A Theory of Endogenous Degrowth and Environmental Sustainability*

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

Abstract

ABSTRACT: We develop and quantify a novel growth theory in which economic activity endogenously shifts from material production to quality improvements. Consumers derive utility from goods with differing environmental footprints: necessities are material-intensive and polluting, while luxuries are more service-based and emit less. Innovation can be directed toward either material productivity or product quality. Because demand for luxuries is more sensitive to quality, the economy gradually becomes "weightless": growth is driven by quality improvements, services become the dominant employment sector, and material production stabilizes at a finite level. This structural transformation enables rising living standards with declining environmental intensity, providing an endogenous path to degrowth in material output without compromising economic progress. Policy can accelerate the transition, but its burden is uneven, falling more heavily on the poor than on the rich.

*We thank Zhang Chen, Maarten De Ridder, David Hemous, Peter Klenow, Morten Olsen, Kjetil Storesletten, Chris Tonetti, and audiences at the Conference on the Economics of Innovation in Memory of Zvi Griliches, the SED, and seminars at the Chinese University of Hong Kong, Danmarks Nationalbank, EIEF, London Business School, LSE, Swiss National Bank, University of Copenhagen, University of Lausanne, University of Oslo, University of St. Gallen, the Geneva Graduate Institute, and University of Zurich for helpful comments. We are grateful to Lorenzo Caliendo and Matthew Murillo for assistance with the GTAP data. Nicolas Fajardo and Youdan Zhang provided outstanding research assistance. Aghion thanks the SCOR Foundation for Science for financial support; Peters and Zilibotti thank the Tobin Center for financial support. For generous hospitality, Zilibotti thanks the University of Copenhagen and Danmarks Nationalbank, while Peters thanks the University of Zurich.

†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 and University of Michigan. schwam@umich.edu

‖Yale University, CEPR, and NBER. fabrizio.zilibotti@yale.edu

1 Introduction

Climate change and environmental degradation are critical challenges of our time. There is broad consensus that economic activity is a major contributor to their severity. In response, many scholars, activists, and policymakers have advocated a slowdown in economic growth (*degrowth*) to achieve carbon neutrality and halt global warming.[1] Mainstream economists have remained skeptical, emphasizing that degrowth would impose large costs on billions, especially in developing countries, and would likely provoke strong opposition and political setbacks. They argue that green innovation offers the best route to addressing climate change without sacrificing prosperity (Acemoglu et al., 2012b).

In this debate, degrowth proponents have also called for a transition from quantity to quality. We take this argument seriously and show that such a transition can contribute to addressing the climate crisis. To that end, we link the notion of quality to the service intensity of different consumption categories. We argue that *weightless* economies (Quah, 1999)—where growth is driven by services and immaterial production—can sustain rising living standards without imposing a heavy environmental toll.

The key mechanism in our theory is that, as incomes rise, economic activity shifts from material production toward services. Aggregate demand endogenously moves from goods toward service- and quality-intensive consumption, while production increasingly relies on service inputs due to Baumol's cost disease. This dual shift gradually redirects R&D and innovation away from material productivity and toward quality enhancement.

In standard models, the distinction between quantity and quality is largely semantic. In our framework, by contrast, it carries important implications. Because material production is the main source of pollution, whether innovation improves productivity or quality has significant consequences for environmental outcomes. Cost-reducing innovation increases material output, which—even with cleaner technologies—tends to raise emissions, while quality-enhancing innovation does not. For example, the environmental footprint of an iPhone 16 is comparable to that of an iPhone 3, just as a gourmet restaurant has a similar footprint as a fast-food outlet.

More formally, we develop a growth theory in which the distinction between quantity and quality takes center stage, and where the direction of technological progress—toward

[1] The intellectual roots of the degrowth movement trace back to the 1970s, as discussed in the literature review below. In recent years, several political leaders have advocated moving beyond GDP growth as the central policy objective. Under President Sarkozy, France commissioned the Stiglitz–Sen–Fitoussi report to explore broader measures of economic performance. More recently, President Macron has called for rethinking capitalism in light of planetary boundaries.

either raising material productivity or enhancing quality—is endogenous. Consumer preferences are defined over a range of final goods that differ in production technology and the quality premium they command.

