



Programme on
Innovation and Diffusion

Working Paper

No.150
September 2026

A theory of endogenous degrowth and environmental sustainability

Philippe Aghion
Timo Boppart
Michael Peters
Matthew Schwartzman
Fabrizio Zilibotti

Abstract

We develop and quantify a theory in which rising incomes and technological change redirect economic activity from material production toward intangible quality, generating an endogenous decline in emissions intensity. We document two motivating empirical patterns: greater service intensity is systematically associated with lower emissions per dollar, and U.S. patenting has shifted from innovations that raise material productivity toward those that enhance final use quality. Nonhomothetic preferences shift expenditure toward quality-intensive goods; physical inputs and producer services are complements; and innovative effort is endogenously allocated between material productivity and final-use quality. In equilibrium, the economy becomes “weightless”: innovation becomes increasingly quality-oriented, services dominate employment, and material production converges to a finite level. At the same time, welfare continues to grow because of rising quality. The calibrated model accounts for the rise of services, the measured productivity slowdown, and the long-run path of CO₂ emissions. We use it to evaluate the aggregate and distributional welfare effects of environmental policies that redirect innovation and accelerate structural transformation.

Keywords: Directed technical change, Endogenous degrowth, Environmental degradation, Quality vs. quantity, Structural change, Weightless economy.

We thank Zhang Chen, Maarten De Ridder, Pete Klenow, Morten Olsen, Kjetil Storesletten, Chris Tonetti, and audiences at numerous conferences and seminars for helpful comments. We are especially grateful to David Hémous and Karthik Sastry for their discussions at the NBER Growth and EFG meetings, respectively. We are indebted to Nicolas Fajardo for outstanding research assistance. For financial support, Aghion thanks the SCOR Foundation for Science; Peters and Zilibotti thank the Tobin Center. For hospitality, Zilibotti thanks the University of Copenhagen and Danmarks Nationalbank; Peters thanks the University of Zurich. The views expressed herein are those of the authors and do not necessarily reflect the views of the Programme on Innovation and Diffusion (ES/V009478/1).

Philippe Aghion, College de France, INSEAD and LSE; Timo Boppart, University of Zurich and IIES, Stockholm University; Michael Peters, Yale University; Matthew Schwartzman, University of Michigan and Fabrizio Zilibotti, Yale University.

Published by
Programme on Innovation and Diffusion
London School of Economics and Political Science
Houghton Street
London WC2A 2AE

This paper was produced as part of the Programme on Innovation Diffusion (POID) financed by the Economic and Social Research Council.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

© P. Aghion, T. Boppart, M. Peters, M. Schwartzman and F. Zilibotti, submitted 2026.

1. INTRODUCTION

Can sustained growth be environmentally sustainable? Whether economic growth can continue indefinitely without degrading the ecological systems on which human welfare depends is the subject of a long-standing debate. Production and consumption require energy, materials, and the capacity of the environment to absorb waste. If rising living standards remain tied to an ever-expanding throughput of matter and energy, economic growth must eventually confront ecological limits.

This concern has led scholars, activists, and policymakers to advocate *degrowth*: a deliberate reduction in material production and consumption aimed at bringing ecological pressures within sustainable limits. While mainstream economists argue that innovation, substitution, and well-designed environmental policies can progressively weaken the link between economic activity and ecological pressure, degrowth proponents question whether absolute decoupling can ever be achieved. They maintain that the benefits of cleaner technologies will ultimately be offset by ever-increasing material throughput and waste.

In this paper, we argue that such decoupling is feasible if prosperity rises through improvements in quality rather than increases in quantity. By quantity, we mean the material component of consumption; by quality, we mean attributes of goods and services such as design, convenience, healthfulness, and user experience that increase utility per unit of material throughput. While in most existing theories of growth, the distinction between quality and quantity is often a matter of semantics, the difference becomes crucial once one considers the environmental dimension. Because the environmental impact of producing two cars, two phones, or two restaurant meals is higher than that of producing one car, one phone, or one meal that delivers twice the enjoyment, sustained improvements in living standards can be environmentally sustainable if they are driven by higher quality. In other words, material production can undergo “degrowth” even as a welfare-relevant measure of GDP continues to expand.

We make three contributions. First, we provide empirical evidence on these forces. Second, we develop a model of endogenous innovation in which nonhomothetic demand and technological forces redirect research along a development path progressively away from increasing material productivity and toward quality improvement, making economies increasingly *weightless* (Quah, 1999). Third, we quantify the model and analyze policy coun-

terfactuals, identifying distributional forces that may generate political opposition to policies that accelerate the transition to quality-led growth.

The key mechanism in our theory is that gains in living standards become progressively less tied to material production along the course of economic development. On the demand side, nonhomothetic preferences shift expenditure toward goods and services for which consumer-facing quality matters more than physical quantity. On the production side, gains in material productivity lower physical-input prices and, because physical inputs and producer services are complements, reduce the material-input cost share, shrinking the market for further material-productivity innovation. Together, these forces progressively redirect innovation toward quality. In the long run, material-productivity innovation ceases to drive rising living standards: material production converges to a finite level, while quality-adjusted consumption and welfare continue to rise. Once material production stops expanding, any sustained improvement in green technology drives flow emissions toward zero. The theory therefore predicts eventual absolute decoupling, but not necessarily in time to avert severe environmental damage: if emissions remain unchecked during the transition, cumulative pollution may cross critical levels.

To quantify the model, we parameterize preferences and technology to match key features of structural change in the U.S. economy. We discipline nonhomotheticity using income elasticities estimated from the Consumer Expenditure Survey (CEX). The model reproduces the secular rise in the service employment share, the recent productivity slowdown, the relative-price trend between goods- and service-intensive industries, and the long-run path of CO₂ emissions. In the calibrated model, the mechanism of endogenous quality-led growth is strong enough that sustained growth can be squared with environmental sustainability.

The theory predicts a hump-shaped path for total emissions: emissions initially rise, then peak and decline, if green technological progress reduces emissions per unit of material output at any positive rate. This eventual decoupling does not amount to a plea for policy inaction: the transition may be too slow to avert severe ecological damage while emissions are rising. Pigouvian taxation remains necessary to correct for an environmental externality and can accelerate the shift from material production to quality-led growth.

We use the calibrated model to evaluate one such policy: a revenue-neutral 20 percent tax on physical inputs that finances subsidies to quality provision. The policy accelerates

quality-led growth through both a static reallocation away from material-intensive activities and a dynamic redirection of innovation toward quality. The short-run costs of the policy fall disproportionately on lower-income households, because material-intensive necessities account for a larger share of their budgets. This unequal incidence may generate political opposition, suggesting a role for targeted transfers or rebates.

Is degrowth necessary? Our analysis suggests that the answer crucially depends on the nature of growth along the quantity-quality dimension. Our calibration makes us somewhat optimistic, because our model predicts that material production is going to plateau, while quality-adjusted consumption and welfare continue to rise. Because in practice quality improvements are imperfectly measured, this transition may appear as slower GDP growth even as living standards continue to improve.

Related literature Our study contributes to several strands of research. We build on the literature examining the macroeconomic consequences of climate change (see, e.g., [Golosov et al. \(2014\)](#)). This literature does not distinguish between quality and quantity growth and typically treats the direction of innovation as exogenous. More closely related to our analysis is the literature on directed technical change. The seminal model of [Acemoglu et al. \(2012\)](#) shows how policy can steer innovation toward cleaner production. Several papers extend this framework; see, e.g., [Acemoglu et al. \(2016\)](#), [Aghion et al. \(2024\)](#), [Hémous \(2016\)](#), and [Hémous and Olsen \(2021\)](#) for a survey. This literature focuses on redirecting innovation across production technologies, whereas we distinguish innovation in material productivity from innovation in quality provision.

Our paper also relates to the debate on degrowth, whose proponents argue that continued expansion of material and energy throughput is incompatible with ecological limits ([Hickel et al., 2022](#)). Foundational contributions to this literature emerged in the 1970s. In an influential report, the Club of Rome ([Meadows et al., 1972](#)) warned that continued industrial expansion could lead to ecological collapse. [Georgescu-Roegen \(1971\)](#) argued that material production irreversibly transforms low-entropy matter and energy into high-entropy waste, thereby challenging the possibility that material growth can ever be fully decoupled from resource depletion and environmental degradation. We contribute to this debate by formalizing an endogenous transition in which material production stabilizes while intangible innovation and welfare can grow unboundedly.

By highlighting the link between the provision of quality and services, we contribute to the literature on structural change and service-led growth. [Boppart \(2014\)](#) features nonhomothetic preferences in the PIGL class, which we adopt. As in [Boppart \(2014\)](#), [Alder et al. \(2022\)](#), and [Fan et al. \(2023\)](#), we fit these preferences to data, and exploit their aggregation properties in our normative analysis. An alternative specification with similar implications for structural change is studied in [Comin et al. \(2021\)](#). We contribute to this literature by linking service-led structural transformation to the debate on environmental sustainability.

Our distinctive departure from canonical models of structural change is to explicitly introduce the notion of quality that is particularly valued by rich households. [Bils and Klenow \(2001\)](#) show that rich households allocate a larger share of spending to higher-quality versions of goods, and [Faber and Fally \(2022\)](#) document that rich consumers buy from larger firms and are more sensitive to quality. Related evidence from international trade also highlights that quality is income-elastic; see, e.g., [Fieler \(2011\)](#), [Fajgelbaum et al. \(2011\)](#), and [Hallak and Schott \(2011\)](#).

Our empirical analysis is related to work on environmental Engel curves, which documents how the pollution embodied in household consumption varies with income ([Caron and Fally, 2022](#), [Levinson and O'Brien, 2019](#)). We link these patterns to service and quality intensity and embed them in a general equilibrium model of technical change.

Finally, our measurement discussion relates to research on productivity measurement and the recent slowdown of growth ([Byrne et al., 2016](#), [Syverson, 2017](#)). We provide a mechanism through which innovation may shift over time toward dimensions of quality that are difficult to capture in standard productivity statistics. [De Ridder and Rachel \(2025\)](#) develop an adjusted measure of total factor productivity that incorporates carbon emissions. Their contribution complements ours: our theory explains how endogenous shifts in innovation can promote sustainable growth that standard growth accounting overlooks.

The paper is organized as follows. Section 2 provides empirical motivation. Section 3 develops the theory. Section 4 contains the quantitative analysis. Section 5 discusses policy implications. Section 6 discusses international aspects. Section 7 concludes. The Online Appendix (henceforth, the “Appendix”) contains data details and technical results.

2. EMPIRICAL MOTIVATION

In this section, we present several facts that motivate our theory. We empirically associate material production with the combined output of agriculture and industry and use service value added as a proxy for intangible activity. Because emissions arise primarily from material inputs and the energy used in their physical transformation, service intensity provides an observable measure of the extent to which value creation is tied to material throughput.¹ Details on the data construction are provided in Section 4.1 and Appendix A.1.


Material Degrowth and Emissions. We begin by documenting a systematic correlation between the importance of tangible production and emissions intensity, i.e., emissions per dollar of value added, over time and across industries, countries, U.S. counties, and households.

Panel (a) of Figure 1 depicts two long-run trends in the U.S. economy: a fivefold decline in emissions intensity (red) since 1917 and a sustained increase in the service employment share (green), from about 30% in the early twentieth century to over 80% in recent years. Total CO₂ emissions (blue) have increased substantially over the past century. Yet, over the last three decades they have plateaued and, in fact, declined since 2008. These aggregate trends are consistent with the structural transformation emphasized by our theory: as economic activity shifts away from material production and toward intangible services, emissions intensity falls, slowing down overall emissions.

An immediate concern is that these patterns might reflect international carbon leakage, with the U.S. outsourcing emissions-intensive material production to other countries. Figure 1 shows why this mechanism is unlikely to be quantitatively important to explain the trends. First, the decline in U.S. production-based emissions intensity is secular: it begins well before the acceleration of global value-chain integration in the late twentieth century and shows no sharp break at the outset of globalization. Second, Panel (a) also reports consumption-based emissions (dashed lines), which adjust for emissions embodied in trade. Although these data are available only from 1990 onward, consumption-based emissions

¹This interpretation is in line with the OECD/Eurostat *Oslo Manual 2018*, which characterizes services as intangible activities that are produced and consumed simultaneously. Statistical categories of final products generically embody both goods and service inputs. Some services are therefore emissions-intensive because they rely heavily on material inputs: air transportation, for example, requires aircraft and fuel. Our theory and empirical analysis therefore classify expenditure categories according to the goods and service shares of their value added.

