A Theory of Endogenous Degrowth and Environmental Sustainability*

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

March 2025

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 services, 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 (mis)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 Maarten De Ridder, David Hemous, Chris Tonetti, and audiences at the SED, the Conference on the Economics of Innovation in Memory of Zvi Griliches, the Chinese University of Hong Kong, and the University of Zurich for their helpful comments. 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. Peters and Zilibotti thank the Tobin Center for financial support. Zilibotti thanks the University of Copenhagen and the Danmarks Nationalbank for generous hospitality.

†College de France and INSEAD. philippe.aghion@insead.edu

‡University of Zurich and IIES, Stockholm University. 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.

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, aggregate demand progressively shifts from items that are intensive in material goods to items that are quality and service intensive. The structural transformation of the U.S. 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 and the economies' emission intensity, that is the amount of emissions per unit of GDP, has peaked in 1917 and declined steadily for the last 100 years. However, the ongoing

[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.

structural transformation may be too slow to resolve the environmental problem.

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 non-homothetic 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 because of its differential environmental footprint: producing a larger quantity of output increases emissions more than producing the same quantity with higher quality. A central tenet of our theory is that market forces can direct innovation along two distinct paths: reducing the cost of material production (*quantity innovation*) or enhancing the

quality of consumer goods (*quality innovation*). 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 as 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 of relative spending on physical goods makes cost-reducing innovation in material productivity less important and less profitable. Secondly, due to non-homothetic 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 productivity-led to quality-led growth. This intervention can also curtail the long-term growth rate of physical production.

Is degrowth indeed necessary to save the planet? The answer to this question hinges on the relative efficiency of quality innovation and how GDP is measured. If it is possible to increase quality sufficiently fast and such quality changes were appropriately measured, 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.

We also extend our theory to an open-economy setting, where the possibility of

trade leads to international 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, 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. Secondly, trade and specialization influence the direction of technological advancement. To analyze these factors more formally, we study 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 Review. 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) and recently developed by (Golosov et al., 2014).[4] However, this literature does not distinguish between quality and quantity based growth, nor does it factor in endogenous directed innovation.

More closely related to our analysis is the literature on the environment and endogenous directed technical change. The seminal paper by Acemoglu et al. (2012a) develops a model of directed technical change that shows how policy can influence innovation toward cleaner forms of production. Several papers have since extended it. Among them, Acemoglu et al. (2016) and, more recently, Aghion et al. (2024), construct models of growth and firm dynamics to analyze the process of energy transition. Hémous (2016) extends Acemoglu et al. (2012a) 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.[5] However, this literature does not differentiate between quality-based and quantity-based growth

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

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

[5] See Hémous and Olsen (2021) for an excellent literature review on green innovation and the energy transition.

for different consumer products but focuses entirely on how things are produced. As such, it also overlooks how endogenous directed innovation interacts with consumer demand and income effects.[6]

Our paper contributes to the ongoing debate on the sustainability of economic growth. Proponents of degrowth argue that the pursuit of unlimited economic expansion is fundamentally incompatible with the preservation of a finite stock of non-renewable resources. The foundational contributions to this literature date back to the 1970s. The Club of Rome published an influential report, *The Limits to Growth* (Meadows et al., 1972), which assessed the long-term consequences of population and economic growth on finite planetary resources. The report warned that, if unchecked, continued economic expansion could lead to environmental and economic collapse by the 21st century. A central intellectual figure in the earlier debate was Georgescu-Roegen (1971, 1974) who expressed the view that modern economic systems transform low-entropy resources, such as raw materials, into high-entropy goods. Since low-entropy resources are limited, the rate at which they are consumed ultimately determines the maximum achievable rate of economic growth.[7]

Degrowth proponents challenge the mainstream belief that green technology can reconcile economic growth with environmental sustainability. Hickel (2020) and Hickel et al. (2022) argue that while technological innovation is often promoted as a means to "decouple" economic growth from environmental damage, such decoupling is largely unrealistic. Similarly, D'Alessandro et al. (2020) argue that achieving environmental sustainability demands a fundamental economic transformation, including shifts in the composition of GDP—a key focus of our analysis.

The need to shift from quantity-driven to quality-focused economic development is a recurring theme in this debate. For instance, Latouche (2009) argues that societies should move beyond mere GDP-centric material expansion and instead prioritize improvements in life quality. Our research underscores the importance of evolving from *productivity-led* to *quality-led* growth as a means to reduce environmental impact while sustaining economic dynamism. Our framework takes a micro-founded, quantitative

[6] Beerli et al. (2020) study such interaction but their study focuses on the demand of consumer durables with no consideration of its effect on environmental sustainability.

[7] Another influential earlier work was Daly (1977) that proposed an alternative to growth-based economies in the form of a stationary economy, in which economic activities remain within ecological limits. More recently, Kallis (2011) has advocated for degrowth as a necessary response to ecological and social crises, while Kallis et al. (2012) argue for new economic paradigms that explicitly recognize biophysical constraints.

approach that goes beyond traditional green technology models, offering a new pathway for directed technological change to drive economic transformation toward sustainability. Shifting from productivity-led to quality-led growth naturally results in a more service-intensive economy—one that pollutes less and relies less on conventional GDP growth as a progress measure.[8] Rather than suppressing GDP growth, we argue that policy should accelerate this structural shift.

Stiglitz et al. (2009) propose new indicators to measure a nation's social and sustainable progress without relying solely on GDP. In this spirit, the ongoing research of De Ridder and Lukasz (2025) construct a measure of total factor productivity (E-TFP) adjusted for carbon emissions. In our model, the transition from quantity-led to quality-led growth naturally reduces emissions by reallocating innovation and production toward less pollution-intensive activities. While De Ridder and Lukasz (2025) adjust productivity measures to reflect environmental costs, our model explains how shifts in innovation and demand drive environmental sustainability—changes that standard growth accounting may overlook.

Our paper also contributes to the literature on structural change and service-led growth. Boppart (2014) developed a precursor model from which we adopt our non-homothetic preferences (PIGL), providing a theoretical foundation for key empirical facts—most notably, the steady increase in the share of household spending devoted to service-intensive goods over time. Alder et al. (2022) propose a generalization of this preference specification that encompasses our specification.[9] Fan et al. (2023) apply this class of preference to a study of service-led growth in India, while Chen et al. (2023) apply it to the Chinese economy. We contribute to this literature by introducing endogenous technical change and the resulting trade-off between quality-enhancing and productivity-enhancing innovation.

An important assumption in our theory is that the demand for quality is non-homothetic. Bils and Klenow (2001) document that higher income households systematically spend more on higher-quality versions of goods, suggesting that income growth is associated with a shift toward quality rather than just increased quantity. Their findings align with the recent trade literature. Fieler (2011) highlights that wealthier consumers have a greater elasticity of substitution for quality, meaning they

[8] Related to this discussion, Easterlin (1974) documented that happiness does not increase proportionally with income beyond a certain national income threshold, underscoring the limitations of GDP as a sole welfare indicator.

[9] An alternative specification, the non-homothetic CES class, is analyzed in Comin et al. (2021).

are more willing to pay for higher-quality goods as their income rises. In a similar vein, Fajgelbaum and Grossman (2011) also argue that richer consumers exhibit a greater willingness to pay for quality.

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 our findings to the debate on degrowth. Section 5 provides a quantitative analysis. Section 6 focuses on an open economy extension of our basic framework. 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.[10]

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.

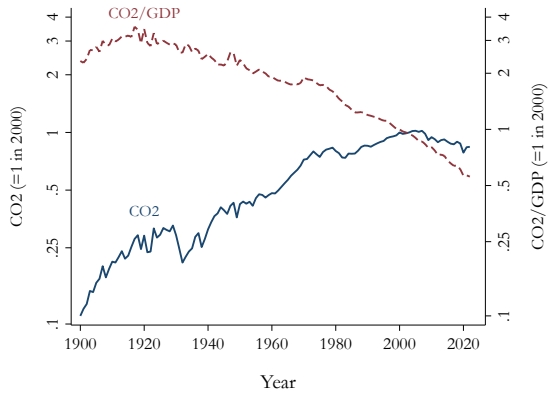
In this paper, we argue that the rise of service activities has played a key role in this shift. In Panel (b) of Figure 1, we show the share of employment in services, which has expanded from around 30% at the start of the 20th century to over 85% today. If the value added by services generates less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

In Panel (c), we show that this is the case. We use data from the National Emissions

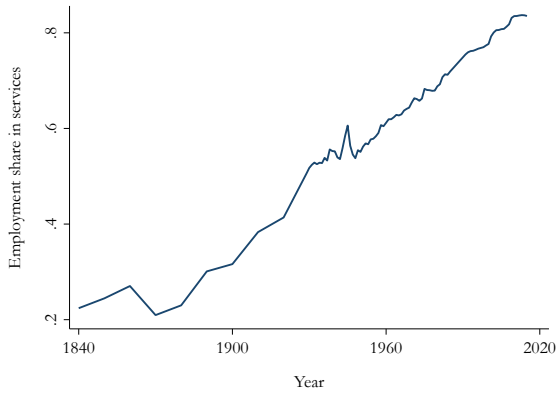
[10] A detailed description of the data used is deferred to Section 5.1.

Figure 1: Services Are Clean And Rising

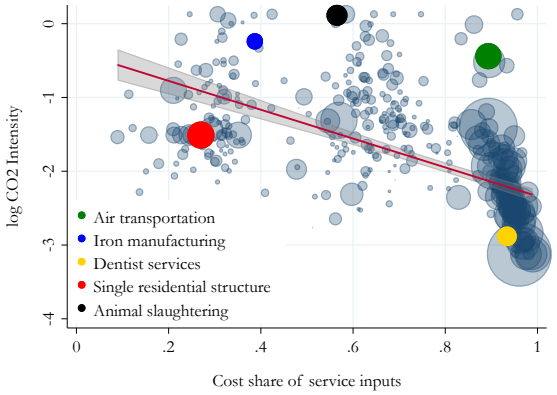
(a) Emissions in the US



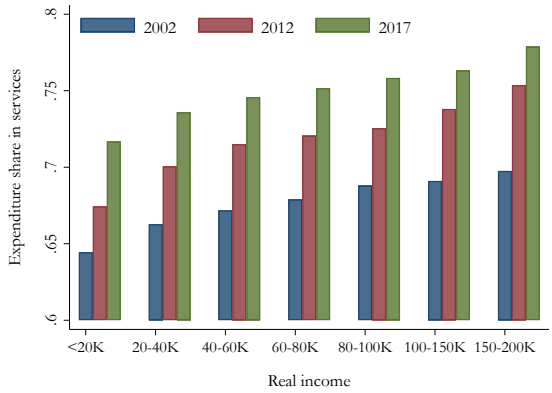
(b) The Rise of Service Employment



(c) Services are Green



(d) Non-homothetic Service Demand



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 U.S. In Panel (c) we display the relationship between the log of CO_2 intensity and the service cost share at the industy level. The CO_2 intensity is taken from the Supply Chain Greenhouse Gas dataset released by the U.S. Environmental Protection Agency. 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.

Inventory (NEI) published by the U.S. Environmental Protection Agency (EPA), which reports total CO_2 emissions for each industry. We calculate each industry's emissions intensity 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 industry i taking all input-output links into account. Likewise, the IO tables enable us to determine the service intensity of each industry, which represents the proportion of service costs in producing a dollar of output in each industry. Panel (c) shows a strong negative correlation between industrial emission intensity and the cost share of services: a 10 percentage point increase in service content corresponds, on average, to a 19.5% 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 are luxuries: the expenditure share on service-intensive goods rises with income. Second, the 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 and changes in the way goods are produced explain why the demand for services rose holding income constant.

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

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

where e/y represents emission intensity (i.e., total CO_2 emissions relative to GDP), s_{ct}^{SERV} is the service 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. Because we control for the agricultural employment share, β is identified from the

variation 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 emission 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 total land area: 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.

In columns 5 to 8, we replicate this analysis focusing on counties *within* the U.S. 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), a service sector (S), and a set of consumption good industries (C) comprising J products. The manufacturing and service sectors provide inputs to produce final goods.

Input Sectors (Goods and Services). The technology of the input sectors is described by the following CES production function:

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

[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.

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

	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 U.S. 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).

where $k \in \{G, S\}$. We assume that individual manufacturing goods are produced with the following linear technology: $y_{iG} = A_i h_{iG}$, where 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 that we will discuss below. We 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 $H_G = \int h_{iG} di$. The service sector operates under a similar technology, but, as in Baumol (1967), we assume there is no material productivity growth in services. Thus, $y_{iS} = h_{iS}$, which implies $Y_S = H_S$ and $H_S = \int h_{iS}\, di$.

Final Goods. Consumers have preferences over J different final products. They value both the quality and the quantity consumed. 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 $j \in \{1, 2, \ldots J\}$ 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_{ij} is a quality index and $\alpha_j \in [0, 1]$ captures the sensitivity of consumers' demand to quality differences for goods. Note that α_j is product specific, indicating

that this sensitivity varies across categories of final 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 input Y_G and of the service input Y_S. More formally, we assume the following technology

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

for $(i,j) \in ([0,1] \times \{1,2,\ldots,J\})$. Here, Y_{ijG} and Y_{ijS} represent the input of manufacturing good and services utilized in the production of y_{ij}, respectively. The parameter $\lambda_j \in [0,1]$ captures the service intensity of product j. Given that all varieties i are produced with the same technology, symmetry allows us to denote the price of each y_{ij} as $\tilde{p}_j$.

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

Representative Household: The consumer side of the economy consists of a large representative household comprising a continuum of individuals engaged in various activities: paid work, innovative entrepreneurship (or *researchers*), and parasitic entrepreneurship. This structure is designed to keep the dynamic aspects of the model simple. In particular, we abstract from savings decisions and introduce assumptions that ensure that all of the output is consumed within each production cycle.

Each period, all household members inelastically supply a fixed number of hours of labor. In addition, some household members draw specific skills—either innovative or parasitic—that enable them to perform entrepreneurial tasks. Researchers discover new technologies and run new firms using a superior technology for just one period. They earn monopoly profits from this activity. Parasitic entrepreneurs replace incumbent firms without making any technological improvement (hence, they produce no social surplus). They also enjoy monopoly rents for one period, after which they are replaced by new randomly drawn parasitic entrepreneurs.[12]

[12] The microfoundation assumes here that parasitic entrepreneurs must incur a fixed cost to start producing. In each period, only one parasitic entrepreneur enters the market, as potential competitors (including the incumbent researcher, if applicable) anticipate that Bertrand competition would

All agents earn the market wage. Researchers and parasitic entrepreneurs earn, in addition, profits that are transferred to the representative household. Note that the household does not make deliberate occupational choices; its only task is to pool income and share it equally to all members for them to consume.

Preferences: The representative household's preferences are parameterized by the following indirect utility function in the PIGL class:

$$\mathcal{V}\left(e,[\tilde{p}_j]_{j=1}^J,[Q_j]_{j=1}^J\right)=\frac{1}{\varepsilon}\left(e\prod_{j=1}^J\left(\frac{Q_j^{\alpha_j}}{\tilde{p}_j}\right)^{\omega_j}\right)^{\varepsilon}-\tilde{\varsigma}\prod_{j=1}^J\left(\tilde{p}_j/Q_j^{\alpha_j}\right)^{\varsigma_j}-v(\mathcal{P})$$

where $\sum_{j=1}^J\varsigma_j=0$ and $\sum_{j=1}^J\omega_j=1$. 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 (hedonic) prices $p_j\equiv\tilde{p}_j/Q_j^{\alpha_j}$. Namely,

$$\mathcal{V}\left(e,[p_j]_{j=1}^J\right)=\frac{1}{\varepsilon}\left(\prod_{j=1}^J\frac{e}{p_j^{\omega_j}}\right)^{\varepsilon}-\tilde{\varsigma}\prod_{j=1}^J p_j^{\varsigma_j}-v(\mathcal{P})$$

Roy's Identity implies that expenditure share on product k for a consumer with spending level e is given by

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

where $\beta_j\equiv\omega_j+\varsigma_j/\epsilon$ and $\phi_k\equiv\varsigma_k\times\tilde{\varsigma}$. Note that $\sum_{k=1}^J\omega_k=1$, $\sum_{k=1}^J\phi_k=0$, and $\sum_{k=1}^J\beta_k=1$.

drive profits to zero and therefore refrain from entering. Finally, we take the limit as the fixed cost approaches zero.