We build on three core assumptions, each supported by empirical evidence. First, goods are ranked along a *sophistication ladder*, with higher sophistication associated with greater service intensity and a stronger role for quality. For example, food prepared at home relies mainly on physical ingredients and is less sensitive to quality, whereas high-end dining involves extensive service provision—chefs, professional waitstaff, ambiance—and commands a higher quality premium. Second, preferences are nonhomothetic: basic goods are necessities, while sophisticated, service-intensive goods are luxuries. This pattern aligns with evidence from Bils and Klenow (2001) and our own findings in Section 2. Third, sophisticated goods generate fewer emissions per unit of expenditure than basic ones. As societies grow wealthier, consumers allocate a larger share of income to these goods, and emissions per dollar spent decline. Our assumption is consistent with cross-sectional evidence from Klenow et al. (2024), who find, using firm-level data from Chile, that higher-quality producers exhibit both greater revenue productivity and lower energy intensity.

Two forces drive the endogenous shift from quantity-led to quality-led growth. First, when goods and services are gross complements, technological progress in manufacturing *lowers* the cost share of physical inputs, weakening incentives for further material productivity gains. Second, nonhomothetic preferences imply that, as income rises, demand shifts toward quality-sensitive, service-intensive goods. Together, these forces redirect innovation and reduce the environmental impact of growth.

To quantify the model's predictions, we parameterize preferences and technology to match key features of structural change in the U.S. economy. The model replicates the secular rise in the service employment share, the recent productivity slowdown, and the trend in relative prices between goods- and service-intensive industries. We discipline income effects by estimating the income elasticity of different consumption categories using the Consumer Expenditure Survey and mapping these estimates to our model. The model also captures the long-run trajectory of CO_2 emissions observed in the data.

Is degrowth necessary? Our analysis highlights that the answer depends on how we measure GDP. Our theory predicts that material production—the main source of environmental damage—converges endogenously to a finite upper bound, not due to resource constraints, but to a realignment of innovation. While material productivity eventually

stagnates, quality-enhancing innovation grows unboundedly. This leads to rising living standards despite a constant amount of physical production. If such quality improvements were fully captured in GDP measurement, the resulting shift toward quality-led growth would not imply degrowth. In practice, however, these gains—especially in services—are poorly measured. Accordingly, our theory predicts a gradual slowdown in measured GDP growth. This perspective offers a new interpretation of the post-2000 decline in TFP growth: rather than reflecting diminished innovation, it may signal a structural shift toward sectors where progress is harder to capture in standard productivity statistics.

The environmental economics literature, e.g., Brock and Taylor (2005) and Shapiro and Walker (2018), highlights three channels through which economic growth affects carbon emissions: scale of production; technology—improvements in production processes that reduce emissions per unit of output; and composition—shifts in the types of products produced, such as moving from more polluting goods to less polluting services. Our theory predicts a role for each of these channels and allows us to decompose the contribution of each component to the observed trajectory of emissions. We show that the time-series variation in emissions is primarily driven by scale, that is, the endogenous decline of material productivity, induced by nonhomothetic preferences and technological substitution toward service inputs. Notably, this material degrowth does not imply a scaling down of the welfare-relevant GDP measure.

Despite its virtues, the transition to a weightless economy may not occur quickly enough to avoid an ecological disaster. In such cases, policy should aim to reinforce—rather than resist—this structural evolution. More generally, measures that accelerate structural change—such as taxing goods production or subsidizing clean services—can be welfare-enhancing by reducing the environmental externalities of polluting activities. But our theory also offers a caution: political support for such policies can vary across the income distribution. Poorer households are generally less supportive, both because of their higher marginal utility of consumption and the heavier tax burden they face on goods-intensive necessities.

Literature Review Our study contributes to several strands of research. We build on the literature examining the macroeconomic and welfare consequences of climate change, pioneered by Nordhaus (1991) and developed by Golosov et al. (2014). This work neither distinguishes between quality and quantity nor endogenizes innovation.

More closely related to our analysis is the literature on directed technical change. The seminal model of Acemoglu et al. (2012a) shows how policy can steer innovation toward

cleaner production. Several papers extend this framework—see, e.g., Acemoglu et al. (2016), Aghion et al. (2024), Hémous (2016), and Aghion et al. (2023). Hémous and Olsen (2021) provide an excellent review of this literature on green innovation and the energy transition. This work focuses on how goods are produced and does not distinguish between quality- and quantity-based growth across consumer products.