FIGURE 1.—MATERIAL DEGROWTH AND EMISSIONS



Notes: Panel (a) shows total annual CO₂ emissions (blue lines), CO₂ emissions relative to GDP (red lines), and the U.S. service employment share (green line). Solid lines report production-based measures, and dashed lines report consumption-based measures; the emissions series are normalized to one in 2000. Panel (b) plots CO₂ emissions per U.S. dollar of GDP against real GDP per capita in PPP-adjusted U.S. dollars for China, India, the United States, the European Union (EU-27), OECD countries, and the world from 1970 to 2024. Panel (c) plots industry-level CO₂ emissions intensity against the cost share of service inputs. Panel (d) reports the service value-added share embodied in household consumption by income group in 2002, 2012, and 2017. Panel (e) plots the relationship between household income and CO₂ emissions per dollar of expenditure, before and after controlling for the service value-added share of expenditure. Panel (f) reports bin-weighted tailpipe CO₂ emissions per mile per dollar of vehicle price in 2022, normalized to one for the lowest-income group. Appendix A.1 provides construction details for Panel (f).

track production-based emissions closely. Third, as shown in Panel (b), declining emissions intensity is not unique to the U.S. but is a robust feature of economic development. Emissions per unit of output have fallen dramatically not only in advanced economies but also in emerging and developing economies such as China and India. This pattern is at odds with the view that falling emissions intensity in rich economies is an artifact of offshoring, because it would imply persistently high or rising emissions intensity in the main destinations of outsourced manufacturing. We discuss global trends in total emissions in more detail in Section 6.

The negative association between emissions intensity and the rise of intangible services appears not only in time-series data but also in the cross section. Panel (c) relates total emissions per dollar of sales to each industry's service intensity, measured as the cost share of service inputs required to produce one dollar of output. For both measures, we account for sectoral linkages through the input–output matrix and hence capture both the service content and emissions embedded in intermediate inputs. A strong relationship emerges: a 10-percentage-point increase in the service share is associated, on average, with a 19.5% reduction in emissions per dollar, consistent with the idea that emissions are generated primarily by the material transformation of physical output rather than by intangible services.

A similar relationship holds geographically, both across countries and across U.S. counties. As detailed in Appendix A.2, across countries a 1-percentage-point employment shift from manufacturing to services is associated with a 5% decline in emissions intensity. The pattern is robust to controlling for overall economic development and country fixed effects, suggesting that the reallocation away from material production contributes to the reduction in emissions intensity beyond other correlates of development, such as richer nations adopting environmentally friendly technologies. A similar relationship also holds across U.S. counties.

Finally, we turn to household consumption. Using Consumer Expenditure Survey (CEX) data, we compute the service value-added share embodied in each dollar of consumer spending across income bins. Panel (d) of Figure 1 highlights two facts. First, service-intensive consumption goods are luxuries: richer households devote a larger share of expenditure to goods with higher service content. Second, holding income fixed, the service content of spending has increased over time. For example, among households earning \$80,000–\$100,000, the service share rose from about 67% in 2002 to 75% in 2017. These

patterns point to two tenets of our theory. First, demand is nonhomothetic: as income rises, expenditure shifts toward products that are less intensive in material production and more intensive in intangible attributes such as quality. Second, as emphasized by [Baumol \(1967\)](#), rising material productivity causes the cost share of services to increase.

Panels (e) and (f) link these consumption patterns to emissions. In Panel (e), we combine industry-level emissions data with households' consumption baskets to compute emissions per dollar of spending. The blue line shows the raw environmental Engel curve, following [Levinson and O'Brien \(2019\)](#): richer households pollute more in levels, but less per dollar of expenditure. The red line shows the residual relationship after controlling for the service value-added share of the consumption basket. The relationship is substantially flatter, indicating that the lower emissions intensity of high-income households is largely accounted for by their greater service intensity.

Because the environmental Engel curves in Panel (e) are estimated using the CEX, they capture only differences in consumers' spending patterns across roughly 300 consumption categories. They might therefore understate the negative relationship between emissions intensity and income because they do not capture spending differences within narrowly defined product categories. In Panel (f), we examine this within-product variation for cars, an important consumption item for which the necessary microdata are available. Using the 2022 InMoment vehicle transaction survey, we observe the income, specific car models, and prices paid by thousands of U.S. households.² Following [Bils and Klenow \(2001\)](#), prices can be interpreted as a measure of quality. Because the EPA reports CO₂ emissions per mile traveled for more than 600 specific car models, we can calculate each car's use-phase emissions per dollar spent and relate this measure to household income. Panel (f) reveals a strong negative relationship even within this narrowly defined product category. Once again, the claim is not that luxury cars are cleaner in absolute terms, but that their emissions are lower per dollar spent. This pattern is also consistent with the firm-level evidence from Chile in [Klenow et al. \(2025\)](#), who find that higher-quality producers have greater revenue productivity and lower energy intensity.

Quality-Led Growth. Our theory posits that the structural change patterns reflect a reallocation from material quantity toward intangible quality. Specifically, as consumers get

²We are very grateful to Ken Gillingham for sharing these data with us.

richer, demand shifts toward luxury goods in which (clean) quality provision accounts for a larger share of value added and (dirty) material content for a smaller one. Examples include chefs who use their expertise to transform ingredients into a meal worthy of a Michelin star or software designers who enhance the functionality of consumer products.

Consider, for example, the evolution of the iPhone, shown in Panel (a) of Figure 2. Each dot represents a particular iPhone model and reports (on the right axis) its total life-cycle emissions, including both production and use, relative to the iPhone 3G introduced in 2009, whose carbon footprint we normalize to 100. Because iPhone models differ in their specifications, emissions vary even among models introduced in the same year; compare, for example, the iPhone 17e with the iPhone 17 Pro Max. Nevertheless, the figure shows only moderate growth in the average carbon footprint of newly introduced models, largely accounted for by the introduction of larger devices in more recent years.


At the same time, quality improvements have been dramatic. The red line plots the BEA's price index for personal consumption expenditures on telephone and related communication equipment. Following Byrne et al. (2019), this index uses hedonic adjustments to account for improvements in processing power, storage capacity, display quality, and camera performance. Between 2010 and 2025, quality-adjusted prices declined by almost 90%, equivalent to an average annual rate of almost 15%.³

For instance, compare the first-generation iPhone SE, introduced in 2016, with the third-generation model, introduced in 2022. The former was launched at a price of \$499, whereas the latter was launched at \$429. Thus, after adjusting for inflation, the market price fell by 29% over this six-year period. However, the quality-adjusted price fell even further because the third-generation model features a faster processor, a better camera, a larger screen, and 5G connectivity. At the same time, both models generated roughly the same level of emissions, illustrating the paper's central claim: growth driven by improvements in quality can deliver welfare gains without a commensurate increase in emissions.

Although quality-led growth may be particularly salient in high-end consumer electronics such as the iPhone, the shift in innovation away from cost reduction and toward quality improvement is pervasive. To document this pattern, we examine patent data and classify

³Byrne et al. (2019) show that this result is robust to a variety of hedonic adjustments. Between 2010 and 2018, the time period of their study, they find annual declines of quality-adjusted prices between 15.5% and 19.6%.

FIGURE 2.—QUALITY-LED GROWTH
 (a) The iPhone: Rising Quality without Emissions (b) Rising Quality: Evidence from Patents



Notes: Panel (a) plots log total life-cycle CO₂ emissions for individual iPhone models (blue dots, right axis) against their introduction dates and the BEA chain-type price index for telephone and related communication equipment (red line, left axis). We normalize emissions of the iPhone 3G, introduced in 2009, and the price index in 2010 to 100. Panel (b) reports the shares of patents classified as productivity-enhancing (*A*), quality-enhancing (*Q*), or ambiguous. Lines report the CPC-based measure, constructed by classifying four-digit CPC subclasses with the production–user utility–value-chain (PUV) prompt and assigning patents fractionally across all listed CPC codes. Bars report the patent-abstract measure, obtained by applying the same prompt to stratified patent samples and reweighting them to represent the underlying patent universe. Appendix A.3 describes the data, sampling procedure, prompt, scoring rule, thresholds, and robustness checks, including the use of only each patent’s main CPC subclass.

innovations according to the economic margin they primarily improve. A patent is classified as productivity-enhancing, or *A*-type, if it primarily increases the quantity of output that can be produced from a given set of inputs. By contrast, a patent is classified as quality-enhancing, or *Q*-type, if it primarily raises the value of final consumption while holding the underlying material flow fixed. Patents for which these two interpretations cannot be clearly distinguished are classified as ambiguous.

Our LLM-based classifier, described in Appendix A.3, operationalizes this distinction using three dimensions: (i) whether the invention improves production efficiency, processes, or costs; (ii) whether it improves user utility, final-use performance, safety, convenience, design, or health outcomes; and (iii) its position in the value chain, distinguishing upstream and infrastructure technologies from final-use and consumer-facing technologies.

We construct two complementary measures. The first is based on patents’ CPC technology classes. We apply the classifier to four-digit CPC subclasses and then use the number of patents assigned to each subclass to construct annual shares of *A*-type, *Q*-type, and

ambiguous patents.⁴ The second is based on patent abstracts and hence uses richer patent-level information. The two measures are therefore constructed from different sources, even though they use the same conceptual prompt. Their correlation is positive, about 0.52, but far from one, suggesting that each measure contains distinct information.

Panel (b) of Figure 2 reports the results. The lines refer to the CPC measure, while the bars are based on the patent abstracts. For either measure, the share of *A*-type patents declines over time, while the share of *Q*-type patents rises.⁵ The ambiguous category remains relatively stable. Hence, patent activity in the U.S. is increasingly directed toward quality, a finding echoed by Fort et al. (2026) and Comin et al. (2025), who document a shift in U.S. patenting from manufacturing toward services over recent decades. Because many quality innovations, such as better restaurants or a larger variety of stores, are not patented, this evidence likely understates the resources devoted to quality improvements.

3. THEORY

In this section, we develop a theory of endogenous degrowth that links declining material and emissions intensities to rising quality provision. We proceed in two steps. We begin by presenting a general environment in which technology is exogenous and impose only minimal assumptions on preferences and production. This framework clarifies how growth can arise either from improvements in material productivity or from increases in quality and shows how environmental sustainability depends on these two sources of growth.

We then specialize the model by parameterizing preferences and technologies and endogenizing the evolution of material productivity and intangible quality through a model of directed technical change. The resulting theory predicts an endogenous transition from material-productivity to quality-led growth and shows how welfare can continue to rise even as material production stabilizes and emissions decline.

⁴When a patent is assigned to multiple CPC subclasses, we use all subclasses with fractional weights; Appendix A.3 shows that the results are essentially unchanged if we assign each patent to its main CPC subclass.

⁵We restrict our analysis to patents from 1976 onward because consistent and reliable digitized textual data are not available for earlier years. In particular, patent abstracts became systematically accessible in the United States Patent and Trademark Office (USPTO) only starting in 1976. For the analysis of abstracts, we considered 15-year time windows in the 80s and 90s, and 10-year windows afterwards. We use 15-year windows in the early part of the sample to reflect the increase in patenting over time.

3.1. Environment

Time is continuous. There is a unit mass of infinitely lived households h with preferences over a unit continuum of goods i . Their lifetime utility is given by

$$\mathcal{V}_0^h = \int_0^\infty \exp(-\varrho t) \left[V(e_t^h, [p_{it}]_{i \in [0,1]}) - \Lambda(\mathcal{P}_t) \right] dt, \quad (1)$$

where V is the indirect utility function, e_t^h is nominal expenditure by household h , p_{it} is the hedonic price of final good i , $\varrho > 0$ is the discount rate, and $\mathcal{P}_t$ is the stock of pollution.⁶

The indirect utility function V satisfies standard regularity conditions, while Λ captures the disutility from pollution.

We begin by characterizing the static equilibrium, taking household expenditure e_t^h as given. In Section 3.6, we derive the intertemporal allocation of expenditure together with the dynamic budget constraint. Consumers value both the tangible quantity and the implemented quality of each final good. Assume that the utility delivered by good i is given by

$$u_i(y_{it}, q_{it}) = \bar{u}_i y_{it}^{1-\alpha_i} q_{it}^{\alpha_i}, \quad (2)$$

where y_{it} denotes the quantity of good i , q_{it} denotes its quality, and $\alpha_i \in [0, 1]$ captures the good's *quality intensity*. We further assume that quantity is produced using a standard constant-returns-to-scale production function,

$$y_{it} = f(x_{it}, \ell_{yit}),$$

which combines material inputs x_{it} and service labor ℓ_{yit} . Material inputs are, in turn, produced through the production function $x_{it} = A_t \ell_{xit}$, where A_t denotes quantity TFP. Empirically, we interpret the production of material inputs x as the manufacturing sector. This specification encompasses standard models of structural change, in which productivity growth is confined to manufacturing (Baumol, 1967, Ngai and Pissarides, 2007).

The novel feature of our theory is that the utility delivered by a given quantity of tangible output can be enhanced through quality. We assume that quality is produced through the

⁶The hedonic price differs from the market price because it is measured per unit of utility delivered by the good.

production function $q_{it} = Q_t \ell_{qit}$, where ℓ_{qit} denotes the labor employed in quality provision for product i , and Q_t is quality TFP. For example, employing more cooks and waitstaff can increase the quality of a restaurant experience while holding the tangible inputs (ingredients, silverware, etc.) fixed. Likewise, employing more engineers can improve the quality of a car or smartphone for a given quantity of physical output.⁷ The index Q , which we endogenize below, captures the productivity of quality provision. In the restaurant example, it reflects the knowledge and expertise that enable a given amount of quality labor to prepare refined dishes and create an appealing experience. We model Q as disembodied knowledge, following the tradition of the growth literature, primarily for reasons of tractability.