Equation (2) highlights the role of the demand parameters ω_k, ϕ_k, and β_k. The parameter ω_k represents the asymptotic expenditure share as spending e gets large. The parameter ϕ_k determines whether product k is income-elastic or income-inelastic: all goods k with $\phi_k < 0$ are classified as luxuries, whereas those with $\phi_k > 0$ are necessities. Finally, the parameter β_k determines the weight of the hedonic price p_k in the pseudo-price index that governs the strength of income effects.

Next, we introduce our key assumption.

Assumption 1 (**The Sophistication Ladder**). *Consumption goods $j \in \{1, 2, \dots J\}$ are ranked on a sophistication ladder, wherein good j' is more sophisticated than good j'' if and only if $j' > j''$. Moreover, $\forall j \in \{1, 2, \dots J-1\}$ $\phi_{j+1} \leq \phi_j$, $\lambda_{j+1} \geq \lambda_j$, 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 material production 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 that we assume is determined by the level of production of goods. More formally:

$$\mathcal{E}_t = \mathcal{E}\left(Y_{Gt}, z_t\right), \tag{4}$$

where $z_t \geq 1$ is a (*green technology*) trend that determines the environmental impact of production activity. Specifically, an increase in z mitigates the environmental impact of production. The function $\mathcal{E}$ is nondecreasing in Y_{Gt} and decreasing in z_t—the larger z_t the less polluting production activity. Our assumption highlights, in a somewhat extreme manner, the differing environmental footprints of goods and services: in line with

empirical evidence, we assume that manufacturing goods produce more emissions than services.[13] Moreover, (4) also highlights the fundamental difference between quantity and quality: using the same quantity of goods production Y_{Gt} with a higher quality would raise welfare and GDP (as long as quality growth is appropriately measured) but not lead to higher emissions.

3.1 Two Consumer Goods: Luxuries and Necessities

In our main analysis, we assume that $J = 2$, namely, there are only two consumer goods. We designate the index N (a mnemonic for *necessity*) for $j = 1$, and the index L (a mnemonic for *luxury*) for $j = 2$. In line with Assumption 1 we assume $\lambda_N < \lambda_L$, indicating that the luxury good is service-intensive. For further simplicity, we assume that consumers are indifferent to quality heterogeneity in the necessity, namely, we set $\alpha_N = 0$, while they exhibit sensitivity to quality heterogeneity in L. We set $\alpha_L = 1$, implying that, as far as the luxury good is concerned, consumers ultimately care about the number of quality units they purchase. In particular, $p_L = \tilde{p}_L/Q$.

3.1.1 Equilibrium Given Technology

The static equilibrium determines the equilibrium allocations given the state of technology (A) and the quality of luxury goods (Q). We proceed to the characterization of the equilibrium by considering first the production side and then the demand side of the economy.

Production: Recall that all goods and services are produced by monopolistically competitive firms. Given the isoelastic demand for different varieties, monopolists set the prices of each variety equal to a constant markup over the marginal cost—see Appendix B-1. Aggregating over the set of varieties yields the following expressions:

$$p_G = \frac{\xi}{\xi - 1}\frac{w}{A}, \quad p_S = \frac{\xi}{\xi - 1}w, \quad \tilde{p}_N = \frac{\xi}{\xi - 1}c_N\left(w, A\right), \quad \tilde{p}_L = \frac{\xi}{\xi - 1}c_L\left(w, A\right), \quad (5)$$

[13] The assumption is stronger than necessary for the main argument of our paper. We could relax it by assuming that $\mathcal{E}_t = \mathcal{E}\left(Y_{Gt}, Y_{St}, a_t\right)$, while still maintaining that the production of services is less polluting than the production of goods. However, our quantitative analysis indicates that our extreme assumption provides a reasonable approximation of the observed distribution of emissions across industries.

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 the goods-producing sector and $c_j(w, A) = \left((1-\lambda_j) p_G^{1-\rho} + \lambda_j p_S^{1-\rho}\right)^{\frac{1}{1-\rho}}$ is the unit cost of production of the $j \in \{N, L\}$ goods. Substituting in the expressions of p_G and p_S, we obtain:

$$c_j(w, A) = \frac{\xi}{\xi - 1} w \times \psi_j(A), \quad \text{where} \quad \psi_j(A) = \left((1-\lambda_j) A^{\rho-1} + \lambda_j\right)^{\frac{1}{1-\rho}}. \tag{6}$$

Note that $\psi_j' < 0$ for $j \in \{N, L\}$, namely, material productivity A reduces the cost of production $c_j(w, A)$ for both goods N and L, however more so for necessities that have a lower service content. For future reference, we note the following asymptotic properties:

$$\lim_{A \to \infty} \psi_j(A) = 0, \qquad \text{if} \quad \rho \geq 1; \tag{7}$$

$$\lim_{A \to \infty} \psi_j(A) = \lambda_j^{\frac{1}{1-\rho}}, \qquad \text{if} \quad \rho < 1. \tag{8}$$

Intuitively, when goods and services are substitutes, productivity growth in manufacturing drives the price of final goods arbitrarily low. Conversely, if goods and services are complements—which we consider the more empirically plausible case—the asymptotic production cost of the final good is determined by the cost share of services and their productivity, which has been normalized to unity.

The market prices in (5) are independent of quality. Alternatively, we can write the *quality-adjusted* prices:

$$p_N = \left(\frac{\xi}{\xi - 1}\right)^2 \psi_N(A) w \qquad \text{and} \qquad p_L = \left(\frac{\xi}{\xi - 1}\right)^2 \frac{\psi_L(A)}{Q} w,$$

where $Q = \left(\int Q_i^{\xi-1} di\right)^{\frac{1}{\xi-1}}$ is the average quality of the varieties of the luxury good. Intuitively, higher quality reduces the price per quality unit of luxury goods. The term $\left(\frac{\xi}{\xi-1}\right)^2$ captures the double marginalization effect arising from monopoly power, which is present in both intermediate and final production stages.

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:

$$\mathcal{V}\left(e, p_N, p_L, Q\right) = \frac{1}{\varepsilon}\left(\frac{e}{p_L^{1-\omega} p_N^{\omega}}\right)^{\varepsilon} + \frac{\phi}{\varsigma}\left(\frac{p_L}{p_N}\right)^{\varsigma} - v(\mathcal{P}).$$

where, in terms of the notation introduced above, $\varsigma \equiv \varsigma_L = -\varsigma_N$, $\omega \equiv \omega_N = 1 - \omega_L$, and $\phi \equiv \phi_N = -\phi_L$. Assumption 1 implies that $\phi > 0$, i.e., good N is a necessity and good L is a luxury.

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

$$\vartheta_N\left(e, p_N, p_L\right) = \omega + \phi\left(\frac{e}{p_L^{1-\beta} p_N^{\beta}}\right)^{-\varepsilon} = \omega + \phi\left(\Upsilon\left(e; A, Q\right)\right)^{-\varepsilon}, \tag{9}$$

$$\vartheta_L\left(e, p_N, p_L\right) = 1 - \omega - \phi\left(\frac{e}{p_L^{1-\beta} p_N^{\beta}}\right)^{-\varepsilon} = 1 - \omega - \phi\left(\Upsilon\left(e; A, Q\right)\right)^{-\varepsilon}, \tag{10}$$

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

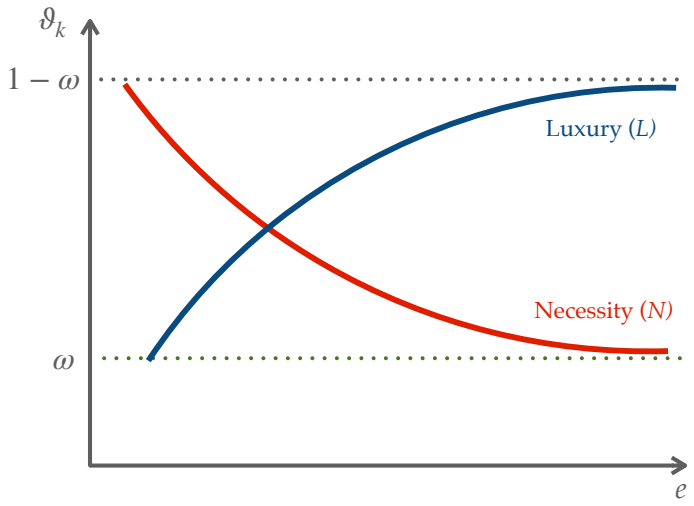
$$\Upsilon\left(e; A, Q\right) = \frac{Q^{1-\beta}}{\psi_N\left(A\right)^{\beta} \psi_L\left(A\right)^{1-\beta}} \times \left(\frac{\xi - 1}{\xi}\right)^{2} e, \tag{11}$$

and we have used the definition $\beta = \omega - \varsigma/\varepsilon$. Conditional on the expenditure level, both productivity growth in manufacturing and enhancements in quality contribute to a shift in expenditure share from the necessity to the luxury good. 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 non-homothetic adjustment. The slope of the Engel curves and the magnitude of income effects are determined by the parameter ε, which we term the *Engel elasticity*. Second, as $e \to \infty$, the expenditure shares ϑ_N and ϑ_L converge to their limiting values, ω and $1 - \omega$, respectively. The spending share on the quality-intensive luxury good approaches $1 - \omega$ from below, while the spending share on the necessity approaches ω from 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

Figure 2: Engel Curves



Notes: The figure shows the expenditure share for necessary goods (red) and luxury goods (blue) as a function of expenditure—see (9).

share on necessary goods while increasing the share on luxury goods for a given level of nominal spending e and fixed prices. Similarly, a rise in productivity, A, also shifts spending toward luxury goods. This occurs because greater productivity lowers the prices of both necessary and luxury goods, effectively raising real income.

Income: In equilibrium, all income from labor and firms' profits accrue to the representative household. Moreover, all household income is allocated toward consumption goods. Therefore, the equilibrium expenditure of the representative consumer can be expressed as:

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

where Π denotes aggregate profits in the economy, and H is the aggregate labor force. These profits accrue part in the intermediate sectors G and S and part in the final good sectors N and L. Because of the constant markup, profits are proportional to

wage income, and more specifically[14]

$$\Pi = \frac{1}{\xi - 1} wH + \frac{1}{\xi - 1}\left(\frac{\xi}{\xi - 1} wH\right) = \left(\left(\frac{\xi}{\xi - 1}\right)^2 - 1\right) wH,$$

which implies, using the normalization $w = 1$, that $e = \left(\frac{\xi}{\xi-1}\right)^2$. Finally, with some slight abuse of notation, we rewrite the term Υ in (11) as

$$\Upsilon\left(A, Q\right) = \left(\left(1 - \lambda_N\right) A^{\rho-1} + \lambda_N\right)^{\frac{\beta}{\rho-1}} \left(\left(1 - \lambda_L\right) A^{\rho-1} + \lambda_L\right)^{\frac{1-\beta}{\rho-1}} Q^{1-\beta}. \tag{12}$$

Υ is fully determined by Q and A, and it increases with both arguments. Importantly, the fact that $\left(\frac{\xi-1}{\xi}\right)^2 e = 1$ highlights that monopolies do not introduce distortions in this economy.[15] Consequently, market power only affects the distribution of income between wages and profits, leaving the overall allocation unaffected due to the assumption of exogenous labor supply. Therefore, the division of labor between goods and services production—our next focus—satisfies the production efficiency criterion.

In the rest of the paper, it is useful to define the expenditure share on the necessary good as

$$\vartheta_N = \vartheta(A, Q) = \omega + \phi\left(\Upsilon\left(A, Q\right)\right)^{-\varepsilon}, \tag{13}$$

where we note that ϑ is decreasing in both its arguments and $\lim_{Q\to\infty} \vartheta(A, Q) = \omega$. Moreover, $\lim_{A\to\infty} \vartheta(A, Q) = \omega$, if $\rho \geq 1$. With this notation, the expenditure share on the luxury good is $\vartheta_L = 1 - \vartheta(A, Q)$.

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 services, the total factor payment (including wages and profits) must equal the associated value added.

Let $\sigma_k(A)$ denote the cost share of good inputs in the production of consumption

[14] The term $\frac{wH}{\xi-1}$ captures the profit generated by intermediate manufacturing and service firms, while the term $\frac{1}{\xi-1}\left(\frac{\xi}{\xi-1} wH\right)$ captures the profit generated by final good firms.

[15] This result arises because firms exhibit uniform market power across all sectors—cf. Epifani and Gancia (2011).

good k. Given the CES production function, we obtain

$$\sigma_k(A) = \frac{(1-\lambda_k)\, p_G^{1-\rho}}{(1-\lambda_k)\, p_G^{1-\rho} + \lambda_k p_S^{1-\rho}} = \frac{(1-\lambda_k)\, A^{\rho-1}}{(1-\lambda_k)\, A^{\rho-1} + \lambda_k} \quad \text{for} \quad k \in \{N, L\}. \tag{14}$$

Given parameters, the cost share is fully determined from the quantity productivity A. σ_k decreases in A if $\rho < 1$, i.e., services and goods are complements, and increases in A if $\rho > 1$, i.e., services and goods are substitutes.

Equating factor payment to value added yields:

$$wH_G = \sigma_N(A)\vartheta\left(A, Q\right) wH + \sigma_L(A)\left(1 - \vartheta\left(A, Q\right)\right) wH, \tag{15}$$

$$wH_S = (1 - \sigma_N(A))\vartheta\left(A, Q\right) wH + (1 - \sigma_L(A))\left(1 - \vartheta\left(A, Q\right)\right) wH. \tag{16}$$

Substituting the expressions of σ_N and σ_L into (15) and (16), allows us to prove the following comparative statics results.

Proposition 1 (Structural Change)**.** *The service employment share H_S/H is:*

1. *increasing in Q;*
2. *increasing in A if $\rho < 1$;*
3. *decreasing in A if $\rho > 1$ and ϕ is small;*
4. *constant if $\rho \to 1$ and $\phi = 0$ (Cobb Douglas).*

Proof of Proposition 1. The four results in the proposition follow then from Equation (16) and from the following properties of our theory: (i) Assumption 1 implies that $\lambda_L > \lambda_N$, which in turn implies that $\sigma_N(A) > \sigma_L(A)$. (ii) $\vartheta\left(A, Q\right)$ is decreasing in both Q and A, if $\phi > 0$, while, $\vartheta\left(A, Q\right) = \omega$, if $\phi = 0$; moreover, $\lim_{\phi \to 0} \vartheta_A\left(A, Q\right) = 0$, where the subscript denotes a partial derivative. (iii) Equation (14) implies that, for $k \in \{N, L\}$, $\sigma_k(A)$ decreases (increases) in A, if $\rho < 1$ ($\rho > 1$), while $\sigma_k(A)$ is independent of A, if $\rho = 1$. In particular, Part 1 follows from (i) and (ii), while Parts 2, 3, and 4 follow from (ii) and (iii). □

Proposition 1 highlights the distinct roles of Q and A in the process of structural change. An increase in Q shifts demand towards luxury goods through an income effect, leaving the factor allocation within products unchanged. As a result, higher quality Q

raises the aggregate employment share of services, given the service-intensive nature of luxury goods. Conversely, an increase in A has two effects. First, similar to an increase in Q, it enriches households, shifting demand towards luxury goods and increasing the service employment share. Second, it impacts the cost structure in the production of final goods. If manufacturing and service inputs are complementary ($\rho < 1$), a rise in A increases the cost share of services in final industries, amplifying the income effect. However, if they are substitutes ($\rho > 1$), a rise in A raises the cost share of manufacturing goods.

3.2 Directed Innovation: Quality versus Quantity

In this section, we examine the determinants of technical progress. We postulate the existence of a mass R of researchers, 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 luxury good (Q_i). We denote by R_Q and R_A, respectively, the research effort directed to increase Q and A, subject to the standard market clearing condition $R_Q + R_A = R$.

The rate at which a unit of research effort directed toward activity $s \in \{A, Q\}$ translates into a successful 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 research efficiency in sector s. A successful innovation augments the quality or productivity of a randomly selected firm by a factor $\gamma > 1$. We assume γ to be sufficiently large to enable the new firm to set the unconstrained monopoly price. Furthermore, we assume that researchers reap profits only for a single period.[16]

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

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

[16] 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 $j = G$ if $s = A$ and $j = L$ if $s = Q$. Here, (i) π_{iG} and π_{iL} denote, respectively, the profits from a productivity innovation in the variety i of sector G and the profits from a productivity innovation in the variety i of sector L and (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 Appendix B-1, we show that the equilibrium profits are equal to:

$$\begin{aligned}\int \pi_{iL} di &= \frac{1}{\xi}\left(1-\vartheta(A,Q)\right) wH \\ \int \pi_{iG} di &= \frac{1}{\xi}\left(\vartheta(A,Q)\sigma_N(A) + \left(1-\vartheta(A,Q)\right)\sigma_L(A)\right) wH.\end{aligned}$$

Note that the demand for quality arises directly from the consumption of luxury goods. In contrast, the demand for manufactured goods is derived from their cost share in producing both necessities and luxury goods.

