Cost Push and the Direction of Endogenous Innovation: Transportation, Regulation, and Scarcity

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

September 30, 2026

Abstract

The current debate over regulation increasingly emphasizes its cost-push effects. We ask what those effects mean when innovation is endogenous and transportation is an integral part of the growth process. We develop a multi-industry model in which firms innovate to reduce the cost of making goods and the cost of shipping them, use labor and fuel in both activities, and choose among multiple transportation modes. The direction of innovation responds to differential cost exposure: a cost-push force common to making and shipping changes the level of costs without changing the allocation of R&D, whereas a shipping-specific force changes the relative return to improving the factory and the fleet. Fuel scarcity provides the first case; cumulative regulation attached to transportation modes provides the second. In the interior equilibrium, manufacturing innovation is affine in the rate of regulatory accumulation, yielding a finite threshold at which manufacturing R&D shuts down while transportation R&D remains active. With substitute modes, industries escape toward the mode with the slower accumulation of regulations; with complementary modes, the link with the more rapid accumulation of regulations can become the bottleneck. Transitional boundary crossing can create persistent level effects, whereas a fully reversed wedge produces full knowledge recovery when adjustment remains everywhere interior.

Keywords: cost push; endogenous innovation; transportation; regulation; scarcity; endogenous growth; modes of transportation

JEL: O31; O40; L51; R41

*Department of Economics, University of South Carolina, USA (e-mail: bentleygcoffey@gmail.com).
†Hoover Institution, Stanford University, USA (e-mail: econpatrick@gmail.com).
‡Department of Economics, Duke University, USA (e-mail: peretto@econ.duke.edu).

1 Introduction

The current debate over regulation increasingly emphasizes its cost-push effects. That language points to a familiar economic idea: regulation can raise the cost of carrying out an activity. In an economy with endogenous innovation, however, the interesting question is not simply whether a regulatory burden raises costs. It is whether it changes the relative return to improving different activities and thereby changes the direction of technical change.

Transportation provides a natural setting for that question. Goods are made before they are delivered, and shipping has its own technology, inputs, modes, and institutional constraints. Taking these properties seriously changes the growth problem. Firms innovate to reduce the cost of making a good and the cost of shipping it, and the allocation of innovation between these activities becomes endogenous. Transportation, therefore, should not be treated as an add-on to an otherwise complete growth model. It is a critical part of the process through which the economy produces and delivers the goods whose expanding abundance defines economic growth.[1]

Our starting point is the distinction between making and shipping developed by Kane and Peretto (2020) in a Schumpeterian growth model. They show that the form of the shipping cost matters for the composition of innovation. An iceberg cost falls with the cost of production because the good that melts in transit must itself be produced, whereas a per-unit shipping cost creates a separate wedge that manufacturing innovation does not directly compress. When firms can innovate in transportation, the two cost-reducing activities become complementary and can sustain a steady state in which their unit costs decline together. We take that mechanism as the benchmark rather than as a contribution of this paper.[2]

We then ask what happens when there is more than one way to ship a good and the institutional environment of transportation evolves over time. We deliberately abstract from quality innovation. This removes an important margin through which firms can redirect innovation effort and lets us isolate the interaction between cost-reducing innovation in making

[1] A large literature documents the role of transportation infrastructure, spatial frictions, and connectivity in economic activity and productivity; see, among others, Fogel (1964), Aschauer (1989), Fernald (1999), Calderón and Servén (2010), Krugman (1991), Fujita and Thisse (2002), Comin, Dmitriev, and Rossi-Hansberg (2012), Allen and Arkolakis (2014), Desmet and Rossi-Hansberg (2014), Donaldson and Hornbeck (2016), and Acheampong et al. (2022). Our focus is different: transportation enters the endogenous innovation mechanism itself.

[2] The distinction between per-unit and iceberg frictions traces back to Alchian and Allen (1964). Hummels and Skiba (2004) provide empirical evidence for the associated quality-composition effect, while Sørensen (2014) and Irarrazabal, Moxnes, and Opromolla (2015) study the consequences of additive versus multiplicative trade costs. The endogenous-growth structure belongs to the broader innovation-driven tradition of Romer (1990), Grossman and Helpman (1991), Aghion and Howitt (1992), and Peretto (1998, 1999).

and shipping. We introduce a second input, fuel, whose price can become endogenous when a fixed resource endowment binds; multiple transportation modes with different technologies; and mode-specific stocks of regulations that accumulate over time. These features are not mere separate extensions; they are the properties of transportation that the growth model must exhibit to capture its real-world integration with production.

The central result is an asymmetry principle. What matters for the direction of innovation is not the absolute cost of transportation, nor even whether that cost is rising, but whether the cost of shipping is rising relative to the cost of making the good. A cost-push force common to the two activities can change the level of costs without changing the allocation of R&D. Under equal factor intensities, the endogenous price of fuel is such a force: it raises the cost of the labor–fuel bundle in both activities and cancels from the ratio of the shipping cost to the manufacturing cost.[3] A shipping-specific force is different. Because it raises shipping cost relative to manufacturing cost, it changes the relative return to transportation and manufacturing innovation and reallocates R&D between the fleet and the factory.

This distinction gives regulation a role that differs from the standard analysis of a fixed wedge. We model regulation as a stock that accumulates rather than as a fixed wedge. The relevant institutional force is therefore a trend in the relative cost of shipping. In the interior equilibrium, manufacturing innovation is affine in the rate of regulatory accumulation. The model consequently yields an explicit finite threshold at which the interior equilibrium meets a regime with inactive manufacturing R&D while transportation R&D remains active. Regulation thus affects not only the level of costs but the direction and timing of endogenous innovation.

Multiple transportation modes make the result sharper. When modes are substitutes, an industry can escape a regulatory burden imposed on one mode by shifting toward another. The long-run constraint is then set by the mode with the slower accumulation of regulations. When modes are complements, by contrast, the modes form links in a transportation chain, so the mode with the more rapid accumulation of regulations can become the bottleneck. A port and the roads that serve it provide a natural example. The technological relation among modes therefore determines which institutional force binds.

The transition adds a further implication. With intratemporal mode choice, firms always choose the current cost-minimizing mode mix, but the relative regulatory stocks that determine

[3]The treatment of fuel also connects the paper to Schumpeterian growth models with a factor in limited supply; see Peretto (2009, 2012), Peretto and Valente (2011, 2015), and Peretto (2027). Here scarcity matters through the cost of an input used in both making and shipping, which makes its effect on the direction of innovation depend on relative factor intensities.

that mix evolve over time. A temporary divergence in regulatory accumulation can therefore redirect innovation during the transition. If the relative shipping cost crosses the boundary at which manufacturing innovation ceases, the industry accumulates less manufacturing knowledge during the shutdown; whether that shortfall survives depends on subsequent catch-up. By contrast, when both innovation margins remain interior, a fully reversed wedge produces full eventual recovery of the knowledge levels. Boundary crossing, rather than temporary cost push by itself, is therefore the source of possible hysteresis.

The paper thus treats transportation as part of the endogenous growth mechanism rather than as a cost appended to it. Its technologies, inputs, modes, and institutional environment determine relative costs and thereby the allocation of innovation. This integration is what generates the threshold, activation, and mode-substitution results.

The empirical motivation for cumulative regulation is that the regulatory stock exhibits persistent net accumulation, interrupted by occasional declines, and that the pace of accumulation differs across transportation activities. RegData provides a way to measure both features. In the 6.0 release, the federal stock of restrictions rises by 175 percent between 1970 and 2025. Our calculations for transportation over 2000–2025 also show substantial heterogeneity in relevance-weighted accumulation across trucking, rail, air, and water transportation. We do not estimate the structural model here. The evidence instead motivates the institutional feature of the model: regulation enters as an evolving stock whose rate of accumulation can differ across transportation modes. Appendix D describes the measurement and reports the underlying descriptive evidence.

The paper proceeds as follows. Section 2 presents the model environment, including household behavior, the production of the consumption good, manufacturing, transportation, and innovation. Section 3 derives firm behavior, entry, and general equilibrium. Section 4 establishes the Kane–Peretto benchmark. Section 5 introduces transportation modes and fuel. Section 6 studies asymmetry and the role of the endogenous price of fuel. Section 7 introduces cumulative regulation and derives the threshold and long-run transport constraints. Section 8 studies activation and ordering along the development path, Section 9 studies transitional dynamics and hysteresis, and Section 10 concludes.

2 The Model Environment

This section sets up the common environment that underlies the analysis. The production side of the economy has two vertically stacked layers. At the bottom, a Cobb–Douglas

aggregator combines differentiated components into a homogeneous consumption good. Each component, in turn, consists of a CES aggregate of differentiated intermediate goods. Thus, an industry consists of an endogenous set of specialized firms that supply the intermediate goods used to produce one differentiated component of the consumption good. Labor moves freely across industries, so all firms face the same wage, and financial markets integrate across industries, so all firms face the same interest rate.

The activities that matter for the paper are making things and shipping them. We first specify the production and innovation technologies for each activity. We then close the environment with household behavior and the demand system. At this stage we use a single transportation technology. Section 5 introduces multiple transportation modes and fuel.

2.1 Technologies

An intermediate firm makes its differentiated good and then ships it to its customers. Each activity has its own technology and its own knowledge stock. Firms can therefore devote R&D to reducing the cost of manufacturing or the cost of transportation.

2.1.1 Manufacturing

Firm i in industry j produces its differentiated intermediate good according to

$$X_{ij} = Z_{ij}^{\theta_j} \left(L_{Xij} - \varphi_j \right), \qquad 0 < \theta_j < 1, \quad \varphi_j > 0, \tag{1}$$

where L_{Xij} denotes labor used in manufacturing, φ_j is overhead labor, and Z_{ij} is the firm's stock of manufacturing knowledge. The specification gives the firm-specific knowledge stock a direct productivity role while preserving increasing returns to the firm through the fixed overhead requirement.

The firm accumulates manufacturing knowledge according to

$$\dot{Z}_{ij} = \alpha_j Z_j L_{Zij}, \qquad \alpha_j > 0, \tag{2}$$

where L_{Zij} is labor devoted to manufacturing R&D and Z_j denotes the knowledge spillover available to firms in industry j. We assume that the spillover equals the average knowledge stock in the industry:

$$Z_j = \frac{1}{N_j} \int_0^{N_j} Z_{sj} \, ds. \tag{3}$$

This specification gives firms access to the knowledge developed by other firms while preserving the private nature of the firm-specific stock Z_{ij}.

2.1.2 Transportation

The firm must ship its output to its customers. We assume in-house distribution and, for now, use labor as the sole transportation input. A separate competitive market for distribution services would leave the results unchanged. The transportation technology is

$$L_{Dij} = \left[\tau_j Z_{ij}^{-\theta_j} + (1+\tau_j)\eta_j S_{ij}^{-\sigma_j}\right] X_{ij}, \qquad \eta_j \geq 0, \quad \sigma_j \geq 0, \tag{4}$$

where L_{Dij} is labor used in shipping, S_{ij} is the firm's transportation-knowledge stock, and η_j and σ_j govern the per-unit transportation technology. The first term represents an *iceberg* component; the second represents a *per-unit* component.

The iceberg term captures the labor cost of producing the additional quantity that melts in transit. Because manufacturing productivity also determines the cost of producing that additional quantity, an improvement in Z_{ij} lowers both manufacturing cost and the iceberg component of shipping cost. The per-unit component instead captures the cost of physically moving the good. An improvement in manufacturing productivity does not directly reduce that component.

This distinction drives the rest of the paper. With $\eta_j = 0$ and $\tau_j > 0$, transportation consists only of the iceberg component, so manufacturing innovation also reduces the shipping cost. With $\eta_j > 0$ and $\tau_j = 0$, transportation consists only of the per-unit component, so manufacturing innovation cannot reduce the shipping cost. More generally, the per-unit component creates an incompressible wedge between the cost of making the good and the cost of delivering it.

The parameter σ_j determines whether firms can improve the per-unit component through transportation innovation. When $\sigma_j = 0$, that component remains exogenous and can prevent manufacturing innovation from generating sustained endogenous growth. When $\sigma_j > 0$, firms can reduce it by accumulating transportation knowledge. Kane and Peretto (2020) establish the resulting dynamic complementarity between manufacturing and transportation innovation; we use that mechanism as the benchmark for the new analysis that follows.

We specify transportation innovation in parallel with manufacturing innovation:

$$\dot{S}_{ij} = \varsigma_j S_j L_{Sij}, \qquad \varsigma_j > 0, \tag{5}$$

where L_{Sij} denotes labor devoted to transportation R&D and S_j is the industry-wide transportation knowledge spillover. We again define the spillover as the industry average,

$$S_j = \frac{1}{N_j} \int_0^{N_j} S_{kj}\, dk. \tag{6}$$

The parallel specification lets us compare the incentives to improve manufacturing and transportation without introducing unnecessary differences in the innovation technologies themselves. The richer environment that matters for this paper enters later through fuel, multiple transportation modes, and cumulative mode-specific regulation.

2.2 Household behavior and demand

The representative household has population

$$L(t) = \Lambda e^{\lambda t}, \qquad \Lambda > 0, \qquad \lambda \in (-\infty, +\infty). \tag{7}$$

We do not restrict population growth to be positive. We also abstract from labor–leisure choice, so the household supplies all of its labor.

The household maximizes

$$U(t) = \int_t^{\infty} e^{-\rho(s-t)} L(s) u(s)\, ds, \qquad \rho > \max\{0, \lambda\}, \tag{8}$$

where instantaneous utility is

$$u = \log\left(\frac{Y}{p_G L}\right). \tag{9}$$

Here Y denotes household expenditure on the consumption good and p_G its price.

The household budget constraint is

$$\dot{A} = rA + \Pi_G + \sum_{j=1}^{J} \Pi_{Gj} + wL - Y, \tag{10}$$

where A denotes asset holdings, r the rate of return on assets, w the wage, and Π_G and Π_{Gj} the profits distributed as dividends by the consumption-good sector and its differentiated-component suppliers. Competitive equilibrium makes these profit terms zero. The household's

optimal expenditure path therefore satisfies

$$r = r_A \equiv \rho + \frac{\dot{Y}}{Y} - \lambda. \tag{11}$$

We interpret r_A as the reservation return at which the household trades current for future consumption.

A representative competitive firm produces the homogeneous consumption good according to

$$G = \prod_{j=1}^{J} G_j^{\phi_j}, \qquad \sum_{j=1}^{J} \phi_j = 1. \tag{12}$$

The firm spends the fraction ϕ_j of its revenue on component j, so

$$\frac{Y_j}{Y} = \phi_j, \qquad j = 1, \ldots, J, \tag{13}$$

and competitive equilibrium implies $\Pi_G = 0$.

A representative competitive firm in industry j produces component G_j according to the CES technology

$$G_j = N_j^{\frac{\chi_j - 1}{\epsilon_j - 1}} \left(\int_0^{N_j} X_{ij}^{\frac{\epsilon_j - 1}{\epsilon_j}} \, di \right)^{\frac{\epsilon_j}{\epsilon_j - 1}}, \qquad \chi_j > 0, \quad \epsilon_j > 1, \tag{14}$$

where χ_j captures increasing returns to variety, ϵ_j is the elasticity of substitution across specialized intermediates, X_{ij} is the quantity of intermediate good i, and N_j is the mass of active suppliers in industry j.

The resulting demand for intermediate good i in industry j is

$$X_{ij} = Y_j \frac{p_{Xij}^{-\epsilon_j}}{\int_0^{N_j} p_{Xkj}^{1-\epsilon_j} \, dk}, \qquad i \in [0, N_j], \tag{15}$$

and competitive equilibrium implies

$$\sum_{j=1}^{J} \Pi_{Gj} = 0. \tag{16}$$

2.3 What the model isolates

The model now contains the essential ingredients for the analysis. Firms make goods with firm-specific manufacturing knowledge and ship those goods with firm-specific transportation knowledge. Manufacturing productivity affects both activities through the iceberg component, whereas transportation knowledge affects only the per-unit component. The resulting asymmetry makes the relative cost of shipping central to the direction of innovation.