The key distinction between quantity and quality is their environmental footprint. We assume that pollution is generated by the usage of materials, but not by intangible activities. Formally, let total material production be $X_t \equiv \int_0^1 x_{it} di$. The stock of pollution then evolves according to

$$\dot{\mathcal{P}}_t = -\varphi \mathcal{P}_t + \mathcal{E}_t, \quad \mathcal{E}_t = \mathcal{E}(X_t, Z_t), \quad (3)$$

where $\mathcal{E}_t$ denotes the flow of emissions at time t , and Z_t is an index of green technology. We assume that $\mathcal{E}_X > 0$, so greater material usage increases emissions, and $\mathcal{E}_Z < 0$, so improvements in green technology—such as higher energy efficiency—reduce emissions for a given level of material production. Our analysis focuses on the endogenous evolution of X_t , taking the path of Z_t as given.⁸

3.2. Static Allocations with Exogenous A_t and Q_t

Consider the static allocation at date t , taking A_t and Q_t as given. Let $\psi(A_t)w_t$ denote the minimum unit cost of producing the composite tangible input $f(x_i, \ell_{yi})$. The total cost of delivering u_i units of utility from good i is linear and given by

⁷For simplicity, our formulation assumes that y_{it} does not depend on q_{it} . An equivalent formulation of our model is $u_i(y_i, q_i) = \bar{u}_i y_i(q_i)^{1-\alpha_i} q_i^{\alpha_i+\nu}$, where $y_i(q_i) = q_i^{-\nu/(1-\alpha_i)} f(x_i, \ell_{yi})$. If $\nu > 0$, higher quality reduces the productivity of quantity production, capturing the idea that providing higher quality may require more costly material implementation. This alternative formulation implies that utility exhibits constant returns to scale in x, ℓ_y , and ℓ_q , yielding allocations that are isomorphic to those in our baseline model.

⁸We do not model the development of green technology because this is the focus of a large literature following Acemoglu et al. (2012).

$$C_i(u_i) = p_{it}u_i, \quad p_{it} = \chi_i w_t \frac{\psi(A_t)^{1-\alpha_i}}{Q_t^{\alpha_i}}, \quad (4)$$

where χ_i is a normalization constant. Thus, the slope of the cost function, p_{it} , is the unit cost of delivering utility from good i , i.e., its hedonic price. An increase in material productivity A_t lowers the unit cost of tangible inputs, $\psi(A_t)$, allowing consumers to purchase more quantity per unit of expenditure. By contrast, an increase in Q_t enables producers to deliver higher quality with a given amount of quality labor, thereby also reducing the hedonic price. Importantly, goods differ in their exposure to these two sources of growth: the higher α_i , the greater the role of quality improvements. In our quantitative example below, we explicitly allow for the fact that quality growth might be imperfectly measured.

To characterize the allocation of labor across material production, producer services, and quality provision, let ℓ_{xt} , ℓ_{yt} , and ℓ_{qt} denote the corresponding labor shares, with aggregate labor normalized to one. Let e_t^h denote the nominal expenditure of household h at time t , and let $\vartheta_{it}(e_t^h)$ denote its expenditure share on good i . Because preferences are nonhomothetic, this expenditure share varies with the household's expenditure level. Let G_t denote the cross-sectional distribution of household expenditure at time t , and define average nominal expenditure by $\bar{e}_t$. The aggregate expenditure share on good i is then

$$\vartheta_{it} \equiv \int \vartheta_{it}(e) \varpi_t(e) dG_t(e), \quad \varpi_t(e) \equiv \frac{e}{\bar{e}_t}. \quad (5)$$

Next, let $\sigma(A_t) = \frac{p_x x_i}{p_x x_i + w \ell_{yi}}$ denote the cost share of material inputs in the production of tangible output y . The equilibrium allocation of labor is then given by

$$\ell_{qt} = \int_0^1 \alpha_i \vartheta_{it} di, \quad \ell_{yt} = (1 - \sigma(A_t))(1 - \ell_{qt}), \quad \ell_{xt} = \sigma(A_t)(1 - \ell_{qt}). \quad (6)$$

The first equation states that the labor share devoted to quality provision equals the expenditure-share-weighted average quality intensity of final demand. The second and third equations divide the remaining labor between producer services and material production. In the data, the service-sector employment share is best interpreted as the empirical counterpart of $\ell_{qt} + \ell_{yt}$: it consists of labor-intensive intangible activities that either provide quality directly or support the production of tangible output.

3.3. *Material Degrowth or Quality Growth?*

Despite its generality, this framework already allows for a tight characterization of the relationship between environmental sustainability and economic growth. For simplicity, suppose that emissions are proportional to material production adjusted by green technology, $\mathcal{E}_t = X_t/Z_t$. Equation (6) then implies a log-linear decomposition of emissions:

$$\ln \mathcal{E}_t = \ln A_t + \ln \ell_{xt} - \ln Z_t = \ln A_t + \ln \sigma(A_t) + \ln(1 - \ell_{qt}) - \ln Z_t. \quad (7)$$

Emissions rise when higher material productivity increases the quantity of material production and fall as green technology improves and manufacturing employment declines, either because technological substitution reduces the material cost share or aggregate demand shifts toward more quality-intensive goods.

The same primitives also determine welfare growth. The growth rate of aggregate consumption, as measured by the Divisia price index, is simply the expenditure-share-weighted average of hedonic price changes, which (as we show in Appendix B.1) is given by

$$g_t^V = - \int_0^1 \vartheta_{it} \frac{d}{dt} \ln p_{it} di = \ell_{qt} g_{Qt} + \ell_{xt} g_{At}. \quad (8)$$

Abstracting from the disutility of pollution, growth in consumption utility is the sum of quality growth and quantity growth. The former is weighted by the labor share in quality provision, the latter by the labor share in material production. Note that a declining material-production share reduces the contribution of quantity TFP growth to aggregate welfare growth, a manifestation of Baumol's cost disease.

Combining (7) and (8) clarifies the *limits to growth*, namely, the conditions under which growth is environmentally sustainable. Suppose that in the long run sectoral employment shares are constant. Equation (7) then shows that the growth rate of emissions hinges on the balance between material productivity growth and improvements in green technology; $g_{\mathcal{E}} = g_A - g_Z$. Requiring emissions to be non-increasing therefore imposes an upper bound on growth: $g^V \leq \ell_q g_Q + \ell_x g_Z$. Thus, as emphasized by proponents of *degrowth*, environmental sustainability indeed constrains economic growth. Yet this constraint operates on material production, the source of emissions, not on quality-led growth. The latter can, in principle,

sustain welfare growth even while emissions decline. The remainder of our theory explains why the economy may, in fact, endogenously transition to such a growth path.

3.4. Functional Forms: PIGL Preferences and CES Technology

Thus far, we have not imposed any restrictions on consumer preferences, summarized by the indirect utility function V and the associated expenditure shares $\vartheta_i(e^h)$, or on the production function for tangible goods, f . In this section, we impose functional forms that are both tractable and suitable for our quantitative analysis.

Preferences: We assume preferences in the PIGL class:

$$V(e, [p_i]) = \frac{1}{\varepsilon} \left(\frac{e}{\exp\left(\int_0^1 \omega_i \ln p_i di\right)} \right)^\varepsilon - \int_0^1 \kappa_i \ln p_i di, \quad (9)$$

where $0 < \varepsilon < 1$, $\int_0^1 \omega_i di = 1$, and $\int_0^1 \kappa_i di = 0$.⁹ Roy's identity implies the expenditure share for good i :

$$\vartheta_i(e, [p_i]) = \omega_i + \kappa_i \Upsilon(e, [p_i])^{-\varepsilon}, \quad \Upsilon(e, [p_i]) \equiv \frac{e}{\exp\left(\int_0^1 \omega_i \ln p_i di\right)}. \quad (10)$$

Equation (10) shows that the demand system resembles Cobb–Douglas preferences with a nonhomothetic adjustment. The term $\Upsilon(\cdot)$, which is common across all goods, captures consumers' real income. Goods with $\kappa_i > 0$ are necessities: their expenditure shares fall with real income. Goods with $\kappa_i < 0$ are luxuries: their expenditure shares increase with real income. The parameter ω_i determines the asymptotic expenditure share as real income tends to infinity; luxuries approach ω_i from below, necessities from above.

An appealing feature of PIGL preferences is their useful aggregation properties, which we exploit in our quantitative analysis below. Using (10), equation (5) yields the aggregate expenditure share as

$$\vartheta_{it} = \omega_i + \kappa_i \mathcal{I}_t \Upsilon(\bar{e}_t, [p_{it}])^{-\varepsilon}, \quad \mathcal{I}_t = \frac{\int e^{1-\varepsilon} dG_t(e)}{\left(\int e dG_t(e)\right)^{1-\varepsilon}}, \quad (11)$$

⁹The restriction $0 < \varepsilon < 1$ ensures that indirect utility is strictly concave in expenditure and that the intertemporal expenditure Euler equation is well defined. We also assume that e is such that the expenditure shares in (10) satisfy $0 \leq \vartheta_i(e, [p_i]) \leq 1$.

where, recall, $\bar{e}_t$ denotes average expenditure. Hence, aggregate expenditure is “as if” it were generated by a representative household with aggregate expenditure $\bar{e}_t$, and a preference parameter $\kappa_i \mathcal{I}_t$. Notably, the entire expenditure distribution is summarized by the scalar $\mathcal{I}_t$, a monotone transformation of the Atkinson index of inequality.¹⁰

Similarly, we can derive the aggregate demand for quality value added. Substituting the aggregate expenditure share in (11) into (6) yields a simple expression for the expenditure and employment share devoted to quality provision:

$$\ell_{qt} = \int_0^1 \alpha_i \vartheta_{it} di = \bar{\omega}_Q + \bar{\kappa}_Q \mathcal{I}_t \Upsilon(\bar{e}_t, [p_{it}])^{-\varepsilon}, \quad \bar{\omega}_Q \equiv \int_0^1 \alpha_i \omega_i di, \quad \bar{\kappa}_Q \equiv \int_0^1 \alpha_i \kappa_i di. \quad (12)$$

Aggregate demand for quality is again in the PIGL class and is conveniently summarized by the three structural parameters $\bar{\omega}_Q \in (0, 1)$, $\bar{\kappa}_Q$, and ε . Interestingly, the income elasticity of aggregate quality demand depends on $\bar{\kappa}_Q$, that is, the *covariance* between goods’ quality intensity, α_i , and the corresponding preference parameters, κ_i . If quality-intensive goods are, on average, income-elastic, then $\bar{\kappa}_Q < 0$, and rising income and greater inequality (lower $\mathcal{I}_t$) increase the expenditure share on quality. Below we show that the service and environmental Engel curves shown in Figure 1 imply that $\bar{\kappa}_Q < 0$.

Moreover, the hedonic prices in (4) imply that Υ takes the simple form

$$\Upsilon_t \equiv \Upsilon(\bar{e}_t, [p_{it}]) = \frac{Q_t^{\bar{\omega}_Q}}{\psi(A_t)^{1-\bar{\omega}_Q}} \frac{\bar{e}_t}{\bar{\chi} w_t}, \quad (13)$$

where $\bar{\chi} = \exp(\int \omega_i \ln \chi_i di)$ is an inconsequential constant. As we show below, in equilibrium $\bar{e}_t \propto w_t$, so that Υ_t is increasing in both Q_t and A_t . Higher Q_t lowers the quality-adjusted price of quality-intensive goods, while higher A_t lowers the tangible-cost component, $\psi(A_t)$, of final-good prices. Thus, growth in either source of growth raises real income and shifts aggregate expenditure toward quality-intensive goods.

Technology: We assume that tangible output is produced using the CES aggregator

$$f(x, \ell_y) = \left[(1 - \lambda)^{\frac{1}{\rho}} x^{\frac{\rho-1}{\rho}} + \lambda^{\frac{1}{\rho}} \ell_y^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}}. \quad (14)$$

¹⁰If $\Phi_{\varepsilon,t}$ denotes the standard Atkinson index with inequality-aversion parameter ε , then $\mathcal{I}_t = (1 - \Phi_{\varepsilon,t})^{1-\varepsilon}$.

Unit costs and material cost shares are then given by

$$\psi(A) = [(1 - \lambda)A^{\rho-1} + \lambda]^{\frac{1}{1-\rho}}, \quad \sigma(A) = \left(1 + \frac{\lambda}{1-\lambda}A^{1-\rho}\right)^{-1}. \quad (15)$$

In the baseline model we focus on $0 < \rho < 1$, so material inputs and producer services are gross complements. Then $\sigma(A)$ falls as A rises. This is the standard Baumol force: productivity growth in material production lowers the cost share of material inputs and raises the relative importance of producer services, which, in our theory, experience no productivity growth. Note that $\lim_{A \rightarrow \infty} \psi(A) = \lambda^{1/(1-\rho)}$, i.e., production costs are bounded away from zero even if material production becomes infinitely productive, reflecting the “bottleneck” logic of Baumol’s cost disease (Jones and Tonetti, 2025).