In equilibrium, the value of the marginal product of researchers will be equalized between quality and productivity improvements. Substituting the expressions for V_Q and V_A, this arbitrage condition can be expressed as:

$$\frac{R_Q}{R_A} = \left(\frac{(1-\tau_Q)\eta_Q}{(1-\tau_A)\eta_A}\right)^{\frac{1}{\zeta}} \times \left(\frac{1-\vartheta(A,Q)}{\sigma_N(A)\vartheta(A,Q) + \sigma_L(A)\left(1-\vartheta(A,Q)\right)}\right)^{\frac{1}{\zeta}} \tag{17}$$

$$= \left(\frac{(1-\tau_Q)\eta_Q}{(1-\tau_A)\eta_A}\right)^{\frac{1}{\zeta}} \times \left(\frac{1-\vartheta(A,Q)}{1-\frac{H_S}{H}}\right)^{\frac{1}{\zeta}} \tag{18}$$

The right-hand side of (17) captures the effects of both technological and demand forces. To isolate these effects, consider the case in which $\phi = 0$, corresponding to Cobb-Douglas homothetic preferences. In this case, only technological forces operate:

$$\frac{R_Q}{R_A} = \left(\frac{(1-\tau_Q)\eta_Q}{(1-\tau_A)\eta_A}\right)^{\frac{1}{\zeta}} \times \left(\frac{1-\omega}{\omega\sigma_N(A) + (1-\omega)\sigma_L(A)}\right)^{\frac{1}{\zeta}},$$

implying that R_Q/R_A increases or decreases in A depending on whether $\rho < 1$ (gross complements) or $\rho > 1$ (gross substitutes). Intuitively, when goods and services are complements (substitutes) in the production functions for final goods, technical progress in manufacturing raises the service (manufacturing good) share in final production—the classical Baumol effect (Baumol, 1967). This shift diminishes (raises) the relative profitability of innovations aimed at enhancing material productivity. As

a result, purely technological forces drive innovation incentives away from (toward) material productivity and toward (away from) quality improvements as A increases. In the particular case of Cobb Douglas production function ($\rho \to 1$), the cost shares σ_N and σ_L are constant, implying that R_Q/R_A is independent of A. Notably, when $\phi = 0$, Q has no effect on the direction of technical progress.

Next, consider the 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: $g_Q = R_Q^{1-\zeta}\eta_Q(\gamma - 1)$ and $g_A = R_A^{1-\zeta}\eta_A(\gamma - 1)$. Note that, under our assumptions, the model results in a backward-looking dynamic system in terms of the technology state vector (A_t, Q_t), which fully characterizes the equilibrium path.

Taking Stock. Our characterization of the static equilibrium has established that, given (A_t, Q_t), the income effect is determined by the term Υ in (12), while the sectoral labor allocation is determined by (15) and (16)). At the same time, (A_t, Q_t) also fully determines the dynamic equilibrium. The allocation of research is given by (17) and the market clearing condition for research skills. This allocation in turn implies the law of motion for both quality and productivity growth.

3.3 Asymptotic Equilibrium Dynamics

In this section, we discuss the asymptotic equilibrium dynamics as both Q and A become arbitrarily large. The properties of such equilibrium depend critically on whether goods and services are gross complements or substitutes in the production of final goods as determined by the technological parameter ρ. When they are gross substitutes ($\rho > 1$), $\sigma_N \to 1$ and $\sigma_L \to 1$, meaning that the share of labor allocated to the production of goods approaches unity. All workers are ultimately employed in manufacturing. In contrast, when they are gross complements ($\rho < 1$), $\sigma_N \to 0$ and $\sigma_L \to 0$, indicating that the share of labor allocated to services approaches unity. In this case, all workers eventually shift to service production. Intuitively, this reflects the situation where goods become increasingly efficient to produce, lowering their relative cost contribution and making labor-intensive services the dominant factor in overall

production costs. In the Cobb Douglas knife-edge case, we have a positive share of employment in both goods and services.

These results have implications for the asymptotic direction of technical change that we summarize in the following proposition.

Proposition 2. *Asymptotically, the equilibrium direction of technical change and the ensuing productivity growth are characterized as follows:*

1. [**Gross Complements**] *If* $\rho < 1$*, then,* $R_Q \to R$ *and* $R_A \to 0$*. Moreover,*

$$g_A \to 0 \quad \textit{and} \quad g_Q \to R^{1-\zeta}\eta_Q(\gamma - 1).$$

2. [**(Weakly) Gross Substitutes**] *If* $\rho \geq 1$*, then* $\frac{R_Q}{R_A} \to \Phi(\rho)$ *and*

$$g_A \to \left(\frac{1}{1+\Phi(\rho)}R\right)^{1-\zeta}\eta_A(\gamma-1) \quad \textit{and} \quad g_Q \to \left(\frac{\Phi(\rho)}{1+\Phi(\rho)}R\right)^{1-\zeta}\eta_Q(\gamma-1),$$

where

$$\Phi(\rho) = \begin{cases} \left((1-\omega)\frac{\eta_Q}{\eta_A}\frac{1-\tau_Q}{1-\tau_A}\right)^{\frac{1}{\zeta}}, & \textit{if } \rho > 1; \\ \left(\frac{1-\omega}{\omega\lambda_N+(1-\omega)\lambda_L}\frac{\eta_Q}{\eta_A}\frac{1-\tau_Q}{1-\tau_A}\right)^{\frac{1}{\zeta}}, & \textit{if } \rho = 1. \end{cases}$$

The long-run trajectory of technical progress depends critically on the substitutability/complementarity between goods and services in the production of final goods. When goods and services are gross complements—the empirically relevant case—the goods-producing sector eventually vanishes and the service employment share approaches unity. As a result, material productivity growth eventually tapers off, and innovation focuses *exclusively* on quality. In this case, the economy becomes "weightless" in the long run.

If goods and services are gross substitutes (and in the Cobb Douglas case), research efforts are split between increasing quality and improving material productivity and material productivity growth is positive in the long-run. In this case, overall material production Y_G will grow at a positive rate even asymptotically.[17]

[17] Note that $\Phi(1) > \Phi(\rho)$ if $\rho > 1$, i.e., the "knife-edge" Cobb-Douglas case features faster quality growth than the case of $\rho > 1$. The reason is that, in the case of $\rho > 1$, all workers are employed in goods production in the long run, while this is generally not true under a Cobb Douglas production function.

In addition to what we discussed above, three exogenous factors are also important. The first is the relative efficiency of quality-enhancing research relative to productivity-enhancing research (i.e., η_Q vs. η_A). The second is policy, namely the relative wedge on quality innovation compared to productivity innovation (i.e., τ_Q vs. τ_A). The third is the preference parameter ω that determines the asymptotic expenditure share on luxury goods relative to necessities. This parameter is only relevant if $\rho \geq 1$.

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 consider the theory's predictions regarding the pollution trajectory. Recall that pollution originates from material production, specifically $Y_{Gt} = A_t H_{Gt}$ (see (24)) where H_{Gt} is given as in (15). Over time, the pollution dynamics hinge on the technical progress in material production and the structural change shifting employment from good-intensive to service-intensive activities. It also depends on the development of green technologies that reduce emissions per unit of production, parameterized by the term z_t.

Elasticities. The elasticity of substitution between goods and services is a key determinant of long-term dynamics. Proposition 2 establishes that, if $\rho \geq 1$, both A and Q grow without bound, leading to an asymptotic flow of new emissions growing at a rate $g_A > 0$. In this scenario, an abatement rate (i.e., the growth rate of green technology z_t) in excess of g_A is a necessary (thought not sufficient) condition to avoid an environmental disaster. Achieving this outcome may require a policy intervention that imposes a sufficiently large tax on quantity innovation.

Conversely, if goods and services are gross complements ($\rho < 1$), our theory predicts that productivity-enhancing technical progress will asymptotically decline to zero.[18] In this case, any positive abatement rate ensures that emissions will eventually begin to decline.[19] While more benign, this scenario does not guarantee environmental sustain-

[18] Note that, even in the case of $\rho < 1$, Equation (15) implies that $\lim_{t\to\infty} Y_{Gt} = \lim_{t\to\infty}(A_t \times H_{Gt}) = \infty$. In other words, Y_G grows unboundedly while its growth rate asymptotically declines to zero.

[19] This conclusion does not rely on the extreme assumption that emissions are independent of service activity. When $\rho < 1$, the model predicts that, in the long run, the growth rate of both goods and

ability, as the transition away from material production could progress too slowly to avert disaster. Thus, policy intervention could be essential even in this case. The silver lining, however, is that such intervention would align with and accelerate an existing trend, rather than counteracting market forces.

Degrowth. How can this discussion inform the debate on degrowth? To address this question, we examine more formally the evolution of economic activity within our model. A complicating factor is that, in models with non-homothetic preferences, defining a deflator for the economy's GDP becomes ambiguous since expenditure shares across different goods vary with income. To circumvent this problem, we focus on the asymptotic economy, where expenditure shares are approximately constant, allowing for a standard real GDP definition. Suppose first that we were to measure GDP at market prices, that is without adjusting for quality. In this case, we obtain:[20]

$$GDP_{market} \approx \frac{e}{\tilde{p}_L^{1-\omega} \times \tilde{p}_N^{\omega}} = \frac{1}{\psi_L(A)^{1-\omega} \times \psi_N(A)^{\omega}},$$

where we recall that $\psi_N' < 0$ and $\psi_L' < 0$. Using the properties of the ψ_j functions as outlined in Equations (7)–(8), it follows that, as $A \to \infty$,

$$GDP_{market} \to \infty, \qquad \text{if} \quad \rho \geq 1,$$

$$GDP_{market} \to \lambda_N^{\frac{\omega}{\rho-1}} \lambda_L^{\frac{1-\omega}{\rho-1}} < \infty, \qquad \text{if} \quad \rho < 1.$$

When goods and services are gross complements, GDP_{market} has an upper bound due to the absence of productivity improvements in the service sector—a phenomenon known as Baumol's cost disease.

GDP_{market} reflects expenditure at market prices. However, it is not a welfare-relevant measure of GDP. We can construct an adjusted measure that accounts for quality improvements. Define

$$GDP_{adjusted} \approx \frac{e}{p_S^{1-\omega} \times p_G^{\omega}},$$

services production approaches zero.

[20] The argument presented in the text does not depend on any specific weights in the GDP price deflator. The same qualitative conclusions hold, for instance, if we replace ω with β. In the numerical analysis below, we employ chained indices with Törnqvist weights rather than assuming constant expenditure shares.

where nominal GDP is deflated using quality-adjusted prices instead of market prices. Our equations imply that

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

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 shows that $GDP_{adjusted}$ can grow unboudedly even in a hypothetical scenario where material production has ceased to grow entirely. This is the case in equilibrium if goods and services are gross complements, i.e., $\rho < 1$. The crux is that Q can increase unboundedly.

In practice, statistical offices aim to account for quality improvements. In our quantitative analysis below, we parameterize the extent to which official statistics reflect these improvements. According to our theory, as long as statistics capture a positive share of quality improvements, GDP growth will remain positive in the long run, eventually being entirely driven by quality improvements in luxury goods.

In conclusion, our theory can offer a new perspective on the degrowth debate. While measured GDP growth is bound to decline and possibly stop, quality-led (and service-led) growth can be self-sustained. Moreover, it can turn growth environmentally sustainable.

5 Quantitative Analysis

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

5.1 Data and Measurement

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 A-1.

To measure 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. To map these

data to our model, we aggregate these goods into two mutually exclusive groups—luxuries and necessities. As in our model, we define luxuries and necessities by their service content. To compute the service content of different 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 the Appendix, we use the IO Table together with the BEA bridge tables to compute the total service content embodied in the output of each industry until it reaches final consumers. 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. We associate all final products with being a necessity if their service content is below the median service share and with being a luxury if their service content is above the median service share (see Appendix Section A-1.3). Given this classification, we can compute the spending share of individual i on luxuries and necessities, ϑ_{it}^N and ϑ_{it}^L, and the overall service cost shares σ_N^S and σ_L^S.

To calculate the environmental footprint of each final product k, we rely on the National Emissions Inventory (NEI) from the EPA, which reports total CO2 emissions for each industry. We calculate the emissions intensity by dividing total emissions by total sales. Sectoral linkages in the IO tables are then used to compute the emission intensity of each final product k, denoted as e_k.

In addition to CO2, we also observe total emissions for five additional pollutants (particulates smaller than 10 microns (PM10), volatile organic compounds (VOC), nitrogen oxides (NOx), sulfur dioxide (SO2), and carbon monoxide (CO)). While we do not use this information to calibrate our model, we show below that the service share is also a key determinant of environmental damages using these measures.

5.2 Calibration Strategy and Estimation Results

Our model is characterized by 13 structural parameters and the emission function $\mathcal{E}$ (see (4)). In addition, as highlighted in Section 4, we explicitly allow for the fact that official price indices might mismeasure the growth of quality. As we explain in detail below, we parametrize the degree of mismeasurement with a single parameter μ and estimate it from our data. As a consequence, the set of structural parameters is given by:

$$\mathcal{P} = \{\underbrace{\varepsilon, \omega, \phi, \beta, \xi}_{\text{Preferences}}, \underbrace{\rho, \lambda_N, \lambda_L}_{\text{Technology}}, \underbrace{R, \zeta, [\eta_s]_{A,Q}, \gamma}_{\text{Innovation}}, \underbrace{\mu}_{Q\text{-measurement}}, \underbrace{\mathcal{E}(.)}_{\text{Environment}} \} \tag{19}$$

Household preferences are described by the Engel elasticity ε, the asymptotic expenditure share on good-intensive products ω, the preference shifter ϕ (which determines whether service-intensive goods are luxuries), the parameter β that governs the importance of necessities in consumer preferences, and the elasticity of substitution across individual varieties ξ. The production side is defined by the elasticity of substitution between goods and services, ρ and the service intensities of luxuries and necessities λ_N and λ_L. 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.

We calibrate the parameters in (19) 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: ω, ε, ϕ, β, and ξ. The Engel elasticity, ε, is an important parameter because it determines the strength of income effects. We follow the strategy of Fan et al. (2023) and estimate ε from the cross-sectional correlation between household expenditure and the expenditure share on necessities. Equation (9) implies that

$$\ln(\vartheta_N(e, p_L, p_N) - \omega) = \ln \phi - \varepsilon \ln e + \varepsilon \ln \left(p_N^{\beta} p_L^{1-\beta} \right) \tag{20}$$

Hence, the elasticity between the distance of the expenditure share on necessities ϑ_N and their asymptotic share ω is constant and given by the Engel elasticity ε.

To implement (20) we need to know the value of ω. Our theory implies that all individuals' expenditure shares on necessities should be bounded below by ω. Empirically, the 1% percentile of the observed distribution is equal to 16% and hence already very small. We thus set the asymptotic share ω equal to zero. Equation (20) then implies that we can estimate ε from the regression

$$\ln(\vartheta_{iNt}) = \delta_t - \varepsilon \ln(e_{it}) + x_{it}' \gamma + u_{it}, \tag{21}$$

where ϑ_{iNt} is the expenditure share on necessities of household i at time t, e_{it} is

Table II: Non-Homothetic Service Demand: Estimating ε

	log (Exp. Share Necessities)				
	(1)	(2)	(3)	(4)	(5)
log (Exp)	-0.069***	-0.318***	-0.254***	-0.316***	-0.367***
	(0.002)	(0.020)	(0.029)	(0.033)	(0.039)
Family Size	Yes	Yes	Yes	Yes	Yes
HH Controls	Yes	Yes	Yes	Yes	Yes
IV	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	No	No	No
Year	All	All	2002	2012	2017
F-Stat First Stage		30	12	11	12
N	30774	25609	9891	7967	7751

Notes: Column 1 reports the OLS relationship between households' expenditure share on luxuries and their expenditure across all years including year fixed effects. Column 2 reports the IV estimate using occupational fixed effects as instruments for household expenditure across all years, while columns 3, 4, and 5 report the IV estimate for each year. All specifications control for a set of fixed effects for the size of the household, geographic location, education, race, and marital status.

total household expenditure, and the time fixed effect δ_t controls for the prices $p_N^{\beta} p_L^{1-\beta}$, which are common across households. In addition, in (21) we also control for additional observable covariates x_{it} which could induce a correlation between household spending and the demand for necessities.