The next sections add the firm's optimization problem, entry, and general equilibrium. We then establish the benchmark dynamics before introducing multiple transportation modes and fuel. That sequence allows us to separate the mechanism taken from Kane and Peretto (2020) from the new mechanism generated by endogenous fuel prices and cumulative, mode-specific regulatory frictions.

3 Firm Behavior, Entry, and General Equilibrium

This section characterizes the decisions of intermediate firms and entrants and then constructs the economy's general equilibrium. The firm chooses its price, allocates labor between production, shipping, and the two R&D activities, and takes the industry-wide knowledge spillovers and the mass of firms as given. Free entry then links firm value to market size and delivers a particularly useful aggregation result: expenditure per capita is constant along the balanced growth path, and the interest rate equals the household's discount rate.

3.1 Firm decisions

At time t, firm i in industry j maximizes the present discounted value of dividends,

$$V_{ij}(t) = \int_t^\infty \exp\Big[-\int_t^s \big(r(v) + \delta_j\big)\, dv\Big]\, \Pi_{ij}(s)\, ds, \qquad \delta_j > 0, \tag{12}$$

where the firm's dividend flow is

$$\Pi_{ij} = \Big[p_{Xij} - w(1+\tau_j)\left(Z_{ij}^{-\theta_j} + \eta_j S_{ij}^{-\sigma_j}\right)\Big] X_{ij} - w\varphi_j - wL_{Zij} - wL_{Sij}. \tag{13}$$

The parameter δ_j represents the firm's exit shock.[4]

[4]The exit shock guarantees symmetric dynamics in a neighborhood of the steady state when population growth is zero. It also provides additional empirical discipline in quantitative exercises; see, e.g., Ferraro,

Because the model contains a continuum of intermediate goods, each intermediate firm is atomistic. The firm therefore takes the demand curve in equation (15), the innovation technologies in equations (2) and (5), the initial knowledge stocks, the industry spillovers Z_j and S_j, and the mass of firms N_j as given. It chooses the product price and the allocation of R&D labor subject to $L_{Zij} \geq 0$ and $L_{Sij} \geq 0$.

To characterize the firm's choices, we write its current-value Hamiltonian as

$$\begin{aligned}\mathcal{H}_{ij} = &\left[p_{Xij} - wZ_{ij}^{-\theta_j} - w\tau_j Z_{ij}^{-\theta_j} - w(1+\tau_j)\eta_j S_{ij}^{-\sigma_j}\right] X_{ij} - w\varphi_j - wL_{Zij} - wL_{Sij} \\ &+ v_{Zij}\alpha_j Z_j L_{Zij} + v_{Sij}\varsigma_j S_j L_{Sij}.\end{aligned}$$

where v_{Zij} and v_{Sij} denote the shadow values of manufacturing-specific and transportation-specific knowledge. The firm treats Z_{ij} and S_{ij} as its state variables and p_{Xij}, L_{Zij}, and L_{Sij} as its controls. Given the properties of the primitives, the first-order conditions characterize the firm's optimum.

The firm's value-maximizing price is a constant markup over the marginal cost of making and shipping the intermediate good:

$$p_{Xij} = \frac{\epsilon_j}{\epsilon_j - 1} w(1+\tau_j)\left(Z_{ij}^{-\theta_j} + \eta_j S_{ij}^{-\sigma_j}\right). \tag{14}$$

Under the restriction $1 > \theta_j(\epsilon_j - 1)$, Peretto (2027) shows that the industry equilibrium is symmetric when firms start with identical manufacturing and transportation knowledge and new entrants enter at the industry-average knowledge stocks. We therefore drop the firm subscript i from firm-specific variables from this point onward, with the understanding that these variables denote the industry average. In particular, L_{Xj} denotes average manufacturing employment, L_{Zj} average manufacturing R&D employment, and similarly for the other activities. Since

$$Y_j = \int_0^{N_j} p_{Xij} X_{ij}\, di,$$

we have $p_{Xj}X_j = Y_j/N_j$, where X_j denotes average firm output. Inverting the manufacturing and transportation technologies and using equation (14) therefore gives

$$L_{Xj} + L_{Dj} = \frac{\epsilon_j - 1}{\epsilon_j}\frac{Y_j}{wN_j} + \varphi_j. \tag{15}$$

The Hamiltonian is linear in L_{Zij} and L_{Sij}, so the two R&D choices have a bang-bang structure. We do not need the full costate system here. The two relevant no-arbitrage conditions equate

Ghazi, and Peretto (2020, 2023).

the market interest rate to the return from investing in firm-specific manufacturing and transportation knowledge:

$$r = r_{Zj} \equiv \alpha_j \left[\frac{Y_j}{wN_j} \frac{\theta_j(\epsilon_j - 1)}{\epsilon_j} \frac{Z_{ij}^{-\theta_j}}{Z_{ij}^{-\theta_j} + \eta_j S_{ij}^{-\sigma_j}} - L_{Zj} \right] + \frac{\dot{w}}{w} - \delta_j, \tag{16}$$

$$r = r_{Sj} \equiv \varsigma_j \left[\frac{Y_j}{wN_j} \frac{\sigma_j(\epsilon_j - 1)}{\epsilon_j} \frac{\eta_j S_{ij}^{-\sigma_j}}{Z_{ij}^{-\theta_j} + \eta_j S_{ij}^{-\sigma_j}} - L_{Sj} \right] + \frac{\dot{w}}{w} - \delta_j. \tag{17}$$

These conditions make the model's endogenous-innovation mechanism transparent. For a given market environment, the return to firm-specific innovation depends linearly on firm size, measured by market size times the firm's market share. The same mechanism makes industrial organization endogenous: entry and exit determine the number of firms, average firm size, and concentration, and those market-structure variables in turn determine the incentives to innovate.

3.2 Entry

An entrepreneur can create a new firm with productivity equal to the industry average by hiring $\beta_j Y_j/(wN_j)$ units of labor. The sunk entry cost therefore depends on industry market size.[5] To keep the cross-industry heterogeneity focused on technologies and demand, we impose from this point onward the simplifying restriction $\beta_j = \beta$ for all j. This restriction is not essential to the mechanism, but it makes the industry dynamics and their aggregation considerably more transparent. Because $\sum_j \phi_j = 1$, the expenditure-weighted entry parameter defined below is then the same common β. In equilibrium, no entry occurs when $V_j < \beta_j Y_j/N_j$, while free entry requires

$$V_j = \frac{\beta_j Y_j}{N_j}.$$

The associated rate of return to entry is

$$r = r_{Nj} \equiv \frac{N_j}{\beta_j Y_j} \left(\frac{Y_j}{\epsilon_j N_j} - w\varphi_j - wL_{Zj} - wL_{Sj} \right) + \frac{\dot{Y}_j}{Y_j} - \frac{\dot{N}_j}{N_j} - \delta_j. \tag{18}$$

[5]See Peretto and Connolly (2007) and Peretto (2027) for an interpretation of this assumption and alternative formulations of the sunk entry cost that leave the fundamental equilibrium properties unchanged. In this class of general-equilibrium models, an entrepreneur is a member of the representative household who chooses to create a firm and issues equity in the financial market to finance the labor required to develop a new good and bring it to market.

Each industry can independently support either a no-entry equilibrium or a free-entry equilibrium. Thus, the economy can in principle contain industries with no firm entry alongside industries with free entry. We focus on the latter configuration and leave the former for future work.[6]

3.3 General equilibrium

We choose labor as the numeraire and set $w \equiv 1$. Asset-market equilibrium requires equal returns across all assets and activities,

$$r = r_A = r_{Zj} = r_{Nj}, \qquad j = 1, \ldots, J,$$

and requires the value of the household's portfolio to equal the value of the securities issued by firms,

$$A = \sum_{j=1}^{J} N_j V_j.$$

If an activity offered a return below the reservation return of savers, investors would withdraw resources from it. We therefore focus on the equilibrium in which all industries satisfy free entry.

Free entry gives $V_j = \beta_j Y_j / N_j$. Since the Cobb–Douglas final-good producer spends the fraction ϕ_j of total expenditure on industry j, we have $Y_j = \phi_j Y$. Hence

$$A = \sum_{j=1}^{J} N_j V_j = \sum_{j=1}^{J} \beta_j \phi_j Y = \beta Y, \qquad \beta \equiv \sum_{j=1}^{J} \beta_j \phi_j.$$

This aggregation result is central to the subsequent analysis. Despite heterogeneity across industries, the economy has a constant wealth-to-expenditure ratio.

Substituting the free-entry aggregation $A = \beta Y$ into the household budget constraint and using the Euler equation yields

$$Y = \left[1 - (\rho - \lambda)\beta\right]^{-1} wL, \qquad 1 > (\rho - \lambda)\beta. \tag{19}$$

Thus consumption expenditure remains proportional to wage income. With $w \equiv 1$, expendi-

[6] Specifically, we plan to study structural slumps and a novel form of propagation in the multi-industry setting, which we label contagion; see Peretto (2027) for the definition and analysis in models without transportation.

ture per capita becomes

$$y = y^* \equiv [1 - (\rho - \lambda)\beta]^{-1}. \tag{20}$$

Because expenditure per capita is constant, the household Euler equation implies

$$r = \rho.$$

Importantly, this aggregation result does not require us to take a stand on whether incumbent firms undertake in-house R&D. The subsequent industry analysis can therefore treat the innovation margins separately while taking y^* and the interest rate as given.

The free-entry equilibrium now gives each industry a self-contained dynamical system. This separation is what makes the model tractable: the economy-wide aggregation fixes expenditure per capita and the interest rate, while each industry determines how its own firm size, knowledge stocks, and transportation costs evolve.

4 The Kane–Peretto Benchmark

We now establish the benchmark that underlies the analysis that follows. We take the core mechanism from Kane and Peretto (2020), and reproduce it in the notation and multi-industry environment developed above. We suppress quality innovation to isolate the interaction between cost reduction in making and shipping. This removes a margin through which firms could otherwise redirect innovation and lets the transportation mechanism operate transparently. Making things and shipping things are distinct activities, each with its own cost-reduction technology, and the form of the shipping cost determines whether innovation in manufacturing can sustain endogenous growth.

We first shut down innovation in transportation and recover the result of Kane and Peretto (2020) that an exogenous per-unit shipping cost eventually stops manufacturing cost reduction. We then allow firms to innovate in transportation and recover the complementary mechanism through which manufacturing and shipping costs can decline together at a common rate. This second case gives us the two activation boundaries that organize the dynamics in the remainder of the paper. The purpose of the section is to put the Kane–Peretto mechanism fully in place before we add the elements that generate our contribution: multiple transportation modes, a second input, and mode-specific regulatory frictions that accumulate over time.

4.1 No innovation in transportation

Set $\sigma_j = 0$ for all industries. Firms then choose $L_{Sj} = 0$, and the per-unit component of shipping costs remains exogenous. The economy still contains the ingredients that generate endogenous growth in a conventional Schumpeterian model, but an incompressible shipping cost changes the result qualitatively: manufacturing innovation alone cannot sustain endogenous growth when $\eta_j > 0$. This is the central transportation constraint identified by Kane and Peretto (2020).

Industry j has expenditure $Y_j = \phi_j Y$, so its employment share equals its expenditure share, $L_j/L = \phi_j$. Define firm size by

$$x_j \equiv \frac{Y_j}{N_j}, \tag{17}$$

and define the relative cost of shipping to making as

$$q_j \equiv \eta_j Z_j^{\theta_j} S_j^{-\sigma_j} = \frac{\text{labor productivity in manufacturing}}{\text{labor productivity in shipping}}. \tag{18}$$

With $\sigma_j = 0$, this reduces to $q_j = \eta_j Z_j^{\theta_j}$. Thus manufacturing innovation raises the relative importance of the exogenous shipping cost.

Solving the return condition for manufacturing R&D gives

$$z_j \equiv \frac{\dot{Z}_j}{Z_j} = \alpha_j L_{Zj} = \alpha_j \frac{\theta_j(\epsilon_j - 1)}{\epsilon_j} \frac{x_j}{1 + q_j} - \rho - \delta_j, \tag{19}$$

so firms innovate if and only if firm size exceeds

$$x_{Zj}(q_j) \equiv \frac{\epsilon_j(\rho + \delta_j)}{\alpha_j \theta_j(\epsilon_j - 1)}(1 + q_j). \tag{20}$$

The boundary is increasing in q_j: as shipping becomes more expensive relative to making, a firm must become larger before improving its factory generates a sufficiently high return.

Substituting (19) into the entry condition gives

$$\frac{\dot{q}_j}{q_j} = \theta_j z_j, \tag{21}$$

$$\frac{\dot{x}_j}{x_j} = \rho + \delta_j - \frac{1}{\beta x_j}\left[\frac{x_j}{\epsilon_j} - \varphi_j - \frac{\theta_j(\epsilon_j - 1)}{\epsilon_j}\frac{x_j}{1 + q_j} + \frac{\rho + \delta_j}{\alpha_j}\right]. \tag{22}$$

Whenever $z_j > 0$, q_j rises. The per-unit shipping cost therefore becomes increasingly important

relative to the cost of making the good, which reduces the return to further manufacturing innovation. For $\eta_j > 0$, the system has no steady state with sustained endogenous growth. Only when $\eta_j = 0$ does q_j remain constant and the standard manufacturing-only growth mechanism re-emerge. This is the Kane–Peretto transportation constraint in its simplest form.

The phase diagram also contains path dependence. For some initial conditions, the economy reaches the region in which manufacturing innovation ceases and then converges to the $\dot{x}_j = 0$ locus. The resulting steady state depends on the initial level of manufacturing knowledge and, more generally, on the point at which the economy crosses the $\dot{q}_j = 0$ boundary. Thus the benchmark already separates a level effect from a growth effect: history can affect the eventual level of production knowledge even though the long-run growth rate is unchanged. The omission of quality innovation makes this mechanism especially transparent.[7]

4.2 Innovation in transportation

Now set $\sigma_j > 0$. Firms can improve both manufacturing and transportation. Kane and Peretto show that the two margins together can sustain endogenous growth: a steady state exists in which the unit costs of making things and shipping things decline at a common rate. We reproduce that mechanism here because it supplies the activation boundaries and state variables used later.

Manufacturing and transportation innovation satisfy

$$z_j = \alpha_j L_{Zj} = \alpha_j \frac{\theta_j(\epsilon_j - 1)}{\epsilon_j} \frac{x_j}{1 + q_j} - \rho - \delta_j, \tag{23}$$

$$s_j \equiv \frac{\dot{S}_j}{S_j} = \varsigma_j L_{Sj} = \varsigma_j \frac{\sigma_j(\epsilon_j - 1)}{\epsilon_j} \frac{q_j}{1 + q_j} x_j - \rho - \delta_j. \tag{24}$$

The two activation boundaries are therefore

$$x_{Zj}(q_j) \equiv \frac{\epsilon_j(\rho + \delta_j)}{\alpha_j \theta_j(\epsilon_j - 1)}(1 + q_j), \tag{25}$$

$$x_{Sj}(q_j) \equiv \frac{\epsilon_j(\rho + \delta_j)}{\varsigma_j \sigma_j(\epsilon_j - 1)} \frac{1 + q_j}{q_j}. \tag{26}$$

The first boundary rises with q_j, whereas the second falls. When shipping becomes expensive

[7] Appendix B develops the corresponding phase-diagram analysis and the continuum of stationary outcomes for the present model.

relative to making, a firm must be larger before improving the factory becomes profitable, but it becomes profitable to improve shipping at a smaller firm size. Manufacturing and shipping innovation therefore become complementary margins.

