A Theory of Endogenous Degrowth and Environmental Sustainability*

Philippe Aghion[†], Timo Boppart[‡], Michael Peters[§], Matthew Schwartzman[¶] and Fabrizio Zilibotti[‖]

October 2024 [PRELIMINARY]

Abstract

We develop and quantify a growth theory where consumers' preferences are defined over products with varying environmental impacts. Preferences are non-homothetic: Necessities are intensive in material inputs whose production leads to high emissions, while luxury goods, being more reliant on service labor, exhibit a comparatively lower environmental footprint. Directed innovation is the focal point of the study: it can be aimed at either enhancing the productivity of material production or refining the "quality" of luxury goods. Over time, innovation increasingly prioritizes quality improvement, consequently reducing the environmental impact of economic growth. The pace of structural transformation and the composition of GDP are both endogenous and susceptible to policy interventions. The shift towards quality-oriented growth may result in a decline in measured GDP growth without a decrease in welfare. Extending the model to a two-country trade scenario reveals that trade barriers could have a detrimental effect on environmental sustainability.

*We thank audiences at the SED and the Conference on the Economics of Innovation in Memory of Zvi Griliches for their helpful comments and especially our discussant Maarten de Ridder. We are grateful to Lorenzo Caliendo and Matthew Murillo for their help in accessing the GTAP data. Nicolas Fajardo and Youdan Zhang provided outstanding research assistance.

[†]INSEAD. philippe.aghion@insead.edu

[‡]University of Zurich and IIES. timo.boppart@iies.su.se.

[§]Yale University and NBER. m.peters@yale.edu

[¶]Yale University. matthew.schwartzman@yale.edu

[‖]Yale University and NBER. fabrizio.zilibotti@yale.edu

1 Introduction

There is broad consensus that climate change and environmental protection are critical priorities and that economic activity is a significant contributing factor to their severity. These observations have led a number of public figures to advocate a growth slowdown ("degrowth"), in order to achieve carbon neutrality and stop the rise in temperature.[1] The main objection to degrowth is that it would entail large costs for billions of people worldwide, especially in developing countries.[2] Furthermore, it would likely trigger fierce opposition leading to political and social disruptions. Mainstream economists have been largely skeptical of this view. Rather, they have advocated for the potential for green innovation to curb climate change without sacrificing long-term economic prosperity (Acemoglu et al., 2012b).

One of the tenets of the degrowth manifesto is that in order to avert an environmental disaster, the emphasis of economic activity should switch from quantity to quality. In this paper, we take this argument seriously. We argue that this can be an important part of the solution of the climate change challenge and one that mainstream economists have so far erroneously neglected. However, we also show that it is misleading to equate the shift from quantity to quality with degrowth.

We associate both theoretically and empirically the abstract notion of quality with the value-added intensity of different consumption items in services relative to material production. We argue that *weightless* economies (Quah, 1999) can grow in a much more environmentally friendly way than traditional economies led by an expansion of material production. The shift from quantity to quality is in part a spontaneous process: as an economy develops and people become wealthier, the demand progressively shifts from items that are intensive in material goods to items that are quality and service intensive. The structural transformation of the US economy offers a good illustration of this idea. In the US, services have been growing rapidly over the past decades, and they currently account for approximately eighty percent of total employment. In addition, total emissions have decreased over the last fifteen years. However, the ongoing structural transformation may be too slow to resolve the environmental problem.

[1] The intellectual roots of the degrowth movement stretch back to the 1970s. We reflect this debate in the literature review below.

[2] A natural experiment of "degrowth" is the first lockdown following the irruption of Covid-19 four years ago. Although indispensable at the time when the new Covid vaccines were not yet operational, the "degrowth" induced by this lockdown resulted in a sharp increase in poverty and famine-driven mortality in less developed countries.

More formally, we develop a novel growth theory in which the distinction between quantity and quality takes center stage and where the direction of technological progress toward increasing the productivity of material production versus improving quality is endogenous. In our theory, consumers' preferences are defined over a a range of final products characterized by variations in both their production technology and the degree to which consumers are willing to pay for enhanced quality.

We make three key assumptions, which we document are borne out in the empirical evidence. The first is that consumption goods are ranked on a *sophistication* ladder, where a higher sophistication is associated with both a higher service intensity in production and a higher importance of quality. For example, compare food at home with gourmet restaurants. Food at home uses mostly physical goods (the meal's ingredients) as inputs and consumers are typically more casual about quality. In contrast, a larger share of the gourmet restaurant's bill comprises payments to service workers (chefs, professional waiters, ambiance) and consumers are willing to pay a higher premium for quality embedded in their services.

The second assumption is that consumers have nonhomothetic preferences: basic goods are necessities, whereas sophisticated goods are luxuries. In the example above, richer consumers spend a higher income share on gourmet restaurants and a lower share on food at home. As society becomes richer, aggregate demand shifts toward gourmet restaurants. The assumption that richer households typically buy higher quality goods is consistent with the evidence in Bils and Klenow (2001).

The third assumption is that the environmental impact of sophisticated goods is lower than that of basic goods per unit of expenditure. Thus, as a society becomes richer, the environmental damage per dollar spent diminishes. This forecast aligns with empirical findings indicating that emissions per unit of GDP are lower in wealthier countries and decrease with the employment share of services.

In most existing theories, the distinction between quality and quantity may seem inconsequential and boils down to alternative interpretations of a given set of equilibrium conditions. However, in our theory, this differentiation has significant implications. Specifically, a central tenet of our theory is that market forces can direct innovation along two distinct paths: reducing the cost of material production or enhancing the quality of consumer goods. Innovation aimed at cost reduction enables firms to expand the production of goods. Even though newer technologies typically boast greater environmental friendliness, the expansion of material production inevitably leads to increased emissions. Conversely, quality-driven innovation does not affect emissions. For instance, an iPhone 16 has a similar environmental footprint to an iPhone 3, and a gourmet restaurant exhibits a comparable environmental impact to a fast-food establishment.

Our theory predicts that economic growth is accompanied by an intrinsic shift of innovation from material production towards quality enhancement. This shift is driven by two complementary forces. Firstly, if goods and services act as complementary inputs in the production of final goods, advancements in manufacturing technology gradually reduce the cost share of material inputs over time. This reduction on total spending on physical goods makes cost-reducing innovation in material productivity less important and less profitable. Secondly, due to nonhomothetic preferences, aggregate demand shifts from basic to sophisticated goods. Both of these dynamics contribute to reducing the environmental impact of economic growth in affluent economies. However, in a laissez-faire setting, this transition may occur too gradually. Policy intervention, such as subsidies towards quality-driven innovation, may be necessary to expedite the shift from quantity-oriented to quality-led growth. This intervention can also curtail the long-term growth rate of physical production.

Is degrowth necessary? The answer to this question hinges on how GDP is measured. A precise measure of real GDP growth should incorporate changes in quality. If quality changes were accurately accounted for, the transition to quality-driven growth would not entail degrowth. However, in practice, quality improvements are often inadequately measured, particularly in service-intensive sectors. Given this imperfect measurement, our theory predicts a gradual decline in GDP growth and, conceivably, long-term stagnation. Consequently, our theory could cast new light on the observed decrease in total factor productivity (TFP) growth since the turn of the millennium. From the perspective of our theory, this decline does not signify a waning technological dynamism, but rather a structural shift towards sectors where improvements in quality are poorly measured. Although this argument is per se not new, its connection with the debate on environmental sustainability is novel.

Next, we consider trade and specialization. In the United States, the phenomenon of deindustrialization could, in part, be attributed to the transfer of production activities to other regions worldwide, particularly China.[3] From an environmental standpoint,

[3] It's noteworthy that the decline of manufacturing and the rise of services in the US began well before significant trade with China emerged.

this relocation opens supplementary questions. For instance, despite growing attention to environmental standards, Chinese firms have frequently adopted technologies that are more polluting than those used by their Western counterparts.

However, our theory also underscores opposing forces. First, the benefits derived from trade contribute to the enrichment of all nations, thereby globally shifting demand toward cleaner, service-intensive goods.[4] Secondly, trade and specialization influence the direction of technological advancement. To analyze these factors more formally, we extend our analysis to a two-country model comprising a higher-income country (the US) and a lower-income country (China). This extension yields further insights. We demonstrate that the net effect of trade liberalization is a reduction in global emissions levels, primarily due to the endogenous response of innovation.

Literature: [**Very preliminary and incomplete**] Our study relates to several strands of literature. It is generally related to the literature on the macroeconomic and welfare implications of climate change pioneered by Nordhaus (1991, 1994, 1997) and recently developed by (Golosov et al., 2014).[5] However, this literature does not distinguish between quality and quantity based growth, nor does it factor in endogenous directed innovation and its relationship with consumers' demand and real income.

More closely related to our analysis is the literature on the environment and endogenous directed technical change (starting with Acemoglu et al. (2012a), henceforth AABH). Several papers have since extended AABH: thus Acemoglu et al. (2016) and more recently Aghion et al. (2024), build models of growth and firm dynamics to analyze the process of energy transition; Hémous (2016) extends AABH to a multi-country model with trade; Aghion et al. (2023) investigate the joint impact of consumers' environmental concerns and market competition on firms' incentives to innovate in clean technologies.[6] However, none of these papers distinguish between quality-based and quantity-based growth, nor do they consider the relationship between development as reflected by real income, consumers' demand for quality, and the resulting choice between quality-enhancing and productivity-enhancing innovation and the ultimate effects of that choice on aggregate pollution.

[4] In a recent study, Chen et al. (2023) observe a structural transformation in the Chinese economy, transitioning from manufacturing to services.

[5] See Hassler et al. (2016) for a comprehensive survey of that line of research.

[6] See Hémous and Olsen (2021) for a good literature review on green innovation and the energy transition.

Our paper also contributes to the long-standing debate on degrowth. The proponents of this radical theory advocate for shrinking rather than growing economies, in order to preserve the world's resources. The philosopher André Gorz is credited with having coined the term degrowth during a debate organized by the Nouvel Observateur in 1972. His view was rooted in the so-called post-work ideology, whereby an increase in productivity leading to more goods being produced with fewer workers is bound to result in societal distress in our work-centered societies.[7] Georgescu-Roegen (1974) argues that modern economic systems transform low-entropy resources, such as raw materials, into high-entropy goods. Because low-entropy resources are limited, the rate at which they are consumed determines the maximum achievable rate of economic growth—see also Georgescu-Roegen (1971, 1979). A common theme in the degrowth manifesto is the petition for a switch from quantity to quality, although this is often phrased as the need to stop GDP growth and replace it by quality growth—see, e.g., Hickel et al. (2022)—while we argue that correctly measured GDP should account for both quality and quantity improvements.[8]

We contribute to this debate on degrowth and the relevance of GDP growth measures by reconciling innovation-based growth with the quest for sustainability, and even more specifically by showing that the endogenous move from quantity to quality-based growth contributes to making growth both, increasingly service-intensive and therefore non-polluting, and increasingly less reflected in measured GDP growth.

Our paper is also related to the literature on structural change and service-led growth. Boppart (2014) develops a precursor model from which we borrowed our non-homothetic preferences, and which rationalizes several significant empirical facts, in particular, that the share of agricultural and manufactured goods in household spending decreases at a constant rate over time. More recently, Fan et al. (2023) argued that India exemplifies the case of an economy that has based its development to a large extent on services and service-led growth, contrary to the dominant view that industrialization is the key to economic development. We contribute to this literature

[7] A related viewpoint is that of Schumacher (1973), who criticizes the use of gigantic technologies that deprive the vast majority of individuals of the ability to make independent decisions regarding their development and use.

[8] Related to this critique, Easterlin (1974) documents that happiness does not increase proportionally with income beyond a certain threshold of national income, highlighting the limitations of using GDP as the sole indicator of welfare. Inspired by these authors, Stiglitz et al. (2009) proposes new indicators that aim to measure the social and sustainable progress of a nation without relying solely on the unidimensional GDP measure.

by introducing endogenous technical change and the resulting choice between quality-enhancing and productivity-enhancing innovation, by analyzing how that choice evolves over time as consumers become richer, and by looking at the implications of the evolution of preferences and of the direction of innovation for the dynamics of aggregate pollution.

The remainder of the paper is organized as follows. Section 2 discusses some empirical motivation. Section 3 presents the theory and characterizes equilibrium. Section 4 discusses the implication of the equilibrium characterization for environmental sustainability and relates the findings to the debate on degrowth. Section 6 focuses on an open economy extension of our basic framework. Section 5 provides a quantitative analysis. Section 7 concludes. An appendix contains details of the data and technical results.

2 Services and Pollution: Empirical Motivation

A core premise of our theory is that services are relatively environmentally friendly. In this section, we present evidence on the relationship between service intensity and pollution levels.[9]

Consider Figure 1. Panel (a) displays the time series of total CO2 emissions in the U.S. (in blue) and the economy's emissions intensity—defined as emissions relative to GDP—in red. For ease of comparison, we normalize both series to a baseline of unity in the year 2000. Two key patterns stand out. First, while emissions grew steadily for much of the 20th century, their growth has significantly slowed, and in the past 15 years, emissions have actually been declining. Second, the U.S. economy's emissions intensity has steadily decreased over the last 100 years. Relative to GDP, emissions peaked around 1920 and have since fallen by over 1 log point—almost a threefold reduction.

In this paper, we argue that the rise of the service activities has played a key role in this shift. In Panel (b) of Figure 1, we show the share of employment in services, which has expanded from around 30% at the start of the 20th century to over 85% today. If the value added by services generates less pollution than that of manufacturing, then the rise of services should have contributed to the observed decline in pollution intensity in the U.S.

[9] A detailed description of the data used is deferred to Section ??.

Figure 1: Services Are Clean Luxuries

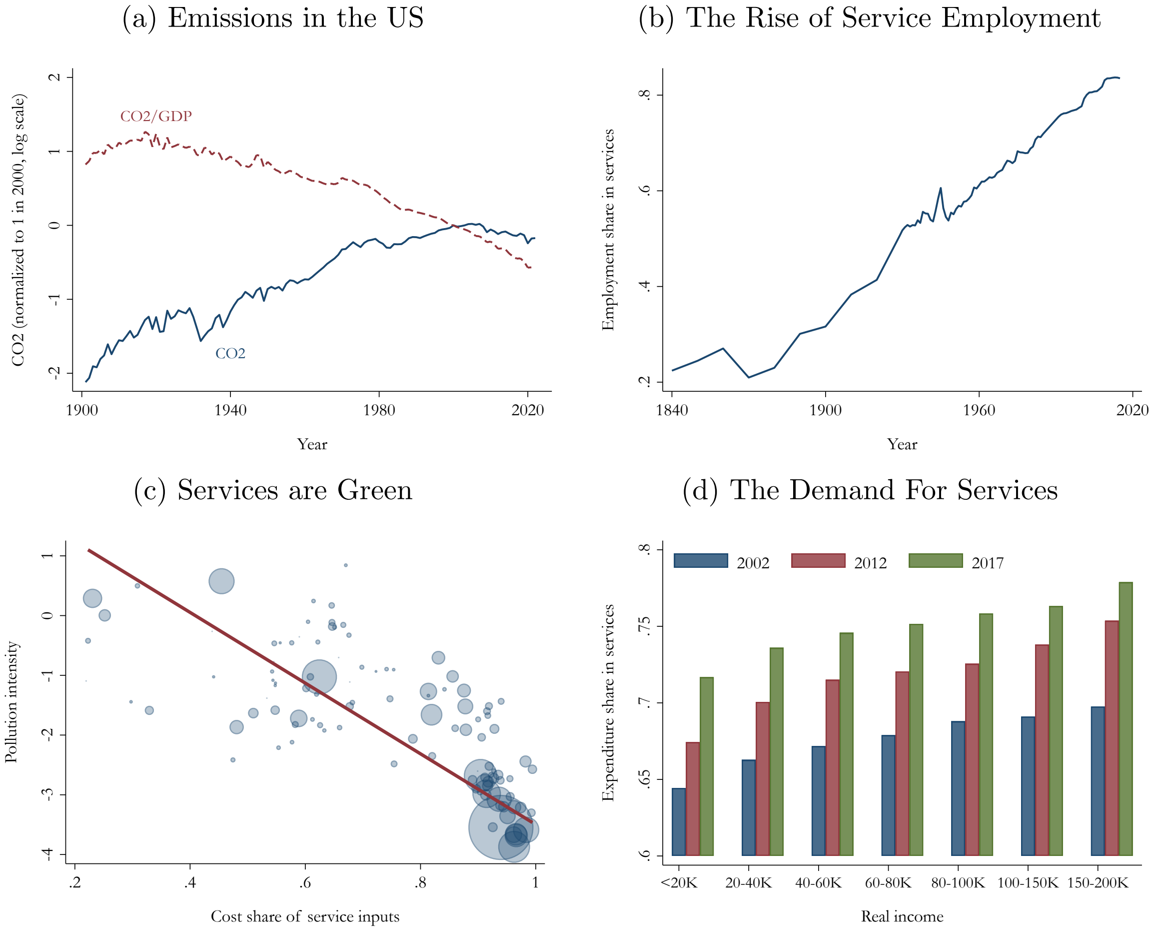

(a) Emissions in the US

(b) The Rise of Service Employment

(c) Services are Green

(d) The Demand For Services

Notes: In Panel (a) we show the total amount of annual CO2 emissions (blue) and total emissions relative to GDP (red). We normalize the respective level in the year 2000 to unity. In Panel (b) we display the service employment share in the US. In Panel (c) we display the relationship between pollution intensity and the service cost share at the industy level. The pollution intensity is taken from Levinson and O'Brien (2019). Panel (d) shows the expenditure share on service value added across consumers of different income and in different time periods. The service cost share in Panel (C) and the service expenditure share in Panel (d) takes sectoral linkages via the Input-Output matrix into account.