We report the results in Table II. In the first column, we estimate (21) using OLS by pooling the CEX data from all years. We control for household size, marital status, race, education, and the geographic location of the household.[21] We find an elasticity of 0.07. In column 2, we implement an IV strategy to estimate ε following the approach of Fan et al. (2023). We do so for two reasons. First, we suspect that measurement error in individual expenditure will bias the coefficient. Second, we aim to capture the variation in permanent income rather than transitory variation as we believe that the former is more informative about non-homothetic demand. We therefore instrument individual log expenditure e_{it} with a full set of occupation fixed effects. Intuitively, we identify the elasticity ε from the systematic variation in spending between high- and low-income occupations.

[21] We observe if the household lives in a urban or rural area, and four broad regions (Northeast, Midwest, South, and West). In terms of the other characteristics (age, education, race, etc.), we associate each household with the characteristics of household member that was assigned as the reference person when the survey was answered.

Column 2 shows that the resulting estimate of ε is given by 0.318 and indeed larger in magnitude. Finally, in column 3–5 we run the same regression separately for each year. We find estimates between 0.25 and 0.37 generally close to the pooled estimate. For our quantitative analysis, we take the estimate in column 2 as our baseline.

To identify the preference parameter ϕ, note that we can, without loss of generality, normalize the level of productivity and quality in our base year 2002 to unity, i.e., $A_{2002} = Q_{2002} = 1$. This implies that $\Upsilon = 1$—see (12)—and that the expenditure share of necessities, $\vartheta(A, Q)$, is simply equal to ϕ—see (13)). Labor market market clearing therefore requires that the employment share of services in 2002—see (16)—is given by

$$\frac{H_{S,2002}}{H_{2002}} = (1 - \sigma_{N,2002})\phi + (1 - \sigma_{L,2002})(1 - \phi), \tag{22}$$

where $\sigma_{N,2002}$ and $\sigma_{L,2002}$ are the cost shares of goods for necessities and luxuries respectively. Equation (22) can be solved for ϕ given data on $\sigma_{j,2002}$.

Finally, we set the elasticity of substitution ξ to 5, a consensus estimate in the literature.

Technology Parameters: ρ, λ_N and λ_L. The parameters λ_k determine the weight of service inputs within the production function for final goods. As a consequence, λ_k directly maps to the cost share, σ_k and we chose (λ_N, λ_L) to match the observed service cost shares for necessities and luxuries in 2002. Using the normalization that $A_{2002} = 1$, (14) directly implies that $\lambda_k = 1 - \sigma_{k,2002}$.

The parameter ρ determines the elasticity of substitution between goods and service workers. We follow Herrendorf et al. (2013), who argue that goods and services are complements and set $\rho = 0.5$. This implies that growth of quantity productivity A *increases* the cost share of services. Below we show that this is the case empirically.

The Innovation Process: R, ζ, γ, η_A, η_Q. Finally, consider the innovation process. The number of researchers R, the step size γ and the efficiency of research labor η_k are not separately identified. We thus set $R = 0.1$, implying that around 10% of the labor force is devoted to research activities, and $\gamma = 1.5$, meaning each successful innovation boosts productivity by 50%. 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 the two R&D efficiencies η_A and η_Q. We calibrate these using three moments: the service

employment share in 1950, GDP per capita growth between 1950 and 2000, and GDP per capita growth between 2000 and 2020. Intuitively, quantity and quality growth both raise GDP pc. At the same time, they have different impacts on the service employment share. While quality growth affects service employment only through the income effect, quantity growth also raises service employment via technological substitution—the Baumol channel. Moreover, the reallocation of research effort toward quality throughout the 20th and 21st century will affect the relative growth rates from 1950–2000 and 2000–2020 depending the relative research efficiencies. As such, using these three moments, we can identify η_Q and η_A separately.

The Measurement of Quality: μ As highlighted in our discussion in Section 4, we explicitly allow for the possibility that quality growth might only be partially measured. More specifically, we assume that the price index for luxury goods as measured from the BLS is given by

$$p_{Lt}^{BLS} = Q_t^{1-\mu} p_{Lt} = \left(\frac{\xi}{\xi - 1}\right)^2 \frac{\psi_L(A)}{Q^\mu} w. \tag{23}$$

If $\mu = 1$, quality is perfectly measured and the BLS price index, p_{Lt}^{BLS}, coincides with the welfare-relevant quality-adjusted prices p_{Lt}. If $\mu = 0$, quality is not measured at all and the BLS price index would understate the decline in the price of luxuries.

For now, we assume that $\mu = 2/3$, indicating that 2/3 of quality growth is captured in the official GDP statistics from the BLS.[22] In the future, we plan to use data on price inflation at the product level to discipline μ. Quantitatively, our results are not overly sensitive to changes in μ.

CO_2 Emissions: $\mathcal{E}_t$. We parameterize the emissions function $\mathcal{E}$ in (4) as follows:

$$\mathcal{E}_t = \kappa_{\mathcal{E}} \frac{Y_{Gt}}{z^t}, \tag{24}$$

where $\kappa_{\mathcal{E}}$ is a scaling parameter to link the production of goods to overall emissions. Note that we assume that only goods production leads to emissions, whereas the weightless service part does not pollute. As discussed above, the term $z \geq 1$ captures the

[22] See Aghion et al. (2019) who estimate about half a percentage point of missing growth per year due to not captured quality improvements from creative destruction.

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. We calibrate z and $\kappa_{\mathcal{E}}$ for our model to match the observed CO2 emission in 1980 and 2000.

Estimated Parameters. We summarize all parameters and corresponding moments in Table III. In line with the observed cost shares of services, we estimate that luxuries are substantially more service-intensive than necessities: $\lambda_L = 0.93 > \lambda_N = 0.41$. We also estimate that the research efficiency for quantity growth (η_A) is larger than for quality growth (η_Q). Finally, we estimate a substantial role for green technological progress: $z = 1.0138$, implying that the environmental footprint of material output declines by 1.4% per year.

Table III: Structural parameters

Parameter	Value	Target	Target value
ε	0.318	Engel curve slope	0.318
λ_L	0.93	Cost share of services in L (IO table)	0.93
λ_N	0.41	Cost share of services in N (IO table)	0.41
η_A	0.3371	$\text{GDP}_{1950}/\text{GDP}_{2000}$	0.413
η_Q	0.1099	$\text{GDP}_{2019}/\text{GDP}_{2000}$	1.269
ϕ	0.3009	1950 U.S. service share	54.5
		2002 U.S. service share	77
z	1.017	1980 U.S. CO2 emissions	$4,721$
$\kappa_{\mathcal{E}}$	144.919	2000 U.S. CO2 emissions	$5,724$
μ	2/3	Set exogenously	-
ω	0	Set exogenously	-
β	0.4	Set exogenously	-
ξ	5	Set exogenously	-
ρ	0.5	Set exogenously	-
ζ	0.5	Set exogenously	-

Notes: The table reports all structural parameters and the corresponding moments. As explained in the text, without loss of generality we normalize $R = 0.1$, $\gamma_A = \gamma_Q = 1.5$, and $A_{2002} = Q_{2002} = 1$.

5.3 Model Fit

Our model provides a good fit to the data. In the upper left panel of Figure 3, we compare the model's predictions with historical data on GDP per worker growth since

1950. The upper right panel focuses specifically on the growth rate of GDP per worker. In both cases, the data is shown in grey, while the model's predictions are depicted in red. Our model successfully captures the overall significant rise in GDP per worker since 1950 as well as the subsequent decline in economic growth. We should note that the average growth rates from 1950 to 2000 and from 2000 to 2019 are specific targets of our calibration.

To illustrate the role of quality measurement, we also present two extreme scenarios: GDP per worker growth assuming all quality improvements are fully measured (dotted lines) and growth assuming none of the quality improvements are captured (dashed lines). Figure 3 shows that our model predicts a decline in growth even if all quality improvements were accounted for. However, the observed decline in measured growth is more pronounced because quality growth accelerates over time. If quality improvements were entirely unmeasured, overall growth between 2000 and 2019 would have been only 0.6% (see the light-red bar).

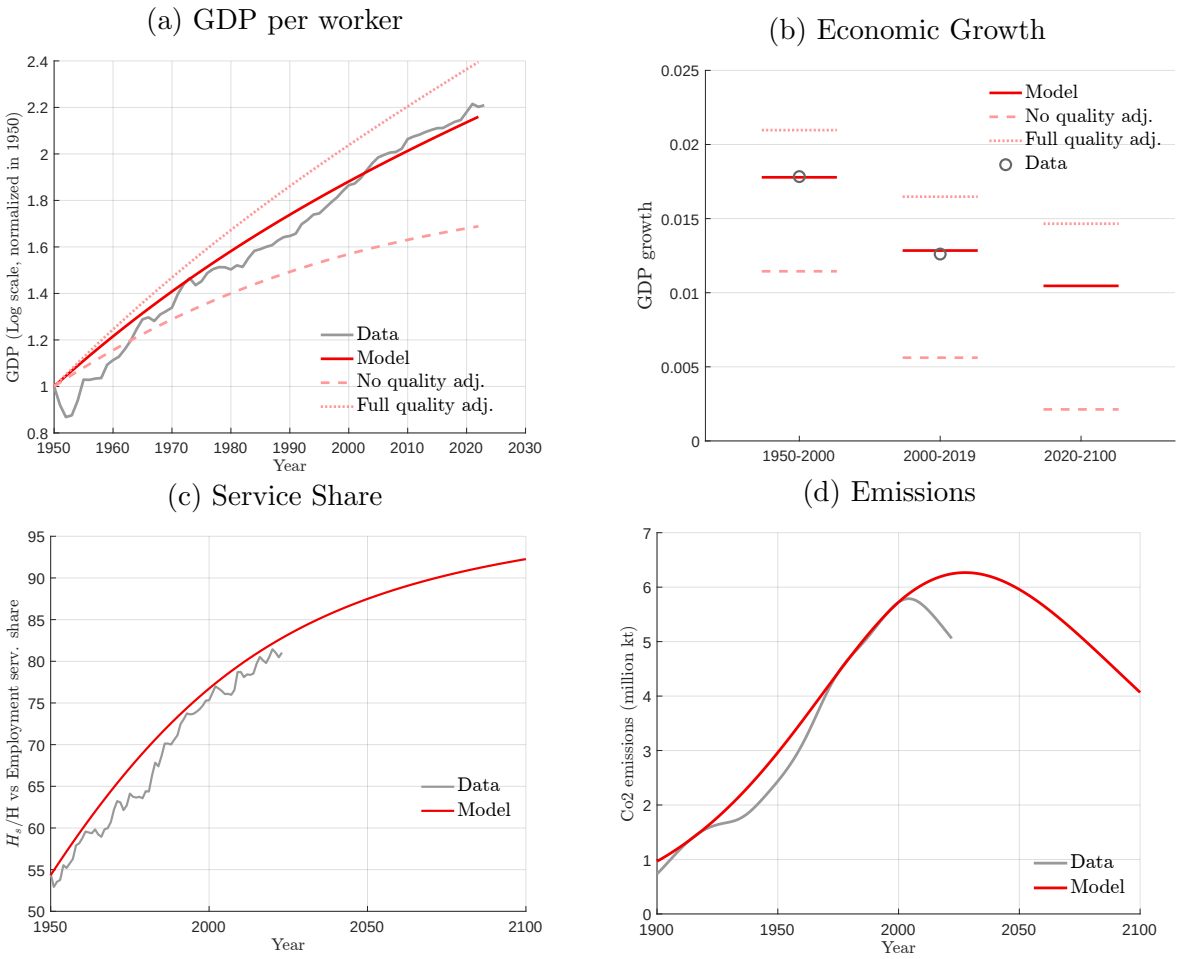
The upper right panel of Figure 3 also includes projected economic growth up to 2100. Even though GDP per capita would still increase by 1.5%, measured growth would be close to zero if quality improvements were completely unaccounted for. This occurs because future growth will be primarily driven by quality improvements rather than increases in quantity.

In the lower panels of Figure 3, we focus on the employment share of services (left panel) and the time series of emissions. The model accurately reflects the rise in the service employment share, partly due to its calibration to match the U.S. service share in 1950 and 2002. Interestingly, the model also approximates the employment share in 1900 reasonably well, even though this data point is not directly used in our estimation. Looking ahead, our model predicts that the shift toward the service sector will continue, albeit at a slower pace, as the majority of the population is already employed in services.[23]

Regarding emissions, the model closely captures the long-term trend since 1900, despite being calibrated only to match total emissions in 1980 and 2000. Notably, the model also replicates the pronounced hump-shaped pattern of CO_2 emissions, even

[23] In our model, the rise in service employment is driven by both a shift towards services within the production of each good, a shift of expenditure toward luxuries. This is in line with the data where both the service cost share and the expenditure share on luxuries increased. Quantitatively, the empirically observed expenditure shift toward luxuries between 2002 and 2017 is stronger than what our model predicts.

Figure 3: Model Fit



Notes: In Panel (a), we show the log GDP per worker in the data (gray line) and the calibrated model (red line). We also depict the respective outcomes if quality was fully measured, i.e., $\mu = 1$, (dotted line) and if quality was not measured at all, $\mu = 0$, (dashed line). In Panel (b), we display the average GDP per worker growth between 1950-2000 and 2000-2019 in the data (gray bar) and the calibrated model (red bar). We also include the outcomes with full quality measurement (light-red bar) and without quality measurement (light-pink bar). Additionally, we include the average GDP growth between 2020-2100 for the three model outcomes. The simulated GDP growth rates use chain-weighted price indices with Tornqvist weights. In Panel (c), we show the service share in the data (gray), corresponding to the GDP share in services (1948–2023), and the calibrated model (red line). In Panel (d), we show the flow of emissions in the data (gray line) and the calibrated model (red line).

though we assume that exogenous green technological progress (z) remains constant over time. This hump-shape arises from a relative decline in material production, Y_G, driven by both a reallocation of labor toward the service sector and a slowdown in quantity productivity growth. Compared to the data, our model suggests that the peak in total emissions occurs somewhat later, implying that the progress of clean technologies has been especially fast in recent decades, arguably due to the sharp reduction in fossil fuel energy production.

Pollution and Service Intensity. The primary mechanism through which our model explains the reduction in pollution is the cleansing effect of service production. In Figure 1, we showed the negative correlation between the cost share of services and pollution intensity at the industry level. In Figure 4, we further explore this relationship by plotting the service cost share against pollution intensity at the product level, comparing our model's predictions with empirical data.

As detailed in Appendix A-1.5, our model makes precise predictions regarding the functional relationship between industrial emissions and service cost share. Specifically, our theory implies that the (log of) emission intensity of industry j—defined as total emissions per dollar of output in industry j—is given by

$$\ln \bar{\mathcal{E}}_j = \delta_t + \ln \sigma_k, = \delta_t + \ln \left(1 \text{ - service share}_k\right), \tag{25}$$

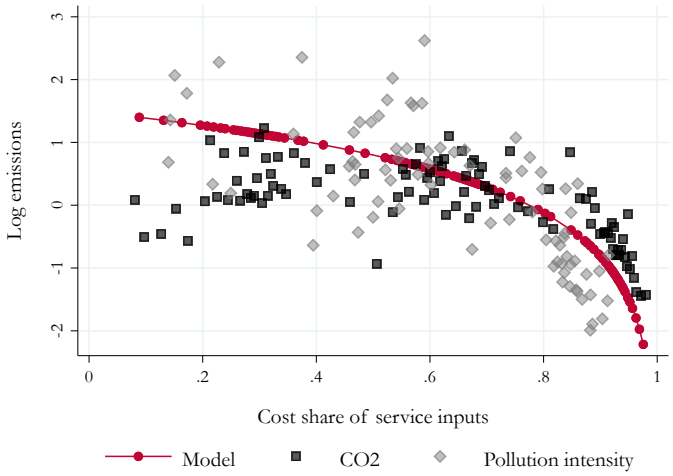
where δ_t is a time-varying effect that depends on the equilibrium price but remains unchanged across industries within a given year. Consequently, the cross-sectional relationship between emission intensity and the service shares of goods should exhibit unitary elasticity.

In Figure 4, we show the relationship between measured CO_2 intensity (black squares) and the industrial service share from the data (the same data presented in Figure 1), superimposing the functional form in (25). We calibrate the scale δ_t so that the CO_2 emissions intensity of the median industry in the model aligns with that of the median industry in the data. Remarkably, the model closely replicates the non-targeted shape of the relationship between service cost share and pollution intensity, supporting our specification of emissions as a linear function of basic goods production.