Mapping to Empirical Evidence: Our theory maps closely to the evidence in Figure 1. Consider first the household-level spending share on service value added shown in Panel (d). The amount of service value added embodied in product i is $s_{it}^S = \alpha_i + (1 - \alpha_i)(1 - \sigma(A_t))$. Using the demand system in (10), the service content of household h ’s consumption is

$$s_t^{S,h} = \int_i s_{it}^S \vartheta_i(e^h, [p_i]) di = 1 - \sigma(A_t) + \sigma(A_t) \left(\bar{\omega}_Q + \bar{\kappa}_Q \Upsilon(e^h, [p_{it}])^{-\varepsilon} \right). \quad (16)$$

Cross-sectionally, $s_t^{S,h}$ is increasing in income. Equation (16) shows that this holds if and only if $\bar{\kappa}_Q < 0$, that is, if income-elastic products are, on average, quality-intensive. Moreover, Figure 1 shows a pronounced upward shift in $s_t^{S,h}$, holding household income fixed. Equation (16) shows that such shifts reflect the Baumolian substitution channel whereby rising material productivity reduces the material cost share $\sigma(A_t)$.

Similarly, our theory delivers the cross-sectional Engel curves shown in Panels (e) and (f) of Figure 1. Household h ’s embedded consumption of physical materials for product i is $(1 - \alpha_i)\sigma(A_t)\frac{\vartheta_i^h e^h}{p_t^x}$. If emissions per unit of material input are Z_t^{-1} , consumer h ’s environmental footprint per unit of spending, $\bar{\mathcal{E}}_t^h$, is given by

$$\ln \bar{\mathcal{E}}_t^h := \ln \left(\frac{1}{Z_t} \frac{\int_i (1 - \alpha_i)\sigma(A_t)\frac{\vartheta_i^h e^h}{p_t^x} di}{e^h} \right) = \ln \left(\frac{A_t \sigma(A_t)}{Z_t} \right) + \ln \left(1 - \bar{\omega}_Q - \bar{\kappa}_Q \Upsilon(e^h, [p_{it}])^{-\varepsilon} \right),$$

where we used $p_t^x = A_t^{-1}$ with $w_t = 1$. Hence, the cross-sectional environmental Engel curves are downward sloping if and only if luxury products are, on average, more quality intensive, that is, if $\bar{\kappa}_Q < 0$.

3.5. Endogenous Technology and Directed Technical Change

So far, we have taken Q_t and A_t as given. We now endogenize their evolution through a standard Schumpeterian model of directed technical change; see Appendix Section B.1.2 for details. We assume that the inputs x_{it} , y_{it} , and q_{it} are CES aggregates of differentiated varieties with elasticity of substitution $\xi > 1$ in a monopolistic-competition environment. In equilibrium, all inputs are subject to the common markup $\xi/(\xi - 1)$, which we absorb into χ_i in equation (4).¹¹

The differentiated varieties entering the CES bundles for x and q are ordered on a vertical quality ladder, and their productivity can be improved through the entry of new firms that creatively destroy the incumbent producer. Successful A -innovations raise the productivity of material-input varieties, thereby increasing A_t , while successful Q -innovations raise the productivity of quality-provision varieties, thereby increasing Q_t .¹²

There is a constant measure R of researchers who can be allocated to material productivity or quality: $R_{At} + R_{Qt} = R$. A researcher allocated to sector $s \in \{A, Q\}$ generates a successful innovation at rate $\eta_s R_{st}^{-\zeta}$, where η_s denotes research efficiency and $0 < \zeta < 1$ captures the usual “fishing-out” congestion externality in research. Research is directed across sectors but not toward individual varieties. Conditional on success, the affected variety is drawn at random within the chosen sector. Its productivity (or quality) is multiplied by $\tilde{\gamma} > 1$, and the innovator replaces the incumbent producer. Moreover, the productivity of every A - and Q -variety depreciates at the common proportional rate $\delta \geq 0$. At the aggregate level, this can be viewed as a reduced form for maintenance (or overhead) requirements in research. The growth rates of the aggregate state variables are therefore

¹¹We impose the same monopolistic-competition structure on y so that, abstracting from the environmental externality, markups do not introduce relative distortions across sectors.

¹²We do not allow innovation to target individual final goods. Instead, innovation generates general knowledge that raises the productivity of quality provision across all goods. Allowing for product-specific innovation would be descriptively appealing, but would substantially complicate the analysis without affecting the main mechanism.

$$g_{st} = (\gamma - 1)\eta_s R_{st}^{1-\zeta} - \delta, \quad s \in \{A, Q\}, \quad (17)$$

where $\gamma \equiv 1 + (\tilde{\gamma}^{\xi-1} - 1)/(\xi - 1) > 1$, so that $\gamma - 1$ measures the contribution of innovation intensity to aggregate productivity growth.

A successful innovator sells the property rights to the innovation to a mutual fund, which owns all firms and receives each variety's profit flow until it is displaced through creative destruction. Let $z_A \equiv a/A_t$ and $z_Q \equiv q/Q_t$ denote a firm's relative productivity in the two sectors. The value function of an incumbent, $V_s(z_s, t)$, satisfies

$$r_t V_s(z_s, t) = \pi_s(z_s, t) - \tau_{st} V_s(z_s, t) + \dot{V}_s(z_s, t), \quad s \in \{A, Q\}, \quad (18)$$

where r_t is the interest rate implied by consumers' Euler equation, $\pi_s(z_s, t)$ denotes flow profits, and $\tau_{st} = \eta_s R_{st}^{1-\zeta}$ is the creative-destruction rate.

Under the CES aggregator for differentiated varieties, we show in Appendix B.1.2 that firm values are homogeneous in relative productivity; $V_s(z, t) = z^{\xi-1} v_s(t)$, where $v_s(t) \equiv V_s(1, t)$ denotes the value of a firm with average quantity TFP or quality, respectively. Conditional on survival, a and q depreciate at rate δ , so that the growth rate of z_s is given by $-(g_{st} + \delta)$. Hence the average values, $v_s(t)$, solve the differential equation

$$[r_t + \tau_{st} + (\xi - 1)(g_{st} + \delta)] v_s(t) = \Pi_s(t) + \dot{v}_s(t), \quad s \in \{A, Q\}, \quad (19)$$

where $\Pi_A(t) = \pi_A(1, t)$ and $\Pi_Q(t) = \pi_Q(1, t)$ are the profits of the "average" firm.

Letting w_t^R denote the equilibrium wage for researchers, the allocation of researchers equalizes the expected return to directing research toward A or Q :

$$w_t^R = \eta_A R_{At}^{-\zeta} \tilde{\gamma}^{\xi-1} v_A(t) = \eta_Q R_{Qt}^{-\zeta} \tilde{\gamma}^{\xi-1} v_Q(t). \quad (20)$$

The term $\tilde{\gamma}^{\xi-1}$ appears because the value of an entrant scales with its relative productivity to the power $\xi - 1$. The equilibrium allocation of researchers thus satisfies

$$\frac{R_{Qt}}{R_{At}} = \left(\frac{\eta_Q v_Q(t)}{\eta_A v_A(t)} \right)^{1/\zeta}. \quad (21)$$

Equations (19) and (21) are the key equations of the innovation block because they determine the direction of technical change. The allocation of research between A - and Q -innovation depends on the relative values of producing physical-input varieties, $v_A(t)$, and intangible quality, $v_Q(t)$. These values are, in turn, driven by the relative profitability of the two activities:

$$\frac{\Pi_Q(t)}{\Pi_A(t)} = \frac{\ell_{qt}}{\ell_{xt}} = \frac{\ell_{qt}}{\sigma(A_t)(1 - \ell_{qt})}. \quad (22)$$

Equation (22) shows that, under CES demand, relative profitability is determined by relative market size, which is exactly given by the employment ratio ℓ_{qt}/ℓ_{xt} . The second equality highlights why quality provision becomes increasingly profitable as the economy grows. First, nonhomothetic preferences shift demand toward quality-intensive final goods, raising ℓ_{qt} . Second, as material productivity improves, the physical-input cost share $\sigma(A_t)$ declines. Both forces increase the current and expected future profitability of quality provision relative to physical-input production. Consequently, research is endogenously redirected toward quality improvements, and growth becomes increasingly quality-led.

3.6. Heterogeneity, Income, and Savings in General Equilibrium

To close the model, we now specify the household side. Households differ in skill. Each household of type h comprises a unit measure of production workers and a measure R of researchers; each worker supplies h efficiency units of labor. The skill index h is distributed according to Γ , normalized so that $\int h d\Gamma(h) = 1$. A household of type h therefore maximizes lifetime utility (1) subject to the dynamic budget constraint

$$\dot{b}_t^h = h \left(w_t + R w_t^R \right) + r_t b_t^h - e_t^h, \quad (23)$$

where b_t^h denotes financial wealth, measured as the market value of the household's position in the mutual fund (the only financial asset in our economy), and r_t denotes the rate of return on that position. Let V_{Mt} denote the total market value of the mutual fund and define s_t^h as the density of mutual-fund ownership with respect to Γ . Thus, $b_t^h = s_t^h V_{Mt}$. Asset-market clearing requires $\int s_t^h d\Gamma(h) = 1$, or, equivalently, $\int b_t^h d\Gamma(h) = V_{Mt}$. The household is also subject to the standard no-Ponzi and transversality conditions.

We use the wage of production workers as the numeraire, retaining w_t momentarily to make the normalization explicit. Aggregate expenditure is determined jointly by goods- and asset-market clearing. As shown in Appendix B.1.3, consolidating the household and mutual-fund budget constraints implies that aggregate expenditure equals production-labor income plus operating monopoly profits. The mutual fund purchases successful innovations and, by free entry, its expected acquisition expenditures equal the aggregate research wage bill. Research income and innovation acquisition expenditures therefore cancel in the consolidated budget constraint. Thus,

$$\bar{e}_t = w_t(\ell_{xt} + \ell_{yt} + \ell_{qt}) + \Pi_Q(t) + \Pi_A(t) + \Pi_Y(t) = \frac{\xi}{\xi - 1} w_t. \quad (24)$$

Normalizing $w_t = 1$ gives $\bar{e}_t = \xi/(\xi - 1)$. Expenditure in units of the production wage is constant over time, whereas real expenditure rises as quality-adjusted prices decline.

We impose that initial mutual-fund ownership is proportional to skill and labor income, $s_0^h = h$. Under this condition, Appendix B.1.3 shows that the invariant skill distribution $\Gamma(h)$ fully determines the cross-sectional distribution of expenditure at every point in time: $e_t^h = h\bar{e}_t = h\frac{\xi}{\xi - 1}$. Intuitively, all households face the same return and price path, so their Euler equations imply the same expenditure growth rate. Since initial wealth is proportional to lifetime labor income, wealth and expenditure remain proportional to skill throughout the transition. Consequently, the spending distribution is time-invariant and entirely determined by $\Gamma(h)$. It follows that the inequality index $\mathcal{I}_t$ in (11) is also constant and given by $\mathcal{I} = \int h^{1-\varepsilon} d\Gamma(h)$.

DEFINITION 1: *An equilibrium is a collection of trajectories for the state variables $\{A_t, Q_t, P_t\}_{t \geq 0}$, flow emissions $\{\mathcal{E}_t\}_{t \geq 0}$, prices $\{r_t, w_t, w_t^R, [p_{it}]_{i \in [0,1]}\}_{t \geq 0}$, labor allocations $\{\ell_{qt}, \ell_{xt}, \ell_{yt}, R_{Qt}, R_{At}\}_{t \geq 0}$, household asset holdings, expenditures, and spending shares $\{b_t^h, e_t^h, [\vartheta_{it}^h]_{i \in [0,1]}\}_{h \in \text{supp}(\Gamma), t \geq 0}$, aggregate expenditure and spending shares $\{\bar{e}_t, [\vartheta_{it}]_{i \in [0,1]}\}_{t \geq 0}$, and average values $\{v_A(t), v_Q(t)\}_{t \geq 0}$, such that, given a set of initial conditions, households maximize utility; firms maximize profits; researchers choose their sector of activity optimally; the goods, production-labor, research-labor, and asset markets clear; aggregate expenditure and spending shares satisfy $\bar{e}_t = \int e_t^h d\Gamma(h)$ and*

$\vartheta_{it} = \int \vartheta_{it}^h \frac{e_t^h}{e_t} d\Gamma(h)$; the value functions satisfy (19); the laws of motion for A_t , Q_t , and $\mathcal{P}_t$ are satisfied; and flow emissions are given by $\mathcal{E}_t = \mathcal{E}(X_t, Z_t)$.

3.7. The ABGP: Quality-Led Growth with Finite Material Usage

The full equilibrium transition must be characterized numerically. We can, however, analytically characterize the unique asymptotic balanced growth path (ABGP) and show that it is locally saddle-path stable.¹³ Along equilibrium paths converging to this ABGP, the economy approaches a “weightless” economy in which sustained growth is driven entirely by improvements in quality, while physical input use converges to a finite level. As a consequence, emissions converge to zero if Z_t grows at any positive asymptotic rate.