Our paper relates to the debate on degrowth, whose proponents argue that the pursuit of unlimited economic expansion is incompatible with a finite stock of non-renewable resources. Foundational contributions to this literature emerged in the 1970s. The Club of Rome's influential report *The Limits to Growth* (Meadows et al. (1972)) assessed the long-term effects of population and economic growth on planetary boundaries. A central figure in this debate is Georgescu-Roegen (1971), who argued that modern economies irreversibly transform low-entropy resources, such as raw materials, into high-entropy waste. Because this transformation is inherently one-way, he concluded that unlimited material growth is physically impossible.[2]

Stiglitz et al. (2009) propose alternative indicators of social and environmental progress beyond GDP. In this spirit, De Ridder and Lukasz (2025) develop an adjusted measure of total factor productivity that incorporates carbon emissions. They show that recent years appear more virtuous when productivity is measured net of environmental costs, despite the slowdown in traditional TFP. Their contribution complements ours: our theory explains how endogenous shifts in innovation can promote sustainable growth patterns that standard growth accounting overlooks.

Our paper also contributes to the literature on structural change and service-led growth. Boppart (2014) features nonhomothetic preferences in the PIGL class, which we adopt. Alder et al. (2022) and Fan et al. (2023) fit these preferences to data, and we exploit their aggregation properties in our normative analysis. An alternative specification is studied in Comin et al. (2021). We contribute to this literature by introducing endogenous technical change and linking service-led structural transformation to sustainable growth.

A key assumption of our framework relative to previous models of structural change is that demand for quality is nonhomothetic. Bils and Klenow (2001) show that higher-income households allocate a larger share of spending to higher-quality versions of goods, suggesting that income growth shifts demand toward quality. Fieler (2011) finds that wealthier consumers have a higher elasticity of substitution for quality, making them more

[2] See also Daly (1977). More recent contributions to the degrowth literature include, among others, Kallis (2011), D'Alessandro et al. (2020), and Hickel et al. (2022).

willing to pay for better products—a conclusion echoed by Fajgelbaum and Grossman (2011).

The paper is organized as follows. Section 2 provides empirical motivation. Section 3 develops the theory. Section 4 contains the quantitative analysis. Section 5 discusses policy implications. Section 6 discusses how the model can be extended to an open economy setting. Section 7 concludes. The Online Appendix (henceforth, the "Appendix") contains data details and technical results.

2 Services and Emissions: Empirical Motivation

In this section, we document empirical evidence that motivates our theoretical analysis. A core premise of our theory is that services are relatively environmentally friendly. We show that there is a pronounced negative relationship between emissions and service intensity over time, across industries, across countries, and across the consumption distribution of households of different income. Details on the data construction are deferred to Section 4.1 and Appendix A.

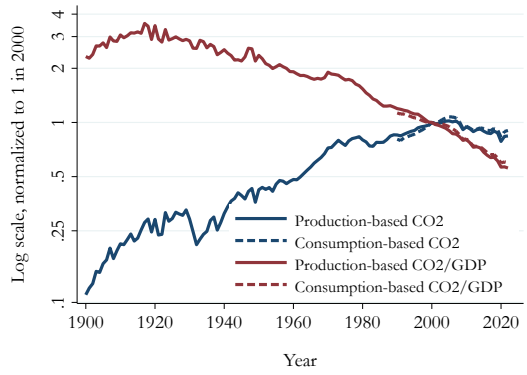
These findings are reported in Figure 1. Panels (a) and (b) show two long-run trends in the U.S. economy: a steady decline in emissions intensity since 1917 (and in total CO_2 emissions since 2008), and a sustained increase in the service employment share, from about 30% in the early 20th century to over 85% today.[3] We argue that these trends are linked: if value added in services generates less pollution than in manufacturing, structural transformation may have contributed to the decline in emissions intensity.