Figure 4 also presents an alternative measure of industrial pollution intensity from Levinson and O'Brien (2019). This measure, described in more detail in Appendix

A-1, is a composite index of five major air pollutants. The figure shows that these pollutants (grey diamonds) are also strongly negatively correlated with the service share and that our functional form provides an equally strong fit to the empirically observed relationship. Thus, service intensity is not only a key determinant of CO_2 emissions intensity but also a crucial factor in broader environmental damages.

Figure 4: Emissions and Service Share: Model vs. Data



Notes: The figure shows the relationship between the cost share of services and the log of emissions for each industry as predicted by the model (red connected dots) and as observed in the data, using two measures: the log of CO_2 emissions per dollar (black squares) and the log of pollution intensity (gray diamonds), taken from the U.S. Environmental Protection Agency and Levinson and O'Brien (2019), respectively. To construct the emissions predicted by the model, we consider a set J of final goods, where each $j \in J$ corresponds to an industry in the data, indexed by its cost share of services $1 - \sigma_j$ and compute total emissions per dollar according to (25). We chose δ_t in (25) to match the median of the distribution of CO_2 emissions.

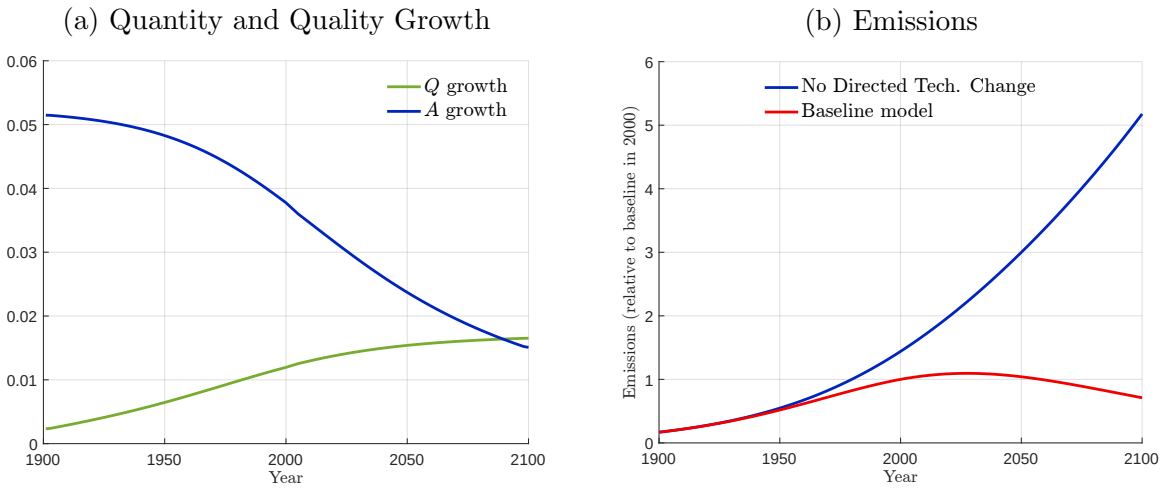
5.4 The Importance of Quality-Led Growth

The central mechanism in our theory is the reallocation of both expenditure and innovation toward product quality. In the right panel of Figure 5, we depict the evolution of quality growth, $\dot{Q}_t/Q_t$, in green, and quantity growth, $\dot{A}_t/A_t$, in blue. Within our model, the increasing share of services reflects a growing expenditure on luxury goods. As a result, the research sector gradually shifts its focus, directing more innovation toward quality rather than quantity.

Quantitatively, our model predicts a substantial realignment of research effort. In 1900, physical productivity grew at nearly 5% per year, while quality growth was

minimal. Over time, research increasingly prioritizes quality improvements. However, the rise in quality growth is smaller than the decline in productivity growth due to our estimate that $\eta_A > \eta_Q$, implying that research efficiency is higher for advancing A than for enhancing quality Q. This greater research efficiency in quantity growth is the primary reason why overall economic growth declines, even if quality growth were fully measured.

Figure 5: The Direction of Technological Change and Clean Growth



Notes: Panel (a) shows the growth rate of Q (green line) and the growth rate of A (blue line). Panel (b) shows the evolution of total emissions in the baseline economy (red line) and in a counterfactual economy without directed technological progress (blue line).

This shift from productivity-led to quality-led growth has significant environmental implications. Emissions primarily originate from the production of necessities. Technical progress enhances the production potential of goods over time. Although this is accompanied by a gradual decline in emissions intensity per unit of production, this decline alone is insufficient to reduce overall emissions. The key mechanism is that the shift in demand and innovation toward quality *slows* the overall quantity of goods produced, ultimately allowing emissions to decrease over time. This dynamic is illustrated in the right panel of Figure 5, which compares emissions under our benchmark model with directed technical change (shown in red) to those under an alternative model with undirected technical change (shown in blue), where the allocation of researchers remains fixed at its 1900 level.[24]

[24] As a consequence, quantity growth $\dot{A}_t/A_t$ and quality growth $\dot{Q}_t/Q_t$ remain constant at approximately 5% and 0.2%, respectively (see left panel of Figure 5).

In the benchmark model, the reallocation of research away from quantity growth leads to a decline in emissions starting in the early 21st century. In contrast, under the alternative model—where the share of research dedicated to quality remains unchanged—sustained productivity growth drives a continuous increase in the quantity of basic goods, resulting in exponential emissions growth. Interestingly, this rise in emissions is driven solely by faster growth in quantity productivity A_t and *not* by a slower reallocation of employment into services (see Appendix A-1.6.) Quantitatively, total CO_2 emissions increased sixfold between 1900 and 2000. In the absence of directed technological change, emissions in the year 2000 would have been 50% higher. By 2100, our baseline model predicts that total emissions will be 20% *lower* than in 2000. In contrast, without a shift toward quality growth, emissions would have been five times larger.

5.5 Welfare and Policy

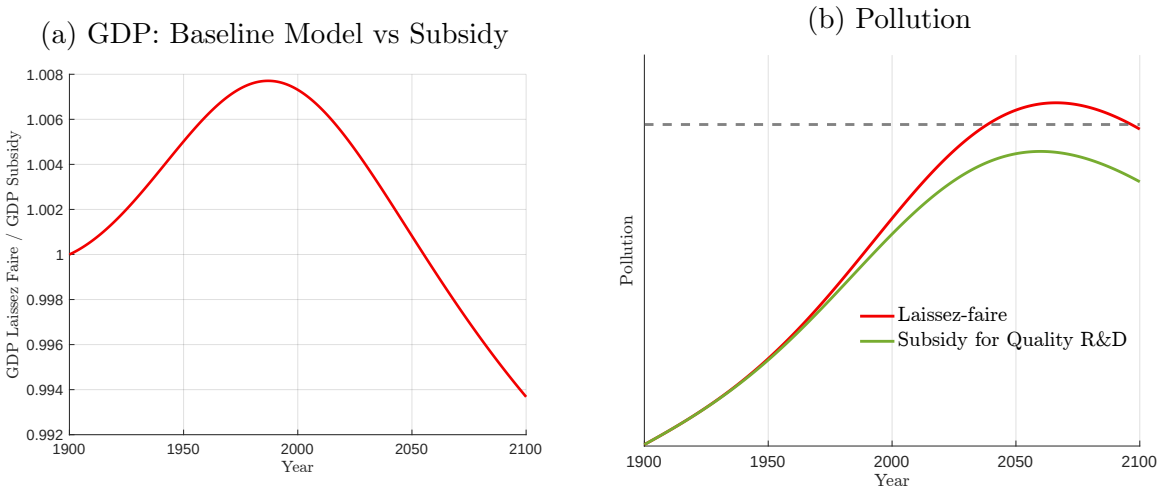
The economy's ability to grow in different ways—either through productivity-led or quality-led growth—implies that research subsidies can play a crucial role in shaping environmental outcomes. To illustrate this, we compare the laissez-faire equilibrium with a counterfactual scenario in which the government subsidizes research directed toward quality by 20% (while financing the subsidy through a lump-sum tax on individuals). In Figure 6, we depict the implications for economic growth (left panel) and environmental outcomes (right panel).

As shown in the left panel, such a policy slightly slows economic growth in the short run, though the quantitative effect is modest. By 2000, GDP per capita would be 1% lower relative to the benchmark economy if quality growth were subsidized. This occurs because research efficiency in quantity growth is relatively high, and expenditure shares on luxury goods remained low in the early 20th century. However, this initial "growth penalty" reverses in the 21st century, when GDP per worker under the subsidy becomes slightly higher than in the baseline case.

The right panel illustrates the environmental benefits of this policy. By shifting research focus toward quality, fewer physical goods are produced, thereby reducing pollution. Since physical goods are the primary source of emissions, reallocating research—and consequently demand and employment—toward quality enables continued economic growth with a slower rise in emissions. Such a policy may have significant

welfare implications, depending on how pollution $\mathcal{P}$ affects utility, $v(\mathcal{P})$. For instance, if emissions surpass a critical threshold leading to an environmental disaster akin to a "tipping point," a quality subsidy could help avert such an outcome. This possibility is illustrated in Figure 6, where the horizontal line indicates the maximum level of emissions the planet can sustain without triggering an environmental breakdown.

Figure 6: Averting an Environmental Disaster: Subsidizing Quality R&D



Notes: Panel (a) shows the GDP in the baseline model relative to the GDP in a counterfactual economy where research directed toward quality innovation is subsidized by $\tau_Q = -20\%$. Panel (b) shows the evolution of total emissions in the baseline economy (red line) and in the counterfactual economy with the subsidy to innovation (green line).

6 International Trade

In this section, we extend our model to an open economy comprising two countries. The main goal is to study the effect of specialization and trade barriers on emissions over the process of structural transformation.

We assume that the two countries produce differentiated manufacturing goods. More formally, we postulate that Y_G is a CES aggregate of domestic and foreign goods:

$$Y_G = \left(\nu^{\frac{1}{\vartheta}} \times \tilde{Y}_G^{\frac{\vartheta-1}{\vartheta}} + (1-\nu)^{\frac{1}{\vartheta}} \times \tilde{Y}_G^{*\,\frac{\vartheta-1}{\vartheta}}\right)^{\frac{\vartheta}{\vartheta-1}}, \text{ with } \vartheta > 1,$$

where $\tilde{Y}_G$ and $\tilde{Y}_G^*$ denote the domestic and foreign manufacturing products, with weight ν given to domestic manufactures.[25] These products are tradable subject to an iceberg

[25] In the rest of this section a star $(*)$ indicates foreign variables while the absence of a star indicates

cost τ. We denote by $\hat{p}_G$ and $\hat{p}_G^*$ the price of the composite good Y_G in the domestic and foreign market, respectively. By contrast, neither the service input, nor final goods are traded, reflecting the fact the service labor has to be provided locally.

We introduce an additional homogeneous tradable good, which is in fixed supply in both countries and can be exchanged against the tradable manufacturing good without any trade cost. We refer to this good as the *endowment*. The sole purpose of such endowment is to allow for the realistic possibility that one economy (e.g., the US) runs a trade deficit in manufacturing goods, i.e., once the economy is open to trade, such an economy exports the endowment and imports the basic good. We denote by E and E^* the supply of the endowment in the two countries and we interpret the endowment in a broad sense as a stand-in for intertemporal trade through capital flows, the export of financial services, royalty payments, or also purchases of domestic real estate by foreign consumers or firms. In our empirical application we will calibrate the relative size of the endowments to match the U.S. trade deficit.

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,

$$\mathcal{V}^{FE}(e, p_N, p_L, p_E) = \frac{1}{\varepsilon}\left(\frac{e}{p_L^{(1-\omega)(1-\varrho)} p_N^{\omega(1-\varrho)} p_E^{\varrho}}\right)^{\varepsilon} + \frac{\phi}{\varsigma}\left(\frac{p_L}{p_N}\right)^{\varsigma(1-\varrho)} - v(\mathcal{P}). \tag{26}$$

Note that the demand for the endowment is homothetic with a unitary price elasticity: consumers spend a constant fraction ϱ of their expenditure 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 labor denoted by H and H^* and research skill R and R^*. Labor is immobile and both quantity productivity A and quality Q are country specific. The state vector is given by $\{A, Q, A^*, Q^*\}$. All parameters are the same in the two economies unless we specify otherwise. Given the symmetric structure, we formally describe only the domestic economy when this is not a source of confusion.

domestic variables.

6.1 Static Equilibrium

Production Costs and Prices. The local production prices of the manufacturing goods in country c are given by $p_G = \frac{\xi}{\xi-1}\frac{w}{A}$. Note that here p_G is defined at the factory and does not include any trade cost.[26] The unit production costs of final goods $j \in \{N, L\}$ is then given by:

$$c_j = \left((1-\lambda_j)\left(\hat{p}_G\right)^{1-\rho} + \lambda_j p_S^{1-\rho}\right)^{\frac{1}{1-\rho}}, \tag{27}$$

where

$$\hat{p}_G = \left(\nu\left(p_G\right)^{1-\theta} + (1-\nu)\left(\tau p_G^*\right)^{1-\theta}\right)^{\frac{1}{1-\theta}}. \tag{28}$$

Note that, absent trade costs, we would have a unique world price $\hat{p}_G = \hat{p}_G^*$.

The prices of manufacturing goods, services, and local prices of final goods incorporate mark-ups like in the closed economy. We assume that mark-ups are identical across the two economies. Substituting in the equilibrium expressions for the prices of manufacturing goods and services yields

$$c\left(p_j, w\right) = \frac{\xi}{\xi-1}\psi_j(A, x)\, w,$$

where we define

$$\psi_j(A, x) = \left((1-\lambda_j)\left(\frac{1}{A}f\left(x\right)\right)^{1-\rho} + \lambda_j\right)^{\frac{1}{1-\rho}}, \tag{29}$$

$$f\left(x\right) \equiv \left(\nu + (1-\nu)\left(x\right)^{1-\theta}\right)^{\frac{1}{1-\theta}}, \tag{30}$$

and

$$x = \tau\pi, \qquad \text{with} \qquad \pi \equiv \frac{w^*/w}{A^*/A}. \tag{31}$$

In plain words, π is the production price of the foreign manufacturing good relative to the domestic good.

[26] Observe that p_G is the production price of the local variety Y_G, while $\hat{p}_G$ is the consumer price of the CES aggregate Y_G in the domestic market.

Expenditure Shares. Under the PIGL preference specification in (26), the expenditure shares of an individual with spending level e are given, respectively, by

$$\begin{aligned}\vartheta_N &= (1-\varrho)\times\left(\omega+\phi\left(\Upsilon\left(A,Q,x,w,e\right)\right)^{-\varepsilon}\right)\\ \vartheta_L &= (1-\varrho)\times\left((1-\omega)-\phi\left(\Upsilon\left(A,Q,x,w,e\right)\right)^{-\varepsilon}\right)\\ \vartheta_E &= \varrho,\end{aligned}$$

where

$$\Upsilon\left(A,Q,x,w,e\right)\equiv\left(\frac{Q^{(1-\beta)}}{\left(\psi_N(A,x)\right)^{\beta}\left(\psi_L(A,x)\right)^{(1-\beta)}}\times\frac{e}{\left(\frac{\xi}{\xi-1}\right)^2 w}\right)^{(1-\varrho)}\left(\frac{e}{p_E}\right)^{\varrho}.$$

As in our baseline model, expenditure shares feature an income effect captured by Υ. This terms depends on nominal income (e), as well as domestic productivity and quality, A and Q, the terms of trade x, the local wage w, and the international price of the endowment p_E. Since p_E is common across countries, we suppress it as an argument of Υ.

Endowment Price. 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. Normalizing domestic labor $H=1$, market clearing for the global supply of the endowment implies that

$$p_E(E+E^*)=\varrho\times(e+e^*H^*). \tag{32}$$

The returns to the endowment are part of total domestic spending, so that $eH = wH+\Pi+p_EE$. In turn, aggregate profits Π can be written as $\Pi=\frac{1}{\xi-1}wH+\frac{1}{\xi}(1-\varrho)eH$, where the first term captures the profits accruing from the sales of intermediate manufacturing and service production, while the second term captures the profit from

the sales of final necessities and luxuries. Standard algebra yields, then:

$$e = \frac{\xi^2}{(\xi-1)^2+\varrho(\xi-1)}\left(w+\frac{\varrho}{1-\varrho}\varpi\left(\frac{\xi}{\xi-1}\right)(w+w^*H^*)\right), \tag{33}$$

$$e^*H^* = \frac{\xi^2}{(\xi-1)^2+\varrho(\xi-1)}\left(w^*H^*+\frac{\varrho}{1-\varrho}(1-\varpi)\left(\frac{\xi}{\xi-1}\right)(w+w^*H^*)\right), \tag{34}$$

where $\varpi \equiv \frac{E}{E+E^*}$ denote the domestic share of the global endowment.