PROPOSITION 1: *For sufficiently small $\delta > 0$, there exists $\underline{\varrho} > 0$, depending only on the model’s primitive parameters, such that, if $\varrho > \underline{\varrho}$, there exists a unique quality-led ABGP, which is locally saddle-path stable in stationary variables and has the following properties:*

1. *Research directed toward material productivity converges to $\bar{R}_A \equiv \left[\frac{\delta}{(\gamma-1)\eta_A} \right]^{\frac{1}{1-\zeta}}$, so that $g_{At} \rightarrow 0$ and $A_t \rightarrow \bar{A} < \infty$.*
2. *Research directed toward quality converges to $\bar{R}_Q \equiv R - \bar{R}_A$, and $g_{Qt} \rightarrow \bar{g}_Q \equiv (\gamma - 1)\eta_Q \bar{R}_Q^{1-\zeta} - \delta > 0$. Consequently, $Q_t \rightarrow \infty$ and $A_t/Q_t \rightarrow 0$.*
3. *The quality and material labor shares satisfy $\ell_{qt} \rightarrow \bar{\omega}_Q$ and $\ell_{xt} \rightarrow \bar{\ell}_x \equiv \sigma(\bar{A})(1 - \bar{\omega}_Q) > 0$. Material production converges to a finite level: $X_t = A_t \ell_{xt} \rightarrow \bar{A} \bar{\ell}_x < \infty$.*
4. *Welfare-relevant consumption growth satisfies $g_t^V \rightarrow \bar{\omega}_Q \bar{g}_Q > 0$.*
5. *If $\mathcal{E}_t = X_t/Z_t$ and Z_t grows at a positive asymptotic rate, then $\mathcal{E}_t \rightarrow 0$.*

Proposition 1 (proved in Appendix B.2) clarifies the sense in which our theory delivers endogenous degrowth. Material production stops growing because research is progressively redirected toward quality innovation. In the long run, only the research required to offset knowledge depreciation remains devoted to material productivity. Together with a finite level of manufacturing employment, this implies that material production converges to a finite level. Growth, however, does not disappear; it changes its nature and becomes fully

¹³Establishing global convergence analytically is generally difficult in nonlinear forward-looking models. In our numerical analysis, equilibrium paths converge to the ABGP from all initial conditions considered.

quality-led. Flow emissions converge to zero if green technology grows at any positive asymptotic rate. Thus, the model features degrowth in physical quantity, but not in welfare.

To clarify the mechanism behind Proposition 1, it is useful to discuss some special cases of our theory that would lead to different outcomes.

Quality-led growth and structural change. Quality-led growth is inherently linked to the structural transformation. To see this, suppose that preferences are homothetic ($\bar{\kappa}_Q = 0$) and production technologies are Cobb–Douglas ($\rho = 1$). In that case, the quality employment share ℓ_q , and the material cost share σ , are both constant. Consequently, the relative value of quality and material innovation, $v_Q(t)/v_A(t)$, remains constant over time, and so does the allocation of research. It follows that the growth rate of quantity TFP, g_A , is positive and constant, implying that material production, X_t , grows without bound. At the same time, relative employment shares remain constant, so the economy exhibits no structural change. Hence, structural change toward services is essential for quantity degrowth, because it is the necessary ingredient to redirect innovation from cost-reducing quantity TFP growth toward rising quality.

Quality and Baumol’s disease. If quality is absent, so that $\alpha_i = 0$ for all goods, technological complementarity can still reduce the growth rate of material production by reallocating labor away from physical-input production. However, it cannot sustain long-run growth. Once material productivity runs into the producer-service bottleneck, further improvements in A_t no longer translate into sustained consumption growth because production costs are bounded away from zero. This is the classical example of Baumol’s cost disease (Jones and Tonetti, 2025). The upshot is that physical output continues to grow without bound, even though additional output contributes progressively less to consumer welfare.¹⁴ Total emissions therefore also continue to grow unless green technological progress is fast enough.

The role of complementarities. Conversely, suppose that quality is present and preferences are nonhomothetic, but physical inputs and producer services are substitutes, i.e., $\rho \geq 1$. Then the economy can sustain long-run growth, but innovation in material productivity remains active asymptotically because the economy converges to a BGP in which a

¹⁴If $\alpha_i = 0$, then $X_t = A_t \sigma(A_t) = \left(A_t^{-1} + \frac{\lambda}{1-\lambda} A_t^{-\rho} \right)^{-1}$, which is increasing in A_t . Moreover, because there is no demand for quality innovation, research remains permanently directed toward quantity, so quantity TFP continues to grow forever. Hence, $\lim_{t \rightarrow \infty} X_t = \infty$.

positive share of research is permanently allocated to material-productivity innovation. As a result, material production grows without bound, implying that emissions can be stabilized only if green technology reduces emissions per unit of material output at least as fast as material output expands, a demanding condition for decoupling growth from emissions.

Quality growth vs. service-led growth. A similar conclusion obtains in a model in which innovation raises the productivity of the producer-service composite, ℓ_y , but there is no distinct quality margin.¹⁵ In that case, even when $\rho < 1$, the economy converges to a BGP in which research is permanently split between material-productivity innovation and producer-service innovation. Growth is therefore self-sustained, but it remains quantity-based, causing material production to grow without bound. Thus, quality is not merely a relabeling of services. Instead, it allows Baumol's cost disease to operate within the production of tangible goods, while still permitting utility from consumption to grow indefinitely.

Finally, Proposition 1 assumes $\delta > 0$. This assumption is realistic and sharpens the asymptotic characterization, but is not essential for the model's main mechanism or quantitative conclusions. When $\delta = 0$, the limiting equilibrium remains quality-led. Although the expected value of a successful material-productivity innovation converges to zero, the Inada property of the idea production function ensures that research directed toward material productivity remains strictly positive at every finite date. Along the corresponding asymptotic path, $R_{At} \rightarrow 0$, $g_{At} \rightarrow 0$, $A_t \rightarrow \infty$, and $\ell_{xt} \rightarrow 0$. Since $A_t \sigma(A_t) \sim [(1 - \lambda)/\lambda] A_t^\rho$, material production $X_t = A_t \ell_{xt}$ grows without bound, but its growth rate converges to zero. Hence, if green technology grows at any positive rate, emissions eventually decline and converge to zero. Growth therefore becomes asymptotically quality-led under either specification. Relative to the case $\delta = 0$, a positive δ sharpens the result by making A_t and X_t converge to finite levels.

4. QUANTITATIVE ANALYSIS

In this section, we structurally estimate our model and evaluate the importance of quality-led growth for the evolution of U.S. emissions and economic growth. The quantitative analysis has three objectives. First, we estimate the extent to which economic growth has shifted

¹⁵Tangible output would be produced according to $y = \left((1 - \lambda)^{1/\rho} (A_t \ell_x)^{(\rho-1)/\rho} + \lambda^{1/\rho} (B_t \ell_y)^{(\rho-1)/\rho} \right)^{\rho/(\rho-1)}$ and research would be directed to A_t and B_t .

from quantity to quality over recent decades. Second, we use the calibrated model to project the future evolution of emissions and economic growth. Third, we evaluate environmental policies, such as a tax on polluting material inputs, and show how their incidence differs across the income distribution, giving rise to conflicting policy preferences.

4.1. Data and Measurement

Our quantitative analysis combines household expenditure data, input–output tables, occupational employment data, emissions data, and macroeconomic time series from the national accounts. Details on data construction are provided in Appendix A.1.

In addition to the aggregate time-series data on GDP per worker, sectoral employment shares, and overall emissions used in Section 2, our theory requires us to distinguish between producer services (y) and quality provision (q). To do so, we use occupational work-content measures from O*NET. We first manually classify 210 of the 968 O*NET occupations as either quality provision (ℓ_q , 85 occupations) or producer services (ℓ_y , 125 occupations). We then classify the remaining occupations using a supervised machine-learning algorithm (LASSO) and interpret the predicted probability that an occupation belongs to the quality group as the share of workers engaged in quality provision. We combine these probabilities with U.S. Census employment data to compute the employment shares of quality provision and producer services by industry and over time.

To estimate household demand for quality value added, we use the 2002 Consumer Expenditure Survey (CEX) to measure spending across final consumption categories and combine it with the BEA’s input–output and bridge tables to compute each household’s derived demand for materials, producer services, and quality labor.

4.2. Calibration Strategy

The model is characterized by a set of structural parameters describing preferences, production and innovation technologies, the extent to which quality growth is measured in official statistics, and the emission function:

$$\Theta = \left\{ \underbrace{\varepsilon, \bar{\omega}_Q, \tilde{\kappa}_Q, \xi, \varrho}_{\text{Preferences}}, \underbrace{\rho, \lambda}_{\text{Technology}}, \underbrace{R, \zeta, \eta_A, \eta_Q, \gamma, \delta}_{\text{Innovation}}, \underbrace{\mu}_{Q\text{-measurement}}, \underbrace{\mathcal{E}(\cdot)}_{\text{Emissions}} \right\}. \quad (25)$$

Because the expenditure distribution is time-invariant in equilibrium, the inequality term $\mathcal{I}$ is constant along the transition. Consequently, only the composite parameter $\tilde{\kappa}_Q \equiv \bar{\kappa}_Q \mathcal{I}$ is identified. For the distributional welfare calculations, we take Γ to be the empirical 2011 U.S. household income distribution and normalize it so that $\int h d\Gamma(h) = 1$. We compute $\mathcal{I} = \int h^{1-\varepsilon} d\Gamma(h)$ and recover the household-level preference parameter as $\bar{\kappa}_Q = \tilde{\kappa}_Q / \mathcal{I}$. The household welfare results are evaluated at the corresponding percentiles of Γ . The calibration does not assume that the economy has reached its asymptotic balanced growth path. We normalize $A_{2000} = Q_{2000} = 1$.

Technology: ρ and λ . The parameter λ governs the relative importance of material inputs and producer services in the tangible component of final consumption. We therefore calibrate λ to match the physical-input cost share, $\sigma(A_t) = \frac{\ell_{xt}}{\ell_{xt} + \ell_{yt}}$, in 2000. This yields $\lambda = 0.611$. The parameter ρ governs the elasticity of substitution between physical inputs and producer services. Following the structural-change literature, we set $\rho = 0.5$, implying that physical inputs and producer services are complements (see, e.g., [Comin et al. \(2021\)](#) and [Buera and Kaboski \(2009\)](#)).

Preferences: ε , $\bar{\omega}_Q$, $\tilde{\kappa}_Q$, ξ , and ϱ . We identify the preference parameters using a combination of time-series and cross-sectional variation. The time-series variation identifies the aggregate demand parameters $\bar{\omega}_Q$ and $\tilde{\kappa}_Q$, whereas the cross-sectional variation identifies the elasticity of nonhomothetic demand ε . Using [\(12\)](#) and the fact that $\mathcal{I}$ is constant in equilibrium, the PIGL demand system implies that the employment shares in quality provision and quantity production satisfy

$$\ell_{qt} = \bar{\omega}_Q + \tilde{\kappa}_Q \Upsilon_t^{-\varepsilon} \quad \text{and} \quad \ell_{xt} = \sigma(A_t)(1 - \bar{\omega}_Q) - \sigma(A_t)\tilde{\kappa}_Q \Upsilon_t^{-\varepsilon}.$$

Moreover, equation [\(10\)](#) implies that the household-level spending share on quantity value added, $\vartheta_{xt}^h = \int_i \sigma(A_t)(1 - \alpha_i)\vartheta_{it}^h di$, satisfies

$$\ln \left(\vartheta_{xt}^h - \sigma(A_t)(1 - \bar{\omega}_Q) \right) = \ln \left(-\sigma(A_t)\bar{\kappa}_Q \Upsilon(e^h, [p_i])^{-\varepsilon} \right) \equiv \delta_t - \varepsilon \ln e^h. \quad (26)$$

Thus, the cross-sectional elasticity of the gap between the quantity expenditure share and its asymptotic level with respect to household expenditure is constant and equal to $-\varepsilon$.

The time fixed effect δ_t absorbs all common price movements as well as the micro-level preference parameter $\bar{\kappa}_Q$.

We estimate $(\varepsilon, \bar{\omega}_Q, \bar{\kappa}_Q)$ jointly through an iterative procedure that combines the cross-sectional regression with the aggregate targets. Specifically, as explained in Appendix C.2, we jointly estimate the triple $(\varepsilon, \bar{\omega}_Q, \bar{\kappa}_Q)$ to target the aggregate quality share, ℓ_q , in 1950 and 2000, the manufacturing employment share (our proxy for ℓ_x) in 1950, and the estimated ε implied by the cross-sectional regression in (26).¹⁶

Finally, we set the elasticity of substitution across differentiated varieties to $\xi = 5$, corresponding to a markup of 25%, and the household discount rate to $\varrho = 0.05$.

Quality measurement: μ . Our theory distinguishes welfare-relevant (hedonic) prices from measured final-good prices. We therefore explicitly allow for the possibility that the statistical agency might capture only a fraction $\mu \in (0, 1)$ of “true” quality growth. Let π_{it} denote the inflation rate of hedonic prices and π_{it}^{BLS} the inflation rate measured by the BLS. The expression for p_{it} in (4), together with our measurement assumption, implies that

$$\pi_{it} = (1 - \alpha_i) \frac{d}{dt} \ln \psi(A_t) - \alpha_i g_{Qt} \quad \text{and} \quad \pi_{it}^{BLS} = (1 - \alpha_i) \frac{d}{dt} \ln \psi(A_t) - \mu \alpha_i g_{Qt}. \quad (27)$$

When $\mu = 1$, measured and hedonic prices coincide. When $\mu = 0$, quality improvements are entirely absent from measured prices. We calibrate μ by targeting the gap between welfare-relevant and measured GDP growth over 2000–2019. In the baseline calibration, we assume that measured GDP growth understates growth by 0.5 percentage points per year. This target is consistent with Aghion et al. (2019), who estimate that missing growth from creative destruction amounts to roughly one-half percentage point per year. This yields $\mu = 0.654$, implying that roughly one third of quality growth is not captured in the official statistics.