In Panel (c), we show that this is the case. We use data from the National Emissions Inventory (NEI) published by the U.S. Environmental Protection Agency (EPA). The NEI offers comprehensive emissions data across five pollutants—particulates smaller than 10 microns (PM10), volatile organic compounds (VOC), nitrogen oxides (NOx), sulfur dioxide (SO2), and carbon monoxide (CO)—for various industries. We calculate each industry's emissions intensity for each pollutant by aggregating total emissions and dividing by total sales, based on data from the 2002 economic and agricultural censuses.

Using sectoral linkages in the Input-Output (IO) tables, we then estimate the pollution intensity for each final product k for pollutant p, denoted by e_k^p. Likewise, the IO tables enable us to determine the service intensity of each product, λ_k, which represents the proportion of service costs in producing each good k. Panel (c) shows a strong negative correlation between total emissions e_k^p and the service share λ_k: a 10 percentage point increase in service content corresponds, on average, to a 25% reduction in emissions per dollar.

Finally, Panel (d) highlights two important features of the demand for services that are central to our theory. Using data from the Consumer Expenditure Survey (CEX), we calculate the service value-added share in consumer spending, showing the average service share both as a function of real income (x-axis) and across three periods: 2007, 2012, and 2017. First, services appear to be as luxuries: the expenditure share on service-intensive goods rises with income. Second, service content in consumer spending, holding income fixed, has grown significantly over time. In 2002, consumers earning \$80,000–\$100,000 allocated around 67% of spending to services; by 2017, this share had risen to 75%. Through the lens of our theory, this pattern has two implications. On the one hand, economic growth has increased service spending by raising income. On the other hand, quality improvements in service-intensive goods explain why the demand for services rose holding income constant.

The patterns in Figure 1 suggest that economic growth reallocates resources toward services, reducing emission intensity. Table I provides correlational evidence supporting this view. We start by analyzing cross-country data using regressions of the following form:

$$\ln\left(\frac{e}{y}\right)_{ct} = \delta_t + \delta_c + \beta s_{ct}^{SERV} + \gamma \ln y_{ct} + \phi s_{ct}^{AG} + x_{ct}'\rho + u_{ct}, \tag{1}$$

where $\frac{e}{y}$ represents emission intensity (i.e., total CO_2 emissions relative to GDP), s_{ct}^{SERV}

is the service sector employment share, $\ln y_{ct}$ denotes log GDP per capita, s_{ct}^{AG} is the agricultural employment share, and x is a vector of other country-specific covariates. Importantly, we control for agricultural employment share, so β is identified by variations in service employment relative to manufacturing.

The results of estimating equation (1) are shown in the first four columns of Table I. Column 1 reveals a significant negative relationship between pollution intensity and a country's share of service employment: a one percentage point increase in service employment is associated with a roughly 4% reduction in emissions per unit of GDP. Columns 2 and 3 add controls for GDP per capita, total population, and country size: the relationship between emissions and service employment remains robust. In column 4, we add country fixed effects, so that β is now identified from within-country changes in service employment and emission intensity over time. While the coefficient size is reduced in absolute terms, a substantial effect remains: a 1 percentage point increase in service employment is associated with a 1.6% decrease in emissions per unit of output.[10]

In columns 5 to 8, we replicate this analysis focusing on counties within the US rather than countries in the international context. Similar to the cross-country findings, there is a significant negative relationship between service employment and emission intensity, with a comparable magnitude (though somewhat smaller).[11]

In conclusion, this section highlights a robust negative empirical correlation between service activity and emissions, even when controlling for standard determinants of emissions. These findings lend empirical support to the theoretical framework developed in the following section.

3 Theory

The production sector of the economy consists of a manufacturing sector (G) and a sector for consumption goods (C) comprising J products. The manufacturing sector's

[10] We also run regressions with the logarithm of emissions (rather than emissions per GDP) as the dependent variable, including GDP as a control. These results are nearly identical.

[11] Due to data availability for only a single year, we cannot estimate the specification with county fixed effects. In column 8, we include state fixed effects, which leaves the coefficient statistically unchanged.

	Across countries				Across counties within US			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Service Empl. Sh.	-4.064***	-4.017***	-3.949***	-1.658***	-2.798***	-1.899***	-1.565***	-1.403***
	(0.722)	(0.708)	(0.658)	(0.159)	(0.534)	(0.554)	(0.403)	(0.268)
ln GDPpc		-0.063	-0.085	-0.257***		-0.991***	-0.651***	-0.401***
		(0.078)	(0.072)	(0.020)		(0.176)	(0.173)	(0.094)
Year FE	Yes	Yes	Yes	Yes				
Ag. Emp. Share	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
log Population			Yes	Yes			Yes	Yes
log Total Land			Yes	Yes			Yes	Yes
Region FE				Yes				Yes
N	4337	4337	4308	4308	3138	3138	3138	3138
R^2	.334	.336	.387	.925	.119	.158	.246	.312

Notes: The table reports the relationship between ln Pollution / GDP with the employment share in services and ln GDP per capita. Columns 1 - 4 focus on the variation across countries for the years 1991 - 2020. Columns 5 - 8 focus on the variation across counties within the US for the emission data in 2017 and 2010 Census data. We always control for the agricultural employment share. Columns 4 (8) control for country FE (state FE).

Table I: Services and Pollution Intensity: Cross-regional evidence

technology is characterized by the following CES production function:

$$Y_G = \left(\int_0^1 y_{iG}^{\frac{\xi-1}{\xi}} di \right)^{\frac{\xi}{\xi-1}},$$

where $y_{iG} = A_i h_{iG}$ represents the production function for individual manufacturing goods, and h_{iG} denotes labor utilized in the production of manufacturing good i. The productivity distribution $\{A_i\}_{i=0}^1$ evolves endogenously over time due to technical change. We will introduce standard assumptions about the microstructure (following Acemoglu (2009), Chap. 14) ensuring that, in equilibrium, $Y_G = AH_G$, where $A \equiv \left(\int A_i^{\xi-1} di \right)^{\frac{1}{\xi-1}}$ and $H_G = \int h_{iG} di$.

Consumers have preferences over J different final products and assign value to them based on utility weights determined by their quality. Each of the J consumption goods is a CES bundle comprising a unit interval of consumption good varieties. More formally, the quality-weighted consumption of good $j \in \{1, 2, ...J\}$ is given by

$$C_j = \left(\int_0^1 \left(Q_{ij}^{\alpha_j} y_{ij} \right)^{\frac{\xi-1}{\xi}} di \right)^{\frac{\xi}{\xi-1}},$$

where Q is a quality index and $\alpha_j \in [0, 1]$ captures the sensitivity of consumers' demand to quality differences for goods j. Note that α_j is product specific, indicating that this sensitivity varies across goods. For example, consumers may be more susceptible to quality differences between restaurants than between pet food brands.

Each consumption good j is produced combining units of the manufacturing good Y_G and labor services. More formally, we assume the following technology

$$y_{ij} = \left((1-\lambda_j)^{\frac{1}{\rho}} Y_{ijG}^{\frac{\rho-1}{\rho}} + \lambda_j^{\frac{1}{\rho}} h_{ijS}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}},$$

for $(i,j) \in ([0,1] \times \{1,2,...,J\})$. Here, Y_{ijG} represents the input of manufacturing good G utilized in the production of y_{ij}. It should be noted that $\lambda_j \in [0,1]$ quantifies the intensity of service of the provision of j: A higher λ_j indicates a higher intensity of service for the good consumption j. Given that all varieties i are produced with the same technology, symmetry allows us to denote the price of each y_{ij} as $\tilde{p}_j$.

The clearing of the labor market implies that $H = H_G + H_S$, where H is the exogenous supply of effective units of labor, $H_G = \int_0^1 h_{iG}\, di$ and $H_S = \sum_{j=1}^J \int_0^1 h_{ijS}\, di$. The clearing of the labor market implies that $Y_G = \sum_{j=1}^J \int_0^1 Y_{ijG}\, di$.

Preferences: Preferences are parameterized by the following indirect utility function in the PIGL class:

$$\mathcal{V}\left(e, [\tilde{p}_j]_{j=1}^J, [Q_j]_{j=1}^J\right) = \frac{1}{\varepsilon}\left(\prod_{j=1}^J \left(\frac{Q_j^{\alpha_j}}{\tilde{p}_j}\right)^{\beta_j} \times e\right)^{\varepsilon} - \sum_{j=1}^J \phi_j \left(\ln \tilde{p}_j - \alpha_j \ln Q_j\right) - v(\mathcal{P})$$

where $\sum_{j=1}^J \phi_j = 0$ and $\sum_{j=1}^J \beta_j = 1$. Here, $\tilde{p}_j$ represents the (non-quality-adjusted) market price of consumption good j, while $Q_j \equiv \left(\int_0^1 Q_{ij}^{\xi-1} di\right)^{\frac{1}{\xi-1}}$ denotes a quality index for the same good. Since the utility derived from consuming good j depends on its quality, the indirect utility $\mathcal{V}(\cdot)$ depends on the prices $\tilde{p}_j$ and the quality indices Q_j.

Finally, the additive-separable term $v(\mathcal{P})$ captures the utility loss associated with pollution, which is a public "bad". Pollution is a state variable whose law of motion we describe below. We assume that $v' > 0$, $v'' > 0$ and $\lim_{\mathcal{P}\to\bar{\mathcal{P}}} v'(\mathcal{P}) = \infty$, for some $\bar{P} < \infty$. We will refer to $\bar{P}$ as the environmental disaster threshold.

In our analysis, it will be useful to rewrite the indirect utility in term of a set of quality-adjusted or "hedonic" prices $p_j \equiv \tilde{p}_j / Q_j^{\alpha_j}$. Namely,

$$\mathcal{V}\left(e, [p_j]_{j=1}^J\right) = \frac{1}{\varepsilon}\left(\prod_{j=1}^J \frac{e}{p_j^{\beta_j}}\right)^{\varepsilon} - \sum_{j=1}^J \phi_j \ln p_j - v(\mathcal{P})$$

As we show in Appendix Section B-1, Roy's Identity implies that expenditure share on product k for a consumer with spending level e is given by

$$\vartheta_k\left(e, [p_j]_{j=1}^J\right) = \beta_k + \phi_k \left(\frac{e}{\prod_{j=1}^J p_j^{\beta_j}}\right)^{-\varepsilon}. \tag{2}$$

Equation (2) highlights the role of the demand parameters β_k and ϕ_k. The parameter ϕ_k determines whether product k is income-elastic or income-inelasitc: all goods k with $\phi_k < 0$ are classified as luxuries, whereas those with $\phi_k > 0$ are necessities. In turn, the parameter β_k represents the asymptotic expenditure share as spending e gets large.

Next, we introduce our key assumption.

Assumption 1 (**The Sophistication Ladder**). *Consumption goods $j \in \{1, 2, ...J\}$ are ranked on a sophistication ladder, wherein good j' is more sophisticated than good j'' if and only if $j' > j''$. Moreover, $\forall j \in \{1, 2, ...J-1\}$ $\phi_{j+1} \leq \phi_j$, $\lambda_{j+1} \leq \lambda_j$, and $\alpha_{j+1} \geq \alpha_j$.*

This important assumption postulates that consumption goods are ranked on a sophistication ladder where growing sophistication is associated with a higher service intensity in production, a higher expenditure elasticity (luxury goods), and a greater salience of the quality aspect. For example, compared to food at home, meals in gourmet restaurants are a luxury good, are more service-intensive, and consumers are willing to pay a higher premium for quality.

Emissions: We assume that the production of manufacturing goods generates a negative externality to consumers, which we call *pollution*. Pollution, denoted by $\mathcal{P}$, is a state variable that evolves according to the following law of motion:

$$\mathcal{P}_t = (1-\delta)\,\mathcal{P}_{t-1} + \mathcal{E}_t \tag{3}$$

where $\mathcal{E}$ denote the flow of new emissions (worldwide in our quantitative two country version). Emissions $\mathcal{E}$ in turn are given by

$$\mathcal{E}_t = \varepsilon_G \frac{Y_{Gt}}{a^t}, \tag{4}$$

where ε_G is a parameter (which in our quantitative two country version we allow to differ across countries). Note that we assume that only goods production leads to emissions, whereas the weightless service part does not pollute. The term $a \geq 1$ captures the effect of green technology on emission reduction that we keep exogenous in this paper, for simplicity. Intuitively, conditional on the level of material production, the emissions fall over time due to abatement or general increases in fuel efficiency.

Recall our assumption that whenever $\mathcal{P}_t$ crosses the threshold $\bar{\mathcal{P}}$, it triggers an environmental disaster. Then, according to the law of motion (3) and the associated equation (4) for emissions, a necessary condition to avoid a disaster is that, asymptotically, the growth rate of Y_G is bounded from above by $a - 1$. If H_{Gt} remains constant in the long run (as we will demonstrate is the case in our equilibrium), then $a - 1$ serves as an upper bound to the growth rate of material productivity A that is environmentally sustainable. However, this condition alone is not sufficient to prevent the disaster. Indeed, it is conceivable that the threshold is crossed during the transition period, even if the necessary condition we discussed is met.

3.1 Two Consumer Goods

In our main analysis, we assume that $J = 2$, namely, there are on;y two consumer goods. We designate the index B (a mnemonic for basic good) for $j = 1$, and the index S (a mnemonic for sophisticated or service-intensive good) for $j = 2$. We assume $\lambda_B = 0$ and $\lambda_S = \lambda > 0$, indicating that the S good is service-intensive. Furthermore, we assume that consumers are indifferent to quality heterogeneity in B, hence we set $\alpha_B = 0$, while they exhibit sensitivity to quality heterogeneity in S. We set $\alpha_S = 1$, implying that, as far as the good S is concerned, consumers ultimately care about the number of quality units they purchase (in particular, $p_S = \tilde{p}_S/Q$).

3.2 Equilibrium Given Technology

The static equilibrium determines the allocation of labor between manufacturing production and services. It also determines the allocation of manufacturing goods between the production of the consumption goods B and S. Since the B good doesn't necessitate any service input, this is tantamount to stating that the equilibrium dictates the division of manufacturing production between final consumption and usage as input in the production of the S good. We proceed to the characterization of the equilibrium by considering first the production side and then the demand side of the economy.

Production: We assume that the different varieties of manufacturing and final goods are produced by monopolists under a regime of fully enforced intellectual property rights. In the appendix, we derive the standard result that, given the isoelastic demand for different varieties, monopolistically competitive firms set the prices of each variety equal to a constant markup over the marginal cost. Aggregating over the set of varieties yields the following expression:

$$\tilde{p}_B = p_G = \frac{\xi}{\xi - 1}\frac{w}{A} \qquad \text{and} \qquad \tilde{p}_S = \frac{\xi}{\xi - 1} c\left(p_G, w\right), \tag{5}$$

where w is the workers' wage, $A \equiv \left(\int A_i^{\xi-1} di\right)^{\frac{1}{\xi-1}}$ is the average productivity in manufacturing and $c\left(p_G, w\right) = \left((1-\lambda)\, p_G^{1-\rho} + \lambda w^{1-\rho}\right)^{\frac{1}{1-\rho}}$ is the unit cost of production of the S goods. Substituting in the expression of p_G in Equation (5), we obtain:

$$c\left(p_G, w\right) = w \times \psi(A), \quad \text{where} \quad \psi(A) = \left((1-\lambda)\left(\frac{\xi}{\xi-1}\right)^{1-\rho} A^{\rho-1} + \lambda\right)^{\frac{1}{1-\rho}}.$$

Note that higher quantity productivity A reduces the prices of both goods B and S. However, the effect on sophisticated goods operates via $\psi(A)$ and is thus mediated by the fact that sophisticated goods also require service-labor as an input. Note also that these market prices are independent of quality. To see the effect of quality, write (5) in terms of *quality-adjusted* prices:

$$p_B = p_G = \frac{\xi}{\xi - 1}\frac{w}{A} \qquad \text{and} \qquad p_S = \frac{\xi}{\xi - 1}\frac{1}{Q}\psi\left(A\right) w,$$

where $Q = \left(\int Q_i^{\xi-1} di\right)^{\frac{1}{\xi-1}}$ is the average quality of S varieties. Hence, higher quality reduces the effective price of sophisticated goods. In the rest of the analysis, we choose the wage as the numéraire, i.e., we set $w = 1$.

Demand: Consider, next, the demand side. We can write the (PIGL) indirect utility function as follows:

$$\mathcal{V}\left(e, p_G, p_S, Q\right) = \frac{1}{\varepsilon}\left(\frac{e}{p_S^{1-\beta} \times p_G^{\beta}}\right)^{\varepsilon} - \phi \times \left(\ln p_S - \ln p_G\right) - v(\mathcal{P}).$$

Note that, for simplicity, we have written $\mathcal{V}(\cdot)$ as a function of p_G rather than p_B. Note also that in case of two goods, the homogeneity restrictions $\sum_{j=1}^{J} \phi_j = 0$ and $\sum_{j=1}^{J} \beta_j = 1$ imply that $\beta_B = \beta = 1 - \beta_S$ and $\phi_B = \phi = -\phi_S$. Together with Assumption 1 this implies that $\phi > 0$, i.e. the basic good is a necessity.