This view implies that a negative relationship between service intensity and emissions should also be observable across both industries and countries. Panel (c) shows the correlation between CO_2 emissions intensity and service cost share across industries, using data from the EPA's National Emissions Inventory (NEI). Emissions intensity is measured as total emissions over total sales, based on the 2002 Economic and Agricultural Censuses. Using sectoral linkages from Input-Output tables, we estimate each industry's emissions intensity by incorporating all upstream and downstream relationships. Likewise, we calculate service intensity as the share of service inputs needed to produce one dollar of

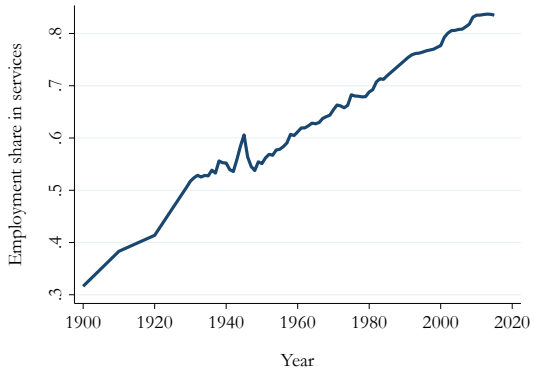
[3] In Panel (a), both total CO_2 emissions (blue) and emissions intensity relative to GDP (red) are normalized to 1 in the year 2000. The dashed line shows consumption-based emissions (available from 1990 onward), adjusted for trade. The two measures are nearly indistinguishable. Some differences emerge post-1990, but correlation is very high. We therefore use the production-based series, which spans a longer period.

Figure 1: The Cleansing Effect of Services

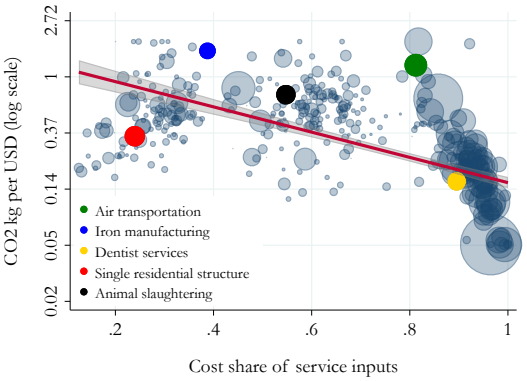
(a) Emissions in the U.S.



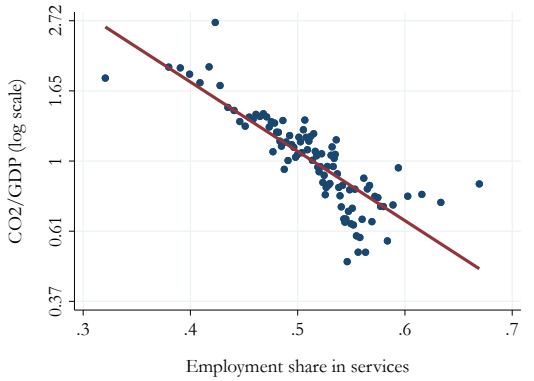
(b) The Rise of Service Employment



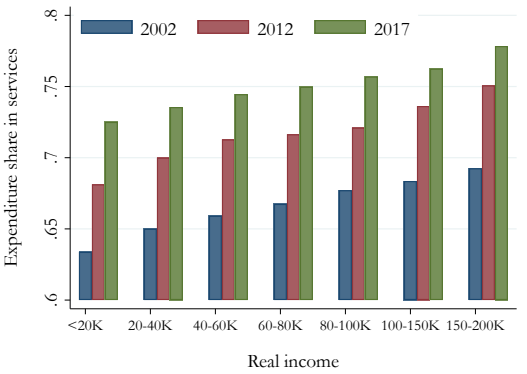
(c) Services are Green: Industry Variation



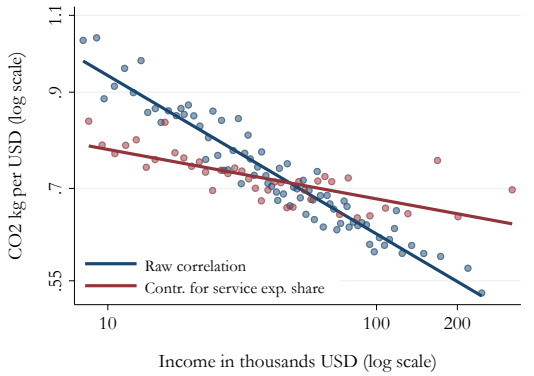
(d) Services are Green: Country Variation