Total spending in each country is determined by domestic and foreign wages as well as the relative supplies of the endowment. It is useful to highlight two particular cases. First, if the endowment has no value, i.e., $\varrho = 0$, expenditure in each country is fully pinned down by local wages. Second, the spending capacity increases in the local endowment. For example, suppose that $\varpi = 1$, i.e., the domestic economy is the only supplier of the endowments. Then, (33)–(34) imply that foreign consumers only earn and spend their labor income, while domestic consumers also earn a rent from the endowment.

Combining these equations with (32) yields the equilibrium price of the endowment:

$$p_E = \left(\frac{\xi}{\xi-1}\right)^2 \frac{\varrho}{1-\varrho}\frac{w+w^*H^*}{E+E^*}. \tag{35}$$

Equations (33),(34), and (35) fully determine total spending (e, e^*) and the endowment price p_E as a function of parameters and the vector of wages (w, w^*).

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

Proposition 3 (Structural Change)**.** *In equilibrium,*

$$\frac{H_S}{H} = \left(\frac{(1-\sigma_N)\vartheta_N}{1-\varrho}+\frac{(1-\sigma_L)\vartheta_L}{1-\varrho}\right)\left(1+\frac{\varrho\xi}{\xi-(1-\varrho)}\left(\frac{\varpi}{\frac{w}{w+w^*H^*}}-1\right)\right) \tag{36}$$

$$\frac{H_G}{H} = 1-\frac{H_S}{H}. \tag{37}$$

where

$$\sigma_j = \frac{(1-\lambda_j)\hat{p}_j^{1-\rho}}{(1-\lambda_j)\hat{p}_G^{1-\rho}+\lambda_j p_S^{1-\rho}} = \frac{(1-\lambda_j)\left(\frac{f(x)}{A}\right)^{1-\rho}}{(1-\lambda_j)\left(\frac{f(x)}{A}\right)^{1-\rho}+\lambda_j}, \tag{38}$$

is the cost share of goods in the production of final items $j \in \{N, L\}$ (hence, $1-\sigma_j$ is the cost share of services in the production of final items j.) Similar expressions hold true for the foreign economy.

Equation (36) highlights the role of the relative abundance of the endowment. When a country is endowment-rich (relative to its relative labor income) its expenditure and employment pattern shift toward services.

6.2 Trade Equilibrium

The previous section characterizes the equilibrium allocation, conditional on the endogenous wages w and w^*. We now break symmetry in the notation by normalizing the domestic wage to unity, i.e., $w = 1$. Consequently, w^* represents the relative wage in the foreign economy compared to the domestic. This relative wage also governs the equilibrium expression of π. In this section, we leverage the trade equilibrium conditions to solve for w^*.

Total domestic manufacturing exports and imports are given, respectively, by

$$\begin{aligned} EX_G &= e^* H^* (1-\chi^*) \left(\vartheta_N^* \sigma_N^* + \vartheta_L^* \sigma_L^*\right), \\ IM_G &= e(1-\chi) \left(\vartheta_N \sigma_N + \vartheta_L \sigma_L\right), \end{aligned}$$

where

$$\chi = \frac{\nu p_G^{1-\theta}}{\nu p_G^{1-\theta} + (1-\nu)\left(\tau p_G^*\right)^{1-\theta}} = \nu \left(f\left(x\right)\right)^{(\theta-1)}. \tag{39}$$

is cost share of domestic manufacturing goods relative to the total cost of manufacturing goods. Note that, conditional on the state vector $\{A, Q, A^*, Q^*\}$ and on the endowments, both EX_G and IM_G are fully determined up to a single endogenous variable, the foreign relative wage w^*.[27]

Market clearing implies that any trade deficit in goods must be paid for by exports of the endowment. More formally,

$$IM_G - EX_G = p_E E - \varrho e. \tag{40}$$

[27] To see why, note that $x = \tau\pi$ and $x^* = \frac{\tau}{\pi}$,where $\pi \equiv \frac{w^*}{A^*/A}$. Thus, conditional on the state vector and normalizations, χ only depends on w^*. Likewise, σ_j and χ can be expressed as functions of w^* using (38) and (39). Next, p_E is fully determined from (35). Moreover, $e = \left(\frac{\xi}{\xi-1}\right)^2 w + p_E E$, which allows us to compute Υ from which ϑ_G and ϑ_S follow (again, as functions of w^*).

Equation (40) is a single equation in a single unknown, w^*. Therefore, it pins down w^* concluding the characterization of the static equilibrium.

6.3 Endogenous Technology

We now determine the equilibrium evolution of the state vector $\{A, Q, A^*, Q^*\}$. The driving forces are the same as in the closed-economy model. In particular, the domestic and foreign country have a mass R and R^* of research skills, respectively, that can be directed to either increase the productivity of the local manufacturing sector or enhance the quality of local luxury goods.

Different from the closed-economy model, the market for local manufacturing goods here depends on *both* the local and foreign demand. More formally:

$$\frac{R_Q}{R_A} = \left(\frac{(1-\tau_Q)\,\eta_Q}{(1-\tau_A)\,\eta_A}\right)^{\frac{1}{\zeta}} \left(\frac{\vartheta_L \times e}{\Theta_G}\right)^{\frac{1}{\zeta}}, \tag{41}$$

$$\frac{R_Q^*}{R_A^*} = \left(\frac{(1-\tau_Q^*)\,\eta_Q^*}{(1-\tau_A^*)\,\eta_A^*}\right)^{\frac{1}{\zeta}} \left(\frac{\vartheta_L^* \times e^* H^*}{\Theta_G^*}\right)^{\frac{1}{\zeta}}, \tag{42}$$

where

$$\begin{aligned}
\Theta_G &\equiv \chi\left(\sigma_N \vartheta_N + \sigma_L \vartheta_L\right) \times e + \left(1-\chi^*\right)\left(\sigma^* \vartheta_N^* + \sigma^* \vartheta_L^*\right) \times e^* H^*, \\
\Theta_G^* &\equiv \chi^*\left(\sigma_N^* \vartheta_N^* + \sigma_L^* \vartheta_L^*\right) \times e^* H^* + \left(1-\chi\right)\left(\sigma_N \vartheta_N + \sigma_L \vartheta_L\right) \times e.
\end{aligned}$$

Note that Θ_G and Θ_G^* comprise both the domestic and foreign demand of the local manufacturing good. In contrast, the market for quality innovation continues to be local.

To highlight the role of the endowment on the direction of innovation in the two economies, note that

$$\begin{aligned}
\frac{\Theta_G}{e} &= \chi\left(\sigma_N \vartheta_N + \sigma_L \vartheta_L\right) + \left(1-\chi^*\right)\left(\sigma^* \vartheta_N^* + \sigma^* \vartheta_L^*\right) \frac{e^* H^*}{e}, \\
\frac{\Theta_G^*}{e^* H^*} &= \chi^*\left(\sigma_N^* \vartheta_N^* + \sigma_L^* \vartheta_L^*\right) + \left(1-\chi\right)\left(\sigma_N \vartheta_N + \sigma_L \vartheta_L\right) \frac{e}{e^* H^*},
\end{aligned}$$

where $\frac{e^* H^*}{e}$ is determined by (33)–(34) plus the normalization $w = 1$.

If the endowment has no value ($\varrho = 0$), then relative foreign demand equals relative

foreign labor income: $\frac{e^*H^*}{e} = \frac{w^*H^*}{w}$. In the general case with $\varrho > 0$, the distribution of the endowment affects relative demands and the direction of technical change in the two economies. In particular, if the domestic economy is relatively endowment-rich, i.e., $\varpi > \frac{w}{w+w^*H^*}$, then $\frac{e^*H^*}{e} < \frac{w^*H^*}{w}$.[28] In this case, the presence of an endowment (which captures the possibility of trade imbalances in goods) strengthens the incentive for the foreign economy, as the net exporter of goods, to direct innovation toward increasing quantity productivity A. Conversely, it strengthens the incentive for the domestic economy to direct innovation toward improving quality. The opposite holds if the foreign economy is relatively endowment-rich.

This prediction lends itself to an application of the theory to the U.S. and Chinese economies in recent years. The U.S. has run persistent trade deficits over time, arguably providing valuable financial services (in the form of safe assets) to China in exchange. To the extent that this deficit is structural, as in Song et al. (2011), our theory predicts that this imbalance has accelerated technical change in the goods-producing sector in China while fostering quality-improving innovation in the U.S.

6.4 Quantitative Analysis

In this section, we study the quantitative implications of the open-economy version of our model. In line with the theory, we consider a two-country model and calibrate the domestic economy to the U.S., with the foreign economy corresponding to China. For clarity, we now specify the two countries by subscripts US and CH.

Our strategy is as follows. We start from the closed-economy calibration between 1900 and 2000. In the year 2000, the U.S. economy then opens up to China, i.e., trade costs fall from a prohibitive level to $\tau < \infty$. We then keep trade costs constant and trace out the transitional dynamics in both the U.S. and China.

To implement this exercise, we require seven additional parameters: (i) the initial level of quality and productivity in China at the time of the trade opening ($A_{2000,CH}$ and $Q_{2000,CH}$), (ii) the U.S. share of the endowment ϖ_{US},[29] (iii) consumers' expenditure share on the endowment ϱ, (iv) the elasticity of substitution of trade goods θ, and (v) the scale and green technology rate of Chinese emissions per unit of Chinese goods

[28] Note, that $\varpi = \frac{w}{w+w^*H^*} \Rightarrow \frac{e^*H^*}{e} = \frac{w^*H^*}{w}$. In this case, the expressions of $\frac{\Theta_G}{e}$ and $\frac{\Theta_G^*}{e^*H^*}$ are the same as in the case in which $\rho = 0$. Standard algebra establishes then the claim in the text.

[29] Since the equilibrium allocation only depends on the relative endowment ϖ_{US}, we normalize the level of the endowment so that $E_{Chn} = 1$.

production, ($\kappa_{\varepsilon,Chn}$ and z_{Chn}).

We calibrate the first three of these parameters to moments from the data. We pick the level of productivity and quality in China in the year 2000, $A_{2000,CH}$ and $Q_{2000,CH}$, to match GDP pc in China relative to the U.S. in 2000 (10.3%) and the employment share of services in China in 2000 (27%).[30] We pick the U.S. share of the global endowment, ϖ_{US}, to match the size of the U.S. trade deficit relative to U.S. 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 3. As a benchmark assumption, we set the parameters of the Chinese emissions function ($\kappa_{\varepsilon,Chn}, z_{Chn}$) to equal those of the U.S., ($\kappa_{\varepsilon,US}, z_{US}$), which we take from our closed-economy simulation. The new parameters and moments are contained in Table IV. We keep all other parameters the same; see Table III.

Table IV: Additional parameters for open-economy simulation

Parameter	Value	Target	Target value
$A_{2000,Chn}$	0.071	CHN/US GDP p.c. (2000)	0.103
$Q_{2000,Chn}$	0.0041	CHN Service emp. share (2000)	0.27
ϖ_{US}	0.773	US Trade deficit rel. to GDP (2019)	0.016
ϱ	0.3	Set exogenously	-
θ	3.0	Set exogenously	-
z_{Chn}	1.017	Set exogenously	-
$\kappa_{\varepsilon,Chn}$	144.919	Set exogenously	-

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

6.4.1 Results

To illustrate the effects of international trade—and conversely, of trade barriers—we plot the path of the U.S. 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 7 we compare the free trade and autarky scenarios for the years 2000–2040. Free trade outcomes are represented by dark solid lines, while autarky corresponds to

[30] 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.

lighter lines of the same color. Panel 7a shows the effect of trade on structural change in the U.S. and Chinese economies. In the U.S., service employment is higher under free trade for two reasons: free trade makes U.S. consumers richer (the income effect) and shifts U.S. production away from goods toward services (the specialization effect).

In China, these two effects push in opposite directions, and the dominant effect varies over time. Free trade increases the income of Chinese consumers (see Panel 7b), pushing them towards demanding a greater share of luxury goods. However, free trade also pushes China to specialize in goods production, so that it can import the endowment good from the U.S. Initially, the income effect dominates, and trade opening induces structural change towards services in China. Over time, technical change in the two countries responds to trade and strengthens the specialization effect, so that by 2013 China has *lower* service employment under the free trade regime than under autarky. This reversal is not a general prediction of our model; under alternative parameter values, it is possible for Chinese service employment to increase or decrease in response to free trade.[31]

In our calibrated model, the short-run impact of free trade is to reduce global emissions through trade-induced structural change (Cravino and Sotelo, 2019). In the longer run, as China begins to specialize in goods production, the overall effect on the environment is ambiguous: U.S. emissions decline throughout, but Chinese emissions are higher under free trade from 2013 onward. On balance, free trade reduces total world emissions for a significant period, as shown in Panel 7c (we revisit the long-run effects below). This result serves as a cautionary note against policy efforts to "reshore" manufacturing in Western economies. Shifting goods production back to the most developed countries may lead to real income declines and greater overall pollution.

Green Technology. Our calibrated model predicts that trade and economic integration reduce global emissions by accelerating the transition to a world of service economies, which are inherently cleaner. Our quantitative analysis relies on the assumption that industrial production is equally polluting (per unit of industrial output) in both the domestic and foreign economies. In the application to China and the United States, this implies that the observed pollution gap per unit of GDP between the two countries is explained by China's specialization in goods production.

[31] In China, the industrial employment share fell from 24% to 21% over the period from 1997 to 2002, before increasing rapidly in subsequent years to reach 32% today.

Interestingly, our baseline calibrated model significantly overpredicts emissions per unit of GDP in China relative to the U.S. Specifically, according to our model, this factor is 2.94, whereas in the data, it is 1.95. Thus, if anything, our stylized model exaggerates China's contribution to global emissions.

Despite this, one might argue that China employs more polluting technologies for the same type of industrial production. This would introduce a mitigating factor to the benefits of international trade. In our model, we can capture such a difference by assuming a larger parameter κ in the foreign economy's pollution technology. We examine the robustness of our results to this modification. In particular, suppose that China uses technologies that are 30% more polluting in the goods-producing sector. Appendix Figure A-3 presents the results. A world economy with trade remains cleaner than one with prohibitive trade barriers, although the difference is significantly smaller. Moreover, the gap narrows over time. By the year 2042, world economies with and without trade reach the same level of emissions.

The underlying mechanism is interesting: innovation in China shifts toward the manufacturing sector due to specialization. However, by assumption, this has no effect on green technology dynamics. Arguably, the outcome would be different if specialization (or other factors) induced some convergence in green technology over time. In reality, China has significantly improved its environmental standards and is committed to a path of declining total emissions after 2030. This trend could mitigate the negative effects of specialization over time—possibly leading to increasing, rather than decreasing, gains from trade. Figure A-3 also illustrates the results of a smoothly declining green technology gap until 2050, after which domestic and foreign economic activities impose the same environmental damage.

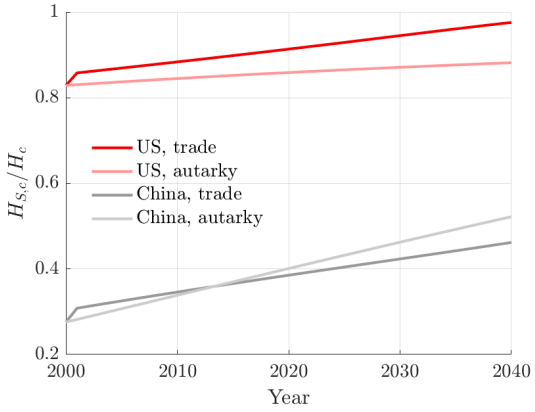
The Long-Run Effect of International Trade. The effect of trade and specialization over time is non-linear. Initially, international trade reduces global emissions, as discussed. This effect is even more pronounced in the long run, as shown in the right panel of Figure A-3.[32] However, there is a significant intermediate period during which emissions are higher in the economy with trade than in the autarkic world. In this phase, the dominant effect is that specialization in innovation enhances efficiency in global goods production. Yet, in the long run, this effect is overtaken by the in-

[32] The results discussed in this paragraph, and more specifically the long-run effects, are independent of the assumption made regarding the relative environmental friendliness of the technologies used in the two countries.

