When Do Firms Clean Up After Themselves? Firm Scale, Environmental Policy, and Endogenous Innovation

Pietro F. Peretto
Department of Economics, Duke University

October 5, 2026

Abstract

When and how does environmental policy induce firms to clean up after themselves? This paper develops a theory of endogenous innovation activation in which firm scale determines when latent innovation opportunities become privately profitable and in what sequence firms exploit them. In a multi-industry endogenous-growth model, firms can undertake productivity or abatement innovation. We distinguish the emissions tax from the *dirtiness* of production, defined as emissions relative to variable production cost. Their product is the relative emissions-tax burden that governs firms' innovation incentives. Rising firm scale can activate initially dormant innovation opportunities under an unchanged tax. Along one development path, productivity innovation activates first and makes production progressively dirtier; abatement innovation begins only later and eventually becomes strong enough to make production cleaner. The theory characterizes analytically the thresholds, ordering, turning points, and timing through which an environmental incentive becomes innovation, innovation becomes cleaner production, and cleaner production eventually becomes lower emissions. The same dynamics generate a delayed abatement response and an environmental Kuznets curve. Alternative uses of tax revenue leave the steady state unchanged but alter transition dynamics. Observed environmental performance therefore reflects both structural characteristics and an industry's position along its development path.

Keywords: environmental policy; endogenous innovation; firm scale; abatement; dirtiness; market structure; environmental Kuznets curve.
JEL Classification: O31, O33, O41, Q55, Q58.

1. Introduction

Environmental policy changes the incentives firms face, but a change in incentives need not produce an immediate technological response. Whether a firm acts on an incentive depends on whether innovation is profitable, and profitability depends in part on the scale over which the resulting knowledge is applied. This observation suggests that the same environmental policy can generate qualitatively different technological responses at different stages of development because endogenous changes in firm scale determine which forms of innovation are economically viable. The same reasoning applies more generally to settings in which economic development changes the profitability of latent investment opportunities and thereby determines when firms act on them.

This paper develops a theory of endogenous innovation activation. Firms face multiple latent innovation opportunities that are not profitable at their current scale. As development changes market structure and

firm scale, the private returns to these opportunities change and different innovation margins become active at different points in time. Development therefore determines not only how much firms innovate but also what they innovate on and when. The environmental application makes this sequencing especially important because productivity and abatement innovation move the environmental state of production in opposite directions.

This paper uses this mechanism to ask a more specific question than whether environmental policy induces innovation: *When and how does environmental policy induce firms to clean up after themselves?* Productivity and environmental innovation take place inside the same firm and apply to the same production activity. Environmental policy changes the relative return to the two activities, but it does not by itself guarantee that either activity is profitable. Rising firm scale can first activate productivity innovation, which makes production dirtier, and only later activate abatement innovation. Even then, firms do not begin to clean up until abatement innovation becomes strong enough to dominate the effect of continuing productivity growth. The model therefore separates three events that need not coincide: the activation of environmental R&D, the point at which production begins becoming cleaner, and the point at which emissions begin falling. More generally, a development process endogenously generates a sequence of technological regimes. The theory characterizes analytically the thresholds, ordering, turning points, and timing that govern this sequence.

A large literature studies whether environmental policy induces technological change. Jaffe, Newell, and Stavins (2002) survey the theoretical and empirical foundations of the induced-innovation mechanism, while Popp (2002) provides early patent evidence that changes in energy prices redirect inventive activity toward energy-saving technologies. More recent firm-level evidence shows that environmental policy can shift innovation toward cleaner technologies; see Calel and Dechezleprêtre (2016), Aghion et al. (2016), and the review by Dechezleprêtre and Sato (2017). A related endogenous growth literature embeds environmental technological change in dynamic general equilibrium. Smulders (1995) is an early contribution, and Smulders and de Nooij (2003) study a setting in which both the rate and direction of technological progress respond endogenously to energy policy. Acemoglu, Aghion, Bursztyn, and Hémous (2012) develop a framework in which innovation can be directed toward clean or dirty sectors. Brock and Taylor (2005) provide a broad synthesis of the growth-and-environment literature.

A recurring empirical finding across different settings motivates this paper's approach: environmental technological responses vary systematically across firms and over time. Qi, Tang, and Xi (2021) show that larger Chinese manufacturing firms are more likely to use clean technology and that distortions to the firm-size distribution materially affect industrial water pollution. De Miguel and Pazó (2017) find that environmental regulation stimulates process innovation in large Spanish manufacturing firms, while the corresponding response among smaller firms takes a different form. Chakraborty (2017) finds that firms' technological response to an environmental standard imposed on Indian exporters is concentrated in the upper half of the firm-size distribution. Brown, Martinsson, and Thomann (2022) find that increases in emissions taxes raise R&D among highly polluting firms and that the response appears after rather than before the tax increase. These findings do not identify the activation mechanism studied here, but they document the empirical regularities for which such a theory is useful: responses vary with firm scale and can occur with a delay.

Shapiro and Walker (2018) show how much of the observed cleaning of U.S. manufacturing can be attributed to changes in environmental regulation. This paper explains the endogenous process through which an environmental incentive becomes technological change and, eventually, actual cleaning. Their quantitative model accounts for the large decline in U.S. manufacturing pollution between 1990 and 2008 primarily through within-product changes in emissions intensity and attributes most of that decline to increased regulatory

stringency. Abatement is a contemporaneous input choice in that exercise. Here firms invest in abatement knowledge, and the profitability of doing so evolves endogenously with firm scale. The model therefore generates the transition itself: an environmental incentive can be present while abatement innovation remains dormant, development can subsequently activate that innovation, and actual cleaning can occur only later.

The model features two innovation activities. Firms can undertake productivity R&D or environmental R&D. Both activities become more attractive as firm scale rises because a technological improvement can be applied to a larger volume of activity. Their effects on the environmental state of production, however, run in opposite directions. We call emissions relative to variable production cost the *dirtiness* of production. Productivity innovation lowers direct production cost and, absent offsetting abatement, makes production dirtier. Environmental innovation lowers emissions and makes production cleaner. The emissions tax prices this dirtiness. The product of the tax rate and dirtiness is the firm's relative emissions-tax burden, which determines the relative return to the two innovation activities. The economy therefore has two endogenous innovation thresholds: one at which productivity R&D becomes profitable and another at which environmental R&D becomes profitable.

The intersection of these thresholds has a useful economic interpretation. It identifies the largest firm scale consistent with dormant firm innovation. While actual R&D is zero in this region, the private returns to productivity and environmental innovation remain well defined. These latent returns rise or fall continuously as population growth and entry jointly determine the scale of the typical firm. The no-innovation dynamics therefore exhibit a counterfactual firm size toward which market structure would converge if firm innovation never activated. When that scale lies beyond the intersection of the two activation loci, at least one latent return reaches the household's reservation return on saving. We call this the *firm-innovation activation condition*: development itself eventually activates investment in firm-specific knowledge even though environmental policy remains unchanged.

Which form of innovation activates first depends on the economy's initial relative emissions-tax burden. If the burden is relatively low, productivity innovation activates first. With the tax rate fixed, productivity innovation then raises dirtiness and hence the relative emissions-tax burden, making environmental innovation progressively more attractive. If the initial burden is relatively high, environmental innovation activates first, reduces dirtiness and the relative emissions-tax burden, and restores the relative return to productivity improvement. The active innovation margin therefore changes the environmental state that governs the return to the other margin. The direction of innovation evolves endogenously even under an unchanged environmental policy.

This structure generates tractable global dynamics. The economy can move from a regime with no firm innovation to one in which only a single innovation margin is active and then to a regime in which both forms of R&D operate. The activation loci determine which margin becomes viable first, while the transition loci determine when the economy moves from rising to falling dirtiness and emissions. These are endogenous economic events, not dates imposed from outside the model. Under public-good provision, the paper characterizes the global dynamics analytically and establishes convergence to the unique interior steady state under economically transparent conditions. Lump-sum redistribution preserves the steady state but changes the transition because it couples industries through the common interest rate.

The activation mechanism suggests a reason not to expect a *universal and immediate* technological response to environmental policy. A particularly revealing experiment is to introduce a permanent emissions tax after productivity innovation has already become active but before abatement innovation is profitable. The tax changes firms' operating environment immediately, yet the response of abatement innovation is delayed

until endogenous firm scale and the relative emissions-tax burden carry the economy to the relevant activation threshold. In other words, a policy change alters the latent returns to innovation immediately, but actual R&D begins only when the dynamics carry one of those returns across its activation threshold. Cross-sectional heterogeneity and delayed responses are therefore not necessarily nuisance features obscuring a common treatment effect. They can be structural implications of endogenous innovation thresholds. The theory predicts that otherwise similar policy incentives can produce different responses across firms, industries, and stages of development because the relevant innovation margins are active in some cases and dormant in others.

The paper builds on Peretto (2008), which studied effluent taxes, endogenous market structure, and firms' allocation of resources between cost-reducing and emission-reducing R&D. That analysis focused primarily on the steady state because of the complexity of the model, but it suggested that studying the transition could produce important insights. This paper makes the transition the central object of analysis.

Two changes make this possible. First, we reformulate the innovation technology so that firm-level research uses labor linearly while industry knowledge spillovers generate declining private returns to a firm's own R&D effort. We can then solve for innovation decisions algebraically at each point in time. Second, we allow the elasticities of production costs and emissions with respect to their respective knowledge stocks to differ. The equal-elasticity assumption used for tractability in the earlier analysis becomes a special case rather than an important restriction. These changes reduce the equilibrium to a low-dimensional autonomous system without eliminating the economic margins of interest.

The paper uses two fiscal closures to separate the innovation mechanism from the general-equilibrium feedback generated by environmental-tax revenue. In the benchmark case, the government allocates the final-good units collected as emissions taxes directly to public consumption, which enters household utility. Tax revenue therefore bypasses the household's consumption–saving decision: private expenditure per capita and the market rate of return remain constant, and each industry has a clean two-state representation in firm scale and the relative emissions-tax burden. We then let the government rebate the same revenue lump sum to households. The technological mechanism and the aggregate resource flow are unchanged, but tax receipts now become household purchasing power and affect private expenditure and the common market rate of return, coupling otherwise self-contained industry dynamics. Comparing the two cases isolates the additional dimensions in which general equilibrium contributes to the innovation mechanism.

The scale mechanism extends beyond emissions. A substantial induced-innovation literature shows that energy prices affect the development, direction, and adoption of energy-saving technologies (Newell, Jaffe, and Stavins, 1999; Popp, 2002; Linn, 2008; Fried, 2018). Peretto (2009) studies a complementary energy-tax margin, while Peretto, Coffey, and McLaughlin (2026) show how transportation costs, regulation, and scarcity affect the rate and direction of endogenous innovation as firm scale evolves. These environments differ in their technological details but share the broader mechanism emphasized here: changes in costs alter the returns to distinct innovation activities, while endogenous firm scale determines when those activities become privately viable.

The same tractability carries over to the multi-industry economy because each industry retains the same low-dimensional innovation dynamics. Industries differ in market size, production technology, environmental technology, and research productivity but face a common emissions tax. The policy maker need not identify which industries are "clean," choose technologies, or determine where environmental R&D should occur. Different industries respond differently because their endogenous scales and technological positions place them in different innovation regimes. The common policy therefore produces endogenous sorting across industries in both the timing and direction of technological change.

The paper studies the economy under the maintained assumption that entry remains active in every industry. This restriction keeps the global analysis focused on the paper's research question—the endogenous activation and direction of firm innovation—rather than adding the secular transition dynamics from a pre-industrial regime to the industrial regime; see, e.g., Peretto (2027). Under active entry, aggregation remains particularly simple. The household asset market and tax rebates couple the industries through a common market rate of return without creating an additional aggregate dynamic state. The resulting interaction has a rank-one structure: a single general-equilibrium rate of return links industry-specific innovation decisions. This preserves industries as meaningful technological objects while allowing common policy and aggregate expenditure to link their dynamics. If entry becomes inactive in some industries, this aggregation property no longer holds and the affected industries must be solved jointly with the aggregate economy, see Peretto (2027) for examples of this mechanism.