(e) Nonhomothetic Service Demand



(f) Environmental Engel Curves



Notes: Panel (a) shows total annual CO_2 emissions (blue) and emissions relative to GDP (red), using production-based (solid lines) and consumption-based (dashed lines) measures. Panel (b) shows the U.S. service employment share. Panel (c) plots the relationship between the CO_2 intensity and the service cost share at the industry level, using data from the EPA. Panel (d) shows the cross-country correlation between CO_2 intensity and the service employment share, controlling for the employment share in agriculture, population, land size, and year fixed-effects, using data from the World Bank. Panel (e) shows the expenditure share on services across consumers by income and period. Panel (f) displays the relation between CO_2 intensity and income (blue dots), and controlling for the expenditure share in services (red dots). The service cost shares in Panel (c) and expenditure shares in Panels (e) and (f) account for sectoral linkages via the Input–Output matrix.

output, accounting for the full production network. Panel (c) reveals a strong negative correlation: a 10 percentage point increase in service share is associated, on average, with a 19.5% reduction in emissions per dollar.

In Panel (d) we show that this negative relationship between emission intensity and the prevalence of the service sector is also present in the cross-country data. Specifically, we estimate cross-country regressions of the form:

$$\ln (e/y)_{ct} = d_t + \beta \; s_{ct}^{SERV} + \gamma \ln y_{ct} + \phi s_{ct}^{AG} + x_{ct}'\rho + u_{ct}. \tag{1}$$

Here, e/y denotes emissions intensity (total emissions relative to GDP); s_{ct}^{SERV} is the service employment share; $\ln y_{ct}$ is log GDP per capita; s_{ct}^{AG} is the agricultural employment share; x is a vector of additional country-specific controls; and d_t are time fixed effects. Since we control for the employment share in agriculture, β captures the correlation between service employment and emissions intensity relative to manufacturing. Panel (d) displays the resulting residual cross-country correlation as a binscatter plot and shows a robust negative relationship: a one percentage point increase in the service employment share is associated with a 5% reduction in emissions per unit of GDP.

In Appendix A.2, we report further details on the estimation of (1). In particular, we show that the estimated coefficient β is robust to alternative sets of controls in x_{ct} and we also estimate a version with country fixed effects, so that β is identified from within-country variation in service employment and emissions. Although the coefficient declines in magnitude, the effect remains sizable: a one-point increase in service employment is associated with a 2.2% decline in emissions intensity. Finally, we replicate the analysis using cross-sectional data on counties *within* the U.S., and again find a significant negative relationship between service employment and emissions intensity.

In the bottom panels, we turn to cross-sectional evidence across individuals. Using Consumer Expenditure Survey (CEX) data, we compute the service value-added share per dollar of consumer spending across income bins for three years: 2007, 2012, and 2017. Panel (e) highlights two salient features. First, services are luxuries: the expenditure share on service-intensive goods rises with income. Second, holding income fixed, the service share of spending has increased over time. In 2002, households earning \$80,000–\$100,000 allocated about 67% of spending to services; by 2017, this had risen to 75%. This pattern has two implications. On the one hand, rising income boosts service demand. On the other, quality improvements and rising service intensity in production might explain why service spending has grown even at constant income.

Panel (f) estimates the environmental footprint of consumers at different income levels. We compute total emissions per dollar of spending by linking industry-level emissions data from the EPA to households' consumption baskets. The blue line shows the raw correlation—sometimes referred to as the *environmental Engel curve* (Levinson and O'Brien (2019))—which is negative: richer consumers pollute less *per dollar of spending*. Our theory explains this pattern as a consequence of the rising service share by income shown in Panel (e). As shown by the red line, the residual environmental Engel curve, after controlling for service share, is indeed substantially flatter.

These findings provide empirical motivation for our theory, to which we turn next.

3 Theory

The production sector of the economy consists of two intermediate inputs, goods (G) and services (S), and a set of consumption good industries comprising J final products.

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

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

where $k \in \{G, S\}$. Within each sector k, each variety i is produced under monopolistic competition. Individual inputs for material goods are produced with the linear technology $y_{iG} = A_i h_{iG}$, where h_{iG} denotes labor used in the production of manufacturing good i. The productivity distribution $\{A_i\}_{i=0}^1$ evolves endogenously over time due to technical change. The service sector operates under a similar technology, but we assume no productivity growth in services, that is, $y_{iS} = h_{iS}$.