Equation (2) thus implies that the expenditure shares that an individual with spending level e allocates to the final goods B and S are:

$$\begin{aligned}
\vartheta_B\left(e, p_G, p_S\right) &= \beta + \phi\left(\frac{e}{p_S^{1-\beta} p_G^{\beta}}\right)^{-\varepsilon} = \beta + \phi\left(\Upsilon\left(e; A, Q\right)\right)^{-\varepsilon}, \qquad (6)\\
\vartheta_S\left(e, p_G, p_S\right) &= 1 - \beta - \phi\left(\frac{e}{p_S^{1-\beta} p_G^{\beta}}\right)^{-\varepsilon} = 1 - \beta - \phi\left(\Upsilon\left(e; A, Q\right)\right)^{-\varepsilon},
\end{aligned}$$

where Υ summarizes the effect of non-homothetic demand and is given by

$$\Upsilon\left(e; A, Q\right) = \frac{A^{\beta} Q^{1-\beta}}{\psi\left(A\right)^{1-\beta}} \times \frac{\xi - 1}{\xi} e. \qquad (7)$$

Conditional on the expenditure level, both productivity growth in manufacturing and enhancements in quality contribute to a shift in expenditure share from the basic good, which is a necessity, to the sophisticated good, which is a luxury. We will demonstrate below that the function Υ increases with both Q and A when evaluated at the equilibrium value of e.

Figure 2 illustrates the fundamental properties of the demand system. First, the demand system, as defined by the above equations, closely resembles a Cobb–Douglas specification with a nonhomothetic adjustment. The slope of the Engel curves and the magnitude of income effects are determined by the parameter ε. This parameter, which we term the *Engel elasticity*, plays a pivotal role in our analysis. Second, as $e \to \infty$, the expenditure shares ϑ_B and ϑ_S converge to their limiting values, β and $1 - \beta$, respectively. The spending share on the quality-intensive good S approaches $1 - \beta$ from below, while the spending share on the basic good B approaches β from

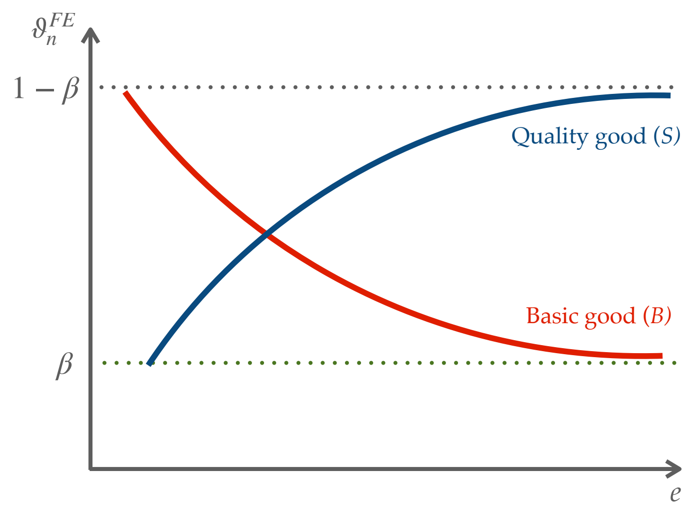

Notes: The figure shows the expenditure share for basic goods (red) and sophisticated goods (blue) as a function of expenditure – see (6).

Figure 2: Engel Curves

above.

As seen in the expression for $\Upsilon(e; A, Q)$, the quality level Q functions as a demand shifter, similar to an increase in real income: higher quality decreases the spending share on basic goods while increasing the share on sophisticated goods for a given level of nominal spending e and fixed prices. Similarly, a rise in productivity, A, also shifts spending toward quality goods. This occurs because greater productivity lowers the prices of both basic and quality goods, effectively raising real income.

Representative Household: We assume the economy is populated by a large, representative household that earns income from labor and firms' profits. In our setting, all household income is allocated toward consumption goods. Therefore, the equilibrium expenditure of the representative household can be expressed as:

$$e = \frac{wH + \Pi}{H},$$

where Π denotes aggregate profits in the economy, and H is the aggregate labor force. Because of the constant markup (see Appendix), profits are proportional to wage income, expressed as $\Pi = \frac{1}{\xi-1} wH$. This implies that

$$e = \frac{wH + \frac{1}{\xi-1} wH}{H} = \frac{\xi}{\xi - 1},$$

where the last equality follows from setting $w = 1$. Thus, with some slight abuse of notation, we can rewrite the real income term Υ in (7) as

$$\Upsilon(A, Q) = A^{\beta} \left((1-\lambda) \left(\frac{\xi}{\xi - 1} \right)^{1-\rho} A^{\rho-1} + \lambda \right)^{\frac{1-\beta}{\rho-1}} Q^{1-\beta}. \quad (8)$$

The demand shifter Υ is fully determined from Q and A and increasing in both arguments.

Labor Market Equilibrium: We now proceed to characterize the equilibrium allocation of labor between manufacturing and services. To do so, we use the market-clearing conditions, which stipulate that, for both goods and service labor, the total factor payment (including wages and profits) must equal the associated value added. First, consider the market-clearing condition for G, where demand arises from the production of both consumption goods B and S:

$$\frac{\xi}{\xi - 1} wH_G = \vartheta_B(e, p_G, p_S) \frac{\xi}{\xi - 1} wH + \sigma_G \left(\vartheta_S(e, p_G, p_S) \frac{\xi}{\xi - 1} wH \right), \quad (9)$$

where σ_G denotes the cost share of good inputs in the production of the sophisticated consumption good S. Given the CES production function, we obtain

$$\sigma_G = \frac{(1-\lambda) p_G^{1-\rho}}{(1-\lambda) p_G^{1-\rho} + \lambda w^{1-\rho}} = \frac{(1-\lambda) \left(\frac{\xi}{\xi-1} \right)^{1-\rho} A^{\rho-1}}{(1-\lambda) \left(\frac{\xi}{\xi-1} \right)^{1-\rho} A^{\rho-1} + \lambda}.$$

The cost share σ_G is fully determined from the quantity productivity A. Moreover, it is decreasing in A if $\rho < 1$, i.e., services and goods are complements, and increasing in A if $\rho > 1$, i.e., services and goods are substitutes.

Next, consider the market for service workers. In this case, the demand stems exclusively from the cost share of services in the production of the sophisticated good

S:

$$\frac{\xi}{\xi-1} w H_S = \frac{\xi}{\xi-1} (1-\sigma_G) \vartheta_S (e, p_G, p_S), \tag{10}$$

where $1-\sigma_G$ is the cost share of services in the production of consumption good S.

Substituting the expressions of ϑ_G and ϑ_S into the market clearing conditions (9) and (10), and recalling the normalization $w=1$, we obtain:

$$H_S = \left(1-\beta-\phi \Upsilon (A,Q)^{-\varepsilon}\right) \frac{\lambda}{(1-\lambda)\left(\frac{\xi}{\xi-1}\right)^{1-\rho} A^{\rho-1} + \lambda} H \tag{11}$$

and

$$H_G = H - H_S. \tag{12}$$

In Appendix B-1, we prove that the following comparative statics hold.

Proposition 2. *Consider the service employment level H_S given in (11). Then:*

1. *H_S is increasing in Q as long as the preferences are nonhomothetic, that is, $\phi > 0$*
2. *H_S is increasing in A if $\rho < 1$*
3. *H_S is decreasing in A if ρ is sufficiently high and β is sufficiently close to 1.*

Proposition 2 highlights the distinct roles of Q and A. An increase in Q shifts demand towards sophisticated goods through an income effect, leaving the factor allocation within products, σ_G, unchanged. As a result, higher quality Q raises the aggregate employment share of services, given the service-intensive nature of sophisticated goods. Conversely, an increase in A has two effects. First, similar to an increase in Q, it enriches households, shifting demand towards sophisticated goods and increasing the service employment share. Second, it impacts the cost structure in the production of sophisticated goods. If service workers and manufacturing inputs are complementary ($\rho < 1$), a rise in A increases the cost share of service workers, amplifying the income effect. However, if they are substitutes ($\rho > 1$), a rise in A raises the cost share of manufacturing goods. When β is large, this substitution effect may dominate, potentially increasing employment in the G sector. Throughout the remainder of this paper, we focus on the empirically-relevant case of complementarity, where $\rho < 1$.

3.3 Directed Innovation: Quality versus Productivity

We now examine the determinants of technical progress. We postulate the existence of a cohort of researchers, with a mass R, capable of directing their research endeavors towards enhancing either the productivity of the varieties of manufacturing goods (A_i) or the quality of the varieties of the S good (Q_i). We denote by R_Q and R_A, respectively, the mass of researchers directing their research efforts to increase Q and A, subject to the standard labor market clearing condition for researchers, $R_Q + R_A = R$.

The probability that a unit of research effort directed toward activity $s \in A, S$ succeeds in generating an innovation is given by $\eta_s R_s^{-\zeta}$, where the parameter ζ quantifies the degree of congestion in research. The parameter η_s denotes the productivity of the innovation technology. A successful innovation augments the quality or productivity of a randomly selected firm in sector s by a factor $\gamma_s > 1$. We assume γ_Q and γ_A to be sufficiently large to enable the new firm to set the unconstrained monopoly price (drastic innovation). Furthermore, we assume that researchers reap profits only for a single period.[12]

Let V_s denote the expected value of directing research towards $s \in \{A, Q\}$. Then:

$$V_s = \underbrace{(1-\tau_s)\left(\eta_s R_s^{-\zeta}\right)}_{\text{Probability of innovation}} \times \underbrace{\int \pi_{ij} di}_{\text{Expected value conditional on innovating}} ,$$

where $j = G$ if $s = A$ and $j = S$ if $s = Q$. Here, (i) π_{iG} and π_{iS} denote, respectively, the profits from a productivity innovation in intermediate sector i of sector G and the profits from a productivity innovation in intermediate sector i of sector S; (ii) τ_s is a wedge (e.g., a tax or subsidy) on $s-$type innovation. These wedges will play a role in the policy analysis because they affect the direction of innovation.

In the Appendix, we show that the equilibrium profits are equal to:

$$\pi_{iS} = \frac{1}{\xi - 1}\frac{1}{\lambda}\frac{(\psi(A))^{1-\rho} H_S w}{Q^{\xi-1}} Q_i^{\xi-1} \text{ and } \pi_{iG} = \frac{1}{\xi-1}\frac{H_G w}{A^{\xi-1}} A_i^{\xi-1} \quad ,$$

[12] This is a simplifying assumption aimed to retain analytical tractability. A rationale for it is that patents confer one-period monopoly rights to innovating firms. Subsequently, a fringe firm, selected randomly from a continuum of firms, attains monopoly power for another period, and so forth. This assumption ensures that each variety's price constitutes a constant markup over marginal cost, averting complications stemming from price disparities between monopolized and competitive varieties. However, the incentive to innovate is determined by a one-period profit rather than the discounted value of future profits.

where, recall, H_S and H_G denote employment in production and services respectively. The terms H_S and H_G reflect market size effects: namely, a higher H_S increases the rents to quality innovations whereas a higher H_G increases the rents of research directed toward quantity innovation. This follows from our assumption that the production of the sophisticated good is intensive in services.

In equilibrium, researchers will be indifferent between directing their efforts toward quality or productivity improvements. Substituting the expressions for V_Q and V_A, this arbitrage condition yields:

$$\frac{R_Q}{R_A} = \left(\frac{(1-\tau_Q)\eta_Q}{(1-\tau_A)\eta_A}\left(\frac{\gamma_Q}{\gamma_A}\right)^{\xi-1}\right)^{\frac{1}{\zeta}} \left(\frac{(1-\lambda)\left(\frac{\xi}{\xi-1}\right)^{1-\rho} A^{\rho-1}+\lambda}{\lambda}\right)^{\frac{1}{\zeta}} \left(\frac{H_S}{H_G}\right)^{\frac{1}{\zeta}}. \tag{13}$$

This expression shows that the relative allocation of researchers to the two activities depends on two endogenous factors: (i) the term H_S/H_G, which reflects the market size effect discussed above. The larger the share of service employment, the more profitable Q innovations become relative to A innovations, leading to more research directed toward quality innovation; (ii) the term $\left((1-\lambda)\left(\frac{\xi}{\xi-1}\right)^{1-\rho} A^{\rho-1}+\lambda\right)$, which reflects technological complementarity or substitution in the S sector. Consider the effect of A while holding employment in goods and services constant. If $\rho < 1$, conditional on H_S/H_G, technical progress in material production discourages innovation directed toward quality relative to material production. This is because the production of sophisticated goods increasingly relies on labor rather than goods, reducing the returns for further quantity innovation.

However, technical progress in manufacturing also impacts the allocation of labor. If goods and services are complementary inputs in the production of S goods, an increase in A shifts employment from goods to services. This shift reverses the comparative statics established above under fixed employment shares, driven by the market size effect. To demonstrate this formally, we substitute the equilibrium expressions for H_S and H_G into the arbitrage equation (13). We obtain:

$$\frac{R_Q}{R_A} = \left(\frac{(1-\tau_Q)\eta_Q}{(1-\tau_A)\eta_A}\left(\frac{\gamma_Q}{\gamma_A}\right)^{\xi-1}\right)^{\frac{1}{\zeta}} \left(\frac{\left(1-\beta-\phi\Upsilon(A,Q)^{-\varepsilon}\right)}{1-\frac{\left(1-\beta-\phi\Upsilon(A,Q)^{-\varepsilon}\right)\lambda}{(1-\lambda)\left(\frac{\xi}{\xi-1}\right)^{1-\rho}A^{\rho-1}+\lambda}}\right)^{\frac{1}{\zeta}}. \tag{14}$$

The right-hand side of (14) captures the effects of both technological forces and demand. To isolate these effects, consider the case in which $\phi = 0$, corresponding to Cobb-Douglas homothetic preferences. Then, equation 14 implies that R_Q/R_A increases in A if and only if $\rho < 1$. The intuition is as follows: when goods and services complement each other in producing S goods, technical progress in manufacturing increases the labor share of services (Baumol effect), thereby enhancing the market size effect for quality-focused innovation. In our model, this effect consistently dominates the partial effect discussed above. Consequently, technological forces alone imply that an increase in A shifts innovation incentives toward quality. Notably, with $\phi = 0$ (Cobb-Douglas preferences), Q has no impact on the direction of technical progress.

Next, consider the more general case where an income effect is present, i.e., $\phi > 0$. Here, increases in both A and Q make consumers wealthier, prompting a shift in expenditure from B to S and thus incentivizing more innovation aimed at enhancing quality.

The equilibrium growth rates of quality and TFP in manufacturing are determined by the allocation of research: $g_Q = R_Q(\gamma_Q - 1)$ and $g_A = R_A(\gamma_A - 1)$. Our analysis in this section demonstrates how economic development increasingly shifts innovation and growth toward quality over time.

3.4 Dynamic Equilibrium

In this section, we characterize the dynamic equilibrium. Under our assumptions, the model results in a straightforward backward-looking dynamic system in terms of the state variables (Q_t, A_t), which fully characterizes the equilibrium path.

Given (Q_t, A_t), we can compute consumer's income Υ from (8) and the sectoral labor allocation (see (11) and (12)). This determines the static equilibrium. At the same time, (Q_t, A_t) also fully determines the allocation os research labor — see (14). This in turn implies the law of motion for both quality and productivity growth.

Asymptotically, A and Q grow without bound, implying that $\vartheta_B \to \beta$ and $\vartheta_S \to 1 - \beta$. The asymptotic direction of technical progress hinges on the technological parameter ρ. If goods and services are gross substitutes in the production of the S good (i.e., $\rho > 1$), then $\sigma_G \to 1$ as time goes infinity. In this case, all workers are eventually employed in the production of material goods as even sophisticated products can be produced with intermediate goods as opposed to service workers. The equilibrium allocation of researchers then converges to

$$\frac{R_Q}{R_A} = \left(\frac{\eta_Q}{\eta_A} \frac{1-\tau_Q}{1-\tau_A} (1-\beta) \left(\frac{\gamma_Q}{\gamma_A} \right)^{\xi-1} \right)^{\frac{1}{\zeta}} \tag{15}$$

However, if $\rho < 1$—which we consider the empirically relevant case—the theory predicts that $\sigma_G \to 0$, $H_G \to \beta H$, and $H_S \to (1-\beta)H$. Intuitively, this implies that, over time, goods become so efficiently produced that service workers dominate the production costs of sophisticated products. In this case, the long-run arbitrage condition in research (cf. Equation (14)) yields:

$$\frac{R_Q}{R_A} = \left(\frac{\eta_Q}{\eta_A} \frac{1-\tau_Q}{1-\tau_A} \frac{1-\beta}{\beta} \left(\frac{\gamma_Q}{\gamma_A} \right)^{\xi-1} \right)^{1/\zeta} \equiv \Phi \tag{16}$$

When β is small (i.e., when the share of basic goods expenditure is low in a highly affluent economy), Φ becomes large, so most innovation is directed toward enhancing quality rather than reducing material production costs. As a result, a large share of researchers focuses on quality innovations—formally, this share is given by $\frac{\Phi}{1+\Phi}$.

The asymptotic equilibrium growth rates of quality Q and quantity productivity A are then given by

$$g_A = \eta_A \left(\frac{1}{1+\Phi} R \right)^{-\zeta} (\gamma_A - 1) \quad \text{and} \quad g_Q = \eta_Q \left(\frac{\Phi}{1+\Phi} R \right)^{-\zeta} (\gamma_Q - 1). \tag{17}$$

In turn, aggregate GDP growth is given by an expenditure-weighted average of these growth rates:

$$g_{GDP} = \beta \cdot g_A + (1-\beta) \cdot g_Q. \tag{18}$$

Equations (17) and (18) underscore three key factors that drive long-run growth increasingly towards quality improvements over material production:

(i) *Innovation*: A higher efficiency in quality-enhancing research relative to productivity-enhancing research (i.e., $\eta_Q > \eta_A$ and $\gamma_Q > \gamma_A$).

(ii) *Policy*: A lower wedge on quality innovation compared to productivity innovation (i.e., $\tau_Q < \tau_A$).

(iii) *Preferences*: A higher asymptotic expenditure share on sophisticated goods relative to basic goods (i.e., a small β).