The central message of the paper is thus simple. Firms face latent innovation opportunities that become privately viable only after development has carried firm scale across endogenous activation thresholds. With multiple innovation margins, deep parameters determine the thresholds, their ordering, and the relevant turning points; the equilibrium transition through this structure determines their timing. In the environmental application, policy creates incentives to reduce the environmental cost of production, but economic scale determines when firms can profitably act on those incentives. An industry with a low relative emissions-tax burden activates productivity innovation first and abatement later; an industry with a high burden activates abatement first and productivity innovation later. The familiar narrative in which productive development precedes environmental innovation is therefore one equilibrium history, not an imposed stage structure. Along that history, the model generates an environmental Kuznets curve analytically: real income per capita rises while emissions per capita first rise with productivity innovation and later fall when abatement becomes sufficiently strong. The distinction between current and fundamental dirtiness follows from the same transition logic: observed environmental performance combines structural characteristics with position along the path.

The paper proceeds as follows. Section 2 develops the multi-industry environment and the two fiscal closures. Section 3 derives firm behavior and entry. Section 4 solves the public-good benchmark and develops the activation and development dynamics in their cleanest form. Section 5 introduces lump-sum redistribution and shows how the resulting expenditure feedback changes transitional dynamics and links industries through the common market return. Section 6 studies endogenous technological sorting across industries under a common environmental policy. Section 7 develops the implications for environmental policy, and Section 8 concludes.

2. The Model Environment

The industrial-organization, household, entry, and asset-market blocks follow the standard structure recently synthesized and expanded in Peretto (2027). We state them compactly here and reserve detailed discussion for assumptions that matter specifically for emissions and abatement. We deliberately narrow the fiscal environment to one instrument—the emissions tax—and compare two uses of its revenue in order to isolate the innovation mechanism and the additional general-equilibrium propagation generated by fiscal recycling.[1]

The economy has J industries. A Cobb–Douglas aggregator combines industry components into the homogeneous consumption good, while a CES aggregator combines the differentiated intermediate goods

[1] Chapter 9 of Peretto (2027) studies fiscal policy in the same class of models with a substantially richer menu of tax instruments, transfers, and public spending arrangements.

supplied within each industry. Labor moves freely across industries and the financial market is fully integrated in no-arbitrage equilibrium across all industries and activities. Firms make goods, generate emissions as a by-product, and can devote R&D to either manufacturing productivity or abatement.

2.1. Technologies

Firm i in industry j produces with the technology

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

where Z_{ij} denotes firm-specific manufacturing knowledge and φ_j overhead labor. Productivity R&D accumulates firm-specific knowledge according to

$$\dot{Z}_{ij} = \alpha_j Z_j L_{Zij}, \qquad Z_j = \frac{1}{N_j}\int_0^{N_j} Z_{sj}\, ds, \qquad \alpha_j > 0. \tag{2}$$

The firm treats the industry-specific average stock Z_j as a knowledge spillover beyond its control.

Production generates emissions according to

$$E_{ij} = B_{ij}^{-\eta_j} X_{ij}, \qquad \eta_j > 0, \tag{3}$$

where B_{ij} denotes firm-specific abatement knowledge. Abatement R&D obeys the parallel technology

$$\dot{B}_{ij} = \varsigma_j B_j L_{Bij}, \qquad B_j = \frac{1}{N_j}\int_0^{N_j} B_{sj}\, ds, \qquad \varsigma_j > 0. \tag{4}$$

We allow the elasticities θ_j and η_j to differ. Their relative magnitude determines interesting aspects of the dynamics.

2.2. Households

Population is an exogenous forcing and follows the law of motion

$$L(t) = \Lambda e^{\lambda t}, \qquad \Lambda > 0, \qquad \lambda \in (-\infty, +\infty), \tag{5}$$

where Λ is a scale parameter and λ is the population growth rate. The household supplies all labor and maximizes

$$U(t) = \int_t^{\infty} e^{-\rho(s-t)} L(s) u(s)\, ds, \qquad \rho > \max\{0, \lambda\}. \tag{6}$$

We specify the instantaneous utility function separately for the two fiscal closures below. In both cases private consumption enters logarithmically, so the household's consumption-saving decision satisfies

$$r = \rho + \frac{\dot{Y}}{Y} - \lambda. \tag{7}$$

This expression defines the household's reservation return on saving at which the household is willing to sacrifice current consumption to finance investment.

2.3. Final production and the intermediate-goods demand system

A competitive final-good producer combines industry components according to

$$G = \prod_{j=1}^{J} G_j^{\phi_j}, \qquad \sum_{j=1}^{J} \phi_j = 1, \tag{8}$$

and therefore spends

$$Y_j = \phi_j Y \tag{9}$$

on component j. Within industry j, a competitive producer uses

$$G_j = N_j^{\chi_j - \frac{1}{\epsilon_j - 1}} \left(\int_0^{N_j} X_{ij}^{\frac{\epsilon_j - 1}{\epsilon_j}} \, di \right)^{\frac{\epsilon_j}{\epsilon_j - 1}}, \qquad \chi_j > 0, \quad \epsilon_j > 1, \tag{10}$$

which yields the demand curve

$$X_{ij} = Y_j \frac{p_{Xij}^{-\epsilon_j}}{\int_0^{N_j} p_{Xkj}^{1-\epsilon_j} \, dk}. \tag{11}$$

This expression characterizes the intermediate-good demand system.

2.4. Environmental policy and two fiscal closures

The government levies the common emissions tax τ. Firm i in industry j pays τE_{ij}. Total tax revenue is

$$\mathcal{T} \equiv \tau \sum_{j=1}^{J} \int_0^{N_j} E_{ij} \, di. \tag{12}$$

We solve the model under two alternative uses of $\mathcal{T}$. Varying only the disposition of emissions-tax revenue lets us identify the propagation mechanism in isolation from other fiscal margins.

Case I: Public-good provision. The government allocates the final-good units collected as emissions taxes directly to public consumption, $P = \mathcal{T}$.[2] No additional government production technology is required. Let

$$u = \log\left(\frac{Y}{p_G L}\right) + v\left(\frac{P}{L}\right), \qquad v' > 0, \quad v'' \leq 0. \tag{13}$$

Public consumption enters utility separably. Because the government, rather than the household, allocates these final-good units to consumption, tax receipts bypass the household's intertemporal consumption–saving decision and never enter disposable income. Hence, the household budget is

$$\dot{A} = rA + wL - Y. \tag{14}$$

[2]This is the standard one-good general-equilibrium convention: the homogeneous final good is the numeraire and can be allocated directly across alternative uses. One unit of final output collected as tax revenue can therefore finance one unit of public consumption, just as one unit of final output allocated to investment becomes one unit of new capital in the Ramsey–Cass–Koopmans model. See Barro and Sala-i-Martin (2004).

This closure isolates the private innovation mechanism from the government's fiscal policy because environmental-tax receipts do not circulate back through private expenditure or asset accumulation.

Case II: Lump-sum redistribution. The government rebates the tax revenue to households as disposable income. The physical resources collected by the government are unchanged, but households now choose how much of the rebate to consume and save. Private utility is

$$u = \log\left(\frac{Y}{p_G L}\right), \tag{15}$$

and the household budget becomes

$$\dot{A} = rA + wL + \mathcal{T} - Y. \tag{16}$$

The rebate therefore places tax receipts inside the household's consumption–saving decision. The tax now affects innovation indirectly because the resulting circulation of purchasing power changes private expenditure and the market return.

3. Firm Behavior and Entry

In this section, we characterize the behavior of firms and then model entry. The details of the derivation are in the Appendix.

3.1. Firm decisions

At time t, firm i in industry j chooses its price and R&D employment to maximize

$$V_{ij}(t) = \int_t^{\infty} \exp\left[-\int_t^s (r(v) + \delta_j)\, dv\right] \Pi_{ij}(s)\, ds, \qquad \delta_j \geq 0, \tag{17}$$

where δ_j is the exit hazard and

$$\Pi_{ij} = \left[p_{Xij} - wZ_{ij}^{-\theta_j} - \tau B_{ij}^{-\eta_j}\right] X_{ij} - w\varphi_j - wL_{Zij} - wL_{Bij}. \tag{18}$$

The firm takes Z_j, B_j, N_j, w, r, and the demand curve as given.

The pricing decision is a constant markup over the marginal cost of production,

$$p_{Xij} = \frac{\epsilon_j}{\epsilon_j - 1}\left(wZ_{ij}^{-\theta_j} + \tau B_{ij}^{-\eta_j}\right). \tag{19}$$

We choose labor as numeraire, $w \equiv 1$. Symmetry within industry j yields $p_{Xj}X_j = Y_j/N_j$. We define firm size, dirtiness, and the relative emissions-tax burden as

$$x_j \equiv \frac{Y_j}{N_j}, \qquad D_j \equiv Z_j^{\theta_j} B_j^{-\eta_j} = \frac{E_j}{L_{Xj} - \varphi_j}, \qquad T_j \equiv \tau D_j = \tau Z_j^{\theta_j} B_j^{-\eta_j}. \tag{20}$$

Then variable manufacturing labor and the emissions-tax payment are

$$L_{Xj} - \varphi_j = \frac{\frac{\epsilon_j - 1}{\epsilon_j} x_j}{1 + T_j}, \qquad \tau E_j = \frac{\frac{\epsilon_j - 1}{\epsilon_j} x_j T_j}{1 + T_j}. \tag{21}$$

Thus x_j is sales per firm. The variable T_j has an equally direct interpretation. Since the marginal production cost of one unit of output is $Z_j^{-\theta_j}$ and the emissions-tax cost per unit is $\tau B_j^{-\eta_j}$, their ratio is

$$T_j = \frac{\tau B_j^{-\eta_j}}{Z_j^{-\theta_j}}. \tag{22}$$

Equivalently, equation (21) implies

$$T_j = \frac{\tau E_j}{L_{Xj} - \varphi_j}, \qquad E_j = \frac{\frac{\epsilon_j - 1}{\epsilon_j}}{\tau} \frac{x_j T_j}{1 + T_j}. \tag{23}$$

The decomposition $T_j = \tau D_j$ separates environmental policy from the environmental state of production. The tax rate τ is the price the government attaches to emissions. Dirtiness D_j is emissions relative to variable manufacturing labor cost. We refer to $D_j(t)$ as the industry's *current dirtiness*; when the industry converges to a steady state, D_j^* denotes its *long-run, or fundamental, dirtiness*: the level of dirtiness selected by the industry's deep technological and economic parameters and the policy environment. This distinction separates where the industry is along its transition from the destination implied by its structural characteristics and maintained policy environment. The relative emissions-tax burden T_j is therefore the firm's emissions-tax bill relative to its variable manufacturing labor cost. This distinction will be useful below. The tax rate is a policy parameter, dirtiness is an endogenous technological state, and their product is the private cost burden that governs innovation incentives. With a common tax rate, ranking industries by T_j is equivalent to ranking them by D_j; with different tax rates, it is not. Since the tax rate is constant along a transition, the two state variables have the same growth rate:

$$\frac{\dot{D}_j}{D_j} = \frac{\dot{T}_j}{T_j} = \theta_j \frac{\dot{Z}_j}{Z_j} - \eta_j \frac{\dot{B}_j}{B_j} = \theta_j \alpha_j L_{Zj} - \eta_j \varsigma_j L_{Bj}. \tag{24}$$

Productivity innovation raises dirtiness by reducing production cost relative to emissions, whereas abatement innovation lowers dirtiness by reducing emissions relative to production cost. Thus the law of motion for T_j below is also the law of motion for dirtiness.