Final Goods: Consumers have preferences over J final products. They value both the quality and the quantity consumed. Each of the J products is a CES bundle comprising a unit interval of varieties. More formally, the quality-weighted consumption of good $j \in \{1, 2, \dots J\}$ 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} denotes the quality of variety i and $\alpha_j \in [0, 1]$ captures the sensitivity of consumers' demand to quality, which is assumed to vary across the J final products. For

example, consumers are more susceptible to quality differences between restaurants than between pet food brands.

Each final good variety is produced with manufacturing (Y_G) and service (Y_S) inputs:

$$y_{ij} = \left((1-\lambda_j)^{\frac{1}{\rho}} Y_{ijG}^{\frac{\rho-1}{\rho}} + \lambda_j^{\frac{1}{\rho}} Y_{ijS}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}, \tag{2}$$

for $(i,j) \in ([0,1] \times \{1,2,\ldots,J\})$. The parameter $\lambda_j \in [0,1]$ captures the final product j's service intensity: the higher λ_j, the more intensive is product j in services.

Because final consumption depends on the quality of the respective consumption varieties, there is an important distinction between the market price and the quality-adjusted, hedonic price index. Symmetry allows us to denote the market price of each variety of final product j, y_{ij}, as $\tilde{p}_j$. The effective price for each unit of final good consumption of product j is then given by $p_j = \tilde{p}_j / Q_j^{\alpha_j}$, where $Q_j \equiv \left(\int_0^1 Q_{ij}^{\alpha_j(\xi-1)} di \right)^{\frac{1}{\alpha_j(\xi-1)}}$ denotes the quality of product j. Like quantity TFP A_i, quality Q_{ij} also evolves endogenously.

Labor Supply: The consumer side of the economy consists of a large representative household comprising a continuum of measure one 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 period.

All household members inelastically supply one unit of labor. In addition, some household members draw specific skills—either innovative or parasitic—that enable them to perform entrepreneurial tasks. Innovators discover new technologies and run firms using a superior technology for 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 either innovating or parasitic firms.[4] This market structure ensures that all output is distributed to households each period and that innovation decisions are static, with no forward-looking component to complicate the analysis. We describe the details of the innovation environment in

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

Section 3.2 below.

All individuals 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 among all members for them to consume.

Preferences: The representative household's preferences are parameterized by the following indirect utility function in the PIGL class (cf. Boppart (2014), Fan et al. (2023))

$$\mathcal{V}\left(e,[p_j]_{j=1}^J\right) = \frac{1}{\varepsilon}\left(\frac{e}{\prod_{j=1}^J p_j^{\omega_j}}\right)^{\varepsilon} - \tilde{\varsigma}\prod_{j=1}^J p_j^{\varsigma_j} - v(\mathcal{P}), \tag{3}$$

where $\sum_{j=1}^J \varsigma_j = 0$ and $\sum_{j=1}^J \omega_j = 1$. Here, e denotes total spending and p_j denotes the hedonic, quality-adjusted price of a unit of consumption of product j. 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 is strictly positive, increasing, and convex.

Roy's Identity implies that the expenditure share on product k for a consumer with spending level e facing prices $[p_j]_j$ is given by

$$\vartheta_k\left(e,[p_j]_{j=1}^J\right) = \omega_k + \phi_k\left(\frac{e}{\prod_{j=1}^J p_j^{\beta_j}}\right)^{-\varepsilon}, \tag{4}$$

where $\beta_j \equiv \omega_j + \varsigma_j/\varepsilon$ and $\phi_k \equiv \varsigma_k \times \tilde{\varsigma}$.[5]

Equation (4) highlights that the demand for product k is fully determined by the two product-specific parameters, ω_k and ϕ_k, and an index of overall spending, $e/\prod_{j=1} p_j^{\beta_j}$. 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. As $e \to \infty$, the expenditure shares ϑ_k converge to their limiting values, ω_k; the share spent on necessities converges from above, the share spent on luxuries from below. Hence, the demand system in our economy resembles a Cobb–Douglas specification with a nonhomothetic adjustment. The slope of the Engel curves and the magnitude of income effects are governed by the parameter ε, which we term the *Engel elasticity*.

Next, we introduce our key assumption on the link between the service-intensity, luxury-sensitivity, and income-elasticity of final goods.

[5] Note that $\sum_{j=1}^J \varsigma_j = 0$ and $\sum_{j=1}^J \omega_j = 1$ imply $\sum_{j=1}^J \phi_j = 0$, and $\sum_{j=1}^J \beta_j = 1$.

