38:1–21, 2020.
- B. Coffey, P. A. McLaughlin, and P. Peretto. Transportation, innovation, and growth. *Working Paper*, 2025.
- P. P. Creanza. Factories of ideas? big business and the golden age of american innovation. *Working Paper*, 2025.
- J. J. Dawson and J. J. Seater. Federal regulation and aggregate economic growth. *Journal of Economic Growth*, 18(2):137–177, 2013.
- J.-P. H. Dube, J. Joo, and K. Kim. Discrete/continuous choice models and representative consumer theory. *Journal of Economic Theory*, 226, 2025.
- J. Ellig and P. A. McLaughlin. The regulatory determinants of railroad safety. *Review of Industrial Organization*, 49(2), 2016.

- C. Makridis and P. A. McLaughlin. Regulatory stringency and industry performance: Evidence from regdata u.s. 6.0. *Working Paper*, 2026.
- P. A. McLaughlin and O. Sherouse. Regdata 2.2: A panel dataset on us federal regulations. *Public Choice*, 180(1):43–55, 2019.
- P. A. McLaughlin and J. Wong. The causal effect of regulations on economic growth: Evidence from the us states. *Mercatus Working Paper*, 2024.
- P. M. Romer. Endogenous technological change. *Journal of Political Economy*, 98(5 (Part 2)):S71–S102, 1990.
- C. Winston. A disaggregate model of the demand for intercity freight transportation. *Econometrica*, 49(4):981–1006, 1981.
- C. Winston, T. M. Corsi, C. M. Grimm, and C. A. Evans. *The Economic Effects of Surface Freight Deregulation*. Brookings Institution Press, 1990.