The decisions about R&D employment yield the two rates of return

$$r + \delta_j = \alpha_j \left[\frac{\theta_j \frac{\epsilon_j - 1}{\epsilon_j} x_j}{1 + T_j} - L_{Zj} \right] \tag{25}$$

and

$$r + \delta_j = \varsigma_j \left[\frac{\eta_j \frac{\epsilon_j - 1}{\epsilon_j} x_j T_j}{1 + T_j} - L_{Bj} \right]. \tag{26}$$

Taking into account the non-negativity constraints, the resulting investment rules have two branches. For

productivity innovation,

$$L_{Zj} = \begin{cases} 0, & \alpha_j \dfrac{\theta_j \frac{\epsilon_j - 1}{\epsilon_j} x_j}{1 + T_j} \leq r + \delta_j, \\ \dfrac{\theta_j \frac{\epsilon_j - 1}{\epsilon_j} x_j}{1 + T_j} - \dfrac{r + \delta_j}{\alpha_j}, & \alpha_j \dfrac{\theta_j \frac{\epsilon_j - 1}{\epsilon_j} x_j}{1 + T_j} > r + \delta_j, \end{cases} \tag{27}$$

and for abatement innovation,

$$L_{Bj} = \begin{cases} 0, & \varsigma_j \dfrac{\eta_j \frac{\epsilon_j - 1}{\epsilon_j} x_j T_j}{1 + T_j} \leq r + \delta_j, \\ \dfrac{\eta_j \frac{\epsilon_j - 1}{\epsilon_j} x_j T_j}{1 + T_j} - \dfrac{r + \delta_j}{\varsigma_j}, & \varsigma_j \dfrac{\eta_j \frac{\epsilon_j - 1}{\epsilon_j} x_j T_j}{1 + T_j} > r + \delta_j. \end{cases} \tag{28}$$

This branch structure is economically important. Each form of firm investment can be inactive even though its latent private return is well defined. Firm size raises the return to both forms of innovation, while the relative emissions-tax burden shifts the return from productivity R&D toward abatement R&D. The two investment decisions therefore generate four possible innovation regimes. Section 4 derives the boundaries in (T_j, x_j) space at which firms switch branches and studies the transitions across those regimes.

3.2. Empirical evidence on R&D activation

The branch structure derived above has a direct empirical motivation. It captures a long-standing regularity in the literature on the Schumpeterian hypotheses concerning firm size, market structure, and innovative activity. That literature, which developed from the 1950s onward, distinguishes two margins that are important here: whether a firm undertakes R&D at all and, conditional on doing so, how much it spends. The probability of performing R&D rises strongly with firm size, while among firms that perform R&D, expenditure tends to rise approximately proportionally with size. Surveys of this literature consequently summarize the relationship, to a first approximation, as one in which R&D becomes active only after a firm-size threshold has been crossed and thereafter rises roughly proportionally with firm size (Cohen and Levin, 1989; Cohen and Klepper, 1996; Symeonidis, 1996; Cohen, 2010). Symeonidis (1996), reviewing a large body of evidence accumulated in the debate over the Schumpeterian hypotheses, summarizes the empirical relationship particularly clearly: R&D expenditure tends to rise approximately proportionally with firm size once it passes a threshold level.

This characterization has a long genealogy and does not depend on any single empirical study. Early contributions by Mansfield, Scherer, Comanor, and others investigated nonlinearities in the relation between firm size and innovative activity, while subsequent work emphasized the importance of distinguishing R&D participation from R&D intensity. Cohen and Levin (1989) survey this evidence as part of the broader empirical literature generated by the Schumpeterian hypotheses. Cohen and Klepper (1996) subsequently emphasize both the strong positive relation between firm size and the propensity to perform R&D and the close relation between firm size and R&D among R&D performers. The evidence does not imply a universal numerical threshold common to firms or industries. Technological opportunity, appropriability, fixed costs, demand conditions, and other industry characteristics affect whether and when R&D becomes profitable. The robust empirical point for present purposes is instead the existence of an important extensive margin: many firms undertake no R&D, the probability of undertaking it increases strongly with firm scale, and R&D

becomes approximately scale-related once the innovation margin is active.

The activation loci derived above provide the model's endogenous counterpart to this empirical regularity. They also sharpen its interpretation and implications. In the model there is not one universal R&D threshold because firms face distinct innovation opportunities. Productivity R&D becomes active when firm scale crosses $x^0_{Zj}(T_j)$, whereas abatement R&D becomes active when it crosses $x^0_{Bj}(T_j)$. Moreover, these thresholds depend on the firm's relative emissions-tax burden T_j. Firm scale therefore determines the profitability of exploiting a latent innovation opportunity, while environmental policy and the firm's dirtiness determine which opportunity becomes profitable and when. The branch structure describes both the extensive margin—whether a particular type of R&D is undertaken—and, once the margin is active, the intensive response of R&D to firm scale. This is the same mechanism that generates an activation threshold followed by R&D that is linear in firm size in the basic endogenous-market-structure model (Peretto, 2027).

Evidence on environmental innovation reinforces this interpretation. Energy prices affect energy-saving innovative activity (Popp, 2002), and changes in energy prices and regulation affect the characteristics and the set of energy-using products that firms choose to offer (Newell, Jaffe, and Stavins, 1999). Plant-level evidence likewise shows that energy prices induce technology adoption, with different responses by incumbents and entrants (Linn, 2008), while firm-level evidence links environmental innovation to accumulated technological capabilities and R&D knowledge capital (Horbach, 2008). More directly, Calel and Dechezleprêtre (2016) find that exposure to the European Union Emissions Trading System caused an increase in low-carbon patenting among regulated firms. This literature does not estimate the particular activation locus $x^0_{Bj}(T_j)$ derived here. It establishes the empirical ingredients behind it: environmental innovation contains an extensive adoption and innovation margin, firms respond to environmental prices and regulation, and those responses depend on firm characteristics and technological capabilities. The model combines these regularities with the older firm-size evidence by making the profitability threshold for abatement innovation depend jointly on firm scale and the environmental incentive.

The model therefore turns a long-established empirical regularity—the extensive margin of R&D participation followed by approximately scale-related R&D once innovation is active—into a theory of endogenous, technology-specific innovation thresholds. In the present application, the existence of two such thresholds and their relative positions determine the sequence in which productivity and abatement innovation become active.

3.3. Entry

An entrepreneur creates a firm in industry j by hiring $\beta_j Y_j/N_j$ units of labor and enters with the industry-average knowledge stocks. This entry technology and its financial-market interpretation are standard in this framework; Peretto (2027) presents the underlying construction and discusses alternative formulations that preserve the equilibrium mechanism.

Free entry requires

$$V_j = \beta_j x_j. \tag{29}$$

The associated return is

$$r + \delta_j = \frac{1}{\beta_j}\left[\frac{1}{\epsilon_j} - \frac{\varphi_j + L_{Zj} + L_{Bj}}{x_j}\right] + \frac{\dot{x}_j}{x_j}. \tag{30}$$

Equivalently,

$$\dot{x}_j = \frac{\varphi_j + L_{Zj} + L_{Bj}}{\beta_j} + x_j\left(r + \delta_j - \frac{1}{\beta_j \epsilon_j}\right). \tag{31}$$

This equation is the general equilibrium law of motion of firm size. In each regime, we specialize it by substituting in the relevant expressions for the firm investment terms L_{Zj} and L_{Bj}.

We maintain active entry in every industry throughout the analysis. Relaxing this restriction would add the secular transition from a pre-industrial regime with inactive entry to the industrial regime studied here and would introduce an additional extensive-margin in state-space. Maintaining active entry lets us isolate how evolving firm scale changes the latent returns to firm innovation and activates investment in firm-specific knowledge.

4. Case I: Public-Good Provision and the Core Mechanism

Case I gives us the clean benchmark. We begin its construction with some general equilibrium results that are the key to the model's tractability. We then characterize the dynamics and the global stability of the system.

4.1. Expenditure and the interest rate

In free entry equilibrium, aggregate financial wealth is

$$A = \sum_{j=1}^{J} N_j V_j = \sum_{j=1}^{J} \beta_j Y_j = \beta Y, \qquad \beta \equiv \sum_{j=1}^{J} \beta_j \phi_j. \tag{32}$$

The two fiscal closures now differ only in what happens to tax revenue after firms pay it.

Substitute the result $A = \beta Y$ and the Euler equation (7) into the household budget (14). The terms involving $\dot{Y}$ cancel, leaving

$$Y = \frac{L}{1 - \beta(\rho - \lambda)}, \qquad 1 > \beta(\rho - \lambda). \tag{33}$$

Expenditure per capita is thus constant, so

$$r = \rho. \tag{34}$$

With this constant interest rate, the expressions in equation (31) simplify substantially.

Recall also that tax revenue affects utility through public-good provision but does not create a private-expenditure state. Together with the constant-interest-rate result that we just obtained, this property yields that each industry is governed by a self-contained two-dimensional dynamical system in the (x_j, T_j) space. We now characterize its properties.

4.2. Dormant firm innovation and activation

The growth rate of the relative emissions-tax burden is

$$\frac{\dot{T}_j}{T_j} = \theta_j \alpha_j L_{Zj} - \eta_j \varsigma_j L_{Bj}. \tag{35}$$

This expression makes transparent what drives the development path in each innovation regime: productivity innovation raises T_j, whereas abatement innovation lowers it. This creates a stabilizing force in the direction of innovation. By raising the relative emissions-tax burden, productivity innovation raises the return to abatement innovation; by lowering it, abatement innovation restores the relative return to productivity innovation. The active innovation margin therefore tends to make activation of the other margin progressively more attractive.

We call the region in which both incumbent R&D margins are inactive *dormant firm innovation*. Entry remains active, so dormancy refers only to investment in firm-specific knowledge, not to firm creation or to aggregate development. For brevity, below we refer to this region as the *dormant regime*. Although actual R&D is zero, the private returns to both innovation activities remain latent variables that evolve with firm scale. The thresholds at which those returns reach the household's reservation return on saving are

$$x_{Zj}^0(T_j) = \frac{(\rho + \delta_j)(1 + T_j)}{\alpha_j \theta_j \frac{\epsilon_j - 1}{\epsilon_j}}, \qquad x_{Bj}^0(T_j) = \frac{(\rho + \delta_j)(1 + T_j)}{\varsigma_j \eta_j \frac{\epsilon_j - 1}{\epsilon_j} T_j}. \tag{36}$$

They intersect at

$$T_{Ij} = \frac{\alpha_j \theta_j}{\varsigma_j \eta_j}, \qquad x_{Ij} = \frac{(\rho + \delta_j)(\alpha_j \theta_j + \varsigma_j \eta_j)}{\alpha_j \varsigma_j \theta_j \eta_j \frac{\epsilon_j - 1}{\epsilon_j}}. \tag{37}$$

The productivity innovation threshold rises with T_j and the abatement innovation threshold falls with it. Hence productivity innovation activates first for $T_j < T_{Ij}$ and abatement innovation activates first for $T_j > T_{Ij}$.

In the dormant region,

$$\dot{x}_j = \frac{\varphi_j}{\beta_j} + x_j \left((\rho + \delta_j) - \frac{1}{\beta_j \epsilon_j} \right). \tag{38}$$

If $1/\epsilon_j > \beta_j(\rho + \delta_j)$, the counterfactual dormant rest point is

$$\bar{x}_j = \frac{\varphi_j}{1/\epsilon_j - \beta_j(\rho + \delta_j)}. \tag{39}$$

Because the lower envelope of the two activation loci reaches its maximum at (T_{Ij}, x_{Ij}), a sufficient firm-activation condition is simply

$$\bar{x}_j > x_{Ij}. \tag{40}$$

Proposition 1 (Firm activation) *If* $1/\epsilon_j > \beta_j(\rho + \delta_j)$ *and* (40) *holds, the dormant firm-innovation region contains no stationary equilibrium. From every admissible initial condition in that region, population growth and active entry change firm scale until the latent return to at least one form of firm innovation reaches the household's reservation return on saving and the corresponding investment margin activates.*

Proof. Equation (38) implies convergence toward the counterfactual dormant rest point $\bar{x}_j$. Under (40), that rest point lies above the maximum of the lower envelope of the two activation loci. Hence a trajectory starting in the dormant region must cross at least one activation locus before it can reach a stationary point. □

The proposition identifies firm scale as the force that moves the economy out of dormant firm innovation. It does not determine which innovation margin activates first; that depends on the relative emissions-tax burden at the activation point. As noted in the Introduction, the mechanism is more general: whenever development changes firm scale, it can activate latent investment opportunities whose profitability depends on the scale over which their returns are realized. We therefore turn next to the two regions in which only one innovation margin is active.