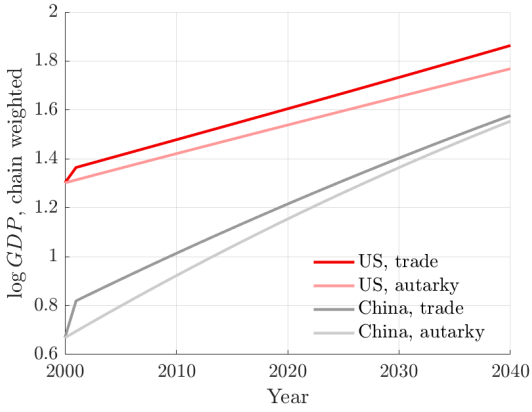
come effect: eventually, both the domestic and foreign economies transition into service economies, directing all innovation efforts toward quality improvements.

Figure 7: Open economy simulations: U.S. and China 2000–2040

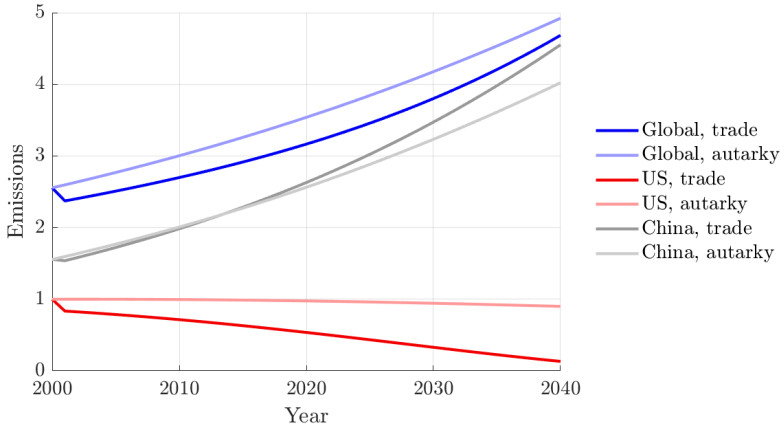
(a) Service share of employment



(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 aggregate emissions (panel c). We always depict the outcomes for the U.S. (China) in blue (red). Emissions are normalized so that U.S. emissions in the year 2000 = 1. The baseline model with international trade is shown with heavy lines, while the closed-economy model is shown with lighter lines. In panel c we also depict total global emissions with blue lines.

7 Conclusion

In this paper, we develop and quantify a growth model where: (i) consumers have non-homothetic preferences between a more basic good (a *necessity*)—for which quality

matters less and whose production is intensive in material inputs—and a "quality" good (a *luxury*) whose production is 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 services has a lower environmental footprint than that of material goods.

The model delivers some novel insights. First, over time as the economy develops, 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 decline or even a stall in measured GDP growth, even though quality-adjusted GDP continues to grow. Third, trade barriers may have a negative effect on global environmental sustainability both in the short and in the long run.

Future research can extend our analysis in several directions. First, it would be useful to consider other factors affecting the environmental footprint of goods of different quality. In the paper, we linked it to service share (which is observable) but this does not capture all the variation in environmental impacts of production across goods. A second extension would be to allow consumers with different income levels to consume different quality versions of the same good.[33]. 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.

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.

[33] For instance, one could consider meals in restaurants as a category of final goods, where poorer consumers purchase services from cheaper restaurants, while wealthier consumers opt for gourmet restaurants. Our theory captures this distinction in a highly stylized manner, but there may be benefits in explicitly incorporating a distinction between non-homotheticity across and within categories of goods. This would require a nontrivial extension of the theory along the lines of Foellmi et al. (2014).

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, T. Boppart, P. J. Klenow, and H. Li: 2019, 'Missing growth from creative destruction'. *American Economic Review* **109**(8), 2795–2822.

Aghion, P., A. Bergeaud, M. De Ridder, and J. Van Reenen: 2024, 'Lost in transition: Financial barriers to green growth'. INSEAD Working Paper No. 2024/16/EPS.

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.

Alder, S., T. Boppart, and A. Müller: 2022, 'A theory of structural change that can fit the data'. *American Economic Journal: Macroeconomics* **14**(2), 160–206.

Baumol, W.: 1967, 'Macroeconomics of Unbalanced Growth: The Anatomy of the Urban Crisis'. *American Economic Review* **57**(3), 415–426.

Beerli, A., F. Weiss, F. Zilibotti, and J. Zweimüller: 2020, 'Demand forces of technical change: Evidence from the Chinese manufacturing industry'. *China Economic Review* **60**, 101157.

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.

Comin, D., D. Lashkari, and M. Mestieri: 2021, 'Structural Change with Long-Run Income and Price Effects'. *Econometrica* **89**(1), 311–374.

Cravino, J. and S. Sotelo: 2019, 'Trade-induced structural change and the skill premium'. *American Economic Journal: Macroeconomics* **11**(3), 289–326.

D'Alessandro, S., A. Cieplinski, T. Distefano, and K. Dittmer: 2020, 'Feasible alternatives to green growth'. *Nature Sustainability* **3**(4), 329–335.

Daly, H. E.: 1977, *Steady-State Economics: The Economics of Biophysical Equilibrium and Moral Growth*. San Francisco: W.H. Freeman.

De Ridder, M. and R. Lukasz: 2025, 'Emissions-Adjusted Total Factor Productivity'. Mimeo, London School of Economics.

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.

Epifani, P. and G. Gancia: 2011, 'Trade, markup heterogeneity and misallocations'. *Journal of International Economics* **83**, 1–13.

Fajgelbaum, P. D. and G. M. Grossman: 2011, 'Income Distribution, Product Quality, and Interna-

tional Trade'. *Journal of Political Economy* **119**(4), 721–765.

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

Fieler, A. C.: 2011, 'Nonhomotheticity and Bilateral Trade: Evidence and a Quantitative Explanation'. *Econometrica* **79**(4), 1069–1101.

Foellmi, R., T. Wuergler, and J. Zweimüller: 2014, 'The macroeconomics of Model T'. *Journal of Economic Theory* **153**, 617–647.

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.

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.: 2020, *Less Is More: How Degrowth Will Save the World*. Penguin Random House.

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.

Kallis, G.: 2011, 'In defence of degrowth'. *Ecological Economics* **70**(5), 873–880.

Kallis, G., C. Kerschner, and J. Martinez-Alier: 2012, 'The economics of degrowth'. *Ecological Economics* **84**, 172–180.

Latouche, S.: 2009, *Farewell to Growth*. Polity Press.

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.

Meadows, D. H., D. L. Meadows, J. Randers, and W. W. B. III: 1972, *The Limits to Growth*. New York: Universe Books.

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, *Managing the Global Commons: The Economics of Climate Change*. Cambridge, MA: MIT Press.

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'. Centre for Economic Performance, London School of Economics and Political, Discussion Paper no. 417.

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.

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.

Song, Z., K. Storesletten, and F. Zilibotti: 2011, 'Growing Like China'. *American Economic Review* **101**(1), 196–233.

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

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.

Additionally, we include data on CO_2 emissions per dollar for each industry from the Supply Chain Greenhouse Gas dataset, also released by the U.S. Environmental Protection Agency. This data includes CO_2-equivalent (CO_2e) emissions in kilograms per dollar for all greenhouse gases, combining direct emissions from input acquisition and production processes with those from distribution and storage at the industry level.

A-1.2 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

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

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 $\hat{S} := S'\hat{T}_{Badj}$.

A-1.3 Classification of Products: Luxuries versus Necessities

In our model, the parameters λ_k directly map 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: luxuries and necessities. We therefore aggregate the data to reflect this distinction. Specifically, we classify all final goods into two mutually exclusive groups: service-intensive ("luxuries") products and goods-intensive ("necessities") products. To do this, we rank final products by their service cost share (which maps one-to-one into the parameter λ_k) categorizing those with a service share above the median, weighted by expenditure, as service-intensive.

Empirically, we account for 218 products classified as "necessities" and 189 as "luxuries", where the mean cost share of services among luxuries is 0.93 and among necessities is 0.41; see Table A-III. In the model, we calibrate λ_L and λ_N to match these observed moments.

Table A-III: Cost Share in Services for Necessities and Luxuries

	Number of products	Mean
Necessities	218	0.409
Luxuries	189	0.935

Notes: The table reports the number of products in the CEX classified as Necessities and Luxuries for 2002, as well as their cost share in services, weighted by total expenditure in each product.

A-1.4 Nonhomothetic Demand

Final Goods

Value Added Content In our main analysis we focus on individuals' expenditure shares of final goods. This allows us to directly estimate the structural parameter ε that governs the consumers' demand system. Alternatively, we can also directly compute individuals' demand for services that are embodied in different final goods. 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

$$\vartheta_{\mathcal{S}}^i = \sum_k \vartheta_k^i s_k. \tag{A-1}$$

In the left panel of Figure A-1 we depict the cross-sectional distribution of $\vartheta_{\mathcal{S}}^i$. In the right panel, we show that this heterogeneity is strongly related to household income using household's income rank, which is directly reported in the CEX.

In Table A-IV we document this positive correlation between individual income and the service content of consumption in a regression format:

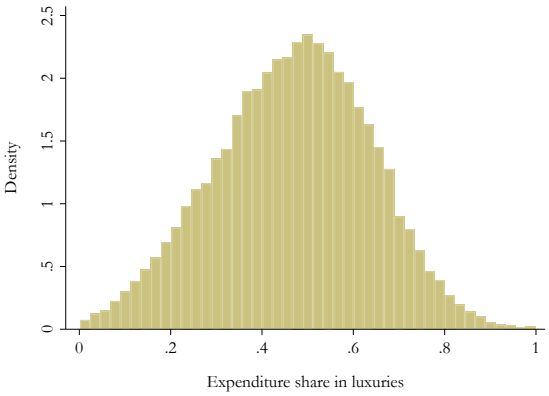
$$\vartheta_{\mathcal{S}}^i = \gamma \ln e_i + x_i'\psi + u_i, \tag{A-2}$$

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 parallel to Table II, 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.

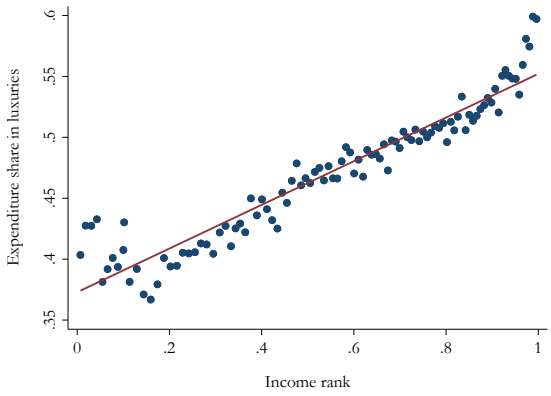
In column 1, we report the simple bivariate correlation, controlling only for household size. In column 2 we add controls for geographic location, education, race, and marital status. In column 3 we report the IV estimate, where, as we did for Table II, we instrument household spending with a full set of occupation fixed effects.

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 (ϑ_L^i). In the right panel, we display a binscatter plot between the expenditure share on service-intensive goods and the income rank of the household.

Table A-IV: Nonhomothetic Service Demand: Estimating ε

	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 $\ln(1-\beta-\vartheta_S^i)$, 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.

A-1.5 Construction of Figure 4

In this section, we describe the details of how we constructed Figure 4. Let σ_k denote the cost share of goods for product k. Because final goods are priced at a constant markup $\frac{\xi}{\xi-1}$, overall spending on goods relative to sales when producing x_k units of product k is given by

$$\frac{p_G Y_G(x_k)}{p_k x_k} = \frac{\xi - 1}{\xi} \times \sigma_k$$

Total emissions per dollar of revenue of product k are thus given by

$$\bar{\mathcal{E}}_{kt} = \kappa_{\mathcal{E}} Y_G(x_k) z^{-t} = \kappa_{\mathcal{E}} z^{-t} \frac{\xi - 1}{\xi} \times \sigma_k p_G^{-1}.$$

The log of the emission intensity of product k is therefore given by

$$\ln \bar{\mathcal{E}}_{kt} = \ln \left(\kappa_{\mathcal{E}} z^{-t} \frac{\xi - 1}{\xi} \times p_G^{-1} \right) + \ln \sigma_k \equiv M_t + \ln(1 \text{ - service share}_k),$$

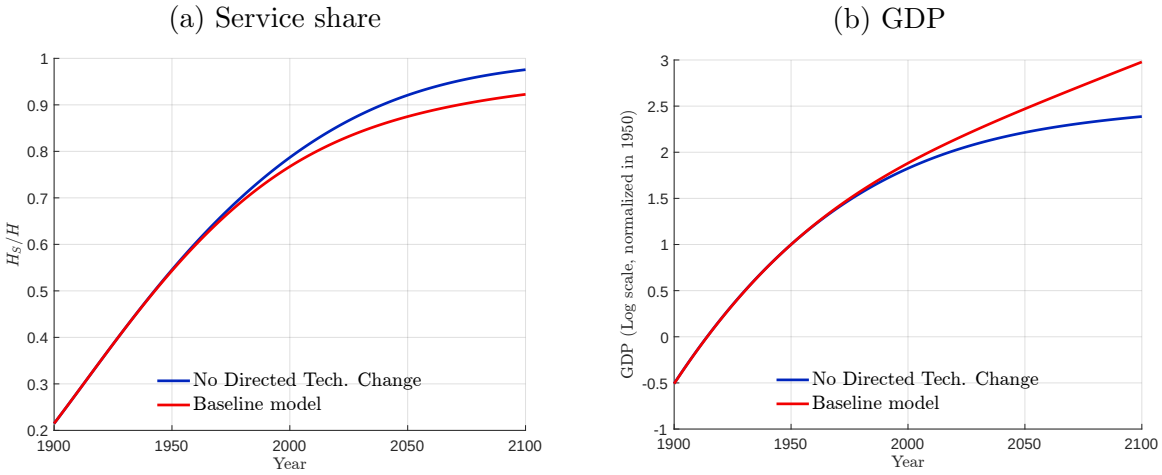
where M_t is an aggregate variable that does not depend on k. Hence, our theory predicts a log-linear structure between emission intensity and the service share at the product level.

A-1.6 Additional Quantitative Results

The Importance of Quality-Led Growth (Section 5.4) In Section 5.4 we analyzed the importance of directed technological change by comparing the predictions of our baseline model with a counterfactual economy, where the allocation of researchers is fixed. In Figure 5 we showed that this would lead to substantially higher emissions. Figure A-2 reports the implications for the service employment share (left panel) and GDP per worker (right panel). Two features are apparent. First, faster quality growth is not a prerogative for services to increase. The service employment share also increases in the absence of directed technical change — if anything the increase is slightly faster. The reason is that faster quantity growth is also a source of rising incomes and that it induces technological substitution (the "Baumol" effect). In the right figure we show that GDP per worker growth is slightly faster when technological change is directed because the direction of innovation responds to changes in the size of the market and hence in the appropriate price index of GDP.

Alternative Assumptions About Chinese Emissions As discussed in the main text, our benchmark specification sets the emissions intensity of Chinese industrial production to equal that of the US. In Figure A-3, we show that our qualitative result that trade reduces global emissions in the short to medium run does not hinge on this assumption. To illustrate that this result is robust, we focus on the relative

Figure A-2: The Impact of Directed Technical Change

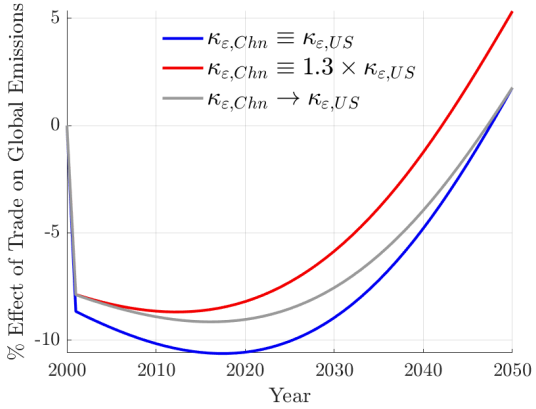


Notes: Panel (a) shows the employment share in services in the baseline model (red line) and in a counterfactual economy without directed technological change (blue line). Panel (b) shows the GDP per worker in in the baseline economy (red line) and in a counterfactual economy without directed technological change (blue line).