Taking logs and time derivatives of q_j gives

$$\frac{\dot{q}_j}{q_j} = \theta_j z_j - \sigma_j s_j = \frac{\epsilon_j - 1}{\epsilon_j} \frac{x_j}{1 + q_j} \left(\alpha_j \theta_j^2 - \varsigma_j \sigma_j^2 q_j \right) - (\theta_j - \sigma_j)(\rho + \delta_j), \tag{27}$$

and the entry condition gives

$$\frac{\dot{x}_j}{x_j} = \rho + \delta_j - \frac{1}{\beta} \left[\frac{1}{\epsilon_j} - \frac{\theta_j + \sigma_j q_j}{1 + q_j} \frac{\epsilon_j - 1}{\epsilon_j} - \frac{\varphi_j - (\rho + \delta_j)(1/\alpha_j + 1/\varsigma_j)}{x_j} \right]. \tag{28}$$

Equations (27)–(28) form an autonomous system in (q_j, x_j). The two activation boundaries partition the state space according to whether firms improve manufacturing, transportation, both, or neither.

In the interior steady state, let

$$\Xi_j \equiv \varphi_j - (\rho + \delta_j) \left(\frac{1}{\alpha_j} + \frac{1}{\varsigma_j} \right). \tag{29}$$

The steady-state conditions imply

$$x_j^* = \frac{(\theta_j - \sigma_j)(\rho + \delta_j)(1 + q_j^*)}{\frac{\epsilon_j - 1}{\epsilon_j} \left(\alpha_j \theta_j^2 - \varsigma_j \sigma_j^2 q_j^* \right)}, \tag{30}$$

with q_j^* determined by the second steady-state condition. We assume $\theta_j > \sigma_j$, so that manufacturing knowledge has the larger direct productivity effect in this benchmark. The parameter η_j does not enter the steady-state ratio q_j^* because firms adjust the two knowledge stocks to maintain that ratio; τ_j likewise does not affect (q_j^*, x_j^*). The equilibrium cost structure therefore depends not on the initial relative cost of making and shipping, but on the relative ability to improve the two activities.

In steady state, production cost, the iceberg component of shipping cost, and the per-unit shipping cost all decline at a common rate while remaining in constant proportions. Since $\dot{q}_j = 0$,

$$\theta_j z_j^* = \sigma_j s_j^*, \qquad s_j^* = \frac{\theta_j}{\sigma_j} z_j^*. \tag{31}$$

This is the key Kane–Peretto complementarity: manufacturing innovation raises the relative importance of shipping, while shipping innovation restores the return to manufacturing R&D.

4.3 TFP and the benchmark growth mechanism

The benchmark also separates transportation's level effect from its growth effect. With no transportation innovation, industry TFP is

$$T_j \equiv \frac{Z_j^{\theta_j} N_j^{\chi_j}}{1 + \eta_j Z_j^{\theta_j}} = \frac{Z_j^{\theta_j} N_j^{\chi_j}}{1 + q_j}, \tag{32}$$

and industry output per worker grows at

$$g_j = \chi_j n_j + \frac{\theta_j z_j}{1 + q_j}. \tag{33}$$

The contribution of manufacturing cost reduction to measured productivity is therefore divided by $1 + q_j$. As q_j rises, that contribution vanishes. In the steady state of this benchmark, $g_j^* = \chi_j \lambda$: transportation prevents manufacturing cost reduction from generating a permanent growth effect.

When firms also innovate in transportation,

$$T_j = \frac{Z_j^{\theta_j} N_j^{\chi_j}}{1 + q_j} = \frac{Z_j^{\theta_j} N_j^{\chi_j} S_j^{\sigma_j}}{S_j^{\sigma_j} + \eta_j Z_j^{\theta_j}}. \tag{34}$$

The per-unit shipping cost makes TFP concave in each knowledge stock separately, so neither innovation margin can sustain growth on its own. Jointly, however, the two knowledge stocks enter homogeneously, which makes endogenous growth possible when firms improve both margins together. Industry growth satisfies

$$g_j = \chi_j n_j + \theta_j z_j - \frac{q_j}{1 + q_j} \frac{\dot{q}_j}{q_j}, \tag{35}$$

so that in the steady state

$$g_j^* = \chi_j \lambda + \theta_j z_j^*, \qquad s_j^* = \frac{\theta_j}{\sigma_j} z_j^*. \tag{36}$$

The endogenous component of growth now survives because firms improve shipping as they improve manufacturing. It depends on the capacity to improve the two technologies, rather than simply on the current cost of transportation.

4.4 What we take from the benchmark

The benchmark establishes the mechanism that we take from Kane and Peretto (2020). Making things and shipping things are distinct activities; a per-unit shipping cost creates a wedge that manufacturing innovation alone cannot eliminate; and when firms can improve shipping as well as manufacturing, the two innovation margins become complementary and can sustain a common decline in their unit costs. The variables q_j and x_j provide the compact state space in which the activation of these margins can be studied.

We do not claim the manufacturing–shipping complementarity, the common-rate steady state, the phase-diagram structure, or the induced-innovation mechanism associated with the per-unit shipping cost as contributions of this paper. Our contribution concerns what happens when transportation uses a second input, operates through multiple modes, and faces mode-specific regulatory frictions that accumulate at different rates.[8]

Suppressing quality innovation serves a second purpose. In addition to isolating the interaction between cost reduction in making and shipping, it makes the present model substantially more tractable than the fuller Kane–Peretto environment. The resulting low-dimensional dynamic system permits analytical comparative statics and a transparent characterization of transitional dynamics. Appendix C shows that the same structure also supports conventional comparative statics, transition dynamics, and welfare analysis.

The distinction is important for the results that follow. The question is no longer whether transportation can constrain growth. Kane and Peretto establish that it can. We instead ask which transportation activity binds, how firms reallocate innovation in response, and what happens when the institutional friction attached to that activity accumulates over time.

5 Transportation Modes and Fuel

The benchmark in Section 4 isolates the complementarity between making things and shipping things with a single transportation technology. We now incorporate two properties of transportation that are central to the paper. First, production and transportation use a second input, which we call fuel. Second, firms can ship through more than one transportation

[8]The appendices are organized to make the relation between the benchmark, the present model, and the additional analytical exercises transparent. Appendix A compares the present model with Kane and Peretto (2020); Appendix B develops the benchmark dynamics of the present model, including the phase-diagram analysis; and Appendix C uses the tractability of the resulting structure to develop standard comparative-statics and transition-dynamics exercises.

mode. We use road and water transportation as the two modes in the analytical specification. Each mode uses the same labor–fuel bundle as manufacturing but has its own productivity factor. Those factors can capture geography, infrastructure, or regulation. This structure gives us a microfoundation for the shipping-cost term that was exogenous in the benchmark.

The specification preserves tractability while making transportation an integral part of the growth mechanism. The key restriction is that manufacturing and transportation use labor and fuel with the same factor intensity within each industry. This restriction allows the cost of the common labor–fuel bundle to factor out of the ratio of the shipping cost to the manufacturing cost. The mode-mix cost remains in that ratio because it applies only to shipping. The distinction is the source of the asymmetry principle developed in Section 6: common cost push changes the level of costs, whereas shipping-specific cost push changes their relative cost and therefore the direction of innovation.

5.1 Production and transportation with fuel and modes

Replace the variable labor input in manufacturing with a composite input produced from labor and fuel. Let $C_j(w, p_f)$ denote the constant-returns unit cost of this composite input in industry j, where w is the wage and p_f is the price of fuel. For the aggregation below we use the constant-share specification

$$C_j(w, p_f) = w^{\omega_j} p_f^{1-\omega_j}, \qquad 0 < \omega_j < 1,$$

up to an inessential normalization constant. The Cobb–Douglas form is convenient because it keeps the labor and fuel cost shares constant; homogeneity alone would not imply that property. The unit cost of making the intermediate good then becomes

$$Z_{ij}^{-\theta_j} C_j(w, p_f). \tag{37}$$

We model transportation as a technology that produces distance covered by one unit of the good. Normalize the distance to one, $D_{ij} = 1$, without loss of generality. The firm can combine road and water transportation according to the constant-returns unit-cost function

$$S_{ij}^{-\sigma_j} C_{Dj}\left(p_{ij}^{rd}, p_{ij}^{wa}\right) D_{ij}, \tag{38}$$

where p_{ij}^{rd} and p_{ij}^{wa} are the internal unit prices of road and water transportation. The transportation knowledge stock S_{ij} therefore scales the productivity of the mode-mix technology

in the same way that Z_{ij} scales manufacturing productivity.

Road and water transportation use the same labor–fuel bundle as manufacturing, but their productivity differs through mode-specific factors. Write their unit costs as

$$p_{ij}^{rd} = R_j^{rd} C_j(w, p_f), \qquad p_{ij}^{wa} = R_j^{wa} C_j(w, p_f). \tag{39}$$

where R_j^{rd} and R_j^{wa} reduce the productivity of the corresponding mode. These factors can represent geographical conditions, such as terrain or access to waterways, or institutional conditions, such as regulation. They can therefore differ across industries and modes.

The mode-mix cost function is homogeneous of degree one in the two mode prices. Consequently,

$$C_{Dj}\Big(R_j^{rd} C_j(w, p_f), R_j^{wa} C_j(w, p_f)\Big) = C_j(w, p_f) C_{Dj}\Big(R_j^{rd}, R_j^{wa}\Big). \tag{40}$$

This factorization separates two objects that will play very different roles. The function $C_j(w, p_f)$ determines the cost of the labor–fuel bundle used in both making and shipping. The function $C_{Dj}(R_j^{rd}, R_j^{wa})$ determines the cost of choosing and combining transportation modes, and therefore applies only to shipping.

The elasticity of substitution between the two modes, denoted by ζ_j, determines how the industry responds when the relative productivity of the modes changes. We have

$$\zeta_j > 1 \quad \text{for substitutes,} \qquad \zeta_j < 1 \quad \text{for complements.}$$

When the modes are substitutes, the firm can shift toward the mode whose relative cost falls. When the modes are complements, the modes represent successive stages of the same logistics chain: for example, water transportation may require road transportation to connect a port to the customer. In that case, a higher cost in one mode cannot be avoided by using the other mode more intensively. The sign of $\zeta_j - 1$ therefore determines whether the industry can escape a mode-specific increase in cost or instead encounters a bottleneck.

For illustration, suppose the mode-mix technology has a CES unit-cost function,

$$C_{Dj}(p^{rd}, p^{wa}) = \left[\varrho_j^{\zeta_j} (p^{rd})^{1-\zeta_j} + (1-\varrho_j)^{\zeta_j} (p^{wa})^{1-\zeta_j}\right]^{1/(1-\zeta_j)}, \qquad 0 < \varrho_j < 1. \tag{41}$$

The corresponding cost shares are the derivatives of the log unit cost with respect to the log price of each mode,

$$c_j^{rd} \equiv \frac{\partial \log C_{Dj}}{\partial \log p^{rd}}, \qquad c_j^{wa} \equiv \frac{\partial \log C_{Dj}}{\partial \log p^{wa}}, \qquad c_j^{rd} + c_j^{wa} = 1. \tag{42}$$

These shares summarize the mode mix without requiring us to carry the individual mode inputs through the rest of the analysis.

5.2 The firm's transportation choice

The firm's price and variable-input decisions are intratemporal. Given Z_{ij}, S_{ij}, the wage, and the fuel price, the firm chooses the least-cost combination of manufacturing and transportation inputs. The resulting current-value Hamiltonian has the same intertemporal terms as in Section 3.1, but its intratemporal profit now includes the composite manufacturing cost and the per-unit shipping cost:

$$\begin{aligned}\mathcal{H}_{ij} = & \left[p_{Xij} - (1+\tau_j) Z_{ij}^{-\theta_j} C_j(w, p_f) - p_{Dij}\right] X_{ij} \\ & - w\varphi_j - wL_{Zij} - wL_{Sij} + v_{Zij}\alpha_j Z_j L_{Zij} + v_{Sij}\varsigma_j S_j L_{Sij}.\end{aligned}$$

where

$$p_{Dij} = S_{ij}^{-\sigma_j} C_j(w, p_f) C_{Dj}(R_j^{rd}, R_j^{wa}) \tag{43}$$

is the least unit cost of shipping the good over the normalized distance.

The firm's pricing equation follows from the same constant-elasticity demand system used in Section 3.1. Solving the firm's pricing equation for p_{Xij} yields

$$p_{Xij} = \frac{\epsilon_j}{\epsilon_j - 1}\left[(1+\tau_j) Z_{ij}^{-\theta_j} C_j(w, p_f) + p_{Dij}\right]. \tag{44}$$

The corresponding gross profit equals $p_{Xij} X_{ij}/\epsilon_j$. More importantly for the dynamics, the new cost structure changes the allocation of gross profit across inputs but does not change the return equations for manufacturing innovation, transportation innovation, or entry. The reason is that the common factor $C_j(w, p_f)$ enters both activities in the same way.

Production and transportation expenditure are, respectively,

$$\begin{aligned}\mathcal{E}_{Xij} &= \frac{Z_{ij}^{-\theta_j} C_j(w, p_f)}{(1+\tau_j) Z_{ij}^{-\theta_j} C_j(w, p_f) + p_{Dij}} \frac{\epsilon_j - 1}{\epsilon_j} p_{Xij} X_{ij}, \\ \mathcal{E}_{Dij} &= \frac{\tau_j Z_{ij}^{-\theta_j} C_j(w, p_f) + p_{Dij}}{(1+\tau_j) Z_{ij}^{-\theta_j} C_j(w, p_f) + p_{Dij}} \frac{\epsilon_j - 1}{\epsilon_j} p_{Xij} X_{ij}.\end{aligned} \tag{45}$$

These two expenditures add to variable cost, $(\epsilon_j - 1) p_{Xij} X_{ij}/\epsilon_j$. The fixed operating cost and R&D expenditure are separate uses of firm revenue, just as in the benchmark.

The cost shares c_j^{rd} and c_j^{wa} determine the allocation of transportation expenditure across the two modes. Because manufacturing and transportation use the labor–fuel bundle with the same factor intensity, the labor and fuel expenditure in each activity follows directly from the common unit-cost function. In particular,

$$w(L_{Xij} - \varphi_j) = \frac{Z_{ij}^{-\theta_j}(1+\tau_j)}{(1+\tau_j)Z_{ij}^{-\theta_j} + S_{ij}^{-\sigma_j} C_{Dj}(R_j^{rd}, R_j^{wa})} \omega_j \frac{\epsilon_j - 1}{\epsilon_j} p_{Xij} X_{ij}, \tag{46}$$

$$p_f F_{Xij} = \frac{Z_{ij}^{-\theta_j}(1+\tau_j)}{(1+\tau_j)Z_{ij}^{-\theta_j} + S_{ij}^{-\sigma_j} C_{Dj}(R_j^{rd}, R_j^{wa})} (1-\omega_j) \frac{\epsilon_j - 1}{\epsilon_j} p_{Xij} X_{ij}, \tag{47}$$

$$wL_{Dij} = \frac{\tau_j Z_{ij}^{-\theta_j} + S_{ij}^{-\sigma_j} C_{Dj}(R_j^{rd}, R_j^{wa})}{(1+\tau_j)Z_{ij}^{-\theta_j} + S_{ij}^{-\sigma_j} C_{Dj}(R_j^{rd}, R_j^{wa})} \omega_j \frac{\epsilon_j - 1}{\epsilon_j} p_{Xij} X_{ij}, \tag{48}$$

$$p_f F_{Dij} = \frac{\tau_j Z_{ij}^{-\theta_j} + S_{ij}^{-\sigma_j} C_{Dj}(R_j^{rd}, R_j^{wa})}{(1+\tau_j)Z_{ij}^{-\theta_j} + S_{ij}^{-\sigma_j} C_{Dj}(R_j^{rd}, R_j^{wa})} (1-\omega_j) \frac{\epsilon_j - 1}{\epsilon_j} p_{Xij} X_{ij}. \tag{49}$$

An analogous decomposition applies to transportation expenditure, with the iceberg component and the per-unit shipping component determining the share of the expenditure devoted to transportation. We will not need the individual labor and fuel demands below; what matters is their aggregation across industries.