4.3. One active innovation margin

If only productivity R&D operates,

$$\left.\frac{\dot{T}_j}{T_j}\right|_Z = \frac{\alpha_j \theta_j^2 \frac{\epsilon_j - 1}{\epsilon_j} x_j}{1 + T_j} - \theta_j(\rho + \delta_j) > 0. \tag{41}$$

Productivity innovation raises T_j and therefore raises the return to abatement innovation, which activates at

$$x_{B|Z,j}(T_j) = \frac{(\rho + \delta_j)(1 + T_j)}{\varsigma_j \eta_j \frac{\epsilon_j - 1}{\epsilon_j} T_j}. \tag{42}$$

If only abatement R&D operates,

$$\left.\frac{\dot{T}_j}{T_j}\right|_B = \eta_j(\rho + \delta_j) - \frac{\varsigma_j \eta_j^2 \frac{\epsilon_j - 1}{\epsilon_j} x_j T_j}{1 + T_j} < 0. \tag{43}$$

Abatement innovation lowers T_j and raises the return to productivity innovation, which activates at

$$x_{Z|B,j}(T_j) = \frac{(\rho + \delta_j)(1 + T_j)}{\alpha_j \theta_j \frac{\epsilon_j - 1}{\epsilon_j}}. \tag{44}$$

Thus the active innovation activity moves the industry toward activation of the dormant innovation activity. This stabilization force is the core of the paper's mechanism.

4.4. Both active innovation margins

When both margins operate,

$$\frac{\dot{T}_j}{T_j} = Q_j(x_j, T_j) \equiv \frac{\frac{\epsilon_j - 1}{\epsilon_j} x_j}{1 + T_j} \left(\alpha_j \theta_j^2 - \varsigma_j \eta_j^2 T_j\right) + (\eta_j - \theta_j)(\rho + \delta_j). \tag{45}$$

The $\dot{T}_j = 0$ locus is

$$x_{Tj}(T_j) = \frac{(\theta_j - \eta_j)(\rho + \delta_j)(1 + T_j)}{\frac{\epsilon_j - 1}{\epsilon_j}(\alpha_j \theta_j^2 - \varsigma_j \eta_j^2 T_j)}. \tag{46}$$

Define

$$T_{0j} \equiv \frac{\alpha_j \theta_j^2}{\varsigma_j \eta_j^2}. \tag{47}$$

The threshold T_{0j} is the relative emissions-tax burden at which the effect of firm scale on the evolution of the relative emissions-tax burden changes sign. For $T_j < T_{0j}$, greater firm scale strengthens the productivity innovation component of $\dot{T}_j/T_j$ relative to the abatement innovation component; for $T_j > T_{0j}$, the abatement innovation component dominates. Since

$$\frac{T_{0j}}{T_{Ij}} = \frac{\theta_j}{\eta_j}, \tag{48}$$

the relative magnitudes of the two knowledge elasticities determine where this change occurs relative to the

intersection of the two innovation-activation loci. The configuration in Figure 1 corresponds to $\theta_j > \eta_j$, and hence $T_{0j} > T_{Ij}$. On the economically relevant branch, $T_{Ij} < T_j < T_{0j}$, the $\dot{T}_j = 0$ locus is strictly increasing. Moreover,

$$\frac{x_{Tj}(T_j)}{x^0_{Zj}(T_j)} = \frac{\alpha_j \theta_j (\theta_j - \eta_j)}{\alpha_j \theta_j^2 - \varsigma_j \eta_j^2 T_j} > 1 \qquad (T_j > T_{Ij}), \tag{49}$$

so the $\dot{T}_j = 0$ locus lies above the activation locus for productivity innovation after the two activation curves meet. This ordering, rather than any global ordering between the $\dot{x}_j = 0$ locus and the activation curves, determines the location of the relevant intersections in the phase diagram.

The entry equation becomes

$$\dot{x}_j = F_{0j}(x_j, T_j) \equiv \frac{\varphi_j - (\rho + \delta_j)\left(\frac{1}{\alpha_j} + \frac{1}{\varsigma_j}\right)}{\beta_j} + x_j \left[(\rho + \delta_j) - \frac{1}{\beta_j \epsilon_j} + \frac{\frac{\epsilon_j - 1}{\epsilon_j}(\theta_j + \eta_j T_j)}{\beta_j (1 + T_j)} \right]. \tag{50}$$

Hence the $\dot{x}_j = 0$ locus is

$$x_{Xj}(T_j) = \frac{\varphi_j - (\rho + \delta_j)\left(\frac{1}{\alpha_j} + \frac{1}{\varsigma_j}\right)}{1/\epsilon_j - \beta_j(\rho + \delta_j) - \frac{\epsilon_j - 1}{\epsilon_j}(\theta_j + \eta_j T_j)/(1 + T_j)}. \tag{51}$$

The two interior steady-state loci have opposite slopes. Under the firm-activation condition and the maintained feasibility restrictions, they intersect once.

Proposition 2 (Interior steady state) *Under the firm-activation, feasibility, and regularity conditions, industry j has a unique steady state with both innovation margins active.*

Proof. The result follows from the geometry derived above. On the economically relevant joint-innovation branch, the $\dot{T}_j = 0$ and $\dot{x}_j = 0$ loci have opposite slopes under the maintained restrictions and intersect once; the firm-activation condition places the relevant intersection in the region in which both innovation margins are active. □

The unique steady state is therefore pinned down by the intersection of two economic forces: the allocation of innovation determines the stationary relative emissions-tax burden, while entry determines stationary firm scale. Uniqueness, however, does not establish that the economy approaches this configuration. We next study its local and global stability.

4.5. Stability and global dynamics

The benchmark dynamics system is

$$\dot{T}_j = T_j Q_j(x_j, T_j), \qquad \dot{x}_j = F_{0j}(x_j, T_j). \tag{52}$$

Let $\mathcal{J}_j^{PG}$ denote its Jacobian at the interior steady state. If $\det \mathcal{J}_j^{PG} > 0$ and $\operatorname{tr} \mathcal{J}_j^{PG} < 0$, the steady state is locally asymptotically stable. If, in addition, the eigenvalues are real, convergence is locally non-oscillatory. In contrast, if

$$(\operatorname{tr} \mathcal{J}_j^{PG})^2 - 4 \det \mathcal{J}_j^{PG} < 0, \tag{53}$$

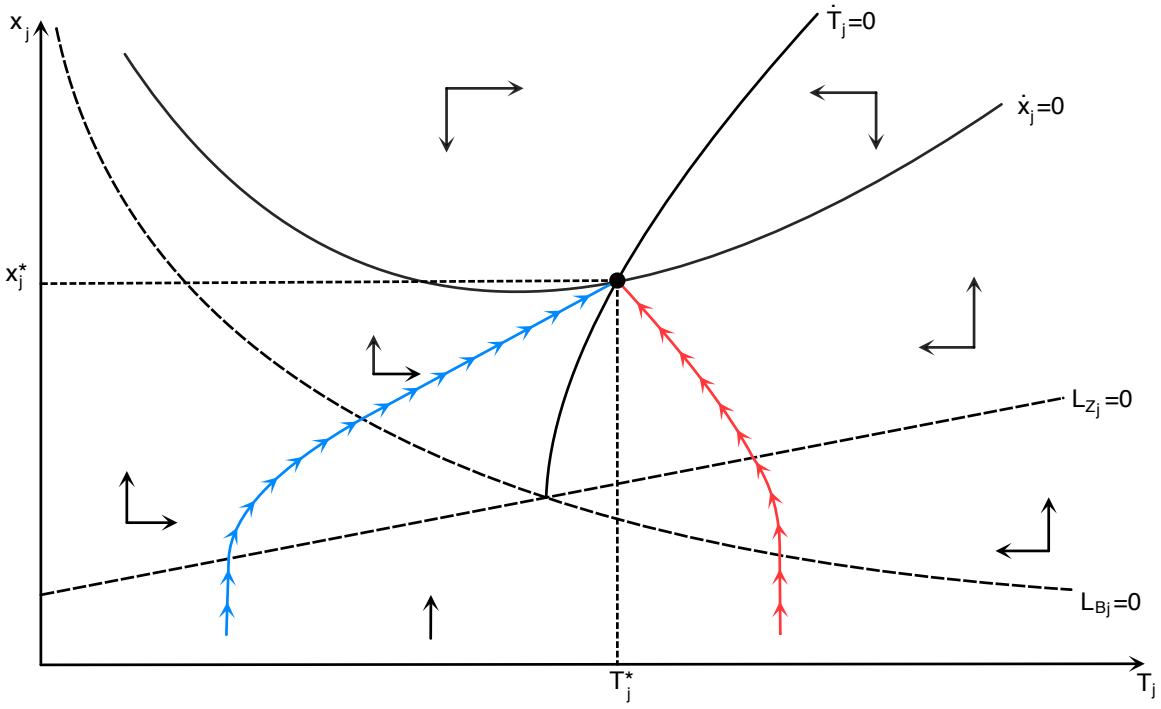

Figure 1: Firm scale, dirtiness, and endogenous activation. The horizontal axis is the relative emissions-tax burden $T_j = \tau D_j$; with a constant common tax, movements in T_j are proportional to movements in dirtiness D_j. The dashed loci mark the thresholds at which productivity and abatement innovation become active; the solid loci are the $\dot{x}_j = 0$ and $\dot{T}_j = 0$ steady-state loci. The two trajectories illustrate how different initial levels of dirtiness generate different sequences of innovation activation.

convergence is oscillatory. Productivity innovation and abatement innovation alternate in relative intensity even under an unchanged emissions tax.

Figure 1 summarizes the global dynamics in the monotonic convergence case. Because $T_j = \tau D_j$ and τ is constant, horizontal motion also records whether the industry is becoming dirtier or cleaner. Below both activation loci, firm innovation is dormant. Entry nevertheless raises x_j, while dirtiness remains fixed because neither knowledge stock changes. The trajectory therefore moves vertically until it reaches one of the two activation boundaries. If initial dirtiness is low, equivalently $T_j < T_{Ij}$, productivity innovation activates first. In that region $\dot{x}_j > 0$ and $\dot{D}_j/D_j = \dot{T}_j/T_j > 0$: productivity innovation makes production progressively dirtier and moves the industry toward activation of abatement innovation. If initial dirtiness is high, equivalently $T_j > T_{Ij}$, abatement innovation activates first. Then $\dot{x}_j > 0$ and $\dot{D}_j/D_j = \dot{T}_j/T_j < 0$: firms clean up production and move the industry toward activation of productivity innovation. Initial conditions therefore determine the sequence of innovation activation, while the active margin itself changes dirtiness in the direction that makes the other margin progressively more attractive.

Once both margins operate, there is convergence to the steady state (T_j^*, x_j^*). For the configuration shown, $\theta_j > \eta_j$ makes the $\dot{T}_j = 0$ locus monotonically increasing on the relevant branch and equation (49) places it above the productivity innovation activation locus. This geometry fixes the economically relevant ordering of intersections. The $\dot{x}_j = 0$ locus need not lie globally above the activation loci; any crossing that occurs beyond the steady state is irrelevant for trajectories converging to (T_j^*, x_j^*). The arrows in the figure illustrate the

direction of motion separately. Together with the activation geometry, they show how trajectories pass from dormant firm innovation through one-margin innovation into the joint-innovation region and then converge to the unique interior steady state.

The public-good closure makes the global analysis especially tractable because $r = \rho$ leaves each industry as a self-contained planar system and, crucially, allows the same Dulac multiplier to be used in every innovation regime. This common multiplier lets the no-cycle argument extend across activation boundaries rather than applying only within each smooth regime. For regime $\mathcal{R} \in \{0, Z, B, ZB\}$, write

$$\left.\frac{\dot{T}_j}{T_j}\right|_{\mathcal{R}} = x_j p_{\mathcal{R}j}(T_j) + d_{\mathcal{R}j}(\rho + \delta_j), \tag{54}$$

where

$$d_{0j} = 0, \qquad d_{Zj} = -\theta_j, \qquad d_{Bj} = \eta_j, \qquad d_{ZBj} = \eta_j - \theta_j.$$

The corresponding functions $p_{\mathcal{R}j}$ are zero in the dormant regime, $\frac{\epsilon_j - 1}{\epsilon_j}\alpha_j\theta_j^2/(1 + T_j)$ in the productivity-innovation-only regime, $-\frac{\epsilon_j - 1}{\epsilon_j}\varsigma_j\eta_j^2 T_j/(1 + T_j)$ in the abatement-innovation-only regime, and $\frac{\epsilon_j - 1}{\epsilon_j}(\alpha_j\theta_j^2 - \varsigma_j\eta_j^2 T_j)/(1 + T_j)$ when both margins operate.