Innovation: $R, \zeta, \gamma, \delta, \eta_A$, and η_Q . The number of researchers, the aggregate innovation-step parameter, and research efficiencies are not separately identified. We set $R = 0.1$ and $\gamma = 1.5$, so that researchers equal 10% of the production workforce and the coefficient multiplying the innovation arrival rate in aggregate productivity growth is $\gamma - 1 = 0.5$. Following Akcigit et al. (2021) and Acemoglu et al. (2018), we set $\zeta = 0.5$, which is consistent

¹⁶In practice we minimize the mean squared error of equally weighted relative deviations from the targeted moments. Note also that having chosen λ to match $\ell_{x2000}/(\ell_{x2000} + \ell_{y2000})$, matching ℓ_{q2000} automatically implies a match for ℓ_{x2000} .

with the empirical literature studying the effects of R&D tax credits on R&D spending (Bloom et al., 2002). We set $\delta = 0.001$, implying slow common depreciation of A - and Q -variety productivities; it determines the asymptotic upper bound on material production but has negligible effects over the historical calibration period, with virtually unchanged results when $\delta = 0$.

The two research-efficiency parameters, η_A and η_Q , are calibrated jointly to match GDP per worker growth over 1950–2000 and 2000–2019, taking into account the quality mis-measurement. Intuitively, both A and Q contribute to growth, but their relative contributions evolve over time as consumption shifts toward quality-intensive goods and research is redirected toward quality provision.

Emissions and green technology. The solution of the model is block-recursive, allowing us to calibrate the emissions block separately from the rest of the model. We parameterize the emissions function as

$$\mathcal{E}_t = \kappa_E e^{-g_Z t} X_t = \kappa_E e^{-g_Z t} A_t \ell_{xt}, \quad (28)$$

so that the efficiency of green technology, Z_t , improves at the constant rate g_Z . We estimate κ_E and g_Z by targeting U.S. production-based emissions per capita in 1980 and 2000. We discipline our model with the data on emissions per capita, because our theory abstracts from population growth.¹⁷ For the welfare analysis in Section 5, we additionally specify the law of motion for the pollution stock and calibrate the depreciation parameter φ to match the assumed atmospheric half-life of emissions.

Table I summarizes the calibration. It reports the parameter values, identifying moments, empirical targets, and model-implied counterparts.

4.3. Model Fit

The calibrated model is designed to match the main features of U.S. structural change, growth, and emissions; Figure 3 displays the fit. Panels (a) and (b) show that the model reproduces the slowdown in measured GDP growth. Welfare-relevant growth is computed

¹⁷In Appendix C.4 we also report the implications of our model for total emissions constructed by multiplying the per-capita series by historical population and, after 2025, by CBO population projections. There we also show that the results are similar if we instead calibrate κ_E and g_Z to total U.S. emissions.

TABLE I
CALIBRATION

Parameter	Value	Target	Target value	Model
ε	0.264	Cross-sectional Engel curve	-0.264	-0.264
η_A	0.332	Measured GDP ₁₉₅₀ /GDP ₂₀₀₀	0.413	0.417
η_Q	0.221	Measured GDP ₂₀₁₉ /GDP ₂₀₀₀	1.27	1.29
$\bar{\omega}_Q$	0.696	1950 U.S. manufacturing employment share	45.5	45.4
$\tilde{\kappa}_Q$	-0.327	1950 U.S. quality-labor employment share	26.5	26.5
		2000 U.S. quality-labor employment share	36.5	37.0
λ	0.611	2000 manufacturing share within tangibles	0.637	0.637
μ	0.654	Welfare-relevant growth – measured growth	0.5 p.p.	0.35 p.p.
$z = e^{g_Z}$	1.022	1980 U.S. CO ₂ emissions (tons per person)	20.41	20.41
κ_E	731.5	2000 U.S. CO ₂ emissions (tons per person)	20.24	20.24

Notes: The table reports calibrated parameters and target moments. The calibration is joint; placing a moment next to a parameter indicates its closest identifying relationship rather than a one-to-one calibration. Employment shares are reported in percentage points, GDP ratios are unit-free, and emissions are in metric tons of CO₂ per person. The parameter $z = e^{g_Z}$ denotes the gross annual improvement in green technology, so that $g_Z = \log z$ is the corresponding continuous-time rate. We normalize $A_{2000} = Q_{2000} = 1$. The following parameters are set exogenously: $\xi = 5$, $\rho = 0.5$, $\varrho = 0.05$, $\delta = 0.001$, $R = 0.1$, $\gamma = 1.5$, $\varphi = 0.023$ (based on IPCC (2007)), and $\zeta = 0.5$ (based on Akcigit et al. (2021)).

using the hedonic prices in (4), while measured growth is computed using the BLS-style prices in (27). The gap between the two widens as quality provision accounts for a larger share of expenditure. Consequently, the endogenous shift toward quality innovation contributes to the slowdown in measured growth.¹⁸

Panels (c) and (d) show that the model also matches the decline in manufacturing employment and the rise of services, decomposed into quality provision, ℓ_q , and producer services, ℓ_y . Service employment rises for two distinct reasons. First, nonhomothetic preferences shift expenditure toward quality-intensive final goods, raising ℓ_q . Second, when physical inputs and producer services are complements, growth in quantity TFP increases the producer-service share, ℓ_y , in the production of the tangible component of final goods.

Panel (e) shows the slowdown and eventual decline in emissions, while Panel (f) reports the underlying redirection of technical change: growth in A slows and growth in Q accelerates. Initially, material production and emissions rise because quantity TFP, A_t , grows rapidly. Over time, however, quantity TFP growth slows and manufacturing employment

¹⁸A second force contributing to the slowdown is that the calibration implies lower research productivity in Q -innovation than in A -innovation, $\eta_Q < \eta_A$. As demand shifts toward quality, research is endogenously reallocated toward the less productive innovation sector, reducing the aggregate innovation rate over time.

shrinks. As material production approaches its upper bound, emissions begin to decline. Our model is conservative in the sense that it underpredicts the post-2000 decline in emissions. This discrepancy likely reflects the recent acceleration in the shift away from fossil fuels, which a model with a constant rate of green technological progress, g_Z , cannot capture. In Appendix C.4 we report the corresponding paths for total emissions, constructed by multiplying the per-capita series by population and the CBO population projections for the time period after 2025. In aggregate terms, the model places the emissions peak around 2020, whereas U.S. emissions peaked in 2008; see Panel (b) of Figure A-6.

4.4. Non-Targeted Moments

We consider two non-targeted quantitative predictions. First, because quantity TFP and quality growth affect measured product prices differently, the model predicts how inflation varies with service-input and quality-provision shares; we compare these predictions with product-level inflation and the price of services relative to goods. Second, it predicts the cross-sectional relationship between service and emissions intensity and expenditure reallocation's contribution to declining aggregate emissions intensity. The model performs very well on both dimensions.


Relative inflation rates. Using the expression for p_{it}^{BLS} in (27) and noting that $\frac{d}{dt} \ln \psi(A_t) = -\sigma(A_t)g_{At}$, measured inflation for product i can be written as

$$\pi_{it}^{BLS} = -g_{At} + \ell_{it}^{SERV} g_{At} - \mu g_{Qt} \ell_{qit}, \quad \ell_{it}^{SERV} \equiv \ell_{yit} + \ell_{qit}. \quad (29)$$

The first two terms together equal $-(1 - \ell_{it}^{SERV})g_{At}$. Products with a larger physical-input share therefore experience larger price declines from quantity TFP growth, while service-intensive products exhibit higher inflation, reflecting Baumol's cost disease. The last term captures the effect of quality growth. Quality improvements lower the inflation rate of quality-intensive products, but only the fraction μ captured by official price statistics enters measured inflation.¹⁹

¹⁹In our model, $\ell_{yit} = (1 - \sigma(A_t))(1 - \alpha_i)$ and $\ell_{qit} = \alpha_i$. Hence, both the service-input share, $\ell_{it}^{SERV} = \ell_{yit} + \ell_{qit}$, and the quality-provision share, ℓ_{qit} , vary across products through α_i .

FIGURE 3.—MODEL FIT



Notes: Panel (a) plots log GDP per worker, normalized in 1950, in the data (gray line) and in the calibrated model (red line), together with two counterfactuals: full quality measurement ($\mu = 1$, dotted line) and no quality measurement ($\mu = 0$, dashed line). Panel (b) reports average annual GDP per worker growth for 1950–2000, 2000–2019, and 2020–2100. Gray circles denote the data, which are available for the first two periods; red horizontal segments denote the calibrated model. Light-red dotted and dashed segments report the corresponding predictions under full and zero quality measurement, respectively. Growth rates are computed using chain-weighted price indices with Törnqvist weights. Panels (c) and (d) report employment shares in percent: Panel (c) shows manufacturing (ℓ_x), and Panel (d) decomposes services into quality provision (ℓ_q) and producer services (ℓ_y), in both the data and the calibrated model. Panel (e) compares the data and model-implied series of emissions in tons of CO₂ per person. Panel (f) reports the model-implied annual growth rates of A and Q .

Given the aggregate paths of g_{At} and g_{Qt} and the calibrated value of μ , equation (29) predicts relative inflation across products without additional parameters. As described in Appendix C.3, we combine input–output tables with our measures of quality workers to construct ℓ_{it}^{SERV} and ℓ_{qit} and compare the predictions with the product-level inflation data of Jaravel (2024). Appendix Figure A-3 shows that the model reproduces the positive relationship between inflation and service intensity. Products with low (below 60%) service-input shares tend to experience price declines (-1.24% , on average), whereas products with high (above 60%) service-input shares exhibit systematically higher inflation ($+1.21\%$).

In addition, we also use equation (29) to test an additional restriction on a product-level inflation regression. For each sample period, the coefficients on ℓ_{it}^{SERV} and ℓ_{qit} should equal the period averages of g_{At} and $-\mu g_{Qt}$, respectively, while time fixed effects absorb the common component $-g_{At}$. Despite wide confidence intervals, the estimates for 1983–2002 and 2003–2019 are broadly consistent with these restrictions. Finally, the price of services relative to goods rises by about 2–2.5% per year in the data, a pattern traditionally associated with low productivity growth in services (Boppart, 2014, Comin et al., 2021, Herrendorf et al., 2013). As we show in Appendix C.3, our model generates a similar pattern and predicts declining relative service-price inflation, in line with the data since 1980.

Emissions intensity and reallocation. The distinction between tangible and intangible inputs also yields sharp predictions for emissions intensity across products. Letting $s_{xi} \equiv (1 - \alpha_i)\sigma(A_t)$ denote the physical-input cost share of final good i , emissions per dollar satisfy

$$\ln \bar{\mathcal{E}}_{it} = \ln \left(\frac{\kappa_E z^{-(t-t_0)}}{p_{x,t}} \right) + \ln s_{xi} \equiv d_t + \ln s_{xi} = d_t + \ln(1 - \ell_{it}^{SERV}), \quad (30)$$

where d_t captures both general-equilibrium effects and the state of green technology that are common across products in a given year. Thus, the cross-sectional variation in model-implied emissions intensity is driven by physical-input intensity and the elasticity of emissions intensity with respect to the physical-input cost share is exactly one.

Panel (a) of Figure 4 compares the model-implied relationship in (30) with EPA data on industry-level CO₂ emissions intensity and the pollution-intensity index of Levinson and O’Brien (2019). Although the calibration only targets aggregate emissions and uses no

industry-level emissions moments, the model nevertheless closely reproduces the negative relationship between service-input shares and emissions intensity. This evidence supports the model’s central environmental assumption that *only* tangible inputs pollute.

These cross-sectional differences also matter for the evolution of aggregate emissions intensity. Even if emissions intensity were unchanged within every industry, aggregate intensity would fall when expenditure shifts toward industries that emit less per dollar. Panel (b) isolates this reallocation channel in the model and the data over 1963–2010.²⁰ In the data, we hold each industry’s emissions intensity fixed at its 2010 level while allowing industry expenditure shares to follow their actual historical paths. Appendix Section C.3 shows that the corresponding model object is exactly $\Delta \ln(1 - \ell_{qt})$.

The blue lines in Panel (b) show the change in aggregate emissions intensity relative to 1963. By 2010, emissions intensity had fallen by nearly 0.8 log points in both the model and the data, in line with the aggregate fit in Figure 3. The red lines isolate the component generated by the reallocation of expenditure. Two results stand out. First, the model matches this non-targeted component extremely well. Second, reallocation accounts for only about 20% of the overall decline. The remaining 80% reflects the net effect of within-industry changes in emissions intensity, which in the model arise from green technological progress and from the scale and substitution effects of material-productivity growth, whose path is shaped by directed innovation. We now turn to a decomposition that isolates the contribution of these forces, together with reallocation, to the evolution of total emissions.