4 GDP, Pollution, and Degrowth

We now return to the core of our motivation. First, does economic growth inevitably lead to unbounded environmental degradation, or are there viable policy interventions that could avert this path? Second, is degrowth necessary to prevent such an outcome?

To address these questions, we revisit the theory's predictions regarding the pollution trajectory. We assume that pollution originates from material production, specifically Y_G (see (4)).[13] Utilizing the employment allocation results in (11) and (12), the production of physical goods, Y_G, is expressed as:

$$Y_G = AH_G = A(H - H_S) = A\left(1 - \frac{\lambda\left(1 - \beta - \phi\Upsilon\left(A, Q\right)^{-\varepsilon}\right)}{(1-\lambda)\left(\frac{\xi}{\xi-1}\right)^{1-\rho} A^{\rho-1} + \lambda}\right) H$$

In the long run, if both A and Q grow unboundedly and $\rho < 1$, then $Y_G \to A\beta H$, as $A^{\rho-1} \to 0$ and $\Upsilon^{-\varepsilon} \to 0$. Consequently, the flow of new emissions increases at the rate of A, g_A, which is determined by equations (16)–(17). To avert an environmental disaster, it is necessary that $g_A < a - 1$. Equation (16) shows that by imposing a sufficiently large wedge on innovation in material production, policy can achieve this target. In standard growth models, such a policy would necessitate drastic restrictions on economic growth. However, in our model, Pigouvian taxation can accomplish this by shifting the focus of innovation from quantity to quality. How does this conclusion inform the debate on degrowth?

To address this question, let us examine the evolution of economic activity within our model. In models with nonhomothetic preferences, defining real GDP becomes ambiguous since expenditure shares across different goods vary with income. To earn some insight, we focus initially on the asymptotic economy, where expenditure shares

[13] Our results do not hinge on this assumption; it is straightforward to generalize to a model where services also contribute to pollution.

are approximately constant, allowing for a standard GDP definition. In this case,[14]

$$GDP_{market} \approx \frac{e}{\tilde{p}_S^{1-\beta} \times \tilde{p}_G^{\beta}}.$$

This measure reflects expenditure at market prices. Using $e = \xi/(\xi - 1)$ along with the equilibrium expressions for $\tilde{p}_G$ and $\tilde{p}_S$ yields:

$$GDP_{market} = \frac{A^{\beta}}{\left((1-\lambda)\left(\frac{\xi}{\xi-1}\right)^{1-\rho} A^{\rho-1} + \lambda\right)^{\frac{1-\beta}{1-\rho}}}.$$

The expression for GDP_{market} is increasing in A and independent of Q—in other words, it does not account for quality improvements. In the long run, as the denominator approaches a constant, the growth rate of GDP_{market} is given by $\beta \times g_A$. Thus, if environmental sustainability imposes a stringent constraint (i.e., if a is a small number), degrowth becomes a necessary condition for preserving the planet.

However, GDP_{market} is not a welfare-relevant measure of GDP. We can construct an adjusted measure that accounts for quality improvements. Define $GDP_{adjusted} \approx \frac{e}{p_S^{1-\beta} \times p_G^{\beta}}$, using hedonic prices instead of market prices. Our equations imply that

$$GDP_{adjusted} = Q^{1-\beta} \times GDP_{market}$$

While $GDP_{adjusted}$ is the theoretically *correct* measure of GDP, properly accounting for quality changes is in practice very difficult, especially in service-related activities (see Bils and Klenow (2001)). Interestingly, our model suggests that $GDP_{adjusted}$ could continue to grow even in a hypothetical scenario where material production growth has ceased entirely.

The optimistic tone of this discussion should not be mistaken as an endorsement of policy inaction. According to our theory, there is no guarantee that a laissez-faire approach would even meet the necessary conditions to avert an environmental catastrophe. Without policy intervention, the growth rate of material production could exceed sustainable levels, potentially leading to escalating pollution, as highlighted by numerous climate science studies.

[14] In the numerical analysis below, we use chained indices with Törnqvist weights rather than assuming constant expenditure shares.

Finally, we recall that our discussion thus far has only addressed the necessary conditions to prevent an environmental crisis. It is possible that, even if the long-term growth rate of A remains below $a - 1$, the critical threshold for environmental disaster could still be crossed at an earlier stage of development, marked by a significant share of industrial production. We will return to this concern below.

5 Quantitative Analysis

In this section, we calibrate our model to assess the role of quality-led growth in shaping the future trajectory of economic growth and environmental sustainability.

5.1 Data

Our analysis relies on three primary data sources: (i) the Consumer Expenditure Survey (CEX), (ii) Input-Output (IO) Tables, and (iii) Environmental Accounts. Below, we provide a brief description of these datasets and our methodology; additional details are available in Appendix Section A-1.

To capture the distribution of individual spending across final goods, we use the 2002 Consumer Expenditure Survey (CEX), which reports consumption expenditures for approximately 12,000 households across 472 final good categories. This allows us to compute household i's expenditure share on each final product k, denoted as ϑ_k^i.

To assess the service content within these final goods, we use data from the 2002 IO Tables, which report intermediate input contents by sector. As we describe in detail in Section A-1 in Appendix, we use the IO Table together with BEA's bridge table to compute the total service content embodied in the output of each industry. We then use the cross-walk from industries to final goods as observed in the CEX to compute the service share of each product k, s_k.

Given the expenditure shares of individuals across products, ϑ_k^i, and the service share of individual products, s_k, we can compute the service content of individual i, as

$$\vartheta_{\mathcal{S}}^i = \sum_k \vartheta_k^i s_k. \tag{19}$$

Finally, we calculate the environmental footprint of each final product k. For this, we rely on the National Emissions Inventory (NEI) from the EPA, which reports total emissions for five pollutants (particulates smaller than 10 microns (PM10), volatile organic compounds (VOC), nitrogen oxides (NOx), sulfur dioxide (SO2), and carbon monoxide (CO)) by industry. We calculate the emissions intensity for each pollutant by aggregating total emissions and dividing by total sales using data from the 2002 economic and agricultural censuses. Sectoral linkages in the IO tables are then used to compute the pollution intensity for pollutant p of each final product k, denoted as e_k^p.

5.2 Calibration Strategy and Estimation Results

Our model economy is characterized by 13 structural parameters

$$\mathcal{P} = \{\underbrace{\phi, \varepsilon, \beta, \xi}_{\text{Preferences}}, \underbrace{\rho, \lambda, a}_{\text{Technology}}, \underbrace{R, \zeta, [\eta_s]_{A,Q}, [\gamma_s]_{A,Q}}_{\text{Innovation}}\} \tag{20}$$

and two initial conditions for the initial level of productivity A_0 and quality Q_0. Household preferences are described by the Engel elasticity ε, the asymptotic expenditure share on good-intensive products β, the preference shifter ϕ (which determines whether service-intensive goods are luxuries), and the elasticity of substitution across individual varieties, ξ. The production side is defined by the elasticity of substitution between goods and services in the production of sophisticated goods, ρ, the service intensity λ, and the rate a at which basic goods production becomes cleaner. Finally, the process of innovation is governed by the mass of researchers R, the decreasing returns of the innovation technology ζ, the sector-specific cost shifter of the R&D technology, η_s, and the sector-specific stepsize parameter γ_s.

We calibrate the parameters in (20) and the initial conditions (A_0, Q_0) by targeting key aspects of the structural transformation of the U.S. economy over the past century. Importantly, our calibration does not assume that the economy has reached its balanced growth path (BGP). Although we calibrate all parameters simultaneously, there remains a clear mapping between specific moments and individual parameters, which we describe in detail as part of our calibration strategy.

Household Preferences: ε, ϕ, β, and ξ The Engel elasticity, ε, is an important parameter because it determines the strength of income effects. We estimate ε from the cross-sectional correlation between household expenditure and the expenditure share on service-intensive goods. Our theory implies a positive relationship between household income and the service share $\vartheta_{\mathcal{S}}^{i}$ (see equation (19)). To test this implication empirically, we estimate regressions of the form

$$\vartheta_{\mathcal{S}}^{i} = \gamma \ln e_i + x_i' \psi + u_i, \tag{21}$$

where e_i represents household spending and x_i includes various observable characteristics that could influence the distribution of household expenditures and may be correlated with spending. In practice, we control for household size, geographic location, education, race, and marital status. Our main parameter of interest, γ, captures the degree to which higher-income households consume goods with a greater service content.

	Service exp. share			log (1-β - service exp. share)
	(1)	(2)	(3)	(4)
Log (Exp)	0.013***	0.082***	0.066***	-0.292***
	(0.001)	(0.005)	(0.008)	(0.037)
Family Size	Yes	Yes	Yes	Yes
HH Controls	No	No	Yes	Yes
IV	No	Yes	Yes	Yes
F-Stat First Stage		2960	1562	913
N	11972	10252	9896	9800
1-β				0.9

Notes: Column 1 reports the OLS relationship between households' expenditure share on services and their expenditure. Columns 2 and 3 report the IV estimate using occupational fixed effects as instruments for household expenditure. Column 4 uses as dependent variable $\ln(1-\beta-\vartheta_{\mathcal{S}}^{i})$, where the asymptotic service share $1-\beta$ is given by 0.9. All specifications control for a set of fixed effects for the size of the household. Columns 3 and 4 control for the geographic location of the household, education, race and marital status.

Table II: Nonhomothetic service demand: Estimating ε

Table II reports the results. In column 1, we report the simple bivariate correlation, controlling only for household size. In column 2, we present results from an instrumental variables strategy to account for measurement error in household expenditure e_i and to identify γ from variation in permanent rather than transitory income. Specifically, we instrument total spending with a full set of occupation fixed effects. These fixed effects strongly predict household spending, and the resulting estimate for γ is 0.082. In column 3, we control for the additional household characteristics mentioned above.

This adjustment lowers the estimate for γ to 0.066, reflecting their correlation with the occupational fixed effects.

In the last column we use the same variation to estimate the structural parameter ε. Equation (2) implies that

$$\ln(1-\beta-\vartheta_{\mathcal{S}}^{i}) = -\varepsilon \ln e_i + \varepsilon \ln\left(\phi \prod_k p_k^{\beta_k}\right) \tag{22}$$

Hence, the elasticity between the distance of the asymptotic expenditure share $1-\beta$ and the actual expenditure share on the quality good $\vartheta_{\mathcal{S}}^{i}$ should be constant and given by the Engel elasticity ε. Note that the second term, $\phi \prod_k p_k^{\beta_k}$, is common across individuals and hence absorbed in the constant of the regression.

To implement (22) we need to know the value of β. Our theory implies that all individuals' expenditure shares on the quality good should be bounded above by β. We set $\beta = 0.1$, which is close to the 99% percentile of the observed distribution of expenditure shares. As seen in column 4 of Table II, we estimate an Engel elasticity of about 0.3.

The preference parameter ϕ is not separately identified from the price level $p_S^{1-\beta} p_B^{\beta}$. We thus normalize $\phi = 1$, making services luxury products as implied by Table II.[15] Finally, we set the elasticity of substitution ξ to 5, a consensus estimate in the literature.

Technology parameters: ρ and λ The parameter λ determines the weight of service inputs within the production function of sophisticated goods. This parameter directly maps to the service share of final commodities observed in the input-output table. To incorporate this information in a model-consistent way, recall that, for simplicity, our theory considers only two final goods: sophisticated and basic goods. We therefore aggregate the data to reflect this distinction. Specifically, we classify all final goods into two mutually exclusive groups: service-intensive ("sophisticated") products and goods-intensive ("basic") products. To do this, we rank final products by their service intensity, λ_k, categorizing those with service intensity above the median as service-intensive. Empirically, the cost share of services among sophisticated goods is 0.93, and we calibrate λ in our model to match this observed moment.

[15] This normalization is without loss of generality. The data identify a nonlinear function of ϕ and the initial condition Q_0, with the estimated value of this function ensuring that quality is a luxury. Thus, setting ϕ serves as a pure normalization. Further details can be found in the appendix.

The parameter ρ determines the elasticity of substitution between goods and service workers. We calibrate ρ following Herrendorf et al. (2013), who argue that goods and services are complements. Specifically, we set $\rho = 0.5$.[16]

Green technological progress a and level of emissions ϵ_G The link between material production and emissions is governed by two parameters: ϵ_G and a. The rate of green technological progress, a, allows goods production to gradually become less polluting over time. For simplicity, we treat this rate of green technological change as exogenous. The parameter ϵ_G simply determines the units in which emissions are measured. We calibrate a and ϵ_G to align our model with observed pollution levels in 1980 and 2000.

The innovation process: R, ζ, $[\eta_s, \gamma_s]_{A,Q}$, and the initial conditions for $[A_0, Q_0]$ Finally, consider the innovation process. We set $R = 0.1$, implying that around 10% of the labor force is devoted to research activities. The innovation stepsize parameter, γ_s, is not separately identified from the efficiency of research labor, η_s (see (17)). We set the stepsize exogenously to $\gamma_A = \gamma_Q = 1.5$, meaning each successful innovation boosts productivity by 50%. This is a normalization without any loss of generality, as we subsequently estimate the research efficiency separately in each of the two sectors. In line with Akcigit et al. (2021), we assume the innovation cost function has an elasticity of 2, setting $\zeta = 0.5$.

This leaves us with two initial conditions for Q and A, along with the two R&D efficiencies η_A and η_Q. We calibrate these using four key moments. First, we target the employment share of services in both 1900 and 2000. Second, we match the average growth rate of GDP per worker, which was 1.94% between 1950 and 2000. Third, we target the IV coefficient of the service expenditure share on household spending reported in column 3 of Table II.

The intuition behind this identification strategy is as follows: by selecting the initial values of Q and A and the R&D efficiencies η_A and η_Q, our model endogenously generates the entire *path* of quality and physical productivity, $\{A_t, Q_t\}_{t=0}^{2002}$. Given the remaining parameters, this path implies a trajectory for the service employment share and income per capita. By targeting the average GDP growth rate from 1950 to 2000

[16] This calibration is preliminary. In future versions, we plan to calibrate ρ based on changes in industry-specific cost shares.

and the service employment share in both 1900 and 2002, we impose three restrictions on this path.

To fully identify our model, we need a fourth restriction. We choose the IV coefficient γ in equation (21) for this purpose. Empirically, we estimate γ to be 0.06 (see Table II, column 3). In our model, this coefficient is not a structural parameter but rather depends on consumers' real income in 2002, $\Upsilon(A_{2002}, Q_{2002})$ (see (8)). Thus, this moment places an additional restriction on the path $\{A_t, Q_t\}$, which is sufficient to fully identify our model.

As highlighted in our discussion in Section 4, our theory draws an important distinction between measured and welfare-relevant GDP. To link GDP in our model to its empirical counterpart, we must determine to what extent quality growth is captured in official GDP statistics. We assume that two-thirds of quality growth is measured and compute GDP growth in our model using a Tornqvist chained index.

We summarize all parameters and corresponding moments in Table III.

Parameter	Value	Target	Target value
ε	0.292	Engel curve slope	0.292
R	0.1	Share of researchers in population	10%
λ	0.937	Serv. share of S (IO table)	0.93
A_{2002}	0.4510	1900 US service emp. share	0.31
Q_{2002}	40230395	2002 US service emp. share	0.80
η_A	0.4451	Q-adjusted GDP p.w. growth	0.0194
η_Q	0.1351	IV coeff	0.066
a	1.0245	1980 US CO2 emissions	$4,721$
ε_G	805.02	2000 US CO2 emissions	$5,724$
ϕ	1	Normalization	-
β	0.1	Set exogenously	-
ξ	5	Set exogenously	-
ρ	0.5	Set exogenously	-
$\gamma_A = \gamma_Q$	1.5	Set exogenously	-
ζ	0.5	Set exogenously	-
δ	0.228	Set exogenously	-

Notes: The table reports all structural parameters and the corresponding moments.

Table III: Structural parameters

5.3 Model Fit

In this section we document that our model provides a good fit to the salient features of the data.

Economic Growth In Figure 3 we compare our model to the data on GDP per worker growth. In the left panel we show GDP growth in the data (black line) and the model (red line). The simulated GDP index are chained Fisher index with Tornqvist weights. Recall that we target the average growth between 1950 and 2000 assuming that 2/3 of overall quality growth is measured. To highlight the importance of the measurement of quality, we also plot the two polar opposites: actual GDP per worker growth assuming all of the quality is measured, and growth if all quality is unmeasured. It is clearly seen that in our calibration a substantive amount of GDP growth would go as unmeasured if quality growth was not taken into account.

This is also seen in the right panel, where we show the annual growth directly. To focus on the long-run frequencies, we report actual GDP growth as the decaded averages, i.e. the growth between 1950 and 1959, 1960 and 1969, etc. Again, we display our calibrated model in red and the models with full and without quality adjustment in the dotted and dashed line. While our model, by construction, matches the observed average growth rate of around 2%, this growth rate would only be 0.5% if quality growth was not taken into account. By contrast, actually GDP growth appropriately measured is in fact 0.5% higher.

Service shares and Aggregate Pollution In Figure 4 we depict our model's implication for the service share (left panel) and the total flow of emissions. Recall that our model is calibrated to match the US service share in 1900 and 2002. As far as the flow of emissions is concerned, our model is unable to capture the rather pronounced decline in emissions starting in the early 2000s. While it matches the flow of emissions between 1900 and 2000 very well (recall that the model is calibrated to match the level of emissions in 1980 and 2000), our theory predicts a further increase in emissions, albeit at a slowing rate. The likeliest reasons are either an increase in the the efficiency of abatement a or international outsourcing whereby emissions no longer are produced in the US. We will explicitly come back to this second explanations in Section 6 when we analyze an open-economy version of our theory.