Crucially, the same Dulac multiplier[3] works in every regime:

$$\mathcal{B}_j(x_j, T_j) = \frac{1}{x_j T_j}. \tag{55}$$

Let $s_{\mathcal{R}j}$ equal 0, $1/\alpha_j$, $1/\varsigma_j$, or $\left(\frac{1}{\alpha_j} + \frac{1}{\varsigma_j}\right)$ according to the active R&D margins. Direct differentiation gives

$$\nabla \cdot (\mathcal{B}_j F_{\mathcal{R}j}) = p'_{\mathcal{R}j}(T_j) - \frac{\varphi_j - (\rho + \delta_j)s_{\mathcal{R}j}}{\beta_j x_j^2 T_j}. \tag{56}$$

For every active regime, $p'_{\mathcal{R}j}(T_j) < 0$; in the dormant regime it equals zero. Hence the simple sufficient restriction

$$\varphi_j > (\rho + \delta_j)\left(\frac{1}{\alpha_j} + \frac{1}{\varsigma_j}\right) \tag{57}$$

makes the divergence strictly negative everywhere in the admissible state space.

Because the multiplier is common, the vector field is continuous when a trajectory crosses an activation boundary. Therefore, since (55) does not change across regimes, the boundary terms generated by applying Green's theorem regime by regime cancel, and the Bendixson–Dulac argument rules out a periodic orbit even if it crosses innovation boundaries.

Proposition 3 (No recurrent cycles under public-good provision) *Suppose* (57) *and the maintained feasibility conditions hold. Then the industry-j piecewise-smooth system admits no periodic orbit, including an orbit that crosses activation boundaries.*

Proof. The common Dulac multiplier in (55) makes the divergence strictly negative in every regime under (57). Because the vector field is continuous at the activation boundaries and the multiplier is the same on both

[3] A Dulac multiplier is a continuously differentiable scalar function that multiplies a planar vector field without changing its phase trajectories, apart from their time parametrization. If the divergence of the resulting vector field has one strict sign on a simply connected region, the Bendixson–Dulac criterion rules out periodic orbits in that region. See Dumortier, Llibre, and Artés (2006).

sides, the interface terms cancel when Green's theorem is applied regime by regime. The Bendixson–Dulac criterion therefore excludes a periodic orbit, including one crossing activation boundaries. □

Combining this result with firm activation yields a global convergence result once the economically admissible region is bounded. In particular, if the admissible state space contains a compact positively invariant region, the interior steady state is the unique attractor in that region, and the firm-activation condition holds, every trajectory in the region converges to that attractor. The boundedness condition is needed because the Dulac argument by itself only rules out periodic orbits. To establish global convergence, we must also ensure that trajectories remain in a bounded economically admissible region rather than approaching its boundary or becoming unbounded.

Proposition 4 (Global convergence in the public-good benchmark) *Under the preceding assumptions, and provided the admissible state space contains a compact positively invariant region with the unique interior steady state, every trajectory starting in that region converges to the steady state.*

Proof. Firm activation rules out a stationary point in the dormant region, Proposition 3 rules out periodic orbits, and the maintained assumptions provide a unique interior steady state in a compact positively invariant region. The planar Poincaré–Bendixson theorem then implies convergence of every trajectory in that region to the unique steady state. □

The result establishes that the activation mechanism does more than determine which form of innovation starts first: under the stated conditions it directs the entire transition toward the same interior steady state. With the dynamics characterized, we can now translate the path of (x_j, T_j) into the observable environmental outcome—the path of emissions.

4.6. Emissions along the development path

The phase diagram allow us to characterize the path of emissions analytically. Equation (21) yields firm emissions as

$$E_j = \frac{\frac{\epsilon_j - 1}{\epsilon_j}}{\tau} \frac{x_j T_j}{1 + T_j}, \tag{58}$$

and therefore

$$\frac{\dot{E}_j}{E_j} = \frac{\dot{x}_j}{x_j} + \frac{1}{1 + T_j} \frac{\dot{T}_j}{T_j}. \tag{59}$$

Thus, firm emissions necessarily rise along the dormant segment of the equilibrium path that lies under the activation boundaries because $\dot{T}_j/T_j = 0$ and $\dot{x}_j > 0$. They also rise in the phase with only productivity innovation, in which both terms in (59) are positive. Once abatement innovation starts, its effect on $\dot{T}_j/T_j$ puts downward pressure on emissions. In particular, firm emissions turn downward when

$$-\frac{1}{1 + T_j} \frac{\dot{T}_j}{T_j} > \frac{\dot{x}_j}{x_j}, \tag{60}$$

which says that abatement innovation generates a reduction in the firm's emissions that dominates the continuing expansion of the firm's production.

Aggregate industry emissions have an even simpler representation. Let $\mathcal{E}_j \equiv N_j E_j$. Since $N_j x_j = Y_j$, we

have

$$\mathcal{E}_j = \frac{\frac{\epsilon_j - 1}{\epsilon_j} Y_j}{\tau} \frac{T_j}{1 + T_j}, \qquad \frac{\dot{\mathcal{E}}_j}{\mathcal{E}_j} = \frac{\dot{Y}_j}{Y_j} + \frac{1}{1 + T_j} \frac{\dot{T}_j}{T_j}. \tag{61}$$

Under public-good provision, $Y_j = \phi_j Y$ and equation (33) gives $\dot{Y}_j / Y_j = \lambda$. Hence, emissions per capita satisfy

$$\frac{d}{dt} \log \left(\frac{\mathcal{E}_j}{L} \right) = \frac{1}{1 + T_j} \frac{\dot{T}_j}{T_j}. \tag{62}$$

Because τ is constant, per-capita industry emissions therefore move in the same direction as dirtiness D_j. Along the low-dirtiness trajectory in Figure 1, production is initially neither becoming cleaner nor dirtier while firm innovation is dormant. Once productivity innovation activates, D_j rises: production becomes progressively dirtier and emissions per capita rise. Abatement innovation subsequently activates, but its activation does not by itself imply that the industry has begun to clean up. As long as $\dot{D}_j > 0$, the dirtiness generated by productivity innovation still dominates the cleaning effect of abatement innovation. The industry starts becoming cleaner only when the trajectory crosses the $\dot{T}_j = 0$, equivalently $\dot{D}_j = 0$, locus. Dirtiness and emissions per capita reach their maximum at that crossing and decline thereafter.

The distinction among dirtiness, firm emissions, and aggregate emissions is economically important. Dirtiness is a technological state: it measures emissions relative to variable production cost. Firm emissions additionally reflect the scale of the firm's activity, so they can rise even when dirtiness is unchanged. Aggregation removes this firm-scale term because entry distributes industry expenditure across a changing number of firms. Under Case I, population growth accounts for the growth of aggregate industry expenditure, leaving dirtiness as a sufficient statistic for the direction of motion of aggregate emissions per capita. The model therefore separates three events that need not coincide: firms can begin abatement R&D, production can begin becoming cleaner, and firm emissions can begin falling. On the low-initial-dirtiness development path, they occur in precisely that order.

The model also characterizes the timing of these events analytically. Let t_{Bj} denote the first date at which the equilibrium trajectory reaches the abatement-innovation activation locus, t_{Dj} the first subsequent date at which it reaches the $\dot{T}_j = 0$ locus, and t_{Ej} the first subsequent date at which condition (60) holds with equality. These dates are endogenous first-passage times. The deep parameters determine the activation and transition loci and their ordering, while the equilibrium path determines when they are reached. Along the low-initial-dirtiness development path, $t_{Bj} < t_{Dj} < t_{Ej}$: abatement R&D begins first, production begins becoming cleaner later, and firm emissions turn down later still. Under Case I, per-capita industry emissions peak at t_{Dj} because equation (62) makes their growth rate change sign exactly with $\dot{T}_j$. The endogenous dates also provide additional empirical implications of the theory. Because the activation and turning times are functions of the model's deep parameters and initial conditions, observed dates and durations—such as the lag between a policy change and the activation of abatement R&D, or between innovation activation and declining emissions—constitute additional data moments that can discipline quantitative implementations of the model.

4.7. A policy experiment: delayed abatement innovation

The phase diagram suggests a natural policy experiment that makes the activation mechanism especially transparent. Consider an economy that initially has no emissions tax. Its market structure and manufacturing knowledge follow the standard one-knowledge-stock dynamics of this class of models, while emissions evolve

in the background without abatement innovation because firms have no incentive to invest in accumulating abatement knowledge. At date $t = 0$, unexpectedly and immediately the government introduces a permanent emissions tax. The policy immediately raises the cost of emissions and therefore affects firms' operating decisions. It does not, however, necessarily activate abatement innovation.

To see this, choose an initial condition such that productivity innovation is already active when the tax is introduced but the post-policy state lies strictly below the activation locus for abatement innovation in Figure 1. The tax causes T_j to jump from zero to the value implied by the inherited knowledge stocks and the new tax rate, while the inherited stocks themselves cannot jump. If the resulting state lies in the productivity-innovation-only region, firms continue to set $L_{Bj} = 0$. Productivity innovation nevertheless raises dirtiness D_j and, with the tax fixed, the relative emissions-tax burden $T_j = \tau D_j$, while population growth and entry continue to change firm scale. The economy therefore moves toward the activation locus for abatement innovation even though the tax itself remains constant.

Proposition 5 (Delayed abatement innovation after an emissions tax) *Suppose Case I applies and a permanent emissions tax is introduced at $t = 0$. If the post-policy initial condition lies in the productivity-innovation-only region, abatement R&D remains inactive until the trajectory first reaches the abatement innovation activation locus. During this interval $\dot{D}_j/D_j = \dot{T}_j/T_j > 0$, so production becomes dirtier and emissions per capita rise according to* (62)*. When the trajectory reaches the activation locus, abatement innovation begins. If the subsequent joint-innovation path crosses the $\dot{T}_j = 0$, equivalently $\dot{D}_j = 0$, steady-state locus from the region in which production is becoming dirtier to the region in which it is becoming cleaner, dirtiness and emissions per capita attain a maximum at that crossing and decline thereafter.*

Proof. This follows directly from the regime dynamics derived above. Below the abatement-innovation activation locus, the branch condition sets $L_{Bj} = 0$ and (41) gives $\dot{T}_j/T_j > 0$. Abatement R&D starts when the trajectory reaches the activation locus. Equation (62) then maps the sign of $\dot{T}_j/T_j$ one-for-one into the sign of per-capita emissions growth, so the emissions peak occurs where the subsequent joint-innovation path crosses the $\dot{T}_j = 0$ locus. □

The proposition isolates an endogenous answer to the question in the title: environmental policy induces firms to clean up after themselves only when firm scale and the relative emissions-tax burden make abatement innovation privately profitable. Nothing prevents firms from observing or responding to the tax at $t = 0$: the tax immediately enters marginal production cost and changes operating decisions. What is delayed is the *abatement-innovation response*. The delay requires neither adjustment costs, imperfect information, policy uncertainty, nor a change in environmental preferences. Moreover, the start of abatement R&D is not yet the point at which production starts becoming cleaner. Initially, productivity innovation can continue to raise dirtiness faster than abatement innovation lowers it. Cleaning begins only when the latter effect dominates and $\dot{D}_j$ turns negative. The mechanism therefore provides a dynamic interpretation of the environmental Kuznets curve that differs from both reduced-form evidence and models that obtain the inverted-U relation from static abatement technology. Grossman and Krueger (1995) provide the classic empirical benchmark, while Stokey (1998) and Andreoni and Levinson (2001) develop influential theoretical accounts. Particularly close to the present emphasis on endogenous technological change, Smulders, Bretschger, and Egli (2011) explain environmental Kuznets curves through the diffusion of clean technologies.[4] Here, the mechanism is

[4] Rubio, García, and Hueso (2009) also obtain a development path along which environmental R&D is initially zero and becomes positive later. In their neoclassical planner problem, environmental deterioration raises the social value of cleaner technology until environmental R&D becomes optimal. The activation mechanism here is different: decentralized firms face latent innovation

the endogenous activation of abatement innovation: the emissions tax can already be in place while firms postpone abatement R&D until firm scale and the relative emissions-tax burden raise its latent return to the household's reservation return on saving.