Assumption 1. *Consumption goods $j \in \{1, 2, \dots J\}$ are ranked on a sophistication ladder: $\phi_{j+1} \leq \phi_j$, $\lambda_{j+1} \geq \lambda_j$, and $\alpha_{j+1} \geq \alpha_j$, for all $j < J$.*

Assumption 1 postulates that consumption goods are ranked on a sophistication ladder where growing sophistication is associated with a higher service intensity in production (λ_j), a higher expenditure elasticity (ϕ_j), and a greater weight given to quality (α_j). 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: Pollution is a state variable that follows the law of motion

$$\mathcal{P}_t = (1 - \varphi)\, \mathcal{P}_{t-1} + \mathcal{E}_t \quad \text{where} \quad \mathcal{E}_t = \mathcal{E}\left(Y_{Gt}, z_t\right). \tag{5}$$

Here $\mathcal{E}$ denotes the flow of new emissions that we assume is determined by the level of production of goods, Y_{Gt}, and an index of *green technology*, z_t, that determines the environmental impact of production activity. The function $\mathcal{E}$ is nondecreasing in Y_{Gt} and strictly decreasing in z_t—the larger z_t, the less polluting production activity.

In line with the empirical evidence of Figure 1, the specification in (5) postulates that material production has a larger environmental footprint than services.[6] In addition, it highlights a fundamental distinction between quantity and quality: for a given level of goods production, higher quality raises welfare and GDP *without* increasing emissions.

3.1 Two Final Products: Luxuries and Necessities

In what follows, we specialize our analysis to two consumer goods ($J = 2$). We label $j = 1$ as N (*necessity*) and $j = 2$ as L (*luxury*). In line with Assumption 1, we assume $\lambda_N < \lambda_L$, indicating that the luxury good is more service-intensive. For simplicity, we assume consumers are indifferent to quality in necessities ($\alpha_N = 0$) and fully sensitive to quality in luxuries ($\alpha_L = 1$). This implies that, for the luxury good, consumers only care about the number of quality units purchased; $p_L = \tilde{p}_L / Q$.

[6] The assumption that services generate no pollution is not essential. It could be relaxed to $\mathcal{E}_t = \mathcal{E}(Y_{Gt}, Y_{St}, a_t)$ while maintaining that service production is less polluting than goods production. Our quantitative analysis suggests that this simplifying assumption provides a reasonable approximation of the observed distribution of emissions across industries shown in Panel (c) of Figure 1.

3.1.1 Static Equilibrium

The static equilibrium determines the equilibrium allocation of labor given the state of technology (A) and the quality of luxury goods (Q).

Production: Goods and services are produced by monopolistically competitive firms. Given the isoelastic demand for different varieties, firms set the price 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 for the prices of inputs (G, S) and final goods (N, L):

$$p_G = \frac{\xi}{\xi - 1}\frac{w}{A}, \quad p_S = \frac{\xi}{\xi - 1}w, \quad \tilde{p}_N = \frac{\xi}{\xi - 1}c_N\left(w, A\right), \quad \tilde{p}_L = \frac{\xi}{\xi - 1}c_L\left(w, A\right), \tag{6}$$

where w is the wage, $A \equiv \left(\int A_i^{\xi-1} di\right)^{1/(\xi-1)}$ is the average productivity in the goods-producing sector and $c_j\left(w, A\right) = \left((1-\lambda_j)\, p_G^{1-\rho} + \lambda_j p_S^{1-\rho}\right)^{1/(1-\rho)}$ is the unit cost of production of $j \in \{N, L\}$. Substituting the expressions of p_G and p_S, we obtain:

$$c_j\left(w, A\right) = \frac{\xi}{\xi - 1} w \times \psi_j(A), \quad \text{where} \quad \psi_j(A) = \left((1-\lambda_j)\, A^{\rho-1} + \lambda_j\right)^{\frac{1}{1-\rho}}. \tag{7}$$

Note that $\psi_j' < 0$, namely, quantity TFP A reduces the cost of production for both goods. However, the impact is stronger for necessities that have a lower service content, $\lambda_N < \lambda_L$.