Figure 3: ECONOMIC GROWTH

(a) Log of real GDP per worker (b) Growth rate of real GDP per worker

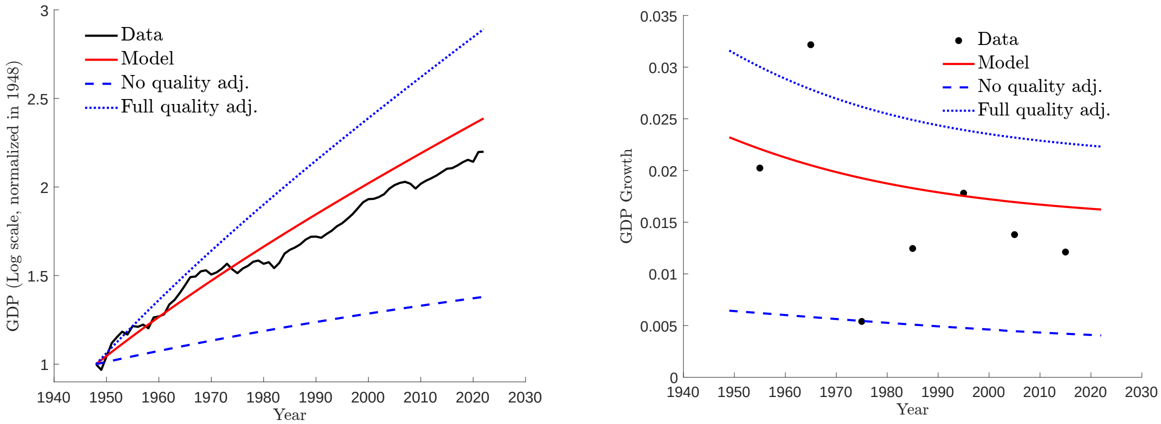

Notes: The figure shows log GDP per worker (left panel) and GDP per work growth (right panel) in the data (black line) and the calibrated model (red line). It also depicts the respective outcomes if quality was fully measured (dotted line) and if quality was not measured at all (dashed line). The simulated GDP growth rates use chain weighted price indices with Tornqvist weights.

Figure 4: SERVICE EMPLOYMENT AND EMISSIONS

(a) Service Employment (b) Emissions

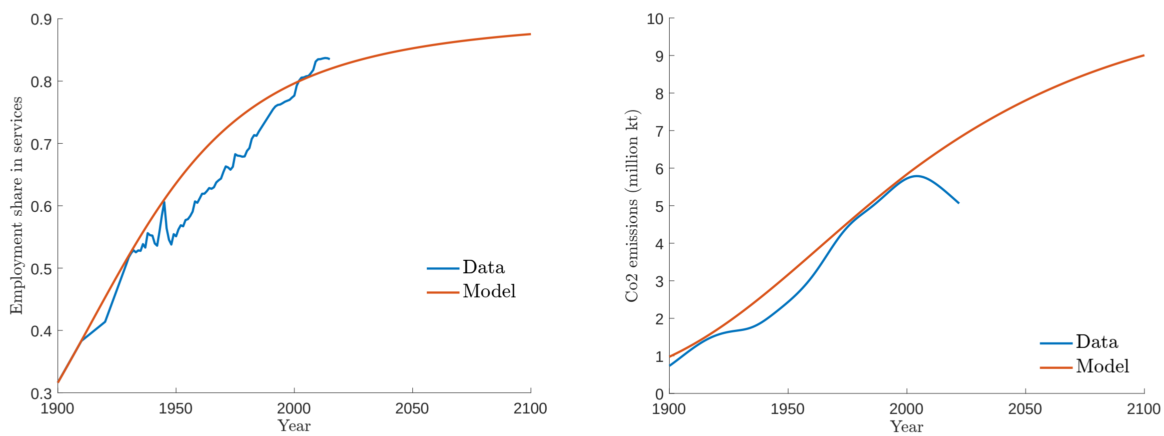

Notes: In the left panel we show the service employment share in the data (blue) and the calibrated model (red). In the right panel we show the flow of emissions in the data (blue) and the model (red).

Pollution and Service Intensity The key mechanism how our model explains a slowdown in pollution is through the "cleansing" effect of service production. In Figure 1 we showed the negative relationship between the cost share of services and the pollution intensity at the industry level. In Figure 5, we plot this relationship between product-level service cost share and pollution intensity in the model and the data. To simulate these relationship in the model, we first generate a grid of λ_j for $j \in \{1, ..., J\}$ products, with λ_j ranging approximately from 0 to 1. Next, for each j, we use the CES production structure to compute the cost share of services, setting A to its simulated value in 2002, the year in which the data were collected. Finally, we compute how many basic goods product j uses per dollar of production costs and use our emissions function to generate "pollution intensity": the quantity of pollution associated with each dollar of production of j. To make the units comparable we set the emissions of the median product to coincide with the pollution intensity of the median product in the data. In this sense, we target the overall level of pollution. The slope and shape of the cross-sectional relationship between service share and pollution intensity are entirely non-targeted. The model matches them well, supporting our specification of emissions as a linear function of basic goods production.

Figure 5: Pollution intensity and service share: model vs. data

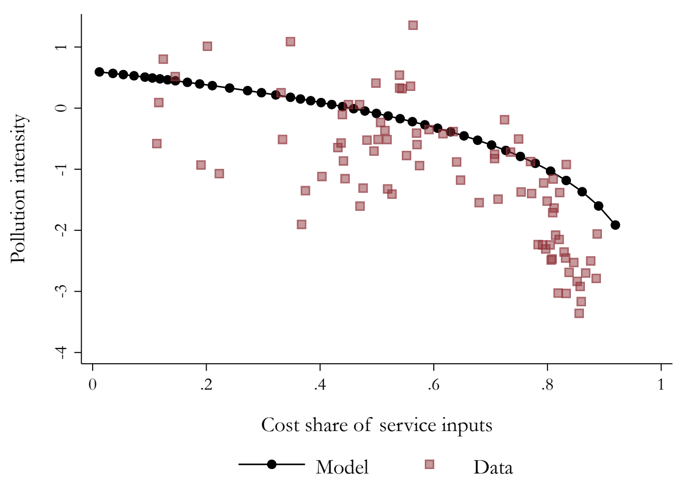

5.4 The Importance of Quality-Led Growth

The central feature in our theory is the reallocation of both expenditure and research inputs toward the purchase and the improvements of product quality. In the right panel of Figure 6, we plot the allocation of employment, both of production and research workers. In blue we depict the employment share of services, in red we show the share of research employment that targets quality improvements as opposed to further increases in the productivity of material production. Interpreted through our model, the rise of the service share of employment indicates a rise in the share of expenditure on the sophisticated good. The research sector responds by shifting innovation towards quality, which augments consumption of the sophisticated good, and away from productivity, which augments production of the basic good. In our calibrated model, the quality share of research rises from around 10% in 1900 to 60% in 2000 and 90% in 2100.

Figure 6: The Direction of Technological Change

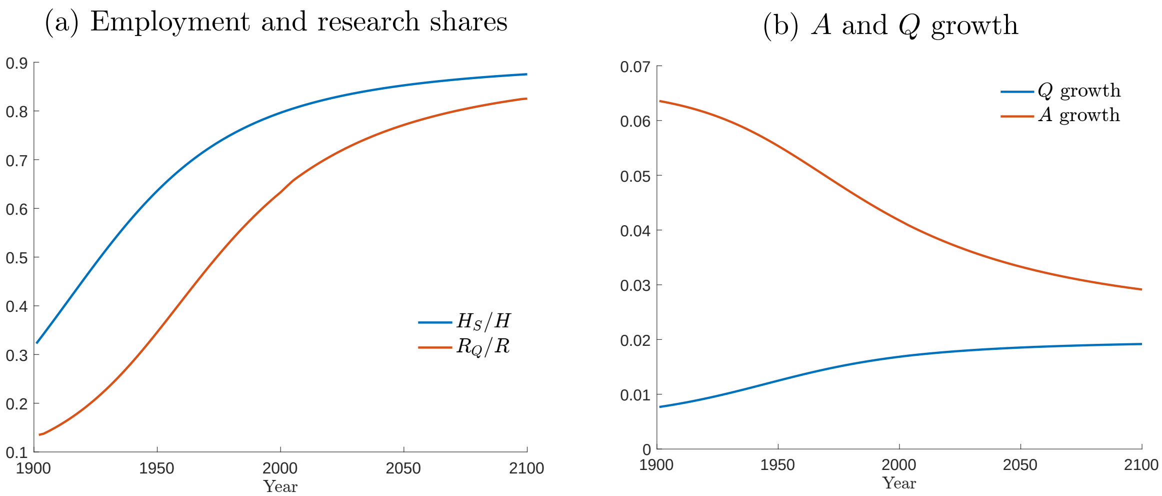

Notes: Panel (a) shows the evolution of the employment share in services (blue line) and the share of researchers working on quality innovations (red line). Panel (b) shows the growth rate of Q (blue line) and the growth rate of A (red line).

In the right panel of Figure 6 we show the implications of this realignment of the direction of research to the growth rates of quality Q and physical productivity A. The nature of directed technological change is very visible. In 1900, physical productivity grows at around 8% per year, while quality growth is is very low. Over time, the direction of research shifts toward quality improvements. This increases quality growth and reduces productivity growth. quantitatively, the growth in quality growth is lower than the decline in productivity growth owing to our estimate that $\eta_A > \eta_Q$, i.e., the efficiency of research to enhance A is higher than to increase quality Q.

Beyond its effects on sectoral quality growth, the shift from productivity-led to quality-led growth has significant implications for the environment. Production of basic goods is what causes emissions, and the emissions impact of each unit of the basic good gradually falls over time. Because quality-led growth leads to a slowdown in the quantity of basic goods, it allows for emissions to decline in the long run. We illustrate this point in Figure 7, which compares emissions in our benchmark model with directed technical change (in blue) against an alternative model (in red) where technical change is undirected. With directed technical change, the change in the composition of growth allows emissions to decline starting in the early 21st century. In the alternative model, we fix the share of research devoted to quality at its value in the initial period and do not allow it to adjust as expenditure shares evolve. Without directed technical change, sustained productivity growth leads to a steadily increasing quantity of basic goods, which in turn leads to continuing exponential growth of emissions.

Figure 7: Emissions with and without directed technical change

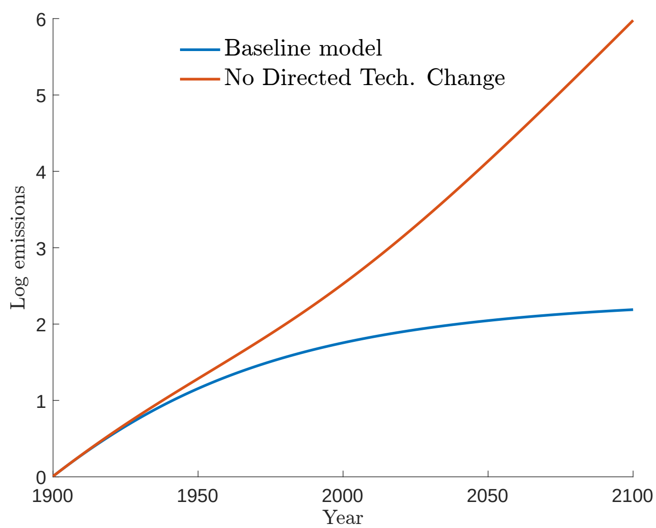

Notes: The figure shows the evolution of total emissions in the baseline economy (blue line) and in a counterfactual economy without directed technological progress (red line).

The economy's ability to grow in different ways – either productivity-led or quality-led growth – means that research subsidies can play a key role in shaping environmental outcomes. As proof of concept, we illustrate the capacity for research policy to avert an environmental disaster in Figure 8. We introduce the notion of an environmental disaster as a threshold level of pollution, to capture the idea of environmental "tipping points". In the laissez-faire equilibrium, pollution surpasses the disaster threshold at around 2030. However, a 20% subsidy to quality research (setting $\tau_Q = -.2$ in our model) allows the economy to avert this disaster. The subsidy causes research to focus more on quality, with the result that fewer physical goods are produced. Because physical goods are the source of pollution, reorienting research, and hence demand

and employment, towards quality allows the economy to continue growing without a catastrophe. While the disaster threshold in the current calibration is meant only to illustrate our point, in future drafts we plan to give the notion of disaster a more serious quantitative treatment.

Figure 8: Averting an Environmental Disaster: Subsidizing Quality R&D

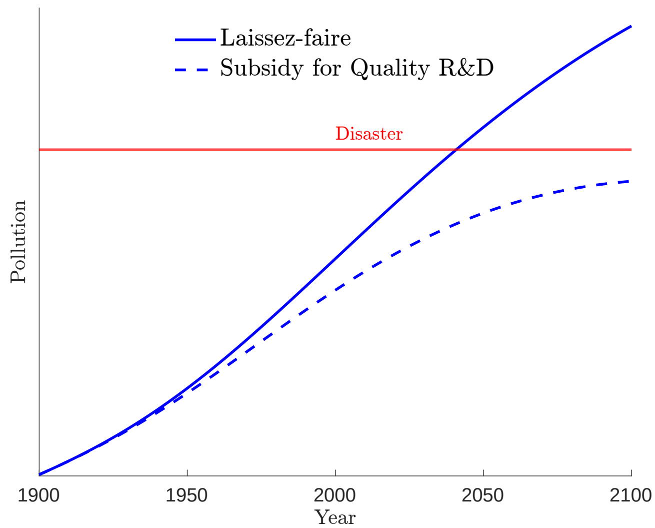

Notes: The figure shows the evolution of total emissions in the baseline economy (blue solid line) and in a counterfactual economy where research directed toward quality innovation is subsidized by $\tau_Q = -20\%$ (blue dashed line).

6 Open economy

In this section, we extend our basic framework to an open economy comprising two countries, which, for concreteness, we refer to as the US and China. The main goal is to study the effect of specialization and trade barriers on emission over the process of structural transformation.

To study this process in the setting of an open economy, we assume that the two countries produce different manufacturing goods ($Y_{G,US}$ and $Y_{G,CH}$) that enter as imperfect substitutes the production function of both basic and sophisticated good. More formally, Y_G is a CES aggregate of $Y_{G,US}$ and $Y_{G,CH}$:

$$Y_G = \left(\left(\frac{1}{2}\right)^{\frac{1}{\theta}} \times Y_{G,US}^{\frac{\theta-1}{\theta}} + \left(\frac{1}{2}\right)^{\frac{1}{\theta}} \times Y_{G,CH}^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}, \text{ with } \theta > 1.$$

By contrast, final goods, are assumed to be non-tradable. This implies that both the value added provided by service workers and the provision of quality cannot be traded internationally.

In addition, we allow for an additional homogenous tradable good, which is in fixed supply in both countries and can be exchanged against the tradable manufacturing good. We refer to this good as the *endowment*. The sole purpose of allowing for such endowment is to allow for the possibility that the US runs a trade deficit in the manufacturing good, i.e. once the economy is open to trade, the US exports the endowment and imports the basic good. In doing so, manufacturing employment in the US can decline in response to a trade liberalization. We denote by E_c the supply of the endowment in country c and we interpret the endowment in a broad sense as a stand-in for inter-temporal trade through capital flows, the export of financial services, royalty payments, or also purchases of US real estate by Chinese consumers or firms. In our empirical application we will calibrate the relative size of the endowments to match US trade deficits.

To generate a positive demand for the endowment, we assume that it directly enters consumers' preferences. In particular, we assume that preferences are given by the same PIGL indirect utility function described above augmented by the endowment good, which is traded at price p_E. Formally,

$$\mathcal{V}^{FE}\left(e_c, p_{G,c}, p_{S,c}, p_E\right) = \frac{1}{\varepsilon}\left(\frac{e_c}{p_{S,c}^{(1-\beta))(1-\varrho)}\left(p_{G,c}\right)^{\beta(1-\phi)} p_E^{\varrho}}\right)^{\varepsilon} - \phi\left(\ln p_{G,c} - \ln p_{S,c}\right) - \mathcal{P}. \tag{23}$$

Hence, the demand for the endowment is homothetic and consumer spend a fraction ϱ of their income on the endowment. If $\varrho = 0$ we are back to a standard model of trade where the endowment is absent.

Each country has a fixed supply of production labor denoted by H_{US} and H_{CH} and research labor R_{US} and R_{CH}. The goods $Y_{G,US}$ and $Y\ _{G,CH}$ are tradable subject to an iceberg cost τ, the homogenous endowment can be traded costlessly. Labor is immobile and service labor must be provided locally. Technology and quality are country specific:

$A_{US,}$ $A_{CH}, Q_{US,}$ and Q_{CH}. All other parameters are assumed to be common to the two economies.

6.1 Equilibrium prices

For $c \in \{US, CH\}$, the production prices of the manufacturing goods are given by $p_{G,c} = \frac{\xi}{\xi-1}\frac{w_c}{A_c}$. Note that here $p_{G,c}$ is defined at the factory and does not include any trade cost. The price of basic goods in the two markets are then given, respectively, by

$$p_{B,US} = \left(\frac{1}{2}p_{G,US}^{1-\theta} + \frac{1}{2}\left(\tau p_{G,CH}\right)^{1-\theta}\right)^{\frac{1}{1-\theta}} \text{ and } p_{B,CH} = \left(\frac{1}{2}(\tau p_{G,US})^{1-\theta} + \frac{1}{2}p_{G,CH}^{1-\theta}\right)^{\frac{1}{1-\theta}}.$$

These prices reflect instead the (iceberg) trade costs associated with importing the foreign manufactiruing goods.