The resulting emissions history has the qualitative properties of an environmental Kuznets curve. Under public-good provision, nominal expenditure per capita is constant, while innovation lowers goods prices. In particular,

$$p_{Xj} = \frac{1}{\frac{\epsilon_j - 1}{\epsilon_j}} \left(Z_j^{-\theta_j} + \tau B_j^{-\eta_j} \right), \tag{63}$$

so either productivity or abatement innovation weakly lowers the industry price. Along a development path with monotonically rising real income per capita, the policy experiment therefore produces an interval in which income and emissions per capita rise together, followed by a decline in emissions after abatement becomes sufficiently strong. The familiar sequence in which productive development precedes environmental innovation is thus an equilibrium response to a policy that can be in place throughout the transition, not an imposed stage structure or the consequence of a later tightening of environmental policy. A high post-policy relative emissions-tax burden, T_j, can instead activate abatement innovation first and generate a different emissions history.

4.8. From industry dynamics to aggregate outcomes

The analysis so far has followed a single industry. The economy, however, consists of industries with different fundamentals, activation thresholds, innovation regimes, and positions along their transition paths. Aggregating these industry dynamics reveals how this heterogeneity maps into economy-wide environmental performance and growth.

Let total emissions be $\mathcal{E} \equiv \sum_{j=1}^{J} \mathcal{E}_j$. Using $Y_j = \phi_j Y$ in equation (61), aggregate emissions per capita under Case I are

$$\frac{\mathcal{E}}{L} = \frac{Y/L}{\tau} \sum_{j=1}^{J} \phi_j \frac{\epsilon_j - 1}{\epsilon_j} \frac{T_j}{1 + T_j}. \tag{64}$$

Since Y/L and τ are constant,

$$\frac{d}{dt}\left(\frac{\mathcal{E}}{L}\right) = \frac{Y/L}{\tau} \sum_{j=1}^{J} \phi_j \frac{\epsilon_j - 1}{\epsilon_j} \frac{T_j}{(1 + T_j)^2} \frac{\dot{T}_j}{T_j}. \tag{65}$$

Aggregate emissions per capita rise or fall according to a positive weighted sum of the industry-specific changes in dirtiness. Industries need not reach their dirtiness peaks simultaneously, and some can be becoming cleaner while others are still becoming dirtier. Whether the resulting aggregate path exhibits an environmental Kuznets curve consequently depends on the collection of industry fundamentals and on where the industries are along their respective transitions. An aggregate hump is an empirical implication that can be evaluated for a given economy, not a shape imposed by the theory.

Growth aggregates differently. The Cobb–Douglas final-good technology in equation (8) implies that the final-good price index is the geometric average of the industry price indexes. Under the CES specification in

opportunities whose private profitability depends on endogenous firm scale, and productivity and abatement innovation have separate activation thresholds. Development therefore determines not only whether firm innovation becomes active but also which innovation margin activates first.

equation (10), symmetry within industry j gives, up to the constant markup,

$$p_{Gj} = Z_j^{-\theta_j}(1+T_j)N_j^{-\chi_j}, \tag{66}$$

so real output per capita is proportional to

$$\frac{G}{L} \propto \frac{Y}{L}\prod_{j=1}^{J}\left[\frac{Z_j^{\theta_j}N_j^{\chi_j}}{1+T_j}\right]^{\phi_j}. \tag{67}$$

Because Y/L is constant in Case I, aggregate growth is

$$g = \sum_{j=1}^{J}\phi_j\left[\theta_j\frac{\dot{Z}_j}{Z_j} + \chi_j\frac{\dot{N}_j}{N_j} - \frac{T_j}{1+T_j}\frac{\dot{T}_j}{T_j}\right]. \tag{68}$$

The first two terms are the productivity innovation and variety innovation contributions generated by each industry's innovation and market-structure dynamics. Since in Case I $\dot{Z}_j/Z_j = \alpha_j L_{Zj}$ and $\dot{N}_j/N_j = \lambda - \dot{x}_j/x_j$, the growth contribution of each industry comes from its own innovation regime (whether productivity innovation is active or not) and transition path due to entry. The last term shows the role of movements in the environmental cost wedges: they introduce a novel policy-induced drag or boost to growth. Industries that are becoming dirtier experience a rising environmental cost wedge that reduces growth, whereas industries that are becoming cleaner experience a declining wedge that raises growth.

This aggregation structure is an extension of the multi-industry growth model in Peretto (2027), where industries with different fundamentals can coexist in different innovation regimes and aggregate growth is the expenditure-share-weighted sum of their contributions. That analysis emphasizes a corresponding policy distinction: policy can operate on the intensive margin by stimulating innovation in industries that already innovate, or on the extensive margin by activating innovation where it is dormant. The environmental setting adds another dimension because productivity and abatement innovation have separate activation thresholds. Environmental policy can therefore affect both the intensity of innovation and its composition across activities and industries.

The environmental side and the growth side moreover aggregate differently. Aggregate emissions depend directly on the levels and movements of the T_j's (equivalently D_j's under a fixed emissions tax), whereas aggregate growth reflects the innovation and market-structure dynamics generated by the individual industries as well as the environmental cost wedges those dynamics generate. Whether these heterogeneous industry transitions produce an aggregate environmental Kuznets curve is consequently an empirical implication of the model, not a restriction imposed by it.

5. Case II: Lump-Sum Redistribution and the General-Equilibrium Feedback

Case II leaves firm technology, pricing, innovation, entry, and the government's resource constraint unchanged. It changes only the route by which tax receipts reach household utility. Because households now receive $\mathcal{T}$ as disposable income, relative emissions-tax burdens affect private expenditure and the common return.

We do not repeat the aggregation exercise of Section 4.8 under this fiscal closure. The purpose of the present section is to isolate how the return of emissions-tax revenues to households changes the transition dynamics through the interest rate and the resulting cross-industry interactions. Aggregate growth and

emissions can be constructed from the resulting industry paths using the same aggregation logic as in Section 4.8. Repeating that exercise would add no new mechanism. We therefore concentrate on the new general-equilibrium propagation mechanism created by the fiscal closure.

5.1. Aggregate expenditure and the common return

Under symmetry within industries, tax revenue from industry j is

$$\mathcal{T}_j = \frac{\epsilon_j - 1}{\epsilon_j} Y_j \frac{T_j}{1 + T_j}. \tag{69}$$

Define the vector of industry relative emissions-tax burdens

$$\mathbf{T} \equiv (T_1, \ldots, T_J). \tag{70}$$

Then

$$\frac{\mathcal{T}}{Y} = \sum_{j=1}^{J} \phi_j \frac{\epsilon_j - 1}{\epsilon_j} \frac{T_j}{1 + T_j}. \tag{71}$$

Substituting $A = \beta Y$ and (7) into (16) gives

$$Y = \frac{L}{D(\mathbf{T})}, \qquad D(\mathbf{T}) \equiv 1 - \beta(\rho - \lambda) - \sum_{j=1}^{J} \phi_j \frac{\epsilon_j - 1}{\epsilon_j} \frac{T_j}{1 + T_j}. \tag{72}$$

Feasibility requires $D(\mathbf{T}) > 0$. Define

$$h_j(\mathbf{T}) \equiv \frac{\phi_j \frac{\epsilon_j - 1}{\epsilon_j} T_j / (1 + T_j)^2}{D(\mathbf{T})}. \tag{73}$$

Since the growth rate of T_j is $\dot{T}_j / T_j$,

$$\frac{\dot{Y}}{Y} = \lambda + \sum_{j=1}^{J} h_j(\mathbf{T}) \frac{\dot{T}_j}{T_j}, \qquad r = \rho + \sum_{j=1}^{J} h_j(\mathbf{T}) \frac{\dot{T}_j}{T_j}. \tag{74}$$

This equation isolates the general-equilibrium feedback. Case I sets the second term to zero because tax receipts never become private expenditure. Case II makes the common return respond to changes in the composition of relative emissions-tax burdens across industries.

5.2. Conditional industry dynamics

To compare the two closures transparently, hold T_k, $k \neq j$, stationary and define

$$(\rho + \delta_j) \equiv \rho + \delta_j, \qquad r + \delta_j = \rho + \delta_j + h_j(\mathbf{T}) \frac{\dot{T}_j}{T_j}. \tag{75}$$

The dormant firm-innovation thresholds and their intersection remain exactly those in (36)–(37), because $\dot{T}_j / T_j = 0$ at activation. The firm-activation condition is therefore unchanged.

Once an innovation margin becomes active, the rebate feedback modifies the speed of adjustment. Productivity-innovation-only dynamics become

$$\left.\frac{\dot{T}_j}{T_j}\right|_Z^{\mathrm{II}} = \frac{\alpha_j \theta_j^2 \frac{\epsilon_j - 1}{\epsilon_j} x_j/(1+T_j) - \theta_j(\rho + \delta_j)}{1 + \theta_j h_j(\mathbf{T})}, \tag{76}$$

whereas abatement-innovation-only dynamics become

$$\left.\frac{\dot{T}_j}{T_j}\right|_B^{\mathrm{II}} = \frac{\eta_j(\rho + \delta_j) - \varsigma_j \eta_j^2 \frac{\epsilon_j - 1}{\epsilon_j} x_j T_j/(1+T_j)}{1 - \eta_j h_j(\mathbf{T})}. \tag{77}$$

With both margins active,

$$\left.\frac{\dot{T}_j}{T_j}\right|_{ZB}^{\mathrm{II}} = \frac{Q_j(x_j, T_j)}{H_j(\mathbf{T})}, \qquad H_j(\mathbf{T}) \equiv 1 - (\eta_j - \theta_j) h_j(\mathbf{T}), \tag{78}$$

where Q_j is the benchmark numerator in (45). The $\dot{T}_j = 0$ locus therefore remains (46): redistribution changes transitional speeds but not the stationary environmental-burden condition.

The firm-size equation becomes

$$\dot{x}_j = F_{0j}(x_j, T_j) + h_j(\mathbf{T}) \left.\frac{\dot{T}_j}{T_j}\right|_{ZB}^{\mathrm{II}} \left(x_j - \frac{\left(\frac{1}{\alpha_j} + \frac{1}{\varsigma_j}\right)}{\beta_j} \right), \tag{79}$$

where F_{0j} is the public-good benchmark in (50). Again the additional term vanishes whenever $\dot{T}_j/T_j = 0$. Thus the two fiscal closures share the same steady-state loci and steady state. They differ in transition dynamics.

This comparison gives the general-equilibrium effect a precise meaning. The rebate does not alter the technological opportunity set or the static incidence of the emissions tax on the firm. It changes household purchasing power. That change moves aggregate expenditure and the return on assets, which in turn changes the hurdle rate for every R&D project and every entry decision.

5.3. Local Stability

Equation (79) can be written

$$\dot{x}_j = F_{0j}(x_j, T_j) + h_j(\mathbf{T}) \left(x_j - \frac{\left(\frac{1}{\alpha_j} + \frac{1}{\varsigma_j}\right)}{\beta_j} \right) \left.\frac{\dot{T}_j}{T_j}\right|_{ZB}^{\mathrm{II}}. \tag{80}$$

At the steady state, the rebate term adds a multiple of the $\dot{T}_j$ row to the $\dot{x}_j$ row of the Jacobian. Relative to a system with the same G_j^B equation and F_{0j} as the firm-size equation, this operation leaves the determinant unchanged and changes the trace. Redistribution can therefore alter the local transition dynamics without moving the steady state.

Let $\mathcal{J}_j^B$ denote the rebate-economy Jacobian. Local stability requires

$$\det \mathcal{J}_j^B > 0, \qquad \operatorname{tr} \mathcal{J}_j^B < 0. \tag{81}$$

Thus the fiscal closure can change the local transition dynamics even though it leaves the long-run steady state unchanged. Appendix B provides a numerical example.