4.5. Quantity Degrowth and Falling Emissions

Combining the emissions function in (28) with $\ell_{xt} = \sigma(A_t)(1 - \ell_{qt})$ yields the exact decomposition of aggregate emissions into four driving forces:

$$d \ln \mathcal{E}_t = \underbrace{d \ln A_t}_{\text{Material scale}} + \underbrace{d \ln \sigma(A_t)}_{\text{Baumolian substitution}} - \underbrace{d \ln Z_t}_{\text{Green technology}} + \underbrace{d \ln(1 - \ell_{qt})}_{\text{Reallocation}}. \quad (31)$$

First, growth in material productivity has two opposing effects. The material-scale term raises emissions by expanding physical production. Yet, because physical inputs and producer services are complements, higher A_t reduces the physical-input cost share $\sigma(A_t)$ and

²⁰We use this period because the BEA provides a consistent industry crosswalk over these years.

lowers emissions through Baumolian substitution. Second, green technological progress reduces emissions per unit of physical output. Finally, rising demand for quality reallocates demand across products and lowers the employment share in the production of quantity. It is this last term that is shown in Panel (b) of Figure 4. Total emissions decline only when the last three forces outweigh the material-scale effect.²¹

Panel (c) of Figure 4 implements the decomposition and shows how the transition from quantity- to quality-led growth eventually turns emissions growth negative. Reallocation and Baumolian substitution each lower annual emissions growth by about 0.5–1 percentage points, but are insufficient by themselves to offset the material-scale effect when quantity TFP grows rapidly. Had quantity TFP continued growing at high early-twentieth-century rates, structural change, even with green-technology growth of about 2% per year, would not have put emissions on a declining path.

Directed technical change is therefore decisive through its effect on the material-scale term. As research shifts from quantity-productivity innovation toward quality innovation, g_{At} declines and green technological progress, Baumolian substitution, and expenditure reallocation make total emissions decline. Thus, the fall is not generated by structural transformation alone, but because material-productivity growth ceases to be the main source of rising living standards.


Panel (d) illustrates this mechanism by holding the allocation of research between A and Q fixed at its 1900 level. Nonhomothetic demand and Baumolian substitution still shift employment toward services (dashed lines), but too slowly to offset expanding material production, so emissions continue rising. Hence, structural change affects emissions mainly by redirecting innovation from material productivity toward quality.

5. ENVIRONMENTAL POLICY

While the endogenous transition toward quality-led growth reduces the material intensity of the economy in the long run, the transition may be too slow given the social cost of pollution. This creates scope for environmental policy. In this section, we study policies that accelerate the reallocation of activity toward service inputs and quality provision. Such

²¹Related decompositions into scale, composition, and technology effects are common in the environmental literature; see, for example, Shapiro and Walker (2018) and Brock and Taylor (2005).

FIGURE 4.—EMISSIONS INTENSITY AND DIRECTED INNOVATION



Notes: Panel (a) plots industry-level log emissions intensity against service-input cost shares. Black squares report EPA CO₂ emissions per dollar, gray diamonds the pollution-intensity index of [Levinson and O'Brien \(2019\)](#), and connected red circles the model relationship in (30). Each of the J final goods is assigned the service-input share of its corresponding industry. Panel (b) reports the change in aggregate emissions intensity since 1963 (blue) and the component attributable to expenditure reallocation (red). The data (solid lines) hold industry emissions intensities at their 2010 levels and use actual expenditure shares; the model counterpart (dashed lines) is $\Delta \ln(1 - \ell_{qt})$. Panel (c) decomposes model emissions growth according to (31) and reports total emissions growth. Panel (d) compares the baseline economy (red) with a counterfactual in which the allocation of researchers between A and Q is held fixed at its 1900 level (blue). Dashed lines report the service employment share on the left axis; solid lines report emissions on the right axis, normalized by their baseline value in 2000.

policies can raise welfare by addressing the pollution externality, but they also change relative prices and the direction of innovation. Because preferences are nonhomothetic, such policies are generically regressive: they raise the prices of material-intensive goods, which occupy larger shares of poorer households' budgets, and favor quality-intensive luxuries consumed disproportionately by richer households.

Taxing Material Inputs. We first consider a counterfactual policy experiment in which the government, starting in 2000, imposes a permanent 20% *ad valorem* tax on purchases of the physical-input composite, x . The revenues finance a time-varying subsidy to quality provision, ℓ_q , set to balance the government budget in every period.²² Appendix D.1 reports the implementation details.

The policy affects pollution through static and dynamic channels. Holding the state variables A_t and Q_t fixed, it raises the relative cost of material throughput and reduces demand for physical-input production—the static reallocation channel. Dynamically, by lowering the profitability of material-input production relative to quality provision, it redirects research from A - to Q -innovation.

Panel (a) of Figure 5 contrasts the pollution stock under the policy (green line) with the laissez-faire benchmark (red line). We also report an alternative counterfactual (blue line) where we impose the same input tax, but keep the allocation of research between A and Q as in the laissez-faire. The difference between the green and blue lines therefore isolates the contribution of the endogenous redirection of innovation. The policy makes pollution peak earlier and at a lower level. Static reallocation away from physical inputs drives the short-run decline; the dynamic effect becomes increasingly important over time.


The model also allows us to compute the welfare consequences of this policy. Because aggregate expenditure is unchanged, the policy affects households through two channels: the paths of hedonic prices, $[p_{it}]_i$, and the pollution stock, $\mathcal{P}_t$.

Let the policy be implemented unexpectedly at time t . Let $[p_{it+s}]_i$ and $\mathcal{P}_{t+s}$ denote the price and pollution paths under laissez-faire, and let $[p_{it+s}^\tau]_i$ and $\mathcal{P}_{t+s}^\tau$ denote the corresponding paths under the policy. We define the equivalent variation for a policy implemented at time t , Δ_t^h , for a household with nominal expenditure e^h by

$$\mathcal{V}_t^h \left(e^h, \{[p_{it+s}^\tau]_i, \mathcal{P}_{t+s}^\tau\}_{s \in [0, \infty)} \right) = \mathcal{V}_t^h \left((1 + \Delta_t^h) e^h, \{[p_{it+s}]_i, \mathcal{P}_{t+s}\}_{s \in [0, \infty)} \right). \quad (32)$$

²²If flow emissions are proportional to material production, $\mathcal{E}_t = X_t/Z_t$, the *ad valorem* input tax can be represented as an emissions tax with per-unit carbon price $\tau p_{x,t} Z_t$, where $p_{x,t}$ is the pretax price of physical inputs.

FIGURE 5.—THE IMPACT OF A TAX ON PHYSICAL INPUTS
(a) Pollution Stock (b) Heterogeneous Welfare Effects



Notes: The intersectoral policy imposes a 20% tax on physical inputs and uses the revenue to subsidize quality provision. Panel (a) plots the pollution stock, normalized to one in 2000, under laissez-faire (red), with the policy (green), and under the assumption that the allocation of research between A and Q is constrained to follow the laissez-faire path (blue). Panel (b) reports the policy's equivalent variation, Δ^h , as a permanent proportional change in nominal expenditure for households at selected percentiles of the income distribution; positive values indicate welfare gains. The horizontal axis reports the social cost of carbon in 2019 U.S. dollars per metric ton of CO_2 ; the vertical dashed line marks \$190 per metric ton.

Thus, Δ_t^h is the percentage change in household h 's annual expenditure for each year $s \geq t$ required under the laissez-faire to attain the same discounted utility as under the policy. A positive value means that the policy raises household h 's welfare.

Calculating Δ_t^h requires specifying pollution damages. We set $\Lambda(\mathcal{P}) = v_0 \mathcal{P}^{\iota+1} / (\iota + 1)$, with $\iota = 0.3$, and calibrate v_0 to a social cost of carbon (SCC) benchmark for 2019, valued using the marginal utility of a household with average expenditure; see Appendix D.2 for details.²³ The net welfare effect trades off the distortionary effects of the tax against the benefit of lower pollution. Its incidence differs across the income distribution because the policy changes relative prices directly through the tax and indirectly through the direction of innovation. The physical-input tax raises the relative price of material-intensive goods, which account for a larger share of poorer households' budgets. At the same time, the induced shift from material-productivity innovation toward quality provision disproportionately benefits richer households, whose expenditure is more quality-intensive.

²³The SCC is the date-2019 wealth equivalent of the discounted welfare loss from one additional model unit of emissions, holding the price path fixed.

Panel (b) of Figure 5 reports the policy's equivalent variation for a range of SCC values used in the literature and policy practice.²⁴ We evaluate the policy for households at the 10th, 25th, 50th, 75th, and 90th percentiles under the invariant skill distribution Γ , which determines the cross-sectional distribution of income and expenditure.²⁵ Welfare gains naturally increase with the SCC. Holding the SCC fixed, richer households benefit more, highlighting the regressive nature of taxing material-intensive necessities. At an SCC of \$190, households at and above the 25th percentile gain, while the poorest quarter lose; bottom-decile households gain only when the SCC approaches \$220 per metric ton of CO₂. The policy's distributional incidence is especially visible when $\text{SCC} = 0$, i.e., when individuals perceive no disutility from pollution. Even then, the top decile gains because the policy lowers the relative cost of quality-intensive goods and redirects innovation toward quality.

Taxing A-Research. We also consider a policy that directly tilts innovation: a 20% tax on A-directed R&D that finances a subsidy to Q-directed R&D. Because the policy operates through technology rather than current input prices, its effects on material production and pollution are more back-loaded than those of taxing material inputs directly: the short-run decline in pollution is smaller although the effect grows over time. This policy is easier to implement from a politico-economic standpoint: at a social cost of carbon of \$190 per metric ton of CO₂, even the poorest decile supports the policy. Appendix D.3 reports the implementation details, and Figure A-7 shows the resulting pollution and welfare effects.

Tipping points. The preceding analysis assumes smooth pollution damages. If damages rise sharply near a threshold $\bar{P}$, for example if $\lim_{P \uparrow \bar{P}} \Lambda'(P) = \infty$, crossing it may cause large, potentially irreversible losses. Although the laissez-faire economy may eventually redirect growth toward quality, it may do so too late to keep the pollution stock below $\bar{P}$. Policy is then valuable both for lowering cumulative pollution and preventing the threshold from being crossed. Timing also affects the choice of instruments: when the threshold is

²⁴Policy benchmarks vary across countries and over time. In the U.S., the 2021 interim federal estimate was \$51 per metric ton of CO₂ for 2020 emissions at a 3% discount rate; EPA's 2023 update increased it to \$190 under a 2.0% discount rate. At present, federal policy no longer maintains a uniform SCC estimate. Germany's Federal Environment Agency recommends a climate-damage cost of €350 (about \$400) per metric ton of 2026 emissions under a 1% time-preference rate and €1,000 when present and future generations are weighted equally. In the academic literature, Ricke et al. (2018) estimate a global SCC of \$417 per metric ton of CO₂.

²⁵We measure these percentiles directly from the observed income distribution in the Current Population Survey (CPS). See Appendix Table A-5.

near, an input tax's immediate reduction in material use is especially valuable, whereas subsidizing research works mainly through future innovation and may be too slow. This issue is especially important globally. Because many developing countries are at an earlier stage of development, their emissions may continue rising long after those of advanced economies have peaked, adding to the global pollution stock and raising the risk of crossing the threshold.

6. DEGROWTH IN THE GLOBAL ECONOMY

Thus far, we have treated the United States as a closed economy. This could overstate the apparent "cleansing effect" of material degrowth if declining U.S. manufacturing mainly reflects the offshoring of carbon-intensive production.


International evidence. Panel (b) of Figure 1 already showed that emissions intensity has declined worldwide, a pattern difficult to reconcile with the view that falling emissions intensity in rich countries primarily reflects the relocation of dirty production to poorer countries. We complement this evidence with trends in emissions levels. Global fossil CO₂ emissions rose from roughly 22.7 Gt in 1990 to 38.1 Gt in 2023, consistent with material productivity remaining an important source of global economic growth. Their growth rate, however, slowed markedly, from an annual average of about 1.9% over 1990–2010 to about 1% over 2010–23 (1.2% before COVID–19).

Figure 6 compares territorial and consumption-based CO₂ emissions in OECD and non-OECD countries, with China shown separately. OECD emissions peaked around 2007 and have since declined. By contrast, non-OECD emissions rose throughout the sample. This shift in the global distribution of emissions is not primarily an artifact of offshoring. In 2023, consumption-based emissions were about 22.9 Gt in non-OECD countries and 13.0 Gt in OECD countries. Thus, even when emissions are assigned to final consumers rather than producers, non-OECD emissions remain roughly three-quarters larger, indicating that their rise is driven mainly by domestic demand rather than by production for OECD consumers. Within the non-OECD group, China is especially informative because, as a major manufacturing exporter, its territorial emissions exceed its consumption-based emissions. The gap widened during the export boom of the early 2000s but subsequently narrowed; by 2023, consumption-based emissions were only about 10 percent lower than territorial emissions. Hence, even in China, most emissions reflect domestic demand rather

than exports. Emissions growth in China has also slowed substantially, with recent estimates suggesting a plateau in 2024–2025.