Similarly, the price of the sophisticated good in country $c \in \{US, CH\}$ are given by

$$p_{S,c} = \frac{\xi}{\xi-1}\frac{1}{Q_c} \times c_Q\left(p_{B,c}, w_c\right),$$

where $c\left(p_{B,c},\, w_c\right) = \left((1-\lambda)\left(p_{B,c}\right)^{1-\rho} + \lambda w_c^{1-\rho}\right)^{\frac{1}{1-\rho}}$ is the unit cost of production of the sophisticated good. Note that $p_{B,c}$ (rather than the price of the manufacturing inputs) enters as an argument of the cost function. The reason is analytical convenience: the price firms pay to buy the manufacturing input in the local market depends on the production prices $p_{G,US}$ and $p_{G,CH}$ as well as on the trade cost. Since in our model this price is equivalent to $p_{B,c}$, we refrain from adding redundant notation. Substituting in the expressions for the prices given above yields

$$c\left(p_{B,c}, w_c\right) = \left((1-\lambda)\left(\frac{\xi}{\xi-1}\frac{1}{A_c}f\left(x_c\right)\right)^{1-\rho} + \lambda\right)^{\frac{1}{1-\rho}} w_c$$

where we define

$$f\left(x\right) \equiv \left(\frac{1}{2} + \frac{1}{2}\left(x\right)^{1-\theta}\right)^{\frac{1}{1-\theta}}, \tag{24}$$

and

$$x_{US} = \tau\pi, \quad x_{CH} = \frac{\tau}{\pi}, \quad \text{and} \quad \pi \equiv \frac{w_{CH}/w_{US}}{A_{CH}/A_{US}}. \tag{25}$$

In plain words, π is the relative production price of the Chinese manufacturing good relative to the US good. Under the assumption that $\theta > 1$, the function f defined in (24) is increasing in the iceberg cost in both countries.

Under the PIGL preference specification in (23), the expenditure shares of an individual with spending level e are given, respectively, by

$$\begin{aligned}
\vartheta_{Gc} &= \beta(1-\varrho) + \phi \times \left(\Upsilon\left(A_c, Q_c, x_c, , w_c, e_c\right)\right)^{-\varepsilon} \\
\vartheta_{Sc} &= \left(1-\beta\right)\left(1-\varrho\right) - \phi \times \left(\Upsilon\left(A_c, Q_c, x_c, , w_c, e_c\right)\right)^{-\varepsilon} \\
\vartheta_{Ec} &= \varrho,
\end{aligned}$$

where

$$\Upsilon\left(A, Q, x, , w_c, e_c\right) \equiv \left(\frac{A^{\beta(1-\varrho)} Q^{(1-\beta)(1-\varrho)}}{f\left(x\right)^{\beta(1-\varrho)} \left(\left(1-\lambda\right)\left(\frac{\xi}{\xi-1}\frac{1}{A} f\left(x\right)\right)^{1-\rho} + \lambda\right)^{\frac{(1-\beta)(1-\varrho)}{1-\rho}} p_E^{\varrho}} \frac{e_c}{w_c^{1-\varrho}} \right).$$

Hence, as in our baseline model, expenditure shares depend on a term capture consumers' "real income" Υ_c. In additional to total spending e_c, this term now also depends on local productivity and quality, A and Q, the terms of trade x, the local wage w_c, and the price of the endowment p_E (which is common across countries and hence we suppress it as an argument of Υ).

In addition to the prices of tradable goods, we can also solve for the price of the endowment, p_E. The introduction of the endowment affects the equilibrium allocations because it transfers resources across countries. Market clearing for the global supply of the endowment implies that

$$p_E(E_{US} + E_{CHN}) = \varrho \times (e_{US} + e_{CHN}), \tag{26}$$

where e_c is total expenditure in country c. At the same time, the returns to the endowment are part of total domestic spending, that is $e_c = \frac{\xi}{\xi-1} w_c + p_E E_c$. Substituting these expressions into (26) and letting $\varpi^c \equiv \frac{E_c}{E_{CHN}+E_{US}}$ denote the share of the global endowment owned by country c, yields

$$e_{US} = \frac{\xi}{\xi-1}\left(w_{US}\left(1+\frac{\varrho}{1-\varrho}\varpi^{US}\right)+\frac{\varrho}{1-\varrho}\varpi^{US}w_{CHN}\right) \tag{27}$$

$$e_{CHN} = \frac{\xi}{\xi-1}\left(w_{CHN}\left(1+\frac{\varrho}{1-\varrho}\varpi^{CHN}\right)+\frac{\varrho}{1-\varrho}\varpi^{CHN}w_{US}\right). \tag{28}$$

Hence, total spending in both countries is fully determined by wages and the relative supplies of the endowment. Note first that if the endowment has no value, i.e. $\varrho = 0$, expenditure in the US is only determined by US wages (and vice versa for China). Second, having an abundance in the endowment allows a country to spend more relative to its wage. For example, suppose that $\varpi^{US} = 1$, i.e. the US was the only supplier of the endowments. In that case, (27) and (28) reduce to $e_{CHN} = \frac{\xi}{\xi-1}w_{CHN}$ and $e_{US} = \frac{\xi}{\xi-1}\left(w_{US}+\frac{\varrho}{1-\varrho}(w_{US}+w_{CHN})\right)$. While consumers in China only spend their labor income, consumers in US spend their labor and capital (i.e. endowment) income.

Combining these equations with (26) allows us to solve for the equilibrium price of the endowment as

$$p_E = \frac{\xi}{\xi-1}\frac{\varrho}{1-\varrho}\frac{w_{CHN}+w_{US}}{E_{US}+E_{CHN}}. \tag{29}$$

Equations (27), (28), and (29), therefore fully determine total spending e_c and the endowment price p_E as a function of parameters and the vector of wages (w_{US}, w_{CHN}). These, in turn, are then fully determined from the trade equilibrium.

6.2 Sectoral employment shares

We set the wage in the US as the numeraire, i.e., $w_{US} = 1$. Hence, w_{CH} is the relative wage in China relative to the US wage; this will be endogenously determined and will also determine the expression of π in equilibrium.

The next proposition establishes the employment shares of goods and services in the two economies.

Proposition 3. *In equilibrium, for $c \in \{US, CH\}$,*

$$\frac{H_{S,c}}{H_c} = \vartheta_{S,c}(1-\sigma_{G,c})\left(1+\frac{\varrho}{1-\varrho}\varpi^c(w_{CHN}+w_{US}).\right) \tag{30}$$

$$\frac{H_{G,c}}{H_c} = 1-\frac{H_{S,c}}{H_c} \tag{31}$$

where $\sigma_{G,c}$ is the cost share of goods relative to services in the production of quality goods.

Proof. The cost share of goods (as opposed to services) in the production of sophisticated goods in country c is equal to

$$\sigma_{G,c} = \frac{(1-\lambda)\,p_{B,c}^{1-\rho}}{(1-\lambda)\,p_{B,c}^{1-\rho}+\lambda w_c^{1-\rho}} = \frac{(1-\lambda)\left(\frac{\xi}{\xi-1}\frac{1}{A_c}\right)^{1-\rho} f\,(x_c)^{1-\rho}}{(1-\lambda)\left(\frac{\xi}{\xi-1}\frac{1}{A_{US}}\right)^{1-\rho} f\,(x_c)^{1-\rho}+\lambda}. \tag{32}$$

where $x_{US} = \tau\pi$ and $x_{CH} = \tau/\pi$. The allocation of employment to services is determined by the following market clearing condition:

$$\frac{\xi}{1-\xi}w_c H_{S,c} = \vartheta_{S,c}\,(1-\sigma_{G,c})\,e_c H_c, \tag{33}$$

where the left hand-side yields the value of the factor payments (including both wages and profits), while the right hand side yields the service value added share in the S sector. Substituing the expressions of ϑ_S and $s_{QG,c}$ into (33) yields the result of the proposition. **QED**

6.3 Trade equilibrium

The equilibrium allocation of labor in Proposition 3 is conditional on endogenous wages w_{US} and w_{CHN}. As we mentioned, the wage in the US is set to unity from the choice of the numeraire. However, the wage of China remains to be determined.

To this aim, we use market clearing for tradable goods, i.e. both the manufacturing good and the endowment. To express these quantities, let $\chi_{G,c}$ denotes the expenditure share on *domestic* manufacturing goods. Given the CES structure, $\chi_{G,c}$ is given by

$$\chi_{G,c} = \frac{p_{G,c}^{1-\theta}}{p_{G,c}^{1-\theta} + (\tau p_{G,\tilde{c}})^{1-\theta}} = \frac{1}{2} \left(f\left(x_c\right)\right)^{(\theta-1)}. \tag{34}$$

Total US manufacturing exports are thus given by

$$EX_{G,US} = e_{CHN}(1 - \chi_{G,CH}) \left(\vartheta_{G,CH} + \vartheta_{S,CH} s_{QG,CH}\right) \tag{35}$$

Similarly, total US manufacturing imports are given by

$$IM_{G,US} = e_{US}(1 - \chi_{G,US}) \left(\vartheta_{G,US} + \vartheta_{S,US} s_{QG,US}\right) \tag{36}$$

Market clearing then implies that any trade deficit is paid for by exports of the endowment. Formally,

$$\text{Deficit}_{US} \equiv IM_{G,US} - EX_{G,US} = p_E E_{US} - \varrho e_{US}. \tag{37}$$

Equation (37) is a single equation in a single unknown, w_{CHN} (given the state variables Q and A). To see why, note first that p_E is fully determined from (29) given w_{CHN}. Similarly, $\sigma_{G,c}$ and $\chi_{G,c}$ can be directly computed - see (32) and (34). Finally, noting that $e_c = \frac{\xi}{\xi-1} w_c + p_E E_c$ and we already computed p_E, e_c can also be computed, given w_{CHN}. This then allows us to compute Υ_c and hence $\vartheta_{G,c}$ and $\vartheta_{S,c}$. Equation (37) therefore pins down w_{CH} and concludes the characterization of the equilibrium.

6.4 Endogenous technologies

We now endogeneize the A_c's and the Q_c's by making them result from innovations. The process of innovation is exactly the same as in our baseline model analyzed above. In particular, there is a mass R_c of researchers in country $c \in \{US, CH\}$, that can direct their research to either increase the productivity of good production (A_c) or the quality of the quality good (Q_c).

As we show in the Appendix, the research allocation takes exactly the same form as in the closed economy (see (13)):

$$\frac{R_{Qc}}{R_{Ac}} = \left(\frac{\eta_{Qc}}{\eta_{Ac}}\left(\frac{\gamma_{Qc}}{\gamma_{Ac}}\right)^{\xi-1}\right)^{1/\zeta}\left(\frac{(1-\lambda)\left(\frac{\xi}{\xi-1}\frac{1}{A_c}f\left(x_c\right)\right)^{1-\rho}+\lambda}{\lambda}\right)^{\frac{1}{\zeta}}\left(\frac{H_{Sc}}{H_{Gc}}\right)^{1/\zeta},$$

where $x_{US} = \tau\pi$ and $x_{CH} = \tau/\pi$ – see (25).

In particular, a higher trade cost will shift $\frac{R_{Q,CH}}{R_{A,CH}}$ downward, i.e. will redirect Chinese innovation away from quality into productivity. To see this, note that:

$$\left((1-\lambda)\left(\frac{\xi}{\xi-1}\frac{1}{A_{CH}}f\left(\tau/\pi\right)\right)^{1-\rho}+\lambda\right)\frac{H_{S,CH}}{1-H_{S,CH}}$$

decreases when τ increases.[17]

6.5 Open economy

We now consider the open-economy extension of our model. As suggested by our theory, we consider a two-country model, with China being the other country.

Our strategy is as follows. We start from the closed-economy calibration between 1900 and 2000. In the year 2000, the US economy then opens up to China, i.e. trade costs fall from ∞ to τ. We then keep trade costs constant and trace out the transitional dynamics in both the US and China.

To implement this exercise, we require five additional parameters: (i) the initial level of quality and productivity in China at the time of the trade opening ($A_{2000,CHN}$ and $Q_{2000,CHN}$), (ii) the relative size of the endowment of the US E_{US},[18] (iii) consumers' expenditure share on the endowment ϱ, and (iv) the elasticity of substitution of trade goods θ.

[17] This in turn follows from the fact that:

$$\frac{\left(1-\beta-\phi\left(\frac{A_c^{\beta}Q_c^{1-\beta}}{\left((f(\tau/\pi))^{\beta}\times\left(\left(\lambda_G\left(\frac{\xi}{\xi-1}\frac{1}{A_{US}}f(\tau/\pi)\right)^{1-\rho}+\lambda_S\right)^{\frac{1}{1-\rho}}\right)^{1-\beta}w_c\right)}\right)^{-\varepsilon}\right)}{1-H_{S,CH}}$$

where $f\left(\tau/\pi\right)$ is increasing in τ and $H_{S,CH}$ is decreasing in τ.

[18] Recall that all allocation only depend on the relative endowment E_{US}/E_{CHN}. We therefore normalize $E_{CHN}=1$ without loss of generality.

We calibrate these parameters to the following moments. We pick the level of productivity and quality in China in the year 2000, $A_{2000,CHN}$ and $Q_{2000,CHN}$, to match GDP pc in China relative to the US in 2000 (10.3%) and the employment share of services in China in 2000 (27%).[19] We pick the size of the endowment in the US, E_{US} to match the size of the US trade deficit relative to US GDP (1.6%). Finally, we set the expenditure share on the endowment, ϱ, to 0.3, and the elasticity of substitution between traded goods, θ, to 5. The new parameters and moments are contained in Table IV. We keep all other parameters the same; see Table III.

Parameter	Value	Target	Target value
$A_{2000,CH}$	1.11	CHN/US GDP p.c. (2000)	0.103
$Q_{2000,CH}$	1.57	CHN Service emp. share (2000)	0.27
E_{US}	21.2	US Trade deficit rel. to GDP (2019)	0.016

Notes: The table reports the additional structural parameters and corresponding moments for the open-economy calibration.

Table IV: Additional parameters for open-economy simulation

6.5.1 Results

To illustrate the effects of free trade – and conversely, of trade barriers – we plot the path of the US and Chinese economies under the benchmark free trade case of $\tau = 1$ and a counterfactual where we set $\tau = 100$, a prohibitively high iceberg trade cost that results in autarky.

In Figure 9 we compare the free trade and autarky scenarios for the years 2000-2040. Free trade outcomes are represented by solid lines, while autarky corresponds to dashed lines. Panel 9a shows that free trade has opposite effects on the sectoral composition of the US and Chinese economies. In the US, service employment is higher under free trade for two reasons: free trade makes the US richer (the income effect) and it allows the US to shift out of basic goods production, instead relying on Chinese imports (the specialization effect). In China, these two effects push in opposite directions. Free trade improves Chinese real incomes (see Panel 9b), pushing consumers towards consuming a greater share of sophisticated goods. However, free trade also pushes

[19] Since we begin our open-economy simulation in 2000, we take our values for $A_{2000,US}$ and $Q_{2000,US}$ directly from the closed-economy simulation.

China to specialize in basic goods production, so that it can import the endowment good from the US. In our simulations, the specialization effect dominates in China and free trade causes a shift towards goods employment. It is worth stressing that this is not a general prediction of our model; under alternative parameter values, it is possible for Chinese service employment to increase in response to free trade.

Because free trade can increase basic goods production in China, its overall effects on the environment are ambiguous: US emissions decline, but Chinese emissions increase. On balance, free trade reduces total world emissions in our model, as displayed in Panel 9c. We view this result as a word of caution against attempts to "reshore" manufacturing in advanced economies. Shifting goods production back to the most developed countries may lead to real income declines and greater overall pollution.

7 Conclusion

In this paper, we developed and quantified a growth model where: (ii) consumers have non-homothetic preferences between a more basic good—for which quality matters less and which are less service-intensive—and a more luxury good – a "quality" good - for which quality matters more and the production of which is more service intensive; (ii) the direction of technological progress—toward increasing the productivity of material production versus improving quality of the luxury good—is endogenous; (iii) the production of the luxury good, which is intensive in service labor, has a lower environmental footprint than that of the basic good.

The model delivered some key insights. First, over time as the economy develops and they become richer, consumers increasingly shift their demand towards the quality good, which in turn tilts the direction of innovation away from increasing material productivity towards increasing quality. Second, the transition to quality-driven growth may translate into a stall or even a decline in measured GDP growth, even though quality-adjusted GDP continues to grow. Third, trade barriers may have a negative effect on global environmental sustainability: by slowing down growth in emerging economies, trade barriers will delay or even reverse local consumer demand's shift toward the higher quality, more service-intensive, good; this, in turn, will result in significantly higher levels of global emissions, as endogenous directed innovation will also move away from quality back to increasing the productivity of material goods.