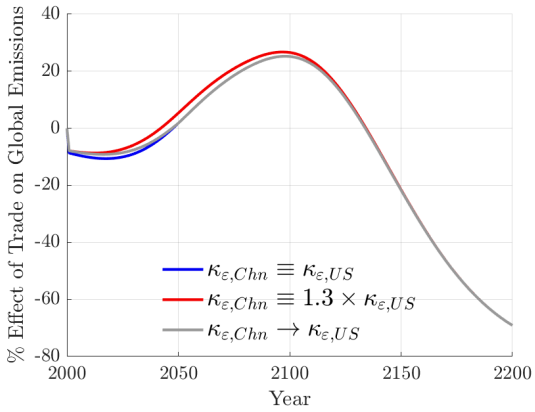
impact of trade on global emissions. In blue, we plot trade's effect in our benchmark specification. In red, we plot a more pessimistic specification where the same physical production generates 30% more emissions when it occurs in China compared to the U.S. In grey, we plot an intermediate specification where Chinese production starts out 30% dirtier than the U.S., but converges to the US's green technology by 2050. In all such specifications, trade reduces global emissions for the first few decades of the 21st century due to trade-induced structural change. However, as China continues to specialize in manufacturing and deepens its physical productive capacity, its relatively dirtier technology begins to take an environmental toll. If Chinese manufacturing remains 30% dirtier than American manufacturing, the emissions benefits of free trade are eliminated by 2042, while this reversal occurs later in more optimistic specifications. Figure A-3b shows the same emissions impact over a longer time horizon, from 2000 to 2200. In the very long run, the income effect of trade once again dominates the specialization effect starting in the 2130s, and emissions significantly decline relative to autarky.

Figure A-3: Robustness: Impact of Trade on Emissions

(a) In the Medium Run



(b) In the Long Run



Notes: The figure displays the change in global emissions under a free trade regime relative to autarky, under different assumptions about the emissions intensity of Chinese industrial production, $\kappa_{\varepsilon,Chn}$. The blue line corresponds to the parameterization in the text, where China's emissions intensity is equal to that of the US. The red line is an alternative specification where Chinese production is 30% more emissions-intensive than the US. The grey line is an intermediate specification where China starts out with 30% more polluting technology, but progressively closes the "green technology gap" from 2000 to 2050. In all specifications, the number plotted is global emissions under free trade minus global emissions under autarky, divided by global emissions under autarky. The two panels plot the same series over different time horizons.

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 our theoretical results.

B-1.1 Consumer Preferences and Expenditure Shares

Consider the indirect utility function

$$\mathcal{V}\left(e,[\tilde{p}_j]_{j=1}^J,[Q_j]_{j=1}^J\right)=\frac{1}{\varepsilon}\left(e\prod_{j=1}^J\left(\frac{Q_j^{\alpha_j}}{\tilde{p}_j}\right)^{\omega_j}\right)^{\varepsilon}-\tilde{\varsigma}\prod_{j=1}^J\left(\tilde{p}_j/Q_j^{\alpha_j}\right)^{\varsigma_j}-v(\mathcal{P})$$

The expenditure share for good k follows from Roy's identity and is given by

$$\begin{aligned}\vartheta_k\left(e,[p_j]_{j=1}^J\right) &= -\frac{\partial\mathcal{V}^{FE}/\partial p_k}{\partial\mathcal{V}^{FE}/\partial e}\frac{p_k}{e}\\ &= -\frac{-\omega_k p_k^{-1}\left(e\prod_{j=1}^J\left(\frac{Q_j^{\alpha_j}}{\tilde{p}_j}\right)^{\omega_j}\right)^{\varepsilon}-\tilde{\varsigma}\varsigma_k p_k^{-1}\prod_{j=1}^J\left(\tilde{p}_j/Q_j^{\alpha_j}\right)^{\varsigma_j}}{e^{-1}\left(e\prod_{j=1}^J\left(\frac{Q_j^{\alpha_j}}{\tilde{p}_j}\right)^{\omega_j}\right)^{\varepsilon}}\frac{p_k}{e}\\ &= \omega_k+\tilde{\varsigma}\varsigma_k\frac{\prod_{j=1}^J\left(\tilde{p}_j/Q_j^{\alpha_j}\right)^{\varsigma_j}}{\left(e\prod_{j=1}^J\left(\frac{Q_j^{\alpha_j}}{\tilde{p}_j}\right)^{\omega_j}\right)^{\varepsilon}}\\ &= \omega_k+\tilde{\varsigma}\varsigma_k\frac{\prod_{j=1}^J p_j^{\varsigma_j}}{\left(e\prod_{j=1}^J\left(p_j\right)^{-\omega_j}\right)^{\varepsilon}}\end{aligned}$$

where the last line uses the expression for hedonic prices $p_j=\tilde{p}_j/Q_j^{\alpha_j}$. Noting that

$$\frac{\prod_{j=1}^J p_j^{\varsigma_j}}{\left(e\prod_{j=1}^J\left(p_j\right)^{-\omega_j}\right)^{\varepsilon}}=\left(\frac{e}{\prod_{j=1}^J p_j^{\varsigma_j/\varepsilon+\omega_j}}\right)^{-\varepsilon}=\left(\frac{e}{\prod_{j=1}^J p_j^{\beta_j}}\right)^{-\varepsilon},$$

where $\beta_j\equiv\varsigma_j/\varepsilon+\omega_j$ yields

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

where $\phi_k = \tilde{\varsigma}\varsigma_k$. This is the expression in (2).

B-1.2 Profits and Prices

In this section, we derive prices and profits earned by firms producing varieties of final goods and intermediate inputs.

Final Goods. Consider the monopolist firm producing variety i in the final good sector $j \in \{1, 2, ...J\}$. Given the production function

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

let p_G and p_S the prices of goods and services. The marginal cost of producing good j are therefore given by

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

Moreover, the cost share of goods, $\sigma_j(p_G, p_S)$ is given by

$$\sigma_j(p_G, p_S) = \frac{(1-\lambda_j) p_G^{1-\rho}}{(1-\lambda_j) p_G^{1-\rho} + \lambda_j p_S^{1-\rho}}. \tag{B-2}$$

Below we show that $p_G = \mu \frac{w}{A}$ and $p_S = \mu w$, where $\mu = \frac{\xi}{\xi - 1}$ is the markup that intermediate producers charge. Using the fact that we take wages as numeraire, that is $w = 1$, we can also express the marginal costs c_j and the costs share of goods, σ_j, as

$$c_j(A) = \mu \, \psi_j(A), \quad \text{where} \quad \psi_j(A) = \left((1-\lambda_j) A^{\rho-1} + \lambda_j \right)^{\frac{1}{1-\rho}}, \tag{B-3}$$

and

$$\sigma_j(A) = \frac{(1-\lambda_j) A^{\rho-1}}{(1-\lambda_j) A^{\rho-1} + \lambda_j}. \tag{B-4}$$

Profit maximization implies that the monopolist producing variety i of product j will set the price

$$\tilde{p}_{ji} = \mu \, c_j(w, A) = \tilde{p}_j.$$

where again $\mu = \frac{\xi}{\xi - 1}$. The (quality-adjusted) price index for good j is therefore given by

$$p_j = \left(\int_0^1 \left(\frac{\tilde{p}_{ji}}{Q_{ij}^{\alpha_j}} \right)^{1-\xi} di \right)^{\frac{1}{1-\xi}} = \mu c_j(A) \left(\int_0^1 Q_{ij}^{\alpha_j(\xi-1)} di \right)^{\frac{1}{1-\xi}} \equiv \frac{1}{Q_j^{\alpha_j}} \mu c_j(A),$$

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

Standard properties of the isoelastic demand for varieties implies that the profits accruing to the monopolist firm are then given by

$$\pi_{ij} = p_j^{\xi} Y_j Q_{ij}^{\alpha_j(\xi-1)} c_j (w, A)^{1-\xi} \frac{(\xi-1)^{\xi-1}}{\xi^{\xi}} = \frac{1}{\xi} \left(\frac{Q_{ij}}{Q_j} \right)^{\alpha_j(\xi-1)} p_j Y_j, \quad \text{(B-5)}$$

where $p_j Y_j$ is aggregate spending on good j.

The aggregate profits generated by firms operating in final good sector $j \in \{1, 2, ...J\}$ are therefore equal to

$$\Pi_j = \int_0^1 \pi_{ij} \, di = \frac{1}{\xi} p_j Y_j. \quad \text{(B-6)}$$

Input Sectors. Denote by p_{ik} and π_{ik}, respectively, the price and profits associated with the monopolist firm producing the input variety i in sector $k \in \{G, S\}$. Profit maximization implies that

$$p_{iG} = \mu \, \frac{w}{A_i} \quad \text{and} \quad p_{iS} = \mu \, w.$$

The aggregate price indices in sector k are therefore given by $p_S = \mu \, w$ and $p_G = \mu \, \frac{w}{A}$, where

$$A \equiv \left(\int_0^1 A_i^{\xi-1} di \right)^{\frac{1}{\xi-1}}.$$

Let $\mathcal{D}_G$ and $\mathcal{D}_S$ denote total spending on goods and services respectively. Overall profits are thus given by

$$\Pi_G = \int_i \pi_{iG} di = \int_i \frac{1}{\xi} \mathcal{D}_G \left(\frac{A_i}{A} \right)^{\xi-1} di = \frac{1}{\xi} \mathcal{D}_G \quad \text{(B-7)}$$

$$\Pi_S = \int_i \pi_{iS} di = \int_i \frac{1}{\xi} \mathcal{D}_S di = \frac{1}{\xi} \mathcal{D}_S. \quad \text{(B-8)}$$

To solve for aggregate demand $\mathcal{D}_k$, note that

$$\mathcal{D}_G \equiv \sum_{j=1}^{J} \sigma_j \, \frac{\xi-1}{\xi} \vartheta_j e H \quad \text{(B-9)}$$

$$\mathcal{D}_S \equiv \sum_{j=1}^{J} (1-\sigma_j) \, \frac{\xi-1}{\xi} \vartheta_j e H. \quad \text{(B-10)}$$

Here, $\vartheta_j e H$ is total spending on product j, a fraction $\frac{\xi-1}{\xi}$ get paid to variable factors G and S, and σ_j is the cost share of goods for product j. Summing over all final goods

j thus yields total spending on goods, $\mathcal{D}_G$. The intuition for $\mathcal{D}_S$ is similar, except that a share $1-\sigma_j$ of spending on goods j goes to services.

B-1.3 Income, Spending and Profits

The representative household is the recipient of both labor income and overall profits.

Overall profits are given by

$$\Pi = \Pi_G + \Pi_S + \sum_j \Pi_j,$$

where Π_G and Π_S are the profits of intermediate producers and Π_j are the profits of all final producers that produce good j. These profits are the returns to researchers, both the innovative types and parasitic types.

Using (B-6), (B-7), and (B-8), it follows that

$$\Pi = \frac{1}{\xi}\left(\sum_j p_j Y_j + \mathcal{D}_G + \mathcal{D}_S\right) = \frac{1}{\xi}\left(eH + \frac{\xi-1}{\xi}eH\right) = \frac{1}{\xi}\left(1+\frac{\xi-1}{\xi}\right)eH \quad \text{(B-11)}$$

Labor income it given by

$$wH = \frac{\xi-1}{\xi}\left(\mathcal{D}_G + \mathcal{D}_S\right) = \left(\frac{\xi-1}{\xi}\right)^2 eH, \quad \text{(B-12)}$$

reflecting the fact that labor receives a share $\frac{\xi-1}{\xi}$ of intermediate spending, which in turn is a share $\frac{\xi-1}{\xi}$ of total final good spending ("double marginalization"). This also implies that overall profits are proportional to wage income, $\Pi = \frac{2\xi-1}{\xi-1}wH$.

Hence, as required total income is equal to total spending

$$\text{Income} = wH + \Pi = \left(\frac{\xi-1}{\xi}\right)^2 eH + \frac{1}{\xi}\left(1+\frac{\xi-1}{\xi}\right)eH = eH. \quad \text{(B-13)}$$

Note also that (B-12) implies that

$$e = \left(\frac{\xi}{\xi-1}\right)^2 w \quad \text{(B-14)}$$

and hence $e = \left(\frac{\xi}{\xi-1}\right)^2$ if w is taken to be the numeraire.

B-1.4 Labor Market Clearing

Now, consider the demand for labor. Labor market clearing for manufacturing workers requires that (see (B-9))

$$wH_G = \frac{\xi-1}{\xi}\mathcal{D}_G = \left(\frac{\xi-1}{\xi}\right)^2 \sum_j \sigma_j \vartheta_j = \left(\frac{\xi-1}{\xi}\right)^2 eH \sum_j \sigma_j \vartheta_j, \tag{B-15}$$

because workers receive a share $\frac{\xi-1}{\xi}$ of overall spending on goods. Using (B-14), this implies that

$$\frac{H_G}{H} = \sum_j \sigma_j \vartheta_j. \tag{B-16}$$

In a similar fashion we have $\frac{H_S}{H} = \sum_j (1-\sigma_j)\vartheta_j$.

B-1.5 Real Income $\Upsilon(A,Q)$

In (B-1) we have shows that expenditure shares can be written as

$$\vartheta_k\left(e,[p_j]_{j=1}^J\right) = \omega_k + \phi_k \Upsilon(e,p_N,p_L)^{-\varepsilon} \quad \text{where} \quad \Upsilon(e,p_N,p_L) = \frac{e}{\prod_{j=1}^J p_j^{\beta_j}}.$$

We now express Υ directly as a function of the state variables A and Q. Using the expression for prices p_j in (B-5) and the expression for marginal costs c_j in (B-3), it follows that

$$\prod_{j=1}^J p_j^{\beta_j} = \mu \prod_{j=1}^J Q_j^{-\alpha_j\beta_j} c_j(A)^{\beta_j} = \mu^2 \prod_{j=1}^J Q_j^{-\alpha_j\beta_j} \psi_j(A)^{\beta_j}.$$

Noting that $\frac{e}{\mu^2} = e\left(\frac{\xi-1}{\xi}\right)^2 = w = 1$, this implies that

$$\Upsilon(A,Q) = \frac{1}{\prod_{j=1}^J Q_j^{-\alpha_j\beta_j}\psi_j(A)^{\beta_j}} = \frac{\prod_{j=1}^J Q_j^{\alpha_j\beta_j}}{\prod_{j=1}^J \psi_j(A)^{\beta_j}}.$$

For the case of two goods and $\alpha_N = 0 < \alpha_L = 1$, this expression reduces to

$$\Upsilon(A,Q) = \frac{Q_L^{1-\beta}}{\psi_L(A)^{\beta}\psi_N(A)^{1-\beta}}.$$

B-1.6 The Optimal Allocation of Research

There are a fixed number of researchers, R, who can direct their research efforts to improve the productivity to produce manufacturing goods (A_i) or improve the quality of the good j, (Q_{ij}). The value of directing research towards improving quality in good j is given by

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

where $\pi_{ij}\left(\gamma_Q Q_{ij}\right)$ denotes the profits of providing variety i for good j at quality $\gamma_Q Q_{ij}$. Using the expression for equilibrium profits in (B-5), we get that

$$\int \pi_{ij}\left(\gamma_Q Q_i\right) di \quad = \quad \frac{1}{\xi} p_j Y_j \gamma_Q^{\alpha_j(\xi-1)}.$$

Hence,

$$V_{Qj} = (1-\tau_Q)\left(\eta_Q R_Q^{-\zeta}\right)\frac{1}{\xi} p_j Y_j \gamma_Q^{\alpha_j(\xi-1)}. \tag{B-17}$$

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

$$V_A = (1-\tau_A)\left(\eta_A R_A^{-\zeta}\right) \times \int \pi_{iG}\left(\gamma_A A_i\right) di = (1-\tau_A)\left(\eta_A R_A^{-\zeta}\right)\frac{1}{\xi}\gamma_A^{\xi-1}\mathcal{D}_G, \tag{B-18}$$

where the last equation uses (B-7). Free entry into innovation implies that $V_A = V_{Qj}$ for all j.

For the case of two goods ($j \in (L,N)$) and $\alpha_L = 0 < \alpha_N = 1$, (B-17) and (B-18) imply that

$$\frac{R_Q}{R_A} \quad = \quad \left(\frac{1-\tau_Q}{1-\tau_A}\frac{\eta_Q}{\eta_A}\left(\frac{\gamma_Q}{\gamma_A}\right)^{\xi-1}\frac{p_L Y_L}{\mathcal{D}_G}\right)^{1/\zeta} \tag{B-19}$$

$$\quad = \quad \left(\frac{1-\tau_Q}{1-\tau_A}\frac{\eta_Q}{\eta_A}\left(\frac{\gamma_Q}{\gamma_A}\right)^{\xi-1}\frac{\xi}{\xi-1}\frac{\vartheta_L}{\sigma_N\vartheta_N+\sigma_L\vartheta_L}\right)^{1/\zeta}, \tag{B-20}$$

where the last equality uses that

$$\frac{p_L Y_L}{\mathcal{D}_G} = \frac{\vartheta_L eH}{\sum_{j\in N,L}\sigma_j\vartheta_j\frac{\xi-1}{\xi}eH} = \frac{\xi}{\xi-1}\frac{\vartheta_L}{\sigma_N\vartheta_N+\sigma_L\vartheta_L}.$$