5.4. Regime-wise recurrent dynamics

For the one-industry rebate benchmark, write

$$\left.\frac{\dot{T}_j}{T_j}\right|_{\mathcal{R}}^{\mathrm{II}} = \frac{x_j p_{\mathcal{R}j}(T_j) + d_{\mathcal{R}j}(\rho + \delta_j)}{1 - d_{\mathcal{R}j} h_j(T_j)}. \tag{82}$$

The Dulac multiplier

$$B_{\mathcal{R}j}(x_j, T_j) = \frac{1 - d_{\mathcal{R}j} h_j(T_j)}{x_j T_j} \tag{83}$$

yields

$$\nabla \cdot (B_{\mathcal{R}j} F_{\mathcal{R}j}) = p'_{\mathcal{R}j}(T_j) + \frac{h_j(T_j)}{T_j} p_{\mathcal{R}j}(T_j) - \frac{\varphi_j [1 - d_{\mathcal{R}j} h_j(T_j)] - (\rho + \delta_j) s_{\mathcal{R}j}}{\beta_j x_j^2 T_j}, \tag{84}$$

where $s_{\mathcal{R}j}$ equals 0, $1/\alpha_j$, $1/\varsigma_j$, or $\left(\frac{1}{\alpha_j} + \frac{1}{\varsigma_j}\right)$ according to the active R&D margins. Under transparent sufficient feasibility restrictions, this divergence is negative in every smooth regime. No periodic orbit can remain inside one regime. A full global-convergence proof still requires control of the interface terms generated when a trajectory crosses activation boundaries.

6. Industry Heterogeneity and Endogenous Technological Sorting

The model starts with industries rather than adding them after solving an aggregate representative-firm economy. The activation conditions therefore describe industry-specific technological responses to a common environmental policy.

6.1. Sorting in the public-good benchmark

In Case I, $r = \rho$. Industry j activates productivity R&D when

$$x_j > \frac{(\rho + \delta_j)(1 + T_j)}{\alpha_j \theta_j \frac{\epsilon_j - 1}{\epsilon_j}}, \tag{85}$$

and abatement R&D when

$$x_j > \frac{(\rho + \delta_j)(1 + T_j)}{\varsigma_j \eta_j \frac{\epsilon_j - 1}{\epsilon_j} T_j}. \tag{86}$$

Industries can therefore occupy different innovation regimes despite facing the same statutory emissions tax. Differences in firm scale, technological position, and R&D productivity generate the sorting.

At any date the economy can contain industries with dormant firm innovation, productivity-innovation industries, abatement-innovation industries, and industries using both forms of R&D. As x_j and $T_j = \tau D_j$

evolve, industries cross their own activation boundaries. Aggregate development therefore contains a sequence of industry-level activation events even under an unchanged common policy.

The decomposition $T_j = \tau D_j$ lets us state the clean–dirty comparison precisely. Under the common emissions tax maintained here, T_j and D_j generate exactly the same cross-industry ranking. We call an industry relatively *dirty* when its D_j, emissions relative to variable production cost, is relatively high, and relatively *clean* when it is relatively low. If all industries have reached their steady states, the economy therefore exhibits a ranking according to $D_j^* = T_j^*/\tau$. Each industry's fundamental dirtiness is determined by its industry-specific deep parameters and the maintained policy environment. Clean and dirty industries are equilibrium outcomes rather than exogenously assigned technological types. But the steady-state ranking cannot in general be inferred from a cross section of current dirtiness $D_j(t)$ because industries can occupy different positions along their transition paths.

6.2. Sorting, propagation, and the rank-one interaction

Case II adds an economy-wide interaction even when every industry remains in free-entry equilibrium. The broader literature on environmental tax reform has long emphasized that the use of environmental-tax revenue can change the macroeconomic effects of policy; see, for example, Bovenberg and de Mooij (1997). Here the recycling rule has a particularly transparent role: it creates a common propagation channel across otherwise technologically distinct industries. Case II preserves the steady-state ranking—indeed, it preserves the same industry steady states as Case I—because all relative emissions-tax-burden growth rates $\dot{T}_j/T_j$ vanish in steady state and the household's reservation return on saving returns to $r = \rho$. The transition, however, is fundamentally different. Away from the steady state, the evolution of each industry's relative emissions-tax burden contains a common macroeconomic term generated by the lump-sum rebate of emissions-tax revenues. Changes in one industry affect household income and saving and thereby the reservation return faced by firms in every industry. Fiscal recycling therefore leaves the destination unchanged but links the paths that industries follow to reach it. For each industry and each innovation regime $\mathcal{R}_j \in \{0, Z, B, ZB\}$, write

$$q_j = M_j^{\mathcal{R}_j}(x_j, T_j) + d_j^{\mathcal{R}_j}(r + \delta_j), \tag{87}$$

with

$$(M_j^0, d_j^0) = (0, 0),$$

$$M_j^Z = \frac{\alpha_j \theta_j^2 \frac{\epsilon_j - 1}{\epsilon_j} x_j}{1 + T_j}, \qquad d_j^Z = -\theta_j,$$

$$M_j^B = -\frac{\varsigma_j \eta_j^2 \frac{\epsilon_j - 1}{\epsilon_j} x_j T_j}{1 + T_j}, \qquad d_j^B = \eta_j,$$

and

$$M_j^{ZB} = \frac{\frac{\epsilon_j - 1}{\epsilon_j} x_j}{1 + T_j}(\alpha_j \theta_j^2 - \varsigma_j \eta_j^2 T_j), \qquad d_j^{ZB} = \eta_j - \theta_j.$$

Thus the representation does not require all industries to operate both R&D margins.

For this cross-industry calculation only, let $\mathbf{q}$ collect the dirtiness (equivalently, relative emissions-tax-

burden) growth rates $q_j \equiv \dot{T}_j/T_j$, and let $\mathbf{d}$ and $\mathbf{h}$ collect $d_j^{\mathcal{R}_j}$ and $h_j(\mathbf{T})$. Define

$$\widetilde{M}_j \equiv M_j^{\mathcal{R}_j} + d_j^{\mathcal{R}_j}\delta_j.$$

Then

$$\mathbf{q} = \widetilde{\mathbf{M}} + \mathbf{d}r, \qquad r = \rho + \mathbf{h}'\mathbf{q}. \tag{88}$$

Substitution gives

$$\mathbf{q} = \widetilde{\mathbf{M}} + \rho\mathbf{d} + \mathbf{d}\mathbf{h}'\mathbf{q}, \tag{89}$$

and therefore

$$r = \frac{\rho + \mathbf{h}'\widetilde{\mathbf{M}}}{1 - \mathbf{h}'\mathbf{d}}. \tag{90}$$

The regularity condition $1 - \mathbf{h}'\mathbf{d} > 0$ keeps this aggregate feedback finite.

Equation (89) gives the interaction a useful structural interpretation. The matrix $\mathbf{d}\mathbf{h}'$ is an outer product and therefore has rank one. Technologically, the economy remains genuinely multi-dimensional: each industry has its own firm scale, dirtiness, relative emissions-tax burden, and innovation regime. General equilibrium links those industry systems through only one propagation channel, the common market return. A disturbance originating in industry k changes aggregate tax revenue and private expenditure; the induced movement in r then changes innovation incentives in every other industry. Low-rank interaction structures and their principal components play a central role in modern network analysis; see, for example, Galeotti, Golub, and Goyal (2020). Here the rank-one structure arises endogenously from the fiscal and asset-market closure rather than from an assumed network of bilateral links.

The formula also makes propagation signs transparent. Holding the state vector fixed,

$$\frac{\partial r}{\partial M_k} = \frac{h_k}{1 - \mathbf{h}'\mathbf{d}}, \qquad \frac{\partial q_j}{\partial M_k} = \mathbf{1}\{j = k\} + \frac{d_j h_k}{1 - \mathbf{h}'\mathbf{d}}. \tag{91}$$

For a receiving industry $j \neq k$, the sign of the directional response is therefore the sign of d_j. A higher common return lowers q_j in a productivity-innovation-only industry, raises it in an abatement-innovation-only industry (making its decline in T_j less rapid), and has sign $\eta_j - \theta_j$ when both margins operate. Industries with dormant firm innovation have $d_j = 0$ for the directional state while the common return changes their latent innovation returns and can move an activation boundary across current firm size.

The effect on innovation intensity is simpler. Conditional on remaining in a given active regime, a higher r lowers each active R&D labor demand: $\partial L_{Zj}/\partial r = -1/\alpha_j$ and $\partial L_{Bj}/\partial r = -1/\varsigma_j$ whenever the corresponding margin is active. Fiscal propagation therefore depresses total R&D when it raises the common return, while its effect on the *direction* of technological change depends on the receiving industry's regime and elasticities. General-equilibrium contagion can consequently alter both how much an industry innovates and the composition of that innovation.

The same endogenous return shifts activation thresholds. Productivity R&D operates when

$$x_j > \frac{(r + \delta_j)(1 + T_j)}{\alpha_j \theta_j \frac{\epsilon_j - 1}{\epsilon_j}}, \tag{92}$$

and abatement R&D operates when

$$x_j > \frac{(r+\delta_j)(1+T_j)}{\varsigma_j \eta_j \frac{\epsilon_j - 1}{\epsilon_j} T_j}. \tag{93}$$

A disturbance in one industry can therefore affect another industry on two margins. It changes the intensity and direction of investment when R&D is already active, and it changes latent innovation returns when firm innovation is dormant. If the induced movement is large enough, it activates or deactivates an investment margin.

This propagation mechanism differs from the secular transition dynamics associated with inactive entry in Peretto (2027). There, the transition from a pre-industrial regime to the industrial regime makes household wealth and expenditure additional channels through which an industry disturbance spreads. Here fiscal recycling creates propagation even while every industry satisfies free entry and aggregate wealth retains the canonical relation $A = \beta Y$. Tax revenue is therefore an additional general-equilibrium transmission mechanism, not a consequence of a breakdown of free entry.

6.3. The role of free entry

The rank-one reduction above exploits the maintained assumption that entry remains active in every industry. If entry becomes inactive in industry j, then $V_j < \beta_j x_j$ and aggregate wealth no longer satisfies $A = \beta Y$. Household wealth becomes an independent state variable, so the common return cannot be recovered from the vector of relative emissions-tax burdens alone. The industry remains a technological object, but its dynamics must then be solved jointly with household wealth and the rest of the economy. Peretto (2027) provides examples of the richer transition dynamics that arise when entry is inactive.

7. Interpretation: Lessons for Environmental Policy

A common emissions tax does not imply a common technological response. Policy sets the price attached to emissions, while firms' productivity and abatement decisions determine the dirtiness of production. Endogenous firm scale determines when an innovation margin becomes viable, while the relative emissions-tax burden $T_j = \tau D_j$ determines its direction. Public-good provision displays that mechanism in isolation. Lump-sum redistribution then shows how the use of tax revenue feeds the same decentralized decisions back through private expenditure and asset prices.

The broader implication is that the steady-state ranking is not itself the main object of interest. Rather, it provides a benchmark for interpreting observed differences across industries. The distinction between current and fundamental dirtiness is analogous to the familiar distinction in growth theory between an economy's current position and the long-run position implied by its fundamentals. The conditional-convergence literature emphasizes that observed income reflects both determinants of the steady state and the economy's position along its transition path; see Barro and Sala-i-Martin (1992). The same logic applies here. Current dirtiness $D_j(t)$ combines structural characteristics with transitional information, whereas fundamental dirtiness D_j^* describes the attractor selected by the industry's deep technological and economic parameters and the maintained policy environment.

In short, *dirtiness is a state, not a type.* Labeling an industry as intrinsically clean or dirty on the basis of its current $D_j(t)$ treats a transitional outcome as a structural characteristic. Under Case I, current dirtiness reflects both industry-specific fundamentals and the industry's position along its transition path. Under Case II, its path also reflects the evolution of other industries through the common macroeconomic effect of

lump-sum recycling. Current and fundamental clean–dirty rankings therefore need not coincide: an industry that is relatively dirty today need not occupy the same relative position once the industries have completed their transitions. Observed dirtiness by itself reveals neither the industry's fundamental dirtiness nor the structural forces responsible for it. This distinction complements approaches that classify technologies or sectors as clean or dirty ex ante, such as Acemoglu et al. (2012): here clean and dirty are endogenous positions along a dynamic path as well as long-run equilibrium outcomes. This interpretation extends beyond the environmental application. More generally, development can activate latent investment opportunities by changing the scale over which their returns are realized.