An important implication of our approach, is that environmental quality and mea-

Figure 9: Open economy simulations: US and China 2000-2040

(a) Service share of employment

(b) Log real GDP per capita

(c) Emissions

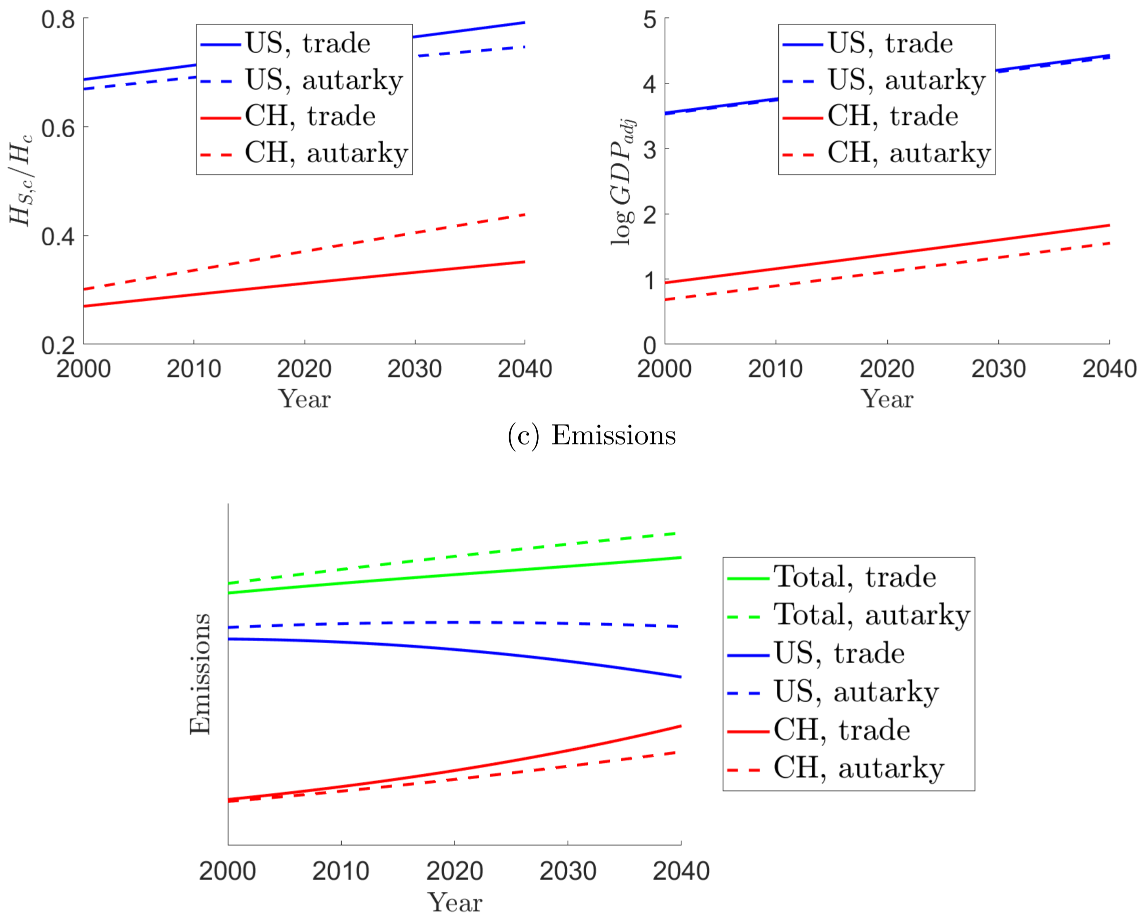

Notes: The figure shows the evolution of service shares (panel a), log GDP per capita (panel b), and log emissions (panel c). We always depict the outcomes for the US (CHN) in blue (red). The baseline model with international trade is shown with solid lines, the closed-economy model is shown with dashed lines. In panel c we also depict total global emissions with green lines.

sured degrowth, are both consistent with the quest for sustained innovation-led growth, insofar as innovation becomes increasingly geared towards quality. And innovation indeed tilts increasingly towards quality as the economy becomes more developed, as it caters to a consumer demand which itself becomes increasingly more quality-oriented as consumers become richer.

Our analysis in this paper can be developed and extended in several interesting directions. First, it would be useful to directly measure the environmental footprint of goods of different quality. In the paper, we linked it to service labor share (which is observable) but this probably does not capture all the variation in environmental impacts of production across goods. A second extension would be to further explore the discrepancy between quality growth and measured GDP growth. For example, focusing on meals and restaurants, one could explore information from gastronomic guides' gradings to capture quality-adjusted GDP as opposed to measured sales. Another extension would be to look at the extent to which the downward sloping cross-country relationship between measured per capita-GDP growth and measured per-capita GDP level, is due to growth becoming increasingly quality-driven as countries become more developed. Similarly, one could try and quantify the extent to which the recent slowdown in measured TFP growth is, at least partly, due to an accelerated shift towards quality-based growth. These and other extensions are left for future research.

References

Acemoglu, D.: 2009, *Introduction to modern economic growth*. Princeton University Press.

Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous: 2012a, 'The environment and directed technical change'. *American economic review* **102**(1), 131–166.

Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr: 2016, 'Transition to clean technology'. *Journal of political economy* **124**(1), 52–104.

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi: 2012b, 'The Network Origins of Aggregate Fluctuations'. *Econometrica* **80**(5), 1977–2016.

Aghion, P., R. Bénabou, R. Martin, and A. Roulet: 2023, 'Environmental preferences and technological choices: Is market competition clean or dirty?'. *American Economic Review: Insights* **5**(1), 1–19.

Aghion, P., A. Bergeaud, M. De Ridder, and J. Van Reenen: 2024, 'Lost in transition: Financial barriers to green growth'.

Akcigit, U., H. Alp, and M. Peters: 2021, 'Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries'. *American Economic Review* **111**(1), 231–275.

Bils, M. and P. J. Klenow: 2001, 'Quantifying quality growth'. *American Economic Review* **91**(4), 1006–1030.

Boppart, T.: 2014, 'Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences'. *Econometrica* **82**(6), 2167–2196.

Chen, X., G. Pei, Z. Song, and F. Zilibotti: 2023, 'Tertiarization Like China'. *Annual Review of Economics* **15**, 485–512.

Dedoussi, I. C., S. D. Eastham, E. Monier, and S. R. Barrett: 2020, 'Premature Mortality related to United States cross-state air pollution'. *Nature* **578**(7794), 261–265.

Easterlin, R. A.: 1974, 'Does economic growth improve the human lot? Some empirical evidence'. In: *Nations and households in economic growth*. Elsevier, pp. 89–125.

Fan, T., M. Peters, and F. Zilibotti: 2023, 'Growing Like India—the Unequal Effects of Service-Led Growth'. *Econometrica* **91**(4), 1457–1494.

Georgescu-Roegen, N.: 1971, *The entropy law and the economic process*. Harvard University Press.

Georgescu-Roegen, N.: 1974, 'Dynamic models and economic growth'. In: *Equilibrium and Disequilibrium in Economic Theory: Proceedings of a Conference Organized by the Institute for Advanced Studies, Vienna, Austria July 3–5, 1974*. pp. 413–449.

Georgescu-Roegen, N.: 1979, 'Demain la décroissance. Entropie–Écologie-Économie'.

Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski: 2014, 'Optimal taxes on fossil fuel in general equilibrium'. *Econometrica* **82**(1), 41–88.

Hassler, J., P. Krusell, and A. A. Smith Jr: 2016, 'Environmental macroeconomics'. In: *Handbook of macroeconomics*, Vol. 2. Elsevier, pp. 1893–2008.

Hémous, D.: 2016, 'The dynamic impact of unilateral environmental policies'. *Journal of International Economics* **103**, 80–95.

Hémous, D. and M. Olsen: 2021, 'Directed technical change in labor and environmental economics'. *Annual Review of Economics* **13**, 571–597.

Herrendorf, B., R. Rogerson, and A. Valentinyi: 2013, 'Two Perspectives on Preferences and Structural Transformation'. *American Economic Review* **103**(7), 2752–89.

Hickel, J., T. Kallis, Giorgosand Jackson, D. W. O'Neill, J. B. Schor, P. A. Steinberger, Julia K.and Victor, and D. Ürge-Vorsatz: 2022, 'Degrowth can work — here's how science can help'. *Nature* **612**, 400–403.

Levinson, A.: 2009, 'Technology, international trade, and pollution from US manufacturing'. *American economic review* **99**(5), 2177–2192.

Levinson, A. and J. O'Brien: 2019, 'Environmental Engel curves: Indirect emissions of common air pollutants'. *Review of Economics and Statistics* **101**(1), 121–133.

Medeiros, G. and T. Howels III: 2017, 'Introducing domestic requirements tables for 1997–2015'. *Bureau of Economic Analysis*.

Nordhaus, W. D.: 1991, 'To Slow or Not to Slow: The Economics of The Greenhouse Effect'. *The Economic Journal* **101**(407), 920–937.

Nordhaus, W. D.: 1994, 'Expert opinion on climatic change'. *American scientist* **82**(1), 45–51.

Nordhaus, W. D.: 1997, 'Discounting in economics and climate change; an editorial comment'. *Climatic Change* **37**(2), 315–328.

Parshall, L., K. Gurney, S. A. Hammer, D. Mendoza, Y. Zhou, and S. Geethakumar: 2010, 'Modeling energy consumption and CO2 emissions at the urban scale: Methodological challenges and insights from the United States'. *Energy policy* **38**(9), 4765–4782.

Quah, D.: 1999, 'The Weightless Economy in Economic Development'.

Reff, A., P. V. Bhave, H. Simon, T. G. Pace, G. A. Pouliot, J. D. Mobley, and M. Houyoux: 2009, 'Emissions inventory of PM2. 5 trace elements across the United States'. *Environmental science & technology* **43**(15), 5790–5796.

Schumacher, E. F.: 1973, *Small is beautiful: economics as if people mattered*. London.

Simon, H., A. Reff, B. Wells, J. Xing, and N. Frank: 2015, 'Ozone trends across the United States over a period of decreasing NOx and VOC emissions'. *Environmental science & technology* **49**(1), 186–195.

Stiglitz, J., A. Sen, and J.-P. Fitoussi: 2009, *Richesse des nations et bien-être des individus.: Performances économiques et progrès social*. Odile Jacob.

APPENDIX A: EMPIRICAL RESULTS

In this section, we discuss our empirical analysis and the construction of the data in more detail.

A-1 Appendix: Empirical Analysis

A-1.1 Data

In this section, we describe the different data sources.

Consumer Expenditure Survey (CEX) The CEX is a nationwide household survey conducted by the U.S. Bureau of Labor Statistics. Its primary aim is to delve into the spending habits of U.S. consumers. The survey comprises two distinct components: the Interview Survey, which captures data on major and/or recurring expenditures, and the Diary Survey, which focuses on more minor or frequently purchased items.

In our empirical analysis, we concentrate primarily on the Interview Survey, as it encompasses approximately 80% to 95% of total household expenditures. To exclude students and retirees, we narrow our sample by restricting the age range of the household head to between 25 and 64 years, excluding those serving in the military. To ensure consistency and relevance, we use all quarter data in the current calendar year's release. This yields a dataset consisting of around 12,000 households for the year 2002.

Consumption and income data in the CEX are organized according to the Universal Classification Codes (UCC) system. To examine expenditure patterns, we exclude all UCC related to assets and gifts. Additionally, individuals may receive reimbursements from government programs, resulting in negative expenditures for certain items. These negative expenditures are also excluded from our analysis.

Input Output Tables The Input-Output Table, a quintennial report generated by the Bureau of Economic Analysis (BEA), offers a comprehensive overview of the U.S. economy. We mainly focus on the detailed Use Table, which includes around 400 industries in 2002. It illuminates the breakdown of value-added components and total intermediate inputs utilized by each industry in their production processes.

The reason for our emphasis on the 2002 dataset is the concordance between the CEX and I-O table provided by Levinson and O'Brien (2019). To ensure precision, we omit scrap and non-comparable imports from the Use Table, given their ambiguous classification as either goods or services. Employing an initial grouping strategy based on the first number of I-O codes, we categorize codes 1 to 3 as goods and 4 to 8, along with government spending, as services. Subsequently, this allows us to ascertain the proportion of services utilized by each industry in their production processes.

Table A-I: Aggregate Industry Code for Input Output Table

Code	Industry
1	Agriculture, Forestry, Fishing and Hunting
2	Mining, Utilities, Construction
3	Manufacturing
4	Wholesale/Retail Trade, Transportation and Warehousing
5	Information, Finance, Real Estate, and Professional Services
6	Educational Services, Health Care and Social Assistance
7	Arts, Recreation, Accommodation and Food Services
8	Other Services except Public Administration
9	Government Industries

Environmental Accounts The National Emissions Inventory (NEI) is a comprehensive air emissions data source, compiled and released every three years by the U.S. Environmental Protection Agency. It has been widely used in environmental science (Dedoussi et al., 2020, Parshall et al., 2010, Reff et al., 2009, Simon et al., 2015). In the economic field, Levinson (2009) demonstrated that while imports in the U.S. trend towards cleaner goods, the lion's share of air pollution reduction stems from technological advancements. In this paper, our emission data primarily relies on the total emission coefficients for each industry, as calculated by Levinson and O'Brien (2019).

We focus on five major air pollutants: particulates smaller than 10 microns (PM10), volatile organic compounds (VOCs), nitrogen oxides (NOx), sulfur dioxide (SO2), and carbon monoxide (CO). Given their varying measurement units, we employ pollutant fixed effects when aggregating them. The total emission coefficient for each pollutant within each industry represents the amount of pollution emitted per dollar of the final product and all associated inputs. Combining with CEX, we can get how much each household emits for their expenditures.

Table A-II: Top 5 Cleanest and Dirtiest Industies

Top 5 Cleanest Industries	
1	Rental of video software/video tapes
2	Contributions to church/religious organization
3	Education tuitions
4	Domestic services
5	Bank service/financial charges
Top 5 Dirtiest Industries	
1	Wood and other fuels, electricity
2	Water and sewerage maintenance
3	Tires - purchased, replaced, installed
4	Materials for patio, walk, fence, etc
5	Gasoline, diesel

A-1.2 Nonhomothetic Service Demand

In the left panel of Figure A-1 we depict the cross-sectional distribution of $\vartheta^i_{\mathcal{S}}$. In the right panel, we show that this heterogeneity is strongly related to household income. The CEX data directly reports each household's income rank. Richer households spend, on average, substantially more on service-intensive goods.

Figure A-1: Expenditure shares on service-intensive goods

(a) Heterogeneity of service shares

(b) Service-intensive goods are luxuries

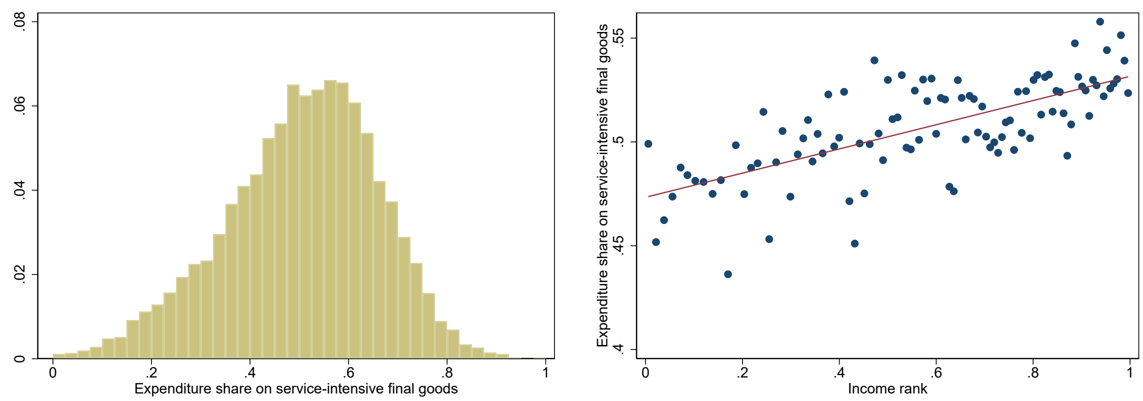

Notes: In the left panel, we display the cross-sectional distribution of households' expenditure share on service-intensive goods ($\vartheta^i_{\mathcal{S}}$). In the right panel, we display a binscatter plot between the expenditure share on service-intensive goods and the income rank of the household.

A-1.3 Regression of $\log(emission/GDP)$ on $\log(GDP\ per\ capita)$

$$\log(emission_{i,t}/GDP_{i,t}) = \alpha \log(GDP\ per\ capita_{i.t}) + \beta services\ employment\ share_{i,t} + X_{i,t}$$

where $X_{i,t}$ include characteristics of each country, e.g. agriculture employment shares, total population and land.

	2015		1991-2016			
	(1)	(2)	(3)	(4)	(5)	(6)
log GDP per capita	0.156***	-0.029	0.224***	-0.090	0.233***	-0.085
	(0.042)	(0.072)	(0.038)	(0.072)	(0.039)	(0.072)
Log of total population	-0.094**	-0.069**	-0.050	-0.056	-0.047	-0.051
	(0.044)	(0.032)	(0.042)	(0.035)	(0.043)	(0.036)
Share of Emp. in Serv.		-3.981***		-4.146***		-3.949***
		(0.614)		(0.619)		(0.658)
Year FE					Yes	Yes
Agriculture Emp. Share		Yes		Yes		Yes
log Total Land	Yes	Yes	Yes	Yes	Yes	Yes
N	188	171	4657	4308	4657	4308
R^2	.102	.377	.135	.38	.148	.387

Notes: The table reports the relationship between emission per GDP and GDP per capita and services employment share. For the year 2015, robust standard errors are in parentheses. For the year 1991-2016, standard errors are clustered at the country level.

Table A-III: Regression of $\log(emission/GDP)$ on $\log(GDP\ per\ capita)$

A-1.4 Service share construction

To construct a measure of the service share for each industry in the Input-Output tables, we use the Total Requirements Tables. These tables provide the value of intermediate inputs used along each industry's supply chain, allowing us to account for the value share of each input allocated to the production of goods and services (Medeiros and Howels III, 2017).

However, the Total Requirements Tables are based on the gross output of each industry. Thus, the value of an intermediate input used by one industry may already include the value of inputs used in its own production process, which can lead to double-counting the value of intermediate inputs. Additionally, the Total Requirements Tables only account for the inputs used in the production process, without considering the wholesale, retail, and transportation costs that an industry incurs to reach the final consumer. To correct for double accountability and consider the costs an industry faces in meeting the consumer, we follow a two-step approach.