5.3 Aggregation and the fuel market

Symmetry within industry j implies

$$p_f F_j = (1-\omega_j) \frac{\epsilon_j - 1}{\epsilon_j} Y_j. \tag{50}$$

Thus aggregate fuel expenditure is proportional to aggregate expenditure on final output,

$$p_f F^d = \kappa Y, \qquad \kappa \equiv \sum_{j=1}^{J} (1-\omega_j) \frac{\epsilon_j - 1}{\epsilon_j} \phi_j < 1. \tag{51}$$

The inequality follows from the maintained restrictions on the expenditure shares and elasticities.

The household owns the raw fuel resource and the labor used to process it. Because processing labor is already included in aggregate labor income wL, the additional income generated by the resource is its scarcity rent, not gross fuel revenue. Let F^d denote current fuel use. Competitive processing therefore gives aggregate resource rent $(p_f - \kappa_f w)F^d$, and the

household budget constraint becomes

$$\dot{A} = rA + wL + (p_f - \kappa_f w)F^d - Y. \tag{52}$$

This accounting distinction matters for aggregation. When the resource constraint does not bind, $p_f = \kappa_f w$ and scarcity rent is zero, so the benchmark result survives unchanged:

$$y^u = \frac{1}{1-(\rho-\lambda)\beta}, \qquad r = \rho. \tag{53}$$

When the fixed resource endowment binds, $F^d = F$ and equation (51) implies $p_f F = \kappa Y$. With $w = 1$, $A = \beta Y$, and the household Euler equation, the budget constraint yields

$$y(t) \equiv \frac{Y(t)}{L(t)} = \frac{1 - \kappa_f F/L(t)}{1-\kappa-(\rho-\lambda)\beta}, \qquad r(t) = \rho + \frac{\dot{y}(t)}{y(t)}. \tag{54}$$

Thus expenditure per capita is not constant during the transition into the resource-constrained regime. If $\lambda > 0$, processing the fixed endowment absorbs a vanishing share of the growing labor force and

$$y(t) \longrightarrow y^c \equiv \frac{1}{1-\kappa-(\rho-\lambda)\beta}, \qquad r(t) \longrightarrow \rho. \tag{55}$$

The earlier benchmark value therefore applies in the unconstrained regime, while the expression with κ is the asymptotic constrained-regime value rather than an exact transition value.

To close the fuel market, let the economy possess an endowment F of a resource that can be processed into fuel. Producing one unit of fuel requires κ_f units of labor. The supply of fuel is therefore perfectly elastic at

$$p_f = \kappa_f w \tag{56}$$

until the economy reaches the resource constraint F. We normalize $w = 1$ as before. Two cases then arise.

When aggregate demand for fuel is sufficiently low, the resource constraint does not bind. Fuel use satisfies

$$F^d = \frac{\kappa y^u L}{\kappa_f} < F, \tag{57}$$

and the equilibrium fuel price is $p_f = \kappa_f$. The per-unit shipping cost in industry j is therefore

$$p_{Dj} = S_j^{-\sigma_j} C_j(1, \kappa_f) C_{Dj}(R_j^{rd}, R_j^{wa}). \tag{58}$$

When fuel demand reaches the resource constraint, the economy uses the entire endowment

and the fuel price adjusts to clear the market:

$$p_f = \frac{\kappa y(t) L}{F}. \tag{59}$$

The per-unit shipping cost then becomes

$$p_{Dj} = S_j^{-\sigma_j} C_j\left(1, \frac{\kappa y(t) L}{F}\right) C_{Dj}(R_j^{rd}, R_j^{wa}). \tag{60}$$

The two cases describe the same intratemporal fuel market at different points along the

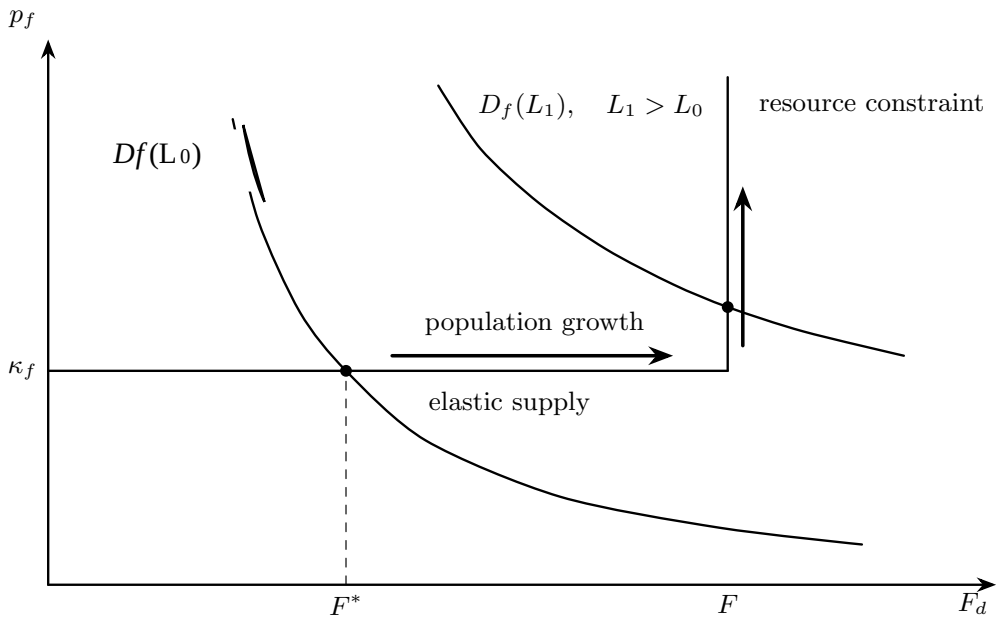

Figure 1: Equilibrium in the fuel market. With an unconstrained fuel supply, elastic supply pins down the fuel price. With a binding aggregate fuel constraint, population growth raises demand relative to the fixed supply and generates continuing growth in the fuel price.

development path. With a constant population, the economy can remain in either regime. With population growth, L/F rises at rate λ, and the economy eventually reaches the resource-constrained regime if fuel demand grows sufficiently rapidly.

The richer model therefore replaces the exogenous parameter η_j of the benchmark with an endogenous shipping cost generated by the cost of the common labor–fuel bundle and the cost of the mode mix. The distinction is important. A change in $C_j(w, p_f)$ affects both making and shipping, whereas a change in $C_{Dj}(R_j^{rd}, R_j^{wa})$ affects shipping alone. The next section shows why those two changes have different implications for the allocation of R&D.

5.4 The resulting cost structure

The richer transportation environment is now part of the growth mechanism rather than a separate block appended to it. The common labor–fuel cost enters manufacturing and shipping symmetrically, while the mode-mix cost enters shipping alone. The distinction survives aggregation and gives the ratio of the shipping cost to the manufacturing cost a particularly simple structure.

The key state variable remains the relative cost of shipping to making. The next section establishes the asymmetry principle directly: a cost-push force common to the two activities does not change the direction of innovation, whereas a shipping-specific force does. The endogenous price of fuel provides the leading neutral case; the mode-specific factors introduced here provide the non-neutral case that becomes the focus of the regulation analysis.

6 Asymmetry and Fuel

The richer transportation environment makes it possible to state the general principle in terms of differential cost exposure: the direction of innovation responds to shocks that change the cost of shipping relative to the cost of making, regardless of the source of the shock. Common and shipping-specific cost push provide two polar cases. Let p_{Dj} denote the per-unit cost of shipping. From the factorization in Section 5,

$$p_{Dj} = S_j^{-\sigma_j} C_j(w, p_f) C_{Dj}(R_j^{rd}, R_j^{wa}), \tag{61}$$

while the productivity-adjusted unit cost of making the good is proportional to $Z_j^{-\theta_j} C_j(w, p_f)$. Their ratio is therefore

$$q_j \equiv \frac{S_j^{-\sigma_j} C_j(w, p_f) C_{Dj}(R_j^{rd}, R_j^{wa})}{Z_j^{-\theta_j} C_j(w, p_f)} = S_j^{-\sigma_j} Z_j^{\theta_j} C_{Dj}(R_j^{rd}, R_j^{wa}). \tag{62}$$

Equation (62) contains the central distinction. The unit cost $C_j(w, p_f)$ of the common labor–fuel bundle appears in both activities and cancels from q_j. The mode-mix cost $C_{Dj}(R_j^{rd}, R_j^{wa})$ applies only to shipping and does not cancel. Thus, what matters for the direction of innovation is not whether shipping is costly, or even whether its cost is rising, but whether its cost is rising relative to the cost of making the good.

Log-differentiating equation (62) gives

$$\frac{\dot{q}_j}{q_j} = \theta_j z_j - \sigma_j s_j + \Psi_j, \qquad \Psi_j \equiv \frac{d \log C_{Dj}(R_j^{rd}, R_j^{wa})}{dt}, \tag{63}$$

where $z_j \equiv \dot{Z}_j / Z_j$ and $s_j \equiv \dot{S}_j / S_j$. When $\Psi_j = 0$, this is the benchmark law of motion. When $\Psi_j > 0$, the relative cost of shipping receives an additional upward push that innovation must offset.

Proposition 1 (Symmetric and asymmetric cost push)**.** *An increase in the price of an input used with the same intensity in making things and shipping things is neutral with respect to the composition of innovation: it changes the level of the two unit costs but cancels from q_j and therefore leaves the allocation of R&D between manufacturing and transportation unchanged. A friction specific to shipping is not neutral: it enters q_j directly, changes the relative return to manufacturing and transportation innovation, and reallocates R&D between the two activities.*

Proof. Under equal factor intensities, the common unit cost $C_j(w, p_f)$ cancels from equation (62), so a change in p_f leaves q_j unchanged. The return equations for manufacturing and transportation innovation depend on the relative cost through q_j, so their allocation is unchanged. A shipping-specific change in C_{Dj} enters q_j directly and therefore changes the two returns. □

This result is the more general form of the distinction in Kane and Peretto (2020). In their benchmark, the iceberg component is symmetric in knowledge because it falls with Z_j just as manufacturing cost does, whereas the per-unit component is asymmetric because manufacturing innovation does not directly compress it. Here the same logic applies to a richer input structure. Under equal factor intensities, fuel is symmetric in factor prices; a mode-specific transportation factor is asymmetric.

The equal-factor-intensity assumption is a tractability restriction, not a claim that transportation and manufacturing use inputs identically. Let their unit costs be $C_j^X(w, p_f)$ and $C_j^S(w, p_f)$, with labor shares ω_j^X and ω_j^S. Then

$$q_j = S_j^{-\sigma_j} Z_j^{\theta_j} \frac{C_j^S(w, p_f)}{C_j^X(w, p_f)} C_{Dj}(R_j^{rd}, R_j^{wa}). \tag{64}$$

With $\dot{p}_f/p_f = \lambda$ in the resource-constrained fuel regime,

$$\frac{d}{dt}\log\left(\frac{C_j^S(w,p_f)}{C_j^X(w,p_f)}\right) = (\omega_j^X - \omega_j^S)\lambda. \tag{65}$$

If shipping is more fuel-intensive, $\omega_j^S < \omega_j^X$, a rising fuel price pushes q_j upward. Fuel neutrality is therefore a consequence of symmetry in factor intensity, not a universal property of fuel.

The endogenous fuel market provides the useful benchmark case. Suppose the economy has an endowment F of a resource that can be processed into fuel, and producing one unit of fuel requires κ_f units of labor. Supply is perfectly elastic at $p_f = \kappa_f w$ until the resource constraint binds. With $w = 1$, aggregate fuel demand satisfies

$$p_f F^d = \kappa y L, \qquad \kappa \equiv \sum_{j=1}^{J}(1-\omega_j)\frac{\epsilon_j - 1}{\epsilon_j}\phi_j < 1. \tag{66}$$

When demand is low, the resource constraint does not bind and $p_f = \kappa_f$. Once it binds,

$$p_f = \frac{\kappa y(t) L}{F}. \tag{67}$$

Along the constrained transition, the fuel price grows at rate $\lambda + \dot{y}/y$ and converges to growth rate λ as $y(t)$ approaches y^c.

The resulting cost push raises the cost of making and shipping goods, but under equal factor intensities it raises both through the same term $C_j(w, p_f)$. It therefore changes delivered prices without changing q_j, the dynamics of (q_j, x_j), or the allocation of R&D between L_{Zj} and L_{Sj}. Scarcity can push costs upward without redirecting innovation.

The mode-specific factors R_j^{rd} and R_j^{wa} are different. An increase in either raises the cost of shipping without raising the cost of making the good. Through $C_{Dj}(R_j^{rd}, R_j^{wa})$ it enters q_j directly and changes the relative return to the two forms of cost reduction. We now let these shipping-specific frictions accumulate over time.

7 Cumulative Regulation

The factors R_j^{rd} and R_j^{wa} have so far been treated as levels. For regulation, however, the empirically relevant feature is accumulation: as documented in Appendix D, the aggregate

regulatory stock exhibits persistent net accumulation, while transportation industries differ substantially in the rate at which their measured exposure grows. We represent that tendency tractably by allowing the mode-specific productivity wedges to trend exponentially upward,

$$R_j^{rd}(t) = R_j^{rd}(0)e^{\psi_j^{rd}t}, \qquad R_j^{wa}(t) = R_j^{wa}(0)e^{\psi_j^{wa}t}, \qquad \psi_j^{rd}, \psi_j^{wa} \geq 0. \tag{68}$$

The levels determine the initial cost of the mode mix. The accumulation rates determine the continuing push on the relative cost of shipping and therefore the direction of innovation.

Using the mode cost shares defined in Section 5, equation (63) becomes

$$\frac{\dot{q}_j}{q_j} = \theta_j z_j - \sigma_j s_j + \Psi_j, \qquad \Psi_j = c_j^{rd}\psi_j^{rd} + c_j^{wa}\psi_j^{wa}, \qquad c_j^{rd} + c_j^{wa} = 1. \tag{69}$$

Regulation therefore enters the growth system through a single object, Ψ_j: the rate at which the mode mix pushes up the relative cost of shipping. If factor intensities differ between making and shipping, the additional term in equation (65) simply enters this forcing as well. We retain equal factor intensities below to keep the regulatory mechanism separate from that qualification.

7.1 Uniform accumulation

Suppose first that the stocks of regulations affecting the two modes accumulate at the same rate, $\psi_j^{rd} = \psi_j^{wa} = \psi_j$. Homogeneity of the transportation unit-cost function gives

$$C_{Dj}\left(R_j^{rd}(0)e^{\psi_j t}, R_j^{wa}(0)e^{\psi_j t}\right) = e^{\psi_j t} C_{Dj}\left(R_j^{rd}(0), R_j^{wa}(0)\right).$$

The cost shares are constant and $\Psi_j = \psi_j$. Hence the system in (q_j, x_j) remains autonomous.

Proposition 2 (Uniform accumulation)**.** *Under uniform regulatory accumulation, the $\dot{x}_j = 0$ locus is unchanged and the $\dot{q}_j = 0$ locus is*

$$x_j = \frac{[(\theta_j - \sigma_j)(\rho + \delta_j) - \psi_j](1 + q_j)}{\frac{\epsilon_j - 1}{\epsilon_j}(\alpha_j \theta_j^2 - \varsigma_j \sigma_j^2 q_j)}. \tag{70}$$

Thus the benchmark expressions carry over with

$$\Omega_j \equiv (\theta_j - \sigma_j)(\rho + \delta_j) - \Psi_j \tag{71}$$

in place of $(\theta_j - \sigma_j)(\rho + \delta_j)$.