Turning to OECD countries, the evidence points in the same direction. Territorial and consumption-based emissions both fell by about 3 Gt between 2007 and 2023, undermining the view that the decline in territorial emissions was driven mainly by offshoring. Net embodied-carbon imports account for the remaining gap between the two measures, but represent only 10–15% of consumption-based emissions and show no clear upward trend. Taken together, these patterns indicate that falling U.S. and OECD emissions intensity is not merely an artifact of offshoring. Global emissions continue to rise mainly because of rapid growth in developing countries whose consumption baskets remain relatively material- and energy-intensive. More generally, many developing economies have yet to pass through the stage in which rising income primarily expands demand for material- and energy-intensive necessities. This makes timely environmental policy especially important, since the endogenous forces that eventually produce material degrowth may operate too slowly to prevent cumulative pollution from crossing critical thresholds.

FIGURE 6.—TOTAL CO₂ EMISSIONS IN THE GLOBAL ECONOMY



Notes: The figure reports annual CO₂ emissions in gigatons. Solid lines show territorial (production-based) emissions, and dashed lines show consumption-based emissions; green, red, and blue denote OECD countries, non-OECD countries, and China, respectively. OECD countries are defined as those that were members of the OECD by 2001. Consumption-based emissions for 1990–2023 are from the Global Carbon Atlas. Territorial emissions through 2024 are from the same source. Territorial emissions for 2025 are estimated by applying country-group-specific growth rates reported by the International Energy Agency (IEA) to the 2024 values.

Extending the model to international trade. Although the preceding evidence indicates that offshoring is not the main driver of declining emissions intensity in advanced economies, trade can still affect where material production takes place and how innovation

is directed. In the working-paper version of this paper ([Aghion et al., 2025](#)), we develop a two-country extension of the model, calibrated to the United States and China. Countries trade differentiated material inputs while producer services and quality are provided locally. We also introduce a tradable endowment that allows the model to reproduce persistent trade imbalances. We calibrate the initial cross-country differences using 2002 data, discipline the subsequent convergence process with 2024 outcomes, and compare the free-trade equilibrium with an autarky counterfactual.

Trade generates opposing forces. Specialization raises Chinese material production and redirects activity toward the material sector, while gains from trade raise income and thus shift expenditure toward quality-intensive goods. In the calibrated model, free trade accelerates China's income convergence but shifts material employment and emissions from the United States to China. While trade increases emissions in China and reduces them in the U.S., the global effect is theoretically ambiguous and depends, among other factors, on relative production technologies. In the calibrated economy, the path of global emissions is very similar under free trade and under autarky, even though Chinese tradable production is assumed to be permanently 40% more emissions-intensive than U.S. production. World emissions are initially higher under free trade, but subsequently fall below the autarky path as the income-induced shift toward quality and the dynamic reallocation of innovation become dominant. This exercise reinforces the interpretation of the international evidence: trade affects the location of emissions, but is unlikely to account for broad global emissions patterns through large-scale carbon leakage.

7. CONCLUSION

In this paper, we develop and quantify a growth model in which consumers have nonhomothetic preferences over quality, the allocation of innovation between material productivity and quality improvement is endogenous, and quality provision has a lower environmental footprint than material production.

The theory delivers several novel insights. First, as economies develop, consumers shift demand toward quality- and service-intensive goods, redirecting innovation away from material productivity and toward quality improvement. Second, this transition to quality-led growth reduces measured GDP growth and ultimately leads material production to converge to a finite level, even as quality-adjusted GDP and welfare continue to rise. Third, environ-

mental policy can accelerate this structural transformation, but its benefits and costs are unevenly distributed across households, creating distributional conflict.

Future research can extend our analysis in several directions. One priority is better product-level measurement of quality's environmental footprint. Our observable service-share proxy misses some variation in emissions across goods; future work could model and measure quality within narrowly defined goods, such as inexpensive versus gourmet restaurants, and connect quality-led growth to relational goods and non-consumption-oriented activities emphasized by degrowth proponents. A second priority is AI, which changed from a hypothetical scenario to a technological revolution while this paper was written. By transforming the production function for many intangible services, AI may raise the embodied emissions of service-intensive industries both via direct energy use and through the construction of data centers. A further extension would introduce demographic change in a semi-endogenous growth model. Standard formulations tie technological progress to population growth, so declining population growth would slow both material-productivity and quality growth. The former would help contain material production and emissions; the latter would slow improvements in living standards through a channel distinct from the redirection of innovation emphasized here. Finally, future work could quantify how much of the recent slowdown in measured growth reflects an accelerated shift toward quality-based innovation. These extensions would deepen our understanding of whether and how weightless growth can support prosperity while reducing environmental pressure.

REFERENCES

- Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hémous: 2012, 'The environment and directed technical change'. *American Economic Review* **102**(1), 131–166. [3, 13]
- Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr: 2018, 'Innovation, reallocation, and growth'. *American Economic Review* **108**(11), 3450–3491. [28]
- Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr: 2016, 'Transition to clean technology'. *Journal of Political Economy* **124**(1), 52–104. [3]
- Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li: 2019, 'Missing Growth from Creative Destruction'. *American Economic Review* **109**(8), 2795–2822. [28]
- Aghion, P., A. Bergeaud, M. De Ridder, and J. Van Reenen: 2024, 'Lost in transition: financial barriers to green growth'. INSEAD Working Paper No. 2024/16/EPS. [3]
- Aghion, P., T. Boppart, M. Peters, M. Schwartzman, and F. Zilibotti: 2025, 'A Theory of Endogenous Degrowth and Environmental Sustainability'. NBER WP 33634. [42]

- 1 Akcigit, U., H. Alp, and M. Peters: 2021, 'Lack of selection and limits to delegation: firm dynamics in developing 1
countries'. *American Economic Review* **111**(1), 231–275. [28, 30] 2
- 2 Alder, S., T. Boppart, and A. Müller: 2022, 'A theory of structural change that can fit the data'. *American* 3
Economic Journal: Macroeconomics **14**(2), 160–206. [4] 3
- 4 Baumol, W.: 1967, 'Macroeconomics of unbalanced growth: the anatomy of the urban crisis'. *American Economic* 4
Review **57**(3), 415–426. [8, 12] 5
- 6 Bils, M. and P. J. Klenow: 2001, 'Quantifying quality growth'. *American Economic Review* **91**(4), 1006–1030. 6
[4, 8]
- 7 Bloom, N., R. Griffith, and J. Van Reenen: 2002, 'Do R&D tax credits work? Evidence from a panel of countries 7
1979–1997'. *Journal of Public Economics* **85**(1), 1–31. [29, A-19] 8
- 9 Boppart, T.: 2014, 'Structural change and the Kaldor facts in a growth model with relative price effects and non- 9
Gorman preferences'. *Econometrica* **82**(6), 2167–2196. [4, 33] 10
- 10 Brock, W. A. and M. S. Taylor: 2005, 'Economic growth and the environment: a review of theory and empirics'. 10
Handbook of Economic Growth **1**, 1749–1821. [35] 11
- 12 Buera, F. J. and J. P. Kaboski: 2009, 'Can traditional theories of structural change fit the data?'. *Journal of the* 12
European Economic Association **7**(2-3), 469–477. [27] 13
- 14 Byrne, D. M., J. G. Fernald, and M. B. Reinsdorf: 2016, 'Does the United States Have a Productivity Slowdown 14
or a Measurement Problem?'. *Brookings Papers on Economic Activity* **47**(1), 109–182. [4] 15
- 15 Byrne, D. M., D. E. Sichel, and A. Aizcorbe: 2019, 'Getting smart about phones: new price indexes and the 15
allocation of spending between devices and services plans in personal consumption expenditures'. Technical 16
report, Finance and Economics Discussion Series 2019-012. Washington: Board of Governors of the Federal 17
Reserve System. [9, A-2] 18
- 18 Caron, J. and T. Fally: 2022, 'Per Capita Income, Consumption Patterns, and CO2 Emissions'. *Journal of the* 18
Association of Environmental and Resource Economists **9**(2), 235–271. [4] 19
- 20 Comin, D., D. Lashkari, and M. Mestieri: 2021, 'Structural change with long-run income and price effects'. 20
Econometrica **89**(1), 311–374. [4, 27, 33] 21
- 22 Comin, D., D. Lashkari, and M. Mestieri: 2025, 'The Structural Transformation of Innovation'. NBER Working 22
Paper 33855, National Bureau of Economic Research. [11] 23
- 23 De Ridder, M. and L. Rachel: 2025, 'Emissions-adjusted total factor productivity'. Discussion Paper CFM- 23
DP2025-09, London School of Economics. [4] 24
- 24 Faber, B. and T. Fally: 2022, 'Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store 24
Scanner Data'. *The Review of Economic Studies* **89**(3), 1420–1459. [4] 25
- 26 Fajgelbaum, P. D., G. M. Grossman, and E. Helpman: 2011, 'Income distribution, product quality, and interna- 26
tional trade'. *Journal of Political Economy* **119**(4), 721–765. [4] 27
- 27 Fan, T., M. Peters, and F. Zilibotti: 2023, 'Growing like India—the unequal effects of service-led growth'. *Econo-* 28
metrica **91**(4), 1457–1494. [4, A-14] 29
- 30 Fieler, A. C.: 2011, 'Nonhomotheticity and bilateral trade: evidence and a quantitative explanation'. *Econometrica* 30
79(4), 1069–1101. [4] 31
- 31 Fort, T. C., N. Goldschlag, J. Liang, P. K. Schott, and N. Zolas: 2026, 'Growth Is Getting Harder to Find, Not 31
Ideas'. NBER Working Paper 35182, National Bureau of Economic Research. [11] 32

- 1 Georgescu-Roegen, N.: 1971, *The entropy law and the economic process*. Harvard University Press. [3] 1
- 2 Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski: 2014, ‘Optimal taxes on fossil fuel in general equilibrium’. 2
Econometrica **82**(1), 41–88. [3] 3
- 3 Hallak, J. C. and P. K. Schott: 2011, ‘Estimating Cross-Country Differences in Product Quality’. *The Quarterly* 3
Journal of Economics **126**(1), 417–474. [4] 4
- 4 Hémous, D.: 2016, ‘The dynamic impact of unilateral environmental policies’. *Journal of International Eco-* 5
nomics **103**, 80–95. [3] 6
- 5 Hémous, D. and M. Olsen: 2021, ‘Directed technical change in labor and environmental economics’. *Annual* 7
Review of Economics **13**, 571–597. [3] 8
- 6 Herrendorf, B., R. Rogerson, and A. Valentinyi: 2013, ‘Two perspectives on preferences and structural transfor-
- 7 mation’. *American Economic Review* **103**(7), 2752–2789. [33, A-15] 9
- 8 Herrendorf, B., R. Rogerson, and A. Valentinyi: 2014, ‘Growth and structural transformation’. In: *Handbook of* 10
Economic Growth, Vol. 2. Elsevier, pp. 855–941. [A-1] 11
- 9 Hickel, J., G. Kallis, T. Jackson, D. W. O’Neill, J. B. Schor, J. K. Steinberger, P. A. Victor, and D. Ürge-Vorsatz:
- 10 2022, ‘Degrowth can work — here’s how science can help’. *Nature* **612**, 400–403. [3] 12
- 11 IPCC: 2007, ‘Chapter 7: couplings between changes in the climate system and biogeochemistry, executive sum-
- 12 mary’. In: S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K. Averyt, M. Tignor, and H. Miller (eds.):
- 13 *Climate change 2007: the physical science basis*. Cambridge University Press. [30] 15
- 14 Jaravel, X.: 2024, ‘Distributional consumer price indices’. CEPR Discussion Paper DP19802. [33, A-13, A-15] 16
- 15 Jones, C. I. and C. Tonetti: 2025, ‘Past Automation and Future AI: How Weak Links Tame the Growth Explosion’.
- 16 Technical report, Working paper, Stanford GSB and NBER. [18, 24] 17
- 17 Klenow, P. J., E. Pastén, and C. Ruane: 2025, ‘Carbon taxes and misallocation in Chile’. Mimeo, Stanford. [8] 18
- 18 Levinson, A. and J. O’Brien: 2019, ‘Environmental Engel curves: indirect emissions of common air pollutants’.
- 19 *Review of Economics and Statistics* **101**(1), 121–133. [4, 8, 33, 36, A-1, A-2, A-13, A-17] 19
- 20 Meadows, D. H., D. L. Meadows, J. Randers, and W. W. Behrens III: 1972, *The limits to growth*. New York:
- 21 Universe Books. [3] 21
- 22 Ngai, L. R. and C. A. Pissarides: 2007, ‘Structural Change in a Multisector Model of Growth’. *American Eco-* 22
nomics Review **97**(1), 429–443. [12] 23
- 23 Quah, D.: 1999, ‘The weightless economy in economic development’. CEP, London School of Economics, Dis-
- 24 cussion Paper no. 417. [1] 24
- 25 Ricke, K., L. Drouet, K. Caldeira, and M. Tavoni: 2018, ‘Country-level social cost of carbon’. *Nature Climate* 25
Change **8**(10), 895–900. [39] 26
- 26 Shapiro, J. S. and R. Walker: 2018, ‘Why is pollution from US manufacturing declining? The roles of environ-
- 27 mental regulation, productivity, and trade’. *American Economic Review* **108**(12), 3814–3854. [35] 27
- 28 Syverson, C.: 2017, ‘Challenges to Mismeasurement Explanations for the US Productivity Slowdown’. *Journal* 28
of Economic Perspectives **31**(2), 165–186. [4] 29
- 29 30
- 30 31
- 31 32