First, let's define our setup based on Levinson and O'Brien (2019). A simple linear production function implies that we can write

$$X = CX + Y,$$

with $Y_{1\times n}$ the vector of aggregate household consumption, $X_{1\times n}$ the vector of total output, and with $C_{n\times n}$ corresponding to the Direct Requirements Table, where an entry c_{ij} is the dollar amount of inputs from industry i that is used to produce one dollar from industry j. Thus, the first term on the right-hand side represents the production used as input, and the second term is the production consumed. Now, we can express the equation as

$$X = [I - C]^{-1}Y.$$

Defining $T_{n\times n} := [I - C]^{-1}$, we have $X = TY$, where T corresponds to the Total Requirements Table. Each column in T represents the total value of production required for domestic industries to supply one dollar of output. Nevertheless, as we explained previously, T can be subject to double accountability across the supply chain of each industry. To correct this problem, we take $T^{-1} = I - C$ and add the columns of T^{-1} to obtain an approximation of the value added per dollar of output in each industry. Then, we construct a diagonal matrix $V_{n\times n}$, where the entries in the diagonal are the value added per dollar of output in each industry (the sum of the columns of T^{-1}). This way, we compute $T_{adj} := VT$, where each entry represents how many units of gross output are embodied in each industry (t_{ij}) times how many units of value-added correspond to each unit of gross output of i (v_i). Notice that once we adjust for the value-added across industries, T_{adj} corrects for the double accountability issue.

Second, we consider the costs that each industry incurs to reach the final consumer. To do this, we incorporate the Bridge PCE Tables into our procedure. These tables identify the value of transactions for each industry at producers' and purchasers' prices, as well as the associated transportation costs and trade margins. However, the Bridge Tables does not make a specific division of how the transportation costs and wholesale and retail margins are divided within its industries (i.e. how transportation costs are divided into air transportation, truck transportation, water transportation, etc.). Therefore, we divide such costs among the different industries, proportional to the size of the total input-output use of that industry relative to the total of the sector reported in the Total Requirements Tables. Then, we compute this information into a matrix $B_{n\times n}$, with the same structure as T. Each column of B corresponds to an industry, and its entries represent the share of the purchase value allocated to transportation costs, retail, and wholesale. Additionally, the diagonal entries represent the share of the purchase value that is allocated to the production of the industry, while the remaining entries are 0. Then, by computing $T_{Badj} := T_{adj}B$, we account for the costs that each industry has to face to reach its final consumer.

Finally, we divide each column of T_{Badj} by the total input-output use of each industry (the sum of each column of T_{Badj}) to obtain $\hat{T}_{Badj}$. Also, we define $S_{n\times 1}$ by assigning a value of 1 if the first digit of the industry's NAICS code is greater than 3, and 0 if it is smaller. We then obtain the service share of each industry ($\hat{S}_{1\times n}$) by computing $\hat{S} := S'\hat{T}_{Badj}$.

APPENDIX B: THEORY

In this section, we discuss the technical material referred to in the text.

B-1 Derivations of theoretical results

This section contains detailed derivations for out theoretical results.

B-1.1 Prices, Profits, Labor Shares and Marginal Costs of Final Good Producers

Consider firm i producing final good j. Let p_G denote the price of physical goods and w the wage rate. Given the production function

$$y_{ij} = \left((1-\lambda_j)^{\frac{1}{\rho}} Y_{ijG}^{\frac{\rho-1}{\rho}} + \lambda_j^{\frac{1}{\rho}} h_{ijS}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}},$$

the marginal cost of production is given by

$$c_j(p_G, w) = \left((1-\lambda_j) p_G^{1-\rho} + \lambda_j w^{1-\rho} \right)^{\frac{1}{1-\rho}}.$$

Furthermore, the relative spending on services (relative to goods) is given by

$$\frac{w h_{ijS}}{p_G Y_{ijG}} = \frac{\lambda_j}{1-\lambda_j} \left(\frac{w}{p_G} \right)^{1-\rho}.$$

The cost share on service inputs is therefore

$$\sigma_j^S \equiv \frac{w h_{ijS}}{w h_{ijS} + p_G Y_{ijG}} = \frac{\lambda_j w^{1-\rho}}{\lambda_j w^{1-\rho} + (1-\lambda_j) p_G^{1-\rho}}. \tag{B-1}$$

Monopolistic competition implies that variety i for product j has a market price of

$$\tilde{p}_{ji} = \frac{\xi}{\xi - 1} c_j(p_G, w) = \tilde{p}_j.$$

The overall quality-adjusted price index for good j is therefore given by

$$\begin{aligned} p_j &= \left(\int_0^1 \left(\frac{\tilde{p}_{ji}}{Q_{ij}^{\alpha_j}} \right)^{1-\xi} di \right)^{\frac{1}{1-\xi}} = \frac{\xi}{\xi-1} c_j(p_G, w) \left(\int_0^1 Q_{ij}^{\alpha_j(\xi-1)} di \right)^{\frac{1}{1-\xi}} \\ &\equiv \frac{1}{Q_j^{\alpha_j}} \frac{\xi}{\xi-1} c_j(p_G, w), \end{aligned} \tag{B-2}$$

where

$$Q_j = \left(\int_0^1 Q_{ij}^{\xi-1} di\right)^{\frac{1}{\xi-1}}.$$

Profits of producer i in sector j are given by

$$\begin{aligned}
\pi_{ij} &= p_j^{\xi} C_j Q_{ij}^{\alpha_j(\xi-1)} c_j \left(p_G, w\right)^{1-\xi} \frac{(\xi-1)^{\xi-1}}{\xi^{\xi}} \\
&= p_j C_j p_j^{\xi-1} Q_{ij}^{\alpha_j(\xi-1)} c_j \left(p_G, w\right)^{1-\xi} \frac{(\xi-1)^{\xi-1}}{\xi^{\xi}} \\
&= \frac{1}{\xi} p_j C_j \left(\frac{Q_{ij}^{\alpha_j(\xi-1)}}{\int_0^1 Q_{\iota j}^{\alpha_j(\xi-1)} d\iota}\right).
\end{aligned} \tag{B-3}$$

B-1.2 Prices and Profits of Manufacturing Firms

Now consider firm i in the manufacturing sector. The optimal price is given by

$$p_{iA} = \frac{\xi}{\xi-1} \frac{w}{A_i}.$$

Letting $\mathcal{D}_G$ denote total spending on manufacturing goods, profits of manufacturing firm i are given by

$$\pi_{iG} = \frac{1}{\xi} \mathcal{D}_G \frac{A_i^{\xi-1}}{\int A_i^{\xi-1} di} = \frac{1}{\xi} \mathcal{D}_G \left(\frac{A_i}{A}\right)^{\xi-1}. \tag{B-4}$$

B-1.3 Consumer Preferences and Expenditure Share

Consider the indirect utility function

$$\mathcal{V}^{FE}\left(e, [p_j]_{j=1}^J\right) = \frac{1}{\varepsilon}\left(\prod_{j=1}^J \frac{e}{p_j^{\beta_j}}\right)^{\varepsilon} - \sum_{j=1}^J \phi_j \ln p_j - \mathcal{P}.$$

The expenditure share for good k is given by

$$\begin{aligned}\vartheta_k\left(e,[p_j]_{j=1}^J\right) &= -\frac{\partial\mathcal{V}^{FE}/\partial p_k}{\partial\mathcal{V}^{FE}/\partial e}\frac{p_k}{e} \\ &= -\frac{\beta_k p_k^{-1}\left(\prod_{j=1}^J \frac{e}{p_j^{\beta_j}}\right)^{\varepsilon}-\phi_k p_k^{-1}}{e^{-1}\left(\prod_{j=1}^J \frac{e}{p_j^{\beta_j}}\right)^{\varepsilon}}\frac{p_k}{e} \\ &= \beta_k+\phi_k\left(\frac{e}{\prod_{j=1}^J p_j^{\beta_j}}\right)^{-\varepsilon}. \\ &= \beta_k+\phi_k\left(\left(\prod_{j=1}^J Q_j^{\beta_j}\right)\frac{\xi-1}{\xi}\frac{e}{\prod_{j=1}^J c_j\left(p_G,w\right)^{\beta_j}}\right)^{-\varepsilon}, \qquad \text{(B-5)}\end{aligned}$$

where the last line uses the expression for sectoral prices p_j (see (B-2)).

B-1.4 The Optimal Allocation of Research

There is a fixed amount of researchers, R, that can direct their research effort towards $J+1$ activities: improving the productivity to produce goods (A_i) or produce the quality of the provision of each of the J final goods (Q_{ij}). Hence, market clearing requires that

$$R = R_A + \sum_j R_{Qj}. \qquad \text{(B-6)}$$

The value of directing research towards improving quality in sector j is given by

$$V_{Qj} = (1-\tau_Q)\left(\eta_Q R_{Qj}^{-\zeta}\right)\times\int \pi_{ij}\left(\gamma_Q Q_{ij}\right)di,$$

where $\pi_{ij}(\gamma_Q Q_{ij})$ denotes to profits of providing variety i in sector j at quality $\gamma_Q Q_{ij}$. Using the expression for equilibrium profits in (B-3), we get that

$$\begin{aligned}\int \pi_{ij}\left(\gamma_Q Q_{ij}\right)di &= \int \frac{1}{\xi}p_j Y_j\left(\frac{\gamma_Q^{\alpha_j(\xi-1)}Q_{ij}^{\alpha_j(\xi-1)}}{\int_0^1 Q_{ij}^{\alpha_j(\xi-1)}di}\right)di \\ &= \frac{1}{\xi}p_j Y_j \gamma_Q^{\alpha_j(\xi-1)}.\end{aligned}$$

Hence,

$$V_{Qj} = (1-\tau_Q)\left(\eta_Q R_{Qj}^{-\zeta}\right)\frac{1}{\xi}p_j Y_j \gamma_Q^{\alpha_j(\xi-1)}. \qquad \text{(B-7)}$$

Similarly, we can solve for the value of directing research towards improving the productivity of manufacturing firms:

$$V_A = (1-\tau_A)\left(\eta_A R_A^{-\zeta}\right) \times \int \pi_{iG}\left(\gamma_A A_i\right) di.$$

Using (B-4), we get that

$$\int \pi_{iG}\left(\gamma_A A_i\right) di = \frac{1}{\xi}\gamma_A^{\xi-1}\mathcal{D}_G.$$

This implies that

$$V_A = (1-\tau_A)\left(\eta_A R_A^{-\zeta}\right)\frac{1}{\xi}\gamma_A^{\xi-1}\mathcal{D}_G. \tag{B-8}$$

Free entry into innovation implies that

$$V_A = V_{Qj} \text{ for all } j = 1, ..., J.$$

Using (B-7), this implies that for all products j and k

$$\frac{R_{Qj}}{R_{Qk}} = \left(\frac{p_j Y_j}{p_k Y_k}\gamma_Q^{(\alpha_j-\alpha_k)(\xi-1)}\right)^{1/\zeta}. \tag{B-9}$$

Hence, relative research effort depends on relative demand $\frac{p_j Y_j}{p_k Y_k}$ and differences in the return to quality $\alpha_j - \alpha_k$. Similarly, (B-7) and (B-8) implies that

$$\frac{R_{Qj}}{R_A} = \left(\frac{1-\tau_Q}{1-\tau_A}\frac{\eta_Q}{\eta_A}\left(\frac{\gamma_Q^{\alpha_j}}{\gamma_A}\right)^{\xi-1}\frac{p_j Y_j}{\mathcal{D}_G}\right)^{1/\zeta} \tag{B-10}$$

Relative research effort depends on (i) relative demand $\frac{p_j Y_j}{\mathcal{D}_G}$, (ii) innovation policy $\frac{1-\tau_Q}{1-\tau_A}$ and (iii) differences in the innovation technology $\frac{\eta_Q}{\eta_A}\left(\frac{\gamma_Q^{\alpha_j}}{\gamma_A}\right)^{\xi-1}$.

Equations (B-9) and (B-10) together with market clearing (B-6) fully determine the allocation of research as a function of relative spending $\left(\mathcal{D}_G, [p_j Y_j]_j\right)$. These spending terms in turn depend on whether or not the economy is open to international trade.

B-1.5 Allocations in the Closed Economy

In the closed economy, the price of the physical good is given by

$$p_G = \frac{\xi}{\xi-1}\frac{w}{A} \quad \text{where} \quad A \equiv \left(\int_0^1 A_i^{\xi-1} di\right)^{\frac{1}{\xi-1}}.$$

Hence,

$$c_j(p_G, w) = \frac{w}{\left((1-\lambda_j)\left(\frac{\xi-1}{\xi}\right)^{\rho-1} A^{\rho-1} + \lambda_j\right)^{\frac{1}{\rho-1}}} \equiv \frac{w}{\psi_j(A)}$$

Using that $e = \frac{\xi}{\xi-1}w$, we have that (see (B-5))

$$\begin{aligned} \frac{e}{\prod_{j=1}^{J} p_j^{\beta_j}} &= \left(\prod_{j=1}^{J} Q_j^{\beta_j}\right) \frac{\frac{\xi-1}{\xi}e}{\prod_{j=1}^{J} c_j(p_G, w)^{\beta_j}} \\ &= \left(\prod_{j=1}^{J} (\psi_j(A) Q_j)^{\beta_j}\right) \equiv \Upsilon(A, [Q_j]) \end{aligned}$$

Hence,

$$\vartheta_k = \beta_k + \phi_k \Upsilon(A, [Q_j])^{-\varepsilon}. \tag{B-11}$$

Now consider the cost share of services (B-1). This cost share is given by

$$\sigma_j^S = \frac{\lambda_j w^{1-\rho}}{\lambda_j w^{1-\rho} + (1-\lambda_j) p_G^{1-\rho}} = \frac{\lambda_j}{\lambda_j + (1-\lambda_j)\left(\frac{\xi-1}{\xi}\right)^{\rho-1} A^{\rho-1}} \tag{B-12}$$

$$= \frac{\lambda_j}{\psi_j(A)^{\rho-1}}. \tag{B-13}$$

Now consider the demand for labor. Labor market clearing for manufacturing workers requires that

$$\frac{\xi}{\xi-1} w H_G = \sum_{j=1}^{J} \vartheta_j \left(1 - \sigma_j^S\right) eH, \tag{B-14}$$

where $1-\sigma_j^S$ denotes the cost share of goods relative to services (see (B-12)). Equation (B-14) states that total income in the goods sector (i.e. wages plus profits) has to equal total spending on goods. By the same token

$$\frac{\xi}{\xi-1} w H_S = \sum_{j=1}^{J} \vartheta_j \sigma_j^S eH.$$

Using that $e = \frac{\xi}{\xi-1} w$, we get

$$\begin{aligned}
\frac{H_S}{H} &= \sum_{j=1}^{J} \vartheta_j \sigma_j^S = \sum_{j=1}^{J} \vartheta_j \sigma_j^S \\
&= \sum_{j=1}^{J} \left(\beta_j + \phi_j \Upsilon\left(A, [Q_j]\right)^{-\varepsilon}\right) \frac{\lambda_j}{\psi_j\left(A\right)^{\rho-1}} \\
&= \left(\sum_{j=1}^{J} \beta_j \lambda_j \psi_j\left(A\right)^{1-\rho}\right) + \left(\sum_{j=1}^{J} \lambda_j \phi_j \psi_j\left(A\right)^{1-\rho}\right) \Upsilon\left(A, [Q_j]\right)^{-\varepsilon}.
\end{aligned}$$

This expression determines the share of labor in services, H_S/H, as a function of parameters and state variables $(A, [Q_j])$.

Now consider the allocation of research. In the closed economy, total spending on final good is given by

$$p_j Y_j = \vartheta_j e = \left(\beta_j + \phi_j \Upsilon\left(A, [Q_j]\right)^{-\varepsilon}\right) \frac{\xi}{\xi - 1} w,$$

where the last equality uses $e = \frac{\xi}{\xi-1} w$ and (B-11). In turn, total demand for manufacturing goods stems from the demand of final good producers as production inputs. Given a level of revenue of $p_j Y_j$, a share $\frac{\xi-1}{\xi}$ is paid to production inputs, i.e. service workers and manufacturing firms. The share that is paid to manufacturing firms is given by $1 - \sigma_j^S$, where σ_j^S is given in(B-12). Hence,

$$\mathcal{D}_G = \sum_j \frac{\xi - 1}{\xi} p_j Y_j \left(1 - \sigma_j^S\right) = w \sum_j \left(\beta_j + \phi_j \Upsilon\left(A, [Q_j]\right)^{-\varepsilon}\right) \frac{\left(1 - \lambda_j\right)\left(\frac{\xi-1}{\xi}\right)^{\rho-1} A^{\rho-1}}{\psi_j\left(A\right)^{\rho-1}}.$$

Using (B-9), (B-10), and (B-6), this determines the allocation of research expenditure as

$$\begin{aligned}
\frac{R_{Qj}}{R_{Qk}} &= \left(\frac{\beta_j + \phi_j \Upsilon\left(A, [Q_j]\right)^{-\varepsilon}}{\beta_k + \phi_k \Upsilon\left(A, [Q_j]\right)^{-\varepsilon}} \gamma_Q^{(\alpha_j - \alpha_k)(\xi-1)}\right)^{1/\zeta} \\
\frac{R_{Qj}}{R_A} &= \left(\frac{1 - \tau_Q}{1 - \tau_A} \frac{\eta_Q}{\eta_A} \left(\frac{\gamma_Q^{\alpha_j}}{\gamma_A}\right)^{\xi-1} \frac{\left(\beta_j + \phi_j \Upsilon\left(A, [Q_j]\right)^{-\varepsilon}\right) \frac{\xi}{\xi-1}}{\sum_j \left(\beta_j + \phi_j \Upsilon\left(A, [Q_j]\right)^{-\varepsilon}\right) \frac{(1-\lambda_j)\left(\frac{\xi-1}{\xi}\right)^{\rho-1} A^{\rho-1}}{\psi_j(A)^{\rho-1}}}\right)^{1/\zeta} \\
R &= R_A + \sum_j R_{Qj}.
\end{aligned}$$

These equations fully determine the allocation of researchers as a function of parameters and state variables $(A, [Q_j])$.