Proof. Substitute $\Psi_j = \psi_j$ into equation (63) and use the innovation-return equations of Section 4. The $\dot{x}_j$ equation contains no regulatory term, while setting $\dot{q}_j = 0$ yields equation (70). □

The steady state therefore satisfies

$$x_j^* = \frac{\Omega_j(1+q_j^*)}{\frac{\epsilon_j-1}{\epsilon_j}(\alpha_j\theta_j^2 - \varsigma_j\sigma_j^2 q_j^*)}, \tag{72}$$

together with the corrected $\dot{x}_j = 0$ condition in equation (28). Rather than solve this pair for a cumbersome expression in q_j^*, the next subsection uses the primitive innovation and entry conditions directly. That route yields a simpler characterization of how the stationary innovation rates depend on the forcing.

7.2 The threshold

The threshold can be expressed directly in terms of the model parameters. Define

$$a_j \equiv \frac{\epsilon_j - 1}{\epsilon_j}, \qquad d_j \equiv \rho + \delta_j, \qquad h_j \equiv \frac{1}{\epsilon_j} - \beta d_j, \tag{73}$$

and

$$\Gamma_{Zj} \equiv \frac{h_j/a_j - \theta_j}{\alpha_j\theta_j}, \qquad \Gamma_{Sj} \equiv \frac{h_j/a_j - \sigma_j}{\varsigma_j\sigma_j}, \qquad \Gamma_{0j} \equiv \varphi_j - \frac{h_j d_j}{a_j}\left(\frac{1}{\alpha_j\theta_j} + \frac{1}{\varsigma_j\sigma_j}\right). \tag{74}$$

The two innovation conditions imply

$$x_j = \frac{1}{a_j}\left[\frac{z_j + d_j}{\alpha_j\theta_j} + \frac{s_j + d_j}{\varsigma_j\sigma_j}\right].$$

Using this expression in the entry condition yields the interior identity

$$\Gamma_{Zj} z_j + \Gamma_{Sj} s_j = \Gamma_{0j} - \beta\dot{x}_j. \tag{75}$$

Under constant uniform forcing $\Psi_j = \psi_j$, an interior stationary state also satisfies

$$\sigma_j s_j^* = \theta_j z_j^* + \psi_j. \tag{76}$$

Hence, with $\Gamma_j \equiv \Gamma_{Zj} + (\theta_j/\sigma_j)\Gamma_{Sj} \neq 0$,

$$z_j^*(\psi_j) = \frac{\Gamma_{0j} - (\Gamma_{Sj}/\sigma_j)\psi_j}{\Gamma_j}, \qquad s_j^*(\psi_j) = \frac{\theta_j z_j^*(\psi_j) + \psi_j}{\sigma_j}. \tag{77}$$

The interior continuation is therefore affine in the regulatory forcing.

Proposition 3 (Regulatory accumulation and the manufacturing threshold)**.** *Suppose* $\Gamma_{0j} > 0$, $\Gamma_{Sj} > 0$, *and* $\Gamma_j > 0$, *and suppose the stationary allocation satisfying the interior first-order conditions is otherwise admissible. Then manufacturing innovation declines linearly with the uniform accumulation rate and reaches zero at the positive finite threshold*

$$\bar{\psi}_j = \frac{\sigma_j \Gamma_{0j}}{\Gamma_{Sj}}. \tag{78}$$

For $0 \leq \psi_j < \bar{\psi}_j$, *the interior continuation has* $z_j^* > 0$. *At the threshold,*

$$z_j^* = 0, \qquad s_j^* = \frac{\bar{\psi}_j}{\sigma_j}, \qquad \frac{x_j^*}{1 + q_j^*} = \frac{\epsilon_j(\rho + \delta_j)}{\alpha_j \theta_j (\epsilon_j - 1)}. \tag{79}$$

For $\psi_j > \bar{\psi}_j$, *the negative continuation of equation* (77) *is not an equilibrium innovation rate; the relevant candidate regime has manufacturing R&D inactive and must satisfy its return inequality.*

Proof. Equation (77) gives $dz_j^*/d\psi_j = -\Gamma_{Sj}/(\sigma_j \Gamma_j) < 0$. Under the stated restrictions, $z_j^*(0) = \Gamma_{0j}/\Gamma_j > 0$ and its unique zero is $\bar{\psi}_j = \sigma_j \Gamma_{0j}/\Gamma_{Sj}$. At that point, stationarity of q_j gives $s_j^* = \bar{\psi}_j/\sigma_j$, while the manufacturing return condition at equality gives the final expression in equation (79). Beyond the threshold, the interior stationary allocation would require $z_j^* < 0$, so the nonnegativity constraint on manufacturing R&D binds. □

The candidate manufacturing-inactive stationary regime can also be stated explicitly. Set $z_j = 0$ and $L_{Zj} = 0$. If transportation R&D remains active, stationarity of q_j requires $s_j = \psi_j/\sigma_j$. The stationary entry condition then gives

$$x_j^0(\psi_j) = \frac{\varphi_j + \psi_j/(\varsigma_j \sigma_j)}{h_j}, \tag{80}$$

provided $h_j > 0$. The transportation return condition determines q_j^0 from

$$\frac{x_j^0 q_j^0}{1 + q_j^0} = \frac{d_j + \psi_j/\sigma_j}{\varsigma_j \sigma_j a_j}. \tag{81}$$

This candidate is an equilibrium whenever $q_j^0 > 0$ and the manufacturing return does not exceed the household return,

$$\alpha_j \theta_j a_j \frac{x_j^0}{1 + q_j^0} \leq d_j. \tag{82}$$

Equality holds at the threshold. Thus, the threshold is the point at which the interior stationary allocation reaches the manufacturing-R&D boundary and meets the manufacturing-inactive regime.

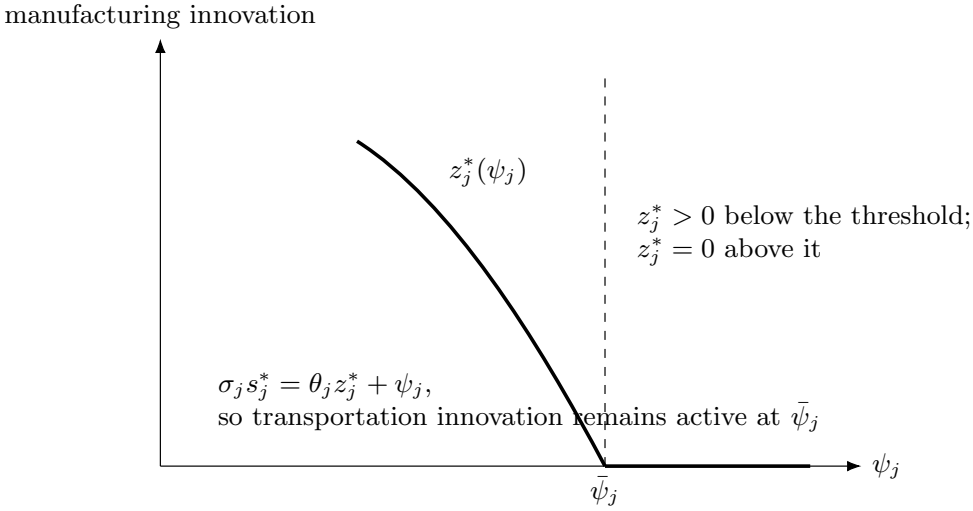

Figure 2: Manufacturing innovation under uniform regulatory accumulation. Under the sufficient conditions in Proposition 3, the interior continuation $z_j^*(\psi_j)$ is affine and decreasing and reaches the manufacturing-R&D boundary at $\bar{\psi}_j$. Beyond the threshold the interior continuation is inadmissible and the manufacturing-inactive regime applies when its return inequality is satisfied.

The economics are transparent. In a stationary relative-cost configuration, transportation innovation must offset both the effect of manufacturing innovation on relative shipping cost and the exogenous regulatory forcing. The linear identity (75) then determines how the two innovation rates can jointly satisfy free entry. Under the sufficient conditions above, stronger forcing shifts that allocation toward the transportation response until the manufacturing margin reaches its boundary. The composite coefficients in equation (74) show directly how entry, innovation productivity, depreciation, fixed operating cost, and demand conditions determine the amount of persistent asymmetric cost push the industry can absorb.

7.3 Differential accumulation and the long-run transport constraint

Uniform accumulation isolates the threshold, but mode-specific regulation need not accumulate uniformly. Let $\psi_j^{rd} \neq \psi_j^{wa}$ and define the relative stock of regulations

$$\mu_j \equiv \frac{R_j^{rd}}{R_j^{wa}}, \qquad \frac{\dot{\mu}_j}{\mu_j} = \psi_j^{rd} - \psi_j^{wa}. \tag{83}$$

By homogeneity,

$$C_{Dj}(R_j^{rd}, R_j^{wa}) = R_j^{wa} c_j(\mu_j), \qquad c_j(\mu) \equiv C_{Dj}(\mu, 1),$$

so we can write the road cost share as

$$c_j^{rd}(\mu_j) = \frac{\mu_j c_j'(\mu_j)}{c_j(\mu_j)} \in (0, 1). \tag{84}$$

The system becomes

$$\frac{\dot{q}_j}{q_j} = \theta_j z_j - \sigma_j s_j + \psi_j^{wa} + c_j^{rd}(\mu_j)(\psi_j^{rd} - \psi_j^{wa}), \qquad \frac{\dot{\mu}_j}{\mu_j} = \psi_j^{rd} - \psi_j^{wa}, \tag{85}$$

together with the benchmark equation for $\dot{x}_j/x_j$. The additional state variable evolves exogenously and monotonically. What matters for the long run is how the mode share responds to the growing relative stock of regulations.

Proposition 4 (Escape and bottlenecks)**.** *Suppose $\psi_j^{rd} > \psi_j^{wa}$ and the elasticity of substitution between modes is not one. If the modes are substitutes, $\zeta_j > 1$, the road cost share converges to zero and*

$$\Psi_j \longrightarrow \psi_j^{wa} = \min\{\psi_j^{rd}, \psi_j^{wa}\}.$$

If the modes are complements, $\zeta_j < 1$, the road cost share converges to one and

$$\Psi_j \longrightarrow \psi_j^{rd} = \max\{\psi_j^{rd}, \psi_j^{wa}\}.$$

In either case the system is asymptotically governed by the corresponding uniform-accumulation system.

Proof. The cost share of a mode falls with its relative price when modes are substitutes and rises with its relative price when they are complements. Since μ_j rises monotonically when $\psi_j^{rd} > \psi_j^{wa}$, the stated limits follow from equation (69). □

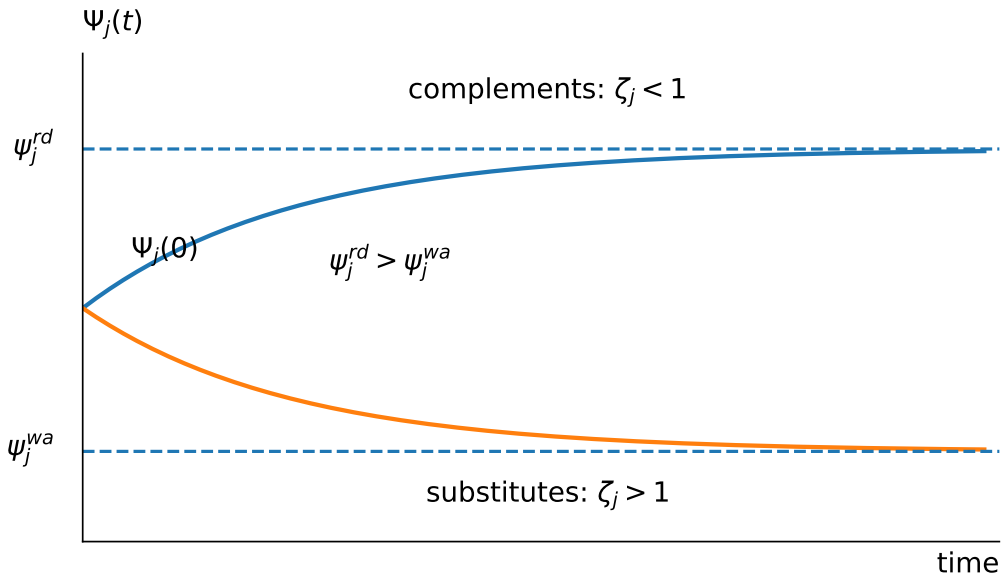

Figure 3: Escape and the network bottleneck when $\psi_j^{rd} > \psi_j^{wa}$. With substitute modes, the cost share of the mode with the more rapid accumulation of regulations vanishes and the regulatory forcing converges to ψ_j^{wa}. With complementary modes, that mode becomes the bottleneck and the forcing converges to ψ_j^{rd}.

This result identifies the long-run transport constraint. With substitutes, an industry can escape the mode with the more rapid accumulation of regulations, so the slower accumulation of regulations governs the long run. With complements, the modes are links in a transportation chain and the more rapid accumulation of regulations becomes the bottleneck. A port and the roads serving it are the natural example: improving access to one link cannot eliminate a worsening constraint on the other.

Combining this result with the threshold gives a direct implication. For substitute modes, continued manufacturing innovation requires the limiting minimum accumulation rate to remain below $\bar{\psi}_j$. For complementary modes, it requires the limiting maximum to remain below the threshold. The same pair of mode-specific regulatory trends can therefore be absorbed by one transportation technology and shut down manufacturing innovation in another.

The long-run result does not describe the path by which the industry escapes the mode with the more rapid accumulation of regulations or becomes increasingly exposed to a bottleneck. That transition matters because knowledge forgone along the path need not be subsequently recovered. Sections 8 and 9 return to those dynamics by studying, respectively, the activation of the innovation margins and the consequences of adjustment along the transition path.

8 Activation along the Development Path

The previous sections studied how transportation changes the composition of innovation once firms are large enough to undertake cost-reducing R&D. We now turn the argument around and ask when those innovation margins become active in the first place. This makes the development path a natural continuation of the asymmetry results. As firm size rises, the economy crosses the activation boundaries derived in Section 4; because those boundaries move in opposite directions with the relative cost of shipping, transportation conditions determine not only how much firms innovate in the factory and the fleet, but also the order in which the two activities begin.

The mechanism is the "stairway to growth" developed in Peretto (2027): firm-size growth activates innovation margins sequentially, and each activation changes the growth process. Here the relevant stairway has three margins—entry, transportation innovation, and manufacturing innovation. The first is active from the outset under free entry. The other two need not be. Abstracting from quality innovation therefore creates a technologically dormant initial regime in which firms expand in number but do not yet reduce either manufacturing or shipping costs. The question is which cost-reducing margin becomes profitable first.

8.1 Three phases

Consider an industry with initial condition $(q_j(0), x_j(0))$ and sufficiently small initial firm size. The first cost-reducing margin to activate is determined by the relative position of the two boundaries below. The transportation-first sequence emphasized in this section has three phases:

1. *Phase I:* variety expansion only, $L_{Nj} > 0$, with $L_{Zj} = L_{Sj} = 0$;
2. *Phase II:* variety expansion and transportation innovation, $L_{Nj} > 0$ and $L_{Sj} > 0$, with $L_{Zj} = 0$;
3. *Phase III:* all three innovation margins are active.

This ordering is not automatic. It applies when the transportation boundary lies below the manufacturing boundary at the relevant crossing. The activation boundaries are

$$x_{Zj}(q_j) = \frac{\epsilon_j(\rho+\delta_j)}{\alpha_j\theta_j(\epsilon_j-1)}(1+q_j), \tag{86}$$

$$x_{Sj}(q_j) = \frac{\epsilon_j(\rho+\delta_j)}{\varsigma_j\sigma_j(\epsilon_j-1)}\frac{1+q_j}{q_j}. \tag{87}$$

The manufacturing boundary rises with q_j while the transportation boundary falls. In particular,

$$x_{Sj}(q_j) < x_{Zj}(q_j) \quad \Longleftrightarrow \quad q_j > \frac{\alpha_j\theta_j}{\varsigma_j\sigma_j}. \tag{88}$$

Thus transportation innovation activates first only when equation (88) holds at the first crossing; below that relative-cost cutoff manufacturing innovation activates first, and equality gives simultaneous activation. Shipping that is expensive relative to making therefore raises the firm size required to make factory innovation profitable but lowers the firm size required to make transportation innovation profitable. The same asymmetry that reallocates R&D in Sections 6–7 thus orders its activation along the development path. The mode-specific factors need not represent institutions. When they are constant, they can instead summarize non-institutional fundamentals—geography, infrastructure, or other persistent determinants of transportation cost. Qualitatively, such factors operate through the same relative-cost state and therefore shift the activation boundaries in the same way as a fixed regulatory wedge; accumulation matters because it makes that state move over time.