This two-closure comparison also clarifies the paper's general-equilibrium channels. The two cases use the same final good, collect the same emissions tax, and differ only in whether the government allocates the receipts directly to public consumption or returns them to households as disposable income. General equilibrium does not mean that every technological state must become an aggregate state. It means that decentralized decisions satisfy the economy-wide budget and asset-market restrictions generated by the fiscal institution. Changing the route followed by the same resources leaves the innovation technologies untouched but changes the circulation of purchasing power and therefore the prices through which industries interact.

8. Conclusion

The central result is about the endogenous sequence through which firms respond to environmental incentives. Economic development changes firm scale and thereby changes which latent innovation opportunities are privately viable. With multiple innovation margins, the model characterizes analytically the thresholds, ordering, turning points, and timing through which an environmental incentive becomes innovation, innovation becomes cleaner production, and cleaner production eventually becomes lower emissions. Environmental policy can therefore be in place well before abatement innovation begins, and abatement innovation can begin well before production actually becomes cleaner.

The two fiscal closures separate this firm-level mechanism from its general-equilibrium propagation. Direct public-good provision isolates the activation dynamics, while lump-sum recycling links industries through household expenditure and the common market return without changing the steady-state destination. Aggregation preserves the distinction between the industry transitions that generate environmental outcomes and those that generate growth. Observed environmental performance consequently reflects both structural fundamentals and position along the transition: dirtiness is a state, not a type.

Although the analysis focuses on environmental innovation, the underlying mechanism is not specific to environmental policy. Development changes the scale over which the returns to investment are realized and can therefore activate opportunities that are dormant at earlier stages. The environmental application makes the mechanism especially transparent because the activation of abatement innovation, the onset of cleaner production, and the decline of emissions are distinct endogenous events.

A. Derivation of Firm Behavior and Entry

This appendix derives the conditions stated in Section 3. Firm i in industry j faces the demand curve (11) and chooses its price and R&D employment taking industry aggregates and factor prices as given.

A.1. Pricing and operating decisions

Using (18), the firm's operating profit before overhead and R&D costs is

$$\left[p_{Xij} - wZ_{ij}^{-\theta_j} - \tau B_{ij}^{-\eta_j}\right] X_{ij}.$$

The demand elasticity is ϵ_j. The first-order condition for price therefore gives the standard constant-markup rule

$$p_{Xij} = \frac{\epsilon_j}{\epsilon_j - 1}\left(wZ_{ij}^{-\theta_j} + \tau B_{ij}^{-\eta_j}\right),$$

which is (19). With $w = 1$ and symmetry, $p_{Xj}X_j = x_j$. Combining the pricing rule with $D_j = Z_j^{\theta_j} B_j^{-\eta_j}$ and $T_j = \tau D_j$ give

$$L_{Xj} - \varphi_j = \frac{\epsilon_j - 1}{\epsilon_j}\frac{x_j}{1 + T_j}, \qquad \tau E_j = \frac{\epsilon_j - 1}{\epsilon_j}\frac{x_j T_j}{1 + T_j},$$

as stated in (21).

A.2. Returns to productivity and abatement R&D

Let the current-value shadow values of Z_{ij} and B_{ij} be μ_{Zij} and μ_{Bij}. Since one unit of R&D labor raises the corresponding knowledge stock at rates $\alpha_j Z_j$ and $\varsigma_j B_j$, the interior first-order conditions are

$$1 = \mu_{Zij}\alpha_j Z_j, \qquad 1 = \mu_{Bij}\varsigma_j B_j.$$

The costate equations equate the capital gain plus the current contribution of knowledge to the required return $r + \delta_j$. Evaluating them in a symmetric equilibrium and using the operating-cost shares above yields

$$r + \delta_j = \alpha_j\left[\theta_j\frac{\epsilon_j - 1}{\epsilon_j}\frac{x_j}{1 + T_j} - L_{Zj}\right]$$

and

$$r + \delta_j = \varsigma_j\left[\eta_j\frac{\epsilon_j - 1}{\epsilon_j}\frac{x_j T_j}{1 + T_j} - L_{Bj}\right].$$

These are (25) and (26). Imposing $L_{Zj} \geq 0$ and $L_{Bj} \geq 0$ gives the two-branch investment rules (27)–(28).

A.3. Entry and the evolution of firm size

An entrant spends $\beta_j x_j$ units of labor and enters with the industry-average knowledge stocks. Active free entry therefore implies

$$V_j = \beta_j x_j.$$

The return on the firm is the dividend yield plus the capital gain. Under the markup rule, operating profit net of overhead and R&D labor is

$$\Pi_j = \frac{x_j}{\epsilon_j} - \varphi_j - L_{Zj} - L_{Bj}.$$

Hence

$$r + \delta_j = \frac{\Pi_j}{V_j} + \frac{\dot{V}_j}{V_j} = \frac{1}{\beta_j}\left[\frac{1}{\epsilon_j} - \frac{\varphi_j + L_{Zj} + L_{Bj}}{x_j}\right] + \frac{\dot{x}_j}{x_j},$$

which is (30). Rearranging gives (31).

B. A Numerical Example of Local Stability

To illustrate the local dynamics under lump-sum redistribution, set $\delta_j = 0$ and

$$\alpha_j = .999, \quad \varsigma_j = .379, \quad \theta_j = .0918, \quad \eta_j = .0723, \quad \rho = .0502,$$

$$\beta_j = 1.509, \quad \epsilon_j = 6.10, \quad \varphi_j = .431.$$

The one-industry rebate benchmark has

$$T_j^* \simeq 3.96079, \qquad x_j^* \simeq 10.15581,$$

with eigenvalues approximately

$$-0.014942 \pm 0.002524i.$$

This parameterization illustrates that the general-equilibrium feedback can affect the local shape of the transition without moving the steady state.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hémous (2012), "The Environment and Directed Technical Change," *American Economic Review* 102(1), 131–166.

Aghion, Philippe, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen (2016), "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry," *Journal of Political Economy* 124(1), 1–51.

Andreoni, James, and Arik Levinson (2001), "The Simple Analytics of the Environmental Kuznets Curve," *Journal of Public Economics* 80(2), 269–286.

Barro, Robert J., and Xavier Sala-i-Martin (1992), "Convergence," *Journal of Political Economy* 100(2), 223–251.

Barro, Robert J., and Xavier Sala-i-Martin (2004), *Economic Growth*, 2nd ed., MIT Press.

Bovenberg, A. Lans, and Ruud A. de Mooij (1997), "Environmental Tax Reform and Endogenous Growth," *Journal of Public Economics* 63(2), 207–237.

Brown, James R., Gustav Martinsson, and Christian Thomann (2022), "Can Environmental Policy Encourage Technical Change? Emissions Taxes and R&D Investment in Polluting Firms," *Review of Financial Studies* 35(10), 4518–4560.

Brock, William A., and M. Scott Taylor (2005), "Economic Growth and the Environment: A Review of Theory and Empirics," in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 1B, Elsevier, 1749–1821.

Calel, Raphael, and Antoine Dechezleprêtre (2016), "Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market," *Review of Economics and Statistics* 98(1), 173–191.

Chakraborty, Pavel (2017), "Environmental Standards, Trade and Innovation: Evidence from a Natural Experiment," *Environment and Development Economics* 22(4), 414–446.

Cohen, Wesley M. (2010), "Fifty Years of Empirical Studies of Innovative Activity and Performance," in Bronwyn H. Hall and Nathan Rosenberg, eds., *Handbook of the Economics of Innovation*, Vol. 1, Elsevier, 129–213.
Cohen, Wesley M., and Steven Klepper (1996), "A Reprise of Size and R&D," *Economic Journal* 106(437), 925–951.
Cohen, Wesley M., and Richard C. Levin (1989), "Empirical Studies of Innovation and Market Structure," in Richard Schmalensee and Robert D. Willig, eds., *Handbook of Industrial Organization*, Vol. 2, North-Holland, 1059–1107.
Dechezleprêtre, Antoine, and Misato Sato (2017), "The Impacts of Environmental Regulations on Competitiveness," *Review of Environmental Economics and Policy* 11(2), 183–206.
de Miguel, Carlos, and Consuelo Pazó (2017), "Environmental Protection, Innovation and Price-Setting Behavior in Spanish Manufacturing Firms," *Energy Economics* 68, 116–124.
Dumortier, Freddy, Jaume Llibre, and Joan C. Artés (2006), *Qualitative Theory of Planar Differential Systems*, Springer.
Galeotti, Andrea, Benjamin Golub, and Sanjeev Goyal (2020), "Targeting Interventions in Networks," *Econometrica* 88(6), 2445–2471.
Grossman, Gene M., and Alan B. Krueger (1995), "Economic Growth and the Environment," *Quarterly Journal of Economics* 110(2), 353–377.
Horbach, Jens (2008), "Determinants of Environmental Innovation—New Evidence from German Panel Data Sources," *Research Policy* 37(1), 163–173.
Jaffe, Adam B., Richard G. Newell, and Robert N. Stavins (2002), "Environmental Policy and Technological Change," *Environmental and Resource Economics* 22, 41–70.
Ma, Clara, Cristina Peñasco, and Laura Díaz Anadón (2025), "Technology Innovation and Environmental Outcomes of Road Transportation Policy Instruments," *Nature Communications* 16, 4467.
Fried, Stephie (2018), "Climate Policy and Innovation: A Quantitative Macroeconomic Analysis," *American Economic Journal: Macroeconomics* 10(1), 90–118.
Linn, Joshua (2008), "Energy Prices and the Adoption of Energy-Saving Technology," *Economic Journal* 118(533), 1986–2012.
Newell, Richard G., Adam B. Jaffe, and Robert N. Stavins (1999), "The Induced Innovation Hypothesis and Energy-Saving Technological Change," *Quarterly Journal of Economics* 114(3), 941–975.
Peretto, Pietro F., Bentley Coffey, and Patrick McLaughlin (2026), "Cost Push and the Direction of Endogenous Innovation: Transportation, Regulation, and Scarcity," working paper.
Proost, Stef, and Kurt Van Dender (2012), "Energy and Environment Challenges in the Transport Sector," *Economics of Transportation* 1(1–2), 77–87.
Peretto, Pietro F. (2008), "Effluent Taxes, Market Structure, and the Rate and Direction of Endogenous Technological Change," *Environmental and Resource Economics* 39(2), 113–138.
Peretto, Pietro F. (2009), "Energy Taxes and Endogenous Technological Change," *Journal of Environmental Economics and Management* 57(3), 269–283.
Peretto, Pietro F. (2027), *Knowledge, Industry, and the Endogenous Growth of Nations*, MIT Press.
Popp, David (2002), "Induced Innovation and Energy Prices," *American Economic Review* 92(1), 160–180.
Qi, Ji, Xin Tang, and Xican Xi (2021), "The Size Distribution of Firms and Industrial Water Pollution: A Quantitative Analysis of China," *American Economic Journal: Macroeconomics* 13(1), 151–183.
Rubio, Santiago J., José L. García, and José L. Hueso (2009), "Neoclassical Growth, Environment and

Technological Change: The Environmental Kuznets Curve," *Energy Journal* 30(2 Suppl.), 143–168.
Shapiro, Joseph S., and Reed Walker (2018), "Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade," *American Economic Review* 108(12), 3814–3854.
Smulders, Sjak (1995), "Environmental Policy and Sustainable Economic Growth: An Endogenous Growth Perspective," *De Economist* 143(2), 163–195.
Smulders, Sjak, Lucas Bretschger, and Hannes Egli (2011), "Economic Growth and the Diffusion of Clean Technologies: Explaining Environmental Kuznets Curves," *Environmental and Resource Economics* 49(1), 79–99.
Smulders, Sjak, and Michiel de Nooij (2003), "The Impact of Energy Conservation on Technology and Economic Growth," *Resource and Energy Economics* 25(1), 59–79.
Symeonidis, George (1996), "Innovation, Firm Size and Market Structure: Schumpeterian Hypotheses and Some New Themes," OECD Economics Department Working Papers No. 161, OECD Publishing.
Stokey, Nancy L. (1998), "Are There Limits to Growth?" *International Economic Review* 39(1), 1–31.