8.2 Leaving the dormant regime

In Phase I, $L_{Zj} = L_{Sj} = 0$. With regulatory factors held fixed, $q_j = q_j(0)$ and the entry condition reduces to

$$\dot{x}_j = \frac{\varphi_j}{\beta} - \left(\frac{1}{\beta\epsilon_j} - \rho - \delta_j\right)x_j. \tag{89}$$

Let

$$\bar{\nu}_j^I \equiv \frac{1}{\beta\epsilon_j} - \rho - \delta_j, \qquad \bar{x}_j^I \equiv \frac{\varphi_j}{1/\epsilon_j - \beta(\rho+\delta_j)}. \tag{90}$$

For $\bar{\nu}_j^I > 0$ and $x_j(0) < \bar{x}_j^I$, firm size rises monotonically toward $\bar{x}_j^I$ according to

$$x_j(t) = x_j(0)e^{-\bar{\nu}_j^I t} + \left(1 - e^{-\bar{\nu}_j^I t}\right)\bar{x}_j^I. \tag{91}$$

The first transition occurs when this path reaches $x_{Sj}(q_j(0))$. Hence, absent regulatory accumulation,

$$t_j^{II} = \frac{1}{\bar{\nu}_j^I} \log\left[\frac{\bar{x}_j^I - x_j(0)}{\bar{x}_j^I - x_{Sj}(q_j(0))}\right]. \tag{92}$$

Because $x_{Sj}(q_j)$ is decreasing in q_j, a higher initial relative shipping cost brings transportation innovation forward. This is induced innovation in its most direct form: the activity that is relatively expensive is the activity firms find profitable to improve first.

Once transportation innovation activates, the industry enters Phase II. Manufacturing innovation remains inactive, but S_j begins to rise and q_j begins to fall. The transportation improvement therefore does more than reduce shipping cost directly. By lowering q_j, it moves the manufacturing activation boundary downward and raises the return to improving the factory. Transportation innovation is thus a catalyst for the second transition. Phase III begins when the trajectory reaches $x_{Zj}(q_j)$, after which the interior dynamics analyzed in the benchmark apply.

8.3 Regulatory accumulation and the ordering of innovation

Section 7 makes this development sequence more revealing. During Phase I there is no cost-reducing innovation, so equation (63) becomes

$$\frac{\dot{q}_j}{q_j} = \Psi_j. \tag{93}$$

Under uniform accumulation, $q_j(t) = q_j(0)e^{\psi_j t}$. Regulation therefore moves the state variable while firm-size growth moves the industry vertically through the activation diagram. Since $x_{Sj}(q_j)$ falls with q_j and $x_{Zj}(q_j)$ rises with it, regulatory accumulation pulls the first transition forward and pushes the second one back. It stretches Phase II from both ends: the fleet arrives earlier and the factory later.

Equation (94) now determines the first transition implicitly:

$$x_j(t) = \frac{\epsilon_j(\rho + \delta_j)}{\varsigma_j \sigma_j(\epsilon_j - 1)} \frac{1 + q_j(0)e^{\psi_j t}}{q_j(0)e^{\psi_j t}}. \tag{94}$$

The left-hand side rises toward $\bar{x}_j^I$, while the right-hand side falls toward the floor

$$\underline{x}_{Sj} \equiv \lim_{q_j \to \infty} x_{Sj}(q_j) = \frac{\epsilon_j(\rho + \delta_j)}{\varsigma_j \sigma_j(\epsilon_j - 1)}. \tag{95}$$

This yields a useful distinction between the threshold for beginning transportation innovation and the threshold derived in Section 7 for sustaining manufacturing innovation.

Proposition 5 (Activation and regulatory accumulation)**.** *Maintain the transportation-first ordering in equation* (88) *along the first crossing. Suppose* $\bar{\nu}_j^I > 0$ *and* $x_j(0) < \bar{x}_j^I$*. Then:*

1. *with no regulatory accumulation, the industry leaves Phase I if and only if*

$$\bar{x}_j^I > x_{Sj}(q_j(0));$$

2. *with any constant* $\psi_j > 0$*, the industry leaves Phase I if and only if*

$$\bar{x}_j^I > \underline{x}_{Sj};$$

Proof. In Phase I, $x_j(t)$ rises monotonically toward $\bar{x}_j^I$. With $\psi_j = 0$, the transportation boundary is constant at $x_{Sj}(q_j(0))$, giving part 1. With $\psi_j > 0$, q_j rises monotonically and $x_{Sj}(q_j)$ falls monotonically toward $\underline{x}_{Sj}$, giving part 2. This proves the two statements. □

The proposition exposes a non-monotonic role for asymmetric cost push along the development path. A rising shipping-specific regulatory forcing can make transportation innovation profitable in an industry that would otherwise never undertake it. What happens after that first activation is a transition question rather than a consequence of the steady-state threshold alone. For parameter configurations satisfying Proposition 3, there are Phase-II paths on which a sufficiently moderate forcing allows the trajectory eventually to cross the manufacturing boundary and enter Phase III, while a sufficiently strong persistent forcing leaves manufacturing innovation inactive in the long run. Thus the friction can help trigger the technology that works against it without necessarily allowing the innovation sequence to reach the factory. This is the asymmetry mechanism viewed dynamically: shipping becomes worth improving because it is relatively costly, while making becomes worth improving only if subsequent transportation improvement restores a sufficiently high return to factory R&D.

Differential accumulation adds the mode-choice result of Section 7. With substitute modes, escape toward the mode with the slower accumulation of regulations causes Ψ_j to decline over time, so the initial push toward transportation innovation weakens as the industry adjusts its mode mix. With complementary modes, the bottleneck works in the opposite direction. The development path therefore carries the long-run escape-versus-bottleneck distinction into the timing of innovation itself. We return in the next section to the transitional consequences

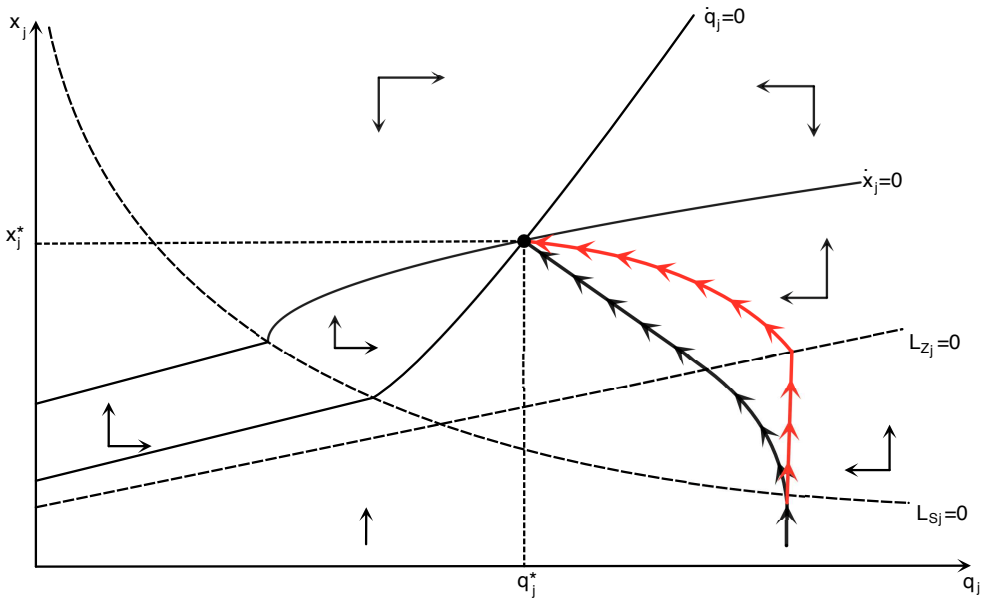

Figure 4: Comparative activation dynamics for two industries that differ in their initial relative shipping cost. A higher initial q_j brings the transportation-innovation boundary closer and the manufacturing-innovation boundary farther away, changing the timing of the two innovation margins even when the industries converge to the same interior steady state.

of these movements, including the possibility that a temporary divergence in regulatory accumulation leaves a permanent level effect on manufacturing knowledge.

9 Transitional Dynamics and Hysteresis

Sections 7 and 8 establish two properties of cumulative regulation. Its long-run effect depends on whether an industry can escape toward the mode with the slower accumulation of regulations, and its effect along the development path depends on how the resulting cost push moves the two activation boundaries. These properties meet during the transition. Mode choice is instantaneous, but the relative regulatory stocks that determine the cost-minimizing mode mix evolve over time. During that evolution, the industry can be exposed to a cost push that differs substantially from the one governing its long-run growth path. Whether the resulting innovation shortfall persists depends on the equilibrium recovery path, especially on whether the trajectory crosses an innovation boundary.

This is the final step in the paper's argument. A shipping-specific cost push changes the direction of innovation; mode substitution determines whether that push persists; and the evolution of relative mode costs determines whether the transition crosses an innovation

boundary and thereby changes the possibility of subsequent knowledge recovery.

9.1 The path of the cost push

Consider differential accumulation with $\psi_j^{rd} > \psi_j^{wa}$. From equation (85),

$$\Psi_j(t) = \psi_j^{wa} + c_j^{rd}(\mu_j(t))(\psi_j^{rd} - \psi_j^{wa}), \qquad \mu_j(t) = \mu_j(0)e^{(\psi_j^{rd} - \psi_j^{wa})t}. \tag{96}$$

For substitute modes, $\zeta_j > 1$, the road cost share falls as the relative stock of road regulations rises, so $\Psi_j(t)$ declines monotonically toward ψ_j^{wa}. The industry eventually escapes toward the mode with the slower accumulation of regulations. For complementary modes, $\zeta_j < 1$, the road share rises and $\Psi_j(t)$ increases toward ψ_j^{rd}. The bottleneck becomes progressively more important.

The transition can therefore be understood as a moving version of the uniform-accumulation system. At each date the current value of $\Psi_j(t)$ determines the position of the $\dot{q}_j = 0$ locus, while the $\dot{x}_j = 0$ locus is unchanged. Differential accumulation makes the former locus sweep monotonically between the positions associated with the initial and limiting cost pushes.

Let $\underline{\Psi}_j$ and $\overline{\Psi}_j$ denote the smaller and larger of the initial and limiting values of $\Psi_j(t)$. The monotonicity of the forcing gives a useful bound on the instantaneous dynamics during the transition.

Proposition 6 (Bracketing the regulatory forcing)**.** *At every date under differential regulatory accumulation,*

$$\underline{\Psi}_j \leq \Psi_j(t) \leq \overline{\Psi}_j. \tag{97}$$

For any given state (q_j, x_j), the resulting value of $\dot{q}_j/q_j$ is therefore bracketed by the values generated by the two constant-forcing systems $\Psi_j = \underline{\Psi}_j$ and $\Psi_j = \overline{\Psi}_j$. Equivalently, the instantaneous $\dot{q}_j = 0$ locus lies between the two corresponding constant-forcing loci.

Proof. Equation (96) writes $\Psi_j(t)$ as a cost-share-weighted average of ψ_j^{rd} and ψ_j^{wa}, with weights in $[0, 1]$, which proves equation (97). Since Ψ_j enters equation (63) additively, the pointwise ordering of $\dot{q}_j/q_j$ follows immediately. Setting $\dot{q}_j = 0$ gives the corresponding ordering of the nullclines. □

The proposition deliberately brackets the forcing and the instantaneous vector field rather than entire trajectories. A global ordering of two-dimensional paths would require additional restrictions on the cross-effects in the (q_j, x_j) system. The pointwise result is all that we need

below: during mode substitution, the innovation incentives move continuously between the incentives associated with the two limiting regulatory environments.

9.2 Temporary divergence, permanent knowledge loss

The substitute case reveals why the transition matters. Suppose

$$\psi_j^{wa} < \bar{\psi}_j < \psi_j^{rd}. \tag{98}$$

The long-run cost push is then below the threshold derived in Section 7: after the industry has escaped toward water transportation, the limiting environment supports both manufacturing and transportation innovation. But early in the transition the road share can be large enough that $\Psi_j(t) > \bar{\psi}_j$. During that interval, the relative cost of shipping rises fast enough to push innovation away from the factory.

The relevant comparison is therefore not between two steady states. It is between two paths with the same long-run cost push. One path experiences the temporary divergence between ψ_j^{rd} and ψ_j^{wa} while the mode mix adjusts; the counterfactual path faces $\Psi_j = \psi_j^{wa}$ from the outset. The two paths eventually confront the same asymptotic institutional trend, but they need not arrive there with the same stock of manufacturing knowledge.

Proposition 7 (Transitional hysteresis)**.** *Let the transportation modes be substitutes, $\zeta_j > 1$, and suppose equation* (98) *holds. Suppose also that, at some point on the manufacturing-innovation boundary, the low-forcing vector field points into the region with active manufacturing R&D while a sufficiently high value of the transitory forcing points into the region with inactive manufacturing R&D. Then there is a nonempty set of nearby initial conditions for which differential regulatory accumulation temporarily switches manufacturing innovation off even though the limiting forcing supports the interior growth regime. Relative to the counterfactual path with $\Psi_j = \psi_j^{wa}$ at all dates, these paths accumulate less manufacturing knowledge during the interval in which manufacturing R&D is inactive. If they subsequently re-enter the interior regime and the cumulative shortfall in manufacturing-knowledge growth is not exactly offset during re-entry, the limiting growth rate is restored but a level gap in manufacturing knowledge remains.*

Proof. Define the signed distance from the manufacturing-innovation boundary by $H_j \equiv x_j - x_{Zj}(q_j)$, so manufacturing R&D is active for $H_j > 0$. Along the boundary,

$$\dot{H}_j = \dot{x}_j - x'_{Zj}(q_j)\dot{q}_j.$$

Because $x'_{Zj}(q_j) > 0$ and the forcing enters $\dot{q}_j/q_j$ additively, raising Ψ_j lowers $\dot{H}_j$ at a given boundary point. By the stated crossing condition, the low-forcing and high-forcing vector fields therefore point to opposite sides of the boundary. Continuity of the vector field gives a nonempty set of nearby initial conditions for which the differential-accumulation path enters $H_j < 0$ while the low-forcing counterfactual remains in $H_j > 0$. Since $\dot{Z}_j = 0$ whenever manufacturing R&D is inactive, the affected path accumulates strictly less Z_j over that interval. For substitute modes, $\Psi_j(t) \to \psi_j^{wa} < \bar{\psi}_j$. If the affected path subsequently re-enters the interior regime, both paths eventually face the same limiting growth system and hence the same limiting growth rate. The initial knowledge shortfall is strictly positive. Exact subsequent offset of that cumulative shortfall is an additional restriction; absent such offset, the two paths converge to the same growth rate with different levels of manufacturing knowledge. □

The result identifies a route to hysteresis in levels rather than in the limiting growth rate, but it also makes clear what remains to be established. Innovation forgone during the shutdown need not be permanently lost if subsequent incentives generate sufficient catch-up. Indeed, when both innovation margins remain active throughout and an exogenous transportation wedge fully returns to its baseline level, the integrated interior conditions imply full recovery of the knowledge levels under the convergence conditions stated below. A persistent level gap therefore requires the boundary crossing, a persistent difference in the exogenous cost index, or another mechanism that prevents exact catch-up. A steady-state comparison alone cannot distinguish these cases.

To see the distinction, consider two paths that remain in the interior regime throughout, have the same primitives and initial normalized states, and converge to the same terminal values of x_j and q_j. Let $I_z \equiv \int_0^\infty (z_j^S - z_j^B)dt$ and $I_s \equiv \int_0^\infty (s_j^S - s_j^B)dt$. Integrating equation (75) across the two paths gives $\Gamma_{Zj} I_z + \Gamma_{Sj} I_s = 0$. If $\Delta \log R_{Dj}$ denotes the terminal-minus-initial difference between their exogenous log transportation-cost indices, integration of equation (63) gives

$$\theta_j I_z - \sigma_j I_s + \Delta \log R_{Dj} = 0.$$

Provided $\sigma_j \Gamma_{Zj} + \theta_j \Gamma_{Sj} \neq 0$,

$$I_z = -\frac{\Gamma_{Sj}}{\sigma_j \Gamma_{Zj} + \theta_j \Gamma_{Sj}} \Delta \log R_{Dj}, \qquad I_s = \frac{\Gamma_{Zj}}{\sigma_j \Gamma_{Zj} + \theta_j \Gamma_{Sj}} \Delta \log R_{Dj}. \tag{99}$$

A fully reversed wedge has $\Delta \log R_{Dj} = 0$ and therefore $I_z = I_s = 0$: the transition can affect discounted welfare without leaving a terminal knowledge gap. Proposition 7 concerns a different experiment because its affected path crosses the manufacturing-R&D boundary. Equation (99) therefore sharpens the interpretation of that proposition: the boundary crossing is potentially substantive precisely because it breaks the everywhere-interior recovery argument.

This also clarifies the role of the development path in Section 8. There, regulatory accumulation can bring transportation innovation forward and manufacturing innovation back. Here the same movement matters after the innovation margins have become relevant: if the manufacturing boundary is crossed in the wrong direction during the period of adjustment, the economy accumulates less productive knowledge during the shutdown even though the force that caused the crossing subsequently disappears. Whether that shortfall survives depends on the catch-up dynamics characterized above. Timing is therefore not a secondary feature of the model. It is part of the mechanism through which asymmetric cost push affects growth.

The analysis in the main text concentrates on regulatory accumulation because it exposes the mechanism that is specific to this paper. Appendix C reports additional steady-state comparative statics and a transition experiment, and shows how the resulting paths map into welfare.

10 Conclusion

We developed an endogenous-growth model in which firms make goods, ship them to their users, and innovate to reduce the cost of both activities. Taking transportation seriously changes the growth problem because making and shipping need not face the same costs. Firms therefore choose not only how much to innovate but where to direct cost-reducing innovation. The model organizes that choice around the cost of shipping relative to making. Manufacturing innovation raises this relative cost by making production cheaper; transportation innovation lowers it by making shipping cheaper. Forces that move the two activities together leave this margin unchanged, whereas forces specific to transportation change the relative return to

innovation in the factory and the fleet.

Fuel scarcity and cumulative regulation illustrate the distinction. When making and shipping use the same labor–fuel bundle with the same factor intensity, an increase in the endogenous price of fuel raises both unit costs proportionally. Scarcity raises delivered costs but does not redirect innovation. Mode-specific regulation is different. Regulatory accumulation raises the cost of shipping relative to making and therefore acts as a forcing term in the relative-cost dynamics. Transportation innovation must work against that forcing. Under the sufficient conditions characterized in Section 7, stronger regulatory forcing reallocates innovation toward the transportation response and away from manufacturing; at a finite threshold, manufacturing innovation reaches zero while transportation innovation remains active.

Multiple transportation modes determine whether that forcing persists. With substitute modes, firms reallocate transportation toward the mode with the slower accumulation of regulations, so the mode with the slower accumulation of regulations governs the long run. With complementary modes, the link with the more rapid accumulation of regulations can become the bottleneck. The transition matters because relative regulatory stocks evolve even though mode choice is intratemporal. Regulatory accumulation can bring transportation innovation forward and delay manufacturing innovation along the development path. If the path crosses the manufacturing-innovation boundary, a temporary period of strong forcing creates a shortfall in manufacturing knowledge. Whether that shortfall persists after the forcing recedes depends on subsequent catch-up; when adjustment remains everywhere interior, a fully reversed wedge instead produces full eventual recovery.

These results answer the question posed at the outset. Transportation is not a cost appended to a growth process generated elsewhere. Once firms must improve the technologies for both making and shipping goods, transportation becomes part of the mechanism that determines the direction and timing of endogenous innovation. The distinction between common and activity-specific cost push is then central: what matters for growth is not simply whether costs rise, but which costs rise relative to which others and how firms can innovate or substitute in response. The framework makes that relative-cost mechanism explicit. In particular, differential exposure reallocates innovation and, when sufficiently strong, reaches a finite threshold at which manufacturing R&D shuts down. Transportation, scarcity, and regulation provide the economic structure that reveals this mechanism; the same logic can be used to study other forms of activity-specific cost push.

A The Kane–Peretto Benchmark and the Present Model

For readers interested in the precise relation between the present model and Kane and Peretto (2020), this appendix reports the main points of continuity and departure. The comparison is not needed to follow the analysis in the paper. Its purpose is to make explicit which elements of the growth mechanism we take from the earlier model and which elements generate the results developed here.

	Kane and Peretto (2020)	**This paper**
Transportation	Making and shipping are distinct activities; shipping contains an iceberg and a per-unit component.	Retains this structure and gives transportation multiple modes, factor inputs, and mode-specific costs.
Innovation	Manufacturing and transportation innovation are complementary and can sustain a common decline in unit costs.	Takes this mechanism as the benchmark and studies how asymmetric cost push changes the direction and activation of the two innovation margins.
Cost push	The per-unit shipping cost creates the relevant asymmetry between making and shipping.	Fuel scarcity can generate symmetric cost push; cumulative mode-specific regulation generates asymmetric forcing.
Implications	Establishes the transportation constraint and the endogenous-innovation response to shipping costs.	Derives regulation thresholds, escape versus bottlenecks across modes, changes in the ordering of innovation, and transitional hysteresis.

B Benchmark Dynamics of the Present Model

Appendix A identifies the elements of the mechanism that the paper takes from Kane and Peretto (2020). This appendix has a different purpose. It reports how that mechanism operates in the particular model used here. The economic logic is inherited, but the phase diagram, activation boundaries, and equilibrium paths below are those generated by the present model after quality innovation has been suppressed. That restriction both isolates the interaction between manufacturing and transportation innovation and reduces the benchmark to a tractable two-state system in (q_j, x_j).

B.1 The benchmark without transportation innovation

When transportation knowledge is fixed, $q_j = \eta_j Z_j^{\theta_j}$ and therefore

$$\frac{\dot{q}_j}{q_j} = \theta_j z_j. \tag{B.1}$$

Manufacturing innovation is active above the boundary $x_{Zj}(q_j)$ in equation (20). Because this boundary rises with q_j, successful manufacturing innovation progressively raises the firm size required to sustain further cost reduction. With a positive exogenous per-unit shipping cost, the economy therefore cannot converge to a steady state with permanently positive manufacturing innovation. The phase diagram makes the transportation constraint transparent: manufacturing innovation raises the relative importance of the shipping cost until its own return is exhausted.

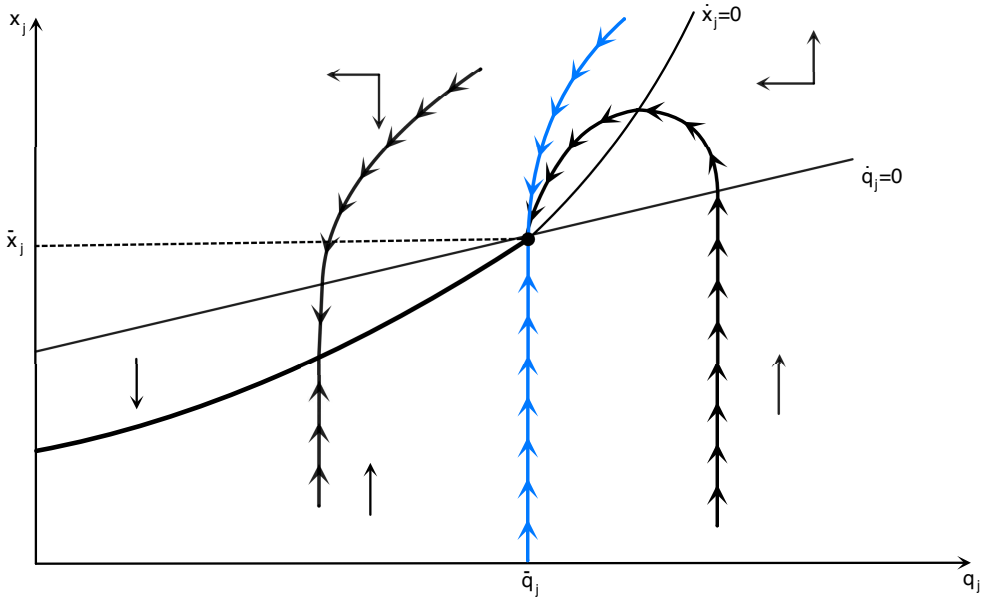

Figure 5: Benchmark dynamics without transportation innovation. The phase diagram shows how manufacturing innovation raises the relative importance of the exogenous shipping cost until the return to further manufacturing cost reduction is exhausted.

B.2 Joint manufacturing and transportation innovation

With both innovation margins active, equations (27) and (28) form an autonomous system in (q_j, x_j). The activation boundaries $x_{Zj}(q_j)$ and $x_{Sj}(q_j)$ move in opposite directions: the manufacturing boundary rises with q_j, while the transportation boundary falls. The state

space is therefore partitioned into regions in which neither, one, or both forms of cost-reducing innovation are active.

In the interior region, the $\dot{q}_j = 0$ locus identifies combinations of relative shipping cost and firm size for which manufacturing and transportation knowledge reduce their respective unit costs at the same rate. The $\dot{x}_j = 0$ locus identifies combinations consistent with stationary firm size under free entry. Their intersection determines the interior steady state (q_j^*, x_j^*). At that point,

$$\theta_j z_j^* = \sigma_j s_j^*, \tag{B.2}$$

so manufacturing and shipping costs decline at a common rate and their ratio remains constant.

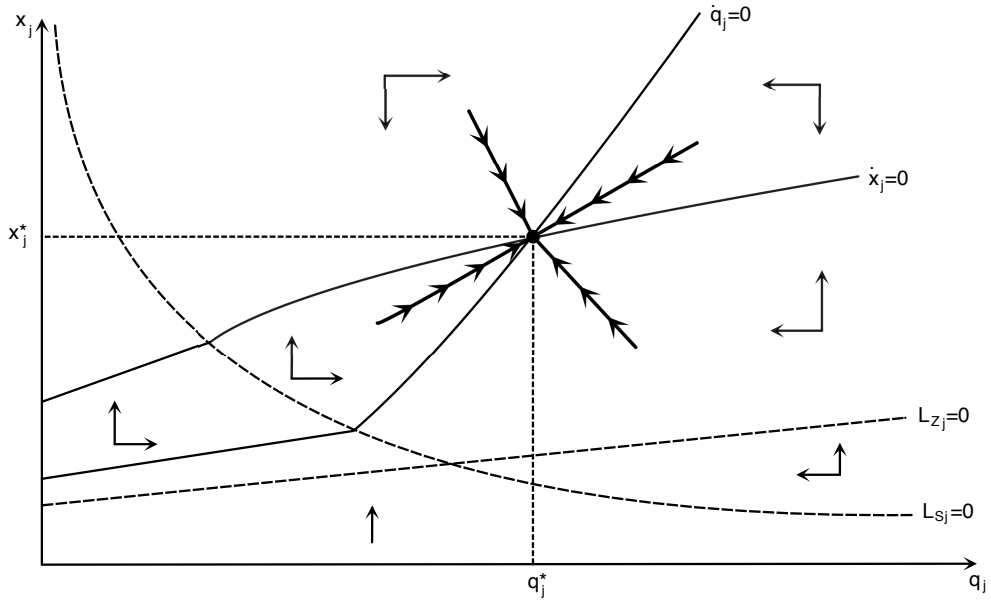

Figure 6: Benchmark dynamics with manufacturing and transportation innovation. The activation boundaries partition the state space, while the $\dot{q}_j = 0$ and $\dot{x}_j = 0$ loci intersect at the interior steady state (q_j^*, x_j^*).

This representation is useful for the analysis in the main text because later forces enter a system whose benchmark geometry is already explicit. A shipping-specific forcing moves the law of motion for q_j without entering the firm-size equation directly, while the activation boundaries continue to identify when the two innovation margins are profitable. Appendix C exploits the same tractability for more conventional parameter changes and transition-dynamics exercises.

C Additional Comparative Statics and Dynamics

The main text uses the model analytically to isolate the mechanisms generated by asymmetric cost push. This appendix reports additional comparative statics, transition dynamics, and welfare implications. Its purpose is analytical rather than quantitative: to show that the tractability used to derive the main results preserves the economically relevant margins needed to study conventional comparative statics, adjustment paths, and welfare. The model therefore obtains its analytical results without sacrificing its usefulness for these standard economic questions.

C.1 Steady-state comparative statics

In the interior regime with both innovation margins active and no regulatory forcing, two loci represent the steady state in (z_j, q_j) space. Combining free entry with the manufacturing-innovation condition gives the free-entry locus

$$q_j = \frac{\alpha_j \theta_j (\epsilon_j - 1)}{1 - \epsilon_j \beta(\rho + \delta_j)} \frac{\varphi_j + \left(\frac{1}{\alpha_j} + \frac{\theta_j}{\varsigma_j \sigma_j}\right) z_j}{z_j + \rho + \delta_j} - 1, \tag{C.1}$$

while the condition that the relative cost q_j be stationary gives the cost-structure locus

$$q_j = \frac{\theta_j \alpha_j}{\varsigma_j \sigma_j} \frac{\frac{\theta_j}{\sigma_j} z_j + \rho + \delta_j}{z_j + \rho + \delta_j}. \tag{C.2}$$

Under the maintained restriction $\theta_j > \sigma_j$, the cost-structure locus is increasing and approaches $\alpha_j \theta_j^2 / (\varsigma_j \sigma_j^2)$. The per-unit shipping-cost parameter η_j is absent from both loci. As in the benchmark, it affects the level of the relative-cost state but not the stationary allocation once transportation knowledge adjusts endogenously.

An increase in the fixed operating cost φ_j shifts the free-entry locus while leaving the cost-structure locus unchanged. The exercise works through firm size and the returns to the two forms of in-house R&D. Changes in the productivity of transportation R&D, ς_j, instead move the cost-structure relation directly as well as changing the R&D term in free entry. The resulting movement in (q_j^*, z_j^*) is therefore a useful illustration of the model's central complementarity: a change in the technology for improving shipping alters the return to improving the factory through the delivered-cost wedge. The exact direction of the movement depends on the parameter configuration.

C.2 The composition margin

The cleanest way to separate technological comparative statics from the regulation experiment is the stationary version of the relative-cost identity. With both innovation margins active,

$$\theta_j z_j^* - \sigma_j s_j^* + \Psi_j = 0. \tag{C.3}$$

When $\Psi_j = 0$, the two rates of cost reduction must satisfy

$$\frac{s_j^*}{z_j^*} = \frac{\theta_j}{\sigma_j}. \tag{C.4}$$

Thus ordinary parameter changes move the steady state subject to a fixed proportionality between the rates at which manufacturing and transportation knowledge reduce costs. A persistent asymmetric cost push is different. A rise in Ψ_j opens a wedge between the two rates: transportation must improve faster relative to manufacturing in order to prevent the relative shipping cost from trending without bound. This is why cumulative regulation has a composition effect that a common input-price shock does not.

The comparison is useful beyond the regulation application. A more productive transportation R&D technology changes how much knowledge effort is required to improve shipping; it does not itself create a trend in the relative cost of shipping versus making. An accumulating shipping-specific friction does create such a trend. The distinction is the same one developed in Sections 6–7, now viewed through standard steady-state comparative statics.

C.3 A transition experiment: a lower per-unit shipping cost

The parameter η_j provides a particularly transparent transition experiment because it does not enter the steady-state loci in equations (C.1) and (C.2). Let η_j^o denote its initial value and suppose that at $t = 0$ the industry experiences an unanticipated, immediate, and permanent fall to $\eta_j < \eta_j^o$. The knowledge stocks cannot jump, so the relative shipping cost jumps from q_j^* to

$$q_j(0) = \eta_j [S_j(0)]^{-\sigma_j} [Z_j(0)]^{\theta_j} < q_j^*, \tag{C.5}$$

while firm size remains at x_j^* on impact. The industry then follows the two-dimensional dynamic system back to the same steady state (q_j^*, x_j^*).

This experiment is useful precisely because the long-run comparison is silent: the shock changes neither steady-state firm size nor the steady-state growth rate. At the instant of

the shock, however, the fall in q_j raises the return to manufacturing innovation and lowers the return to transportation innovation. Along the subsequent transition the components of TFP growth are

$$T_j = \frac{Z_j^{\theta_j} N_j^{\chi_j}}{1+q_j}, \qquad g_j = \chi_j n_j + \theta_j z_j - \frac{q_j}{1+q_j}\frac{\dot{q}_j}{q_j}. \tag{C.6}$$

The response of g_j therefore combines the impulse responses of entry, manufacturing innovation, and the relative shipping cost. Even though the limiting growth rate is unchanged, the transition can cumulate into a permanent difference in the level of TFP. Equations (27) and (28), together with the innovation and entry conditions, trace the responses of all the endogenous variables to the shock.

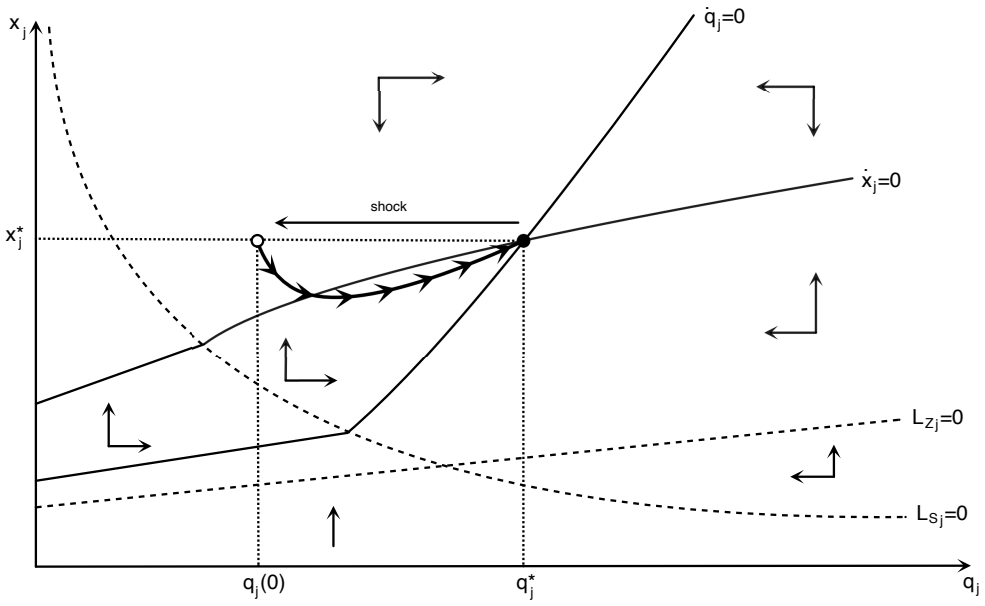

Figure 7: Transition following an unanticipated, immediate, and permanent fall in η_j. The relative shipping cost jumps down on impact while firm size cannot jump; the industry then follows the dynamic system back to the unchanged steady state.

The experiment also clarifies why a transition calculation can contain information absent from comparative statics. A parameter that disappears from the stationary loci can still generate economically important adjustment because it changes the economy's location in the state space. The same observation applies with greater force to the regulatory experiments in the main text, where the forcing itself evolves during the transition.

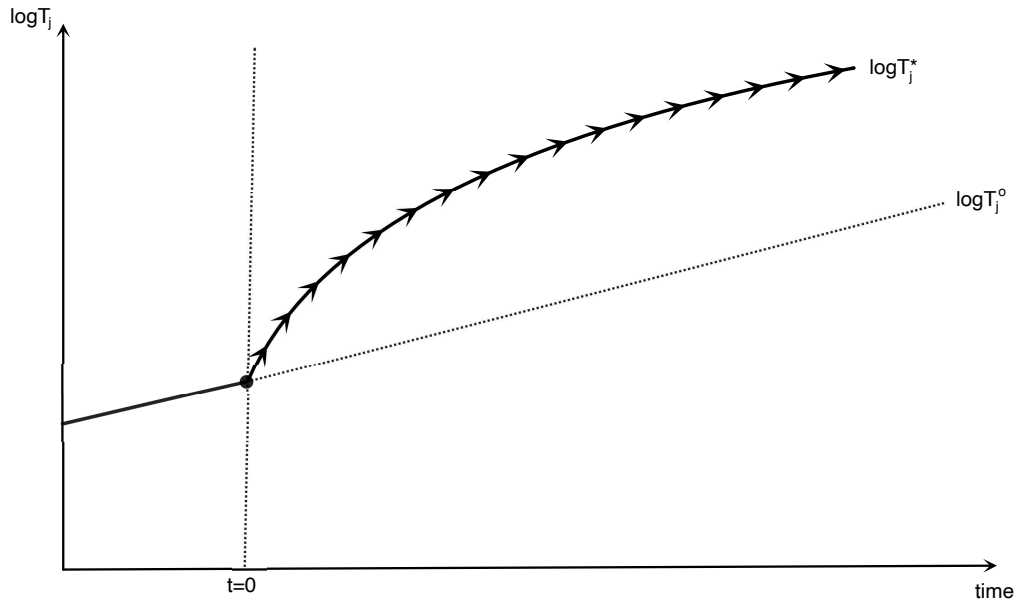

Figure 8: TFP level following the permanent fall in η_j. Although the long-run growth rate returns to its original value, the transitional growth acceleration can cumulate into a permanent gain in the level of TFP.

C.4 Welfare along a transition

We can map the transition paths directly into welfare. Up to constants that are common across experiments, we can write individual flow utility as

$$u(t) = \log y(t) + \sum_{j=1}^{J} \phi_j \log T_j(t), \tag{C.7}$$

where $y(t) = Y(t)/L(t)$. Hence the welfare difference per initial household member between a counterfactual path c and a baseline path b is

$$\frac{\Delta U}{\Lambda} = \int_0^\infty e^{-(\rho-\lambda)t} \left[\log \frac{y^c(t)}{y^b(t)} + \sum_{j=1}^{J} \phi_j \log \frac{T_j^c(t)}{T_j^b(t)} \right] dt. \tag{C.8}$$

Once the transition paths are known, we obtain welfare by discounting and integrating the corresponding consumption and productivity responses. No additional dynamic system is required.

For the experiment in the previous subsection, $y^c(t) = y^b(t) = y^*$ and the shock leaves the

limiting growth rate unchanged. Equation (C.8) therefore reduces to

$$\frac{\Delta U}{\Lambda} = \sum_{j=1}^{J} \phi_j \int_0^{\infty} e^{-(\rho-\lambda)t} \log \frac{T_j^c(t)}{T_j^b(t)}\, dt. \tag{C.9}$$

A transition that raises the TFP path therefore generates a welfare gain even though the before-and-after steady-state growth rates are identical. Conversely, the same calculation evaluates a regulatory transition that temporarily suppresses manufacturing innovation, including the welfare cost of any persistent level loss in productive knowledge.

The point of these exercises is to show that the same low-dimensional structure used to derive the paper's analytical results also determines transition paths and welfare. They are additional implications of the model rather than a separate quantitative exercise.

D RegData and Regulatory Accumulation

The regulation experiment in Section 7 treats the mode-specific wedges as stocks that can accumulate at different rates. This appendix documents the descriptive evidence motivating that formulation. RegData counts binding restrictions in the Code of Federal Regulations and uses estimated industry-relevance probabilities to construct industry-specific measures of regulatory exposure (Al-Ubaydli and McLaughlin 2017; McLaughlin and Sherouse 2019). RegData U.S. 6.0 covers 1970–2025 and reports industry classifications through the six-digit NAICS level (McLaughlin 2026).

At the aggregate level, the federal stock rises from 409,520 restrictions in 1970 to 1,126,959 in 2025, an increase of 175.2 percent. Net declines occur in six of the 55 annual comparisons. The data therefore support persistent net accumulation, not the stronger claim that the stock rises in every year.

For transportation, we construct relevance-weighted exposure for NAICS 481–484. Let r_{pt} denote the restriction count in CFR part p in year t and let π_{pmt} denote RegData's estimated relevance of that part to transportation mode m. Measured exposure is

$$B_{mt} = \sum_{p} \pi_{pmt} r_{pt}. \tag{D.1}$$

The calculation retains all relevance probabilities rather than imposing a probability cutoff. Because a CFR part can be relevant to more than one industry, the mode-specific exposure measures overlap and should not be added as though they were distinct restrictions.

Table 1: Descriptive federal regulatory exposure in transportation, 2000–2025

Mode / NAICS	2000	2025	Change	Annual growth
Truck / 484	930.4	2,815.2	202.6%	4.53%
Rail / 482	7,421.0	15,034.7	102.6%	2.86%
Air / 481	36,454.8	52,567.1	44.2%	1.47%
Water / 483	9,204.5	10,470.8	13.8%	0.52%

Notes: Values are relevance-weighted restriction counts, not estimates of compliance costs. Annual growth is the compound annual growth rate over 25 elapsed years. Calculations use the archived, unthresholded RegData U.S. 6.0 industry-year panel and NAICS 481–484.

Two features matter for the model. First, measured regulatory exposure can accumulate substantially over economically relevant horizons. Second, the rate of accumulation differs sharply across transportation industries. The distinction between levels and growth rates is important: trucking begins this window with much lower measured exposure than water transportation but accumulates restrictions much faster. The least regulated mode at a point in time therefore need not be the mode with the slower accumulation of regulations.

The table is descriptive evidence, not a structural estimate. Restriction counts do not identify the cost wedges R_j^m, their productivity effects, substitution elasticities across modes, or the threshold at which manufacturing innovation becomes inactive. Nor does the ranking in Table 1 hold over every possible sample window: over 1970–2025 rail has the fastest measured growth, and several modal series decline between 2024 and 2025. We therefore use constant nonnegative ψ_j^m as a tractable representation of persistent and heterogeneous accumulation, not as a literal description of every year or transportation mode. The model asks what follows for endogenous innovation when such accumulation changes the relative cost of making and shipping goods.

References

Acheampong, Alex O., Ernest Boateng, Mark Amponsah, and Thomas C. Dzogbewu. 2022. "Unveiling the Effect of Transport Infrastructure and Technological Innovation on Economic Growth, Energy Consumption, and CO2 Emissions." *Technological Forecasting and Social Change* 182: 121843.

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth Through Creative Destruction." *Econometrica* 60: 323–351.

Al-Ubaydli, Omar, and Patrick A. McLaughlin. 2017. "RegData: A Numerical Database on Industry-Specific Regulations for All United States Industries and Federal Regulations, 1997–2012." *Regulation & Governance* 11(1): 109–123.

Alchian, Armen A., and William R. Allen. 1964. *University Economics.* Belmont, CA: Wadsworth.

Allen, Treb, and Costas Arkolakis. 2014. "Trade and the Topography of the Spatial Economy." *Quarterly Journal of Economics* 129(3): 1085–1140.

Aschauer, David A. 1989. "Is Public Expenditure Productive?" *Journal of Monetary Economics* 23(2): 177–200.

Calderón, César, and Luis Servén. 2010. "Infrastructure and Economic Development in Sub-Saharan Africa." *Journal of African Economies* 19(Suppl. 1): i13–i87.

Comin, Diego, Mikhail Dmitriev, and Esteban Rossi-Hansberg. 2012. "The Spatial Diffusion of Technology." NBER Working Paper No. 18534.

Desmet, Klaus, and Esteban Rossi-Hansberg. 2014. "Spatial Development." *American Economic Review* 104(4): 1211–1243.

Donaldson, Dave, and Richard Hornbeck. 2016. "Railroads and American Economic Growth: A 'Market Access' Approach." *Quarterly Journal of Economics* 131(2): 799–858.

Fernald, John G. 1999. "Roads to Prosperity? Assessing the Link between Public Capital and Productivity." *American Economic Review* 89(3): 619–638.

Fogel, Robert W. 1964. *Railroads and American Economic Growth: Essays in Econometric History.* Baltimore: Johns Hopkins University Press.

Fujita, Masahisa, and Jacques-François Thisse. 2002. *Economics of Agglomeration: Cities, Industrial Location, and Regional Growth.* Cambridge: Cambridge University Press.

Grossman, Gene M., and Elhanan Helpman. 1991. *Innovation and Growth in the Global Economy.* Cambridge, MA: MIT Press.

Hummels, David, and Alexandre Skiba. 2004. "Shipping the Good Apples Out? An Empirical Confirmation of the Alchian–Allen Conjecture." *Journal of Political Economy* 112(6): 1384–1402.

Irarrazabal, Alfonso, Andreas Moxnes, and Luca David Opromolla. 2015. "The Tip of the Iceberg: A Quantitative Framework for Estimating Trade Costs." *Review of Economics and*

Statistics 97(4): 777–792.

Kane, Robert F., and Pietro F. Peretto. 2020. "More Apples vs. Better Apples: Distribution and Innovation-Driven Growth." *Journal of Economic Theory* 185: 104964.

Krugman, Paul. 1991. "Increasing Returns and Economic Geography." *Journal of Political Economy* 99(3): 483–499.

McLaughlin, Patrick A. 2026. *RegData U.S. 6.0* [dataset]. QuantGov, Hoover Institution, Stanford University.

McLaughlin, Patrick A., and Oliver Sherouse. 2019. "RegData 2.2: A Panel Dataset on US Federal Regulations." *Public Choice* 180: 43–55.

Peretto, Pietro F. 1998. "Technological Change and Population Growth." *Journal of Economic Growth* 3(4): 283–311.

Peretto, Pietro F. 1999. "Cost Reduction, Entry, and the Interdependence of Market Structure and Economic Growth." *Journal of Monetary Economics* 43: 173–195.

Peretto, Pietro F. 2009. "Energy Taxes and Endogenous Technological Change." *Journal of Environmental Economics and Management* 57: 269–283.

Peretto, Pietro F. 2012. "Resource Abundance, Growth and Welfare: A Schumpeterian Perspective." *Journal of Development Economics* 97: 142–155.

Peretto, Pietro F. 2027. *Knowledge, Industry, and the Endogenous Growth of Nations*. Cambridge, MA: MIT Press.

Peretto, Pietro F., and Michelle Connolly. 2007. "The Manhattan Metaphor." *Journal of Economic Growth* 12: 329–350.

Peretto, Pietro F., and Simone Valente. 2011. "Resources, Innovation and Growth in the Global Economy." *Journal of Monetary Economics* 58: 387–399.

Peretto, Pietro F., and Simone Valente. 2015. "Growth on a Finite Planet: Resources, Technology and Population in the Long Run." *Journal of Economic Growth* 20: 305–331.

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy* 98(5, Part 2): S71–S102.

Sørensen, Anders. 2014. "Additive versus Multiplicative Trade Costs and the Gains from Trade Liberalizations." *Canadian Journal of Economics* 47(3): 1032–1046.