Whether quantity TFP can be the source of long-run growth depends crucially on the substitutability between goods and services. In particular, (7) implies that

$$\lim_{A\to\infty} \psi_j(A) = \begin{cases} 0 & \text{if } \rho \geq 1, \\ \lambda_j^{\frac{1}{1-\rho}} & \text{if } \rho < 1. \end{cases} \tag{8}$$

When goods and services are substitutes ($\rho > 1$), perpetual productivity growth can reduce the cost of production to zero. Conversely, if they are complements, services are essential, and the costs of final goods are bounded by the service cost share λ_j—a manifestation of the classical Baumol (1967) effect.

The elasticity of substitution ρ also plays a crucial role for the relative cost shares of goods and services in the production of final products. Given the CES production

function, the cost share of goods for product k, σ_k, is given by

$$\sigma_k(A) = \frac{(1-\lambda_k)\, p_G^{1-\rho}}{(1-\lambda_k)\, p_G^{1-\rho} + \lambda_k p_S^{1-\rho}} = \frac{(1-\lambda_k)\, A^{\rho-1}}{(1-\lambda_k)\, A^{\rho-1} + \lambda_k}, \quad \text{for} \quad k \in \{N, L\}. \tag{9}$$

For a given level of quantity TFP A, the cost share of goods is decreasing in λ_k. Moreover, σ_k decreases in A if goods and services are complements, and increases in A if they are substitutes. Hence, in the empirically relevant case of goods and services being complements, technological progress in the production of goods endogenously *increases* firms' spending share on services and shifts overall employment into the service sector.

Given the production costs in (7), the *quality-adjusted* prices are given by

$$p_N = \left(\frac{\xi}{\xi-1}\right)^2 \psi_N(A)\, w \qquad \text{and} \qquad p_L = \left(\frac{\xi}{\xi-1}\right)^2 \frac{\psi_L(A)}{Q} w, \tag{10}$$

where $Q = \left(\int Q_i^{\xi-1} di\right)^{1/(\xi-1)}$ is the average quality of the varieties of the luxury good. Equation (10) highlights a crucial distinction between necessities and luxuries: rising quantity TFP *lowers* the relative price of necessities due to their low service content. By contrast, rising quality *increases* the relative price of necessities because luxuries are more sensitive to quality growth.[7] For the remainder of our analysis, we choose the wage as the numéraire, i.e., we set $w = 1$.

Demand: The expression for relative spending shares in (4) implies that aggregate demand for final goods depends on an index of overall spending, Υ, given by

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

where the second equality uses equation (10).

In equilibrium, all labor income and profits accrue to the representative household. Thus, the household's total expenditure is given by $e = w + \Pi$, where Π denotes aggregate

[7] The term $(\xi/(\xi-1))^2$ captures the double marginalization effect arising from monopoly power, which is present in both intermediate and final production stages.

profits. Because of the constant markup, profits are proportional to wage income:[8]

$$\Pi = \frac{1}{\xi - 1} w + \frac{1}{\xi - 1} \left(\frac{\xi}{\xi - 1} w \right) = \left(\left(\frac{\xi}{\xi - 1} \right)^2 - 1 \right) w. \tag{12}$$

The normalization $w = 1$ then implies that $e = \left(\frac{\xi}{\xi - 1} \right)^2$.[9] As a consequence, with some slight abuse of notation, Υ is fully determined by Q and A and given by

$$\Upsilon(A, Q) = \left((1 - \lambda_N) A^{\rho - 1} + \lambda_N \right)^{\frac{\beta}{\rho - 1}} \left((1 - \lambda_L) A^{\rho - 1} + \lambda_L \right)^{\frac{1 - \beta}{\rho - 1}} Q^{1 - \beta}. \tag{13}$$

The expenditure shares on goods N and L are therefore given by

$$\vartheta_N \equiv \vartheta(A, Q) = \omega + \phi \, \Upsilon(A, Q)^{-\varepsilon} \quad \text{and} \quad \vartheta_L = 1 - \vartheta(A, Q), \tag{14}$$

where $\omega \equiv \omega_N = 1 - \omega_L$, $\beta \equiv \beta_N = 1 - \beta_L$, 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. Because Υ increases in both arguments, both quantity and quality growth increase real spending and shift spending from the necessity to the luxury good. Higher quality does so by lowering the quality-adjusted price of luxuries, while higher quantity TFP A reduces the market prices of both goods.

Labor Market Equilibrium: We now characterize the equilibrium allocation of labor. We exploit the fact that, for both goods and services, total factor payment must equal value added. This implies that the sectoral employment shares are given by

$$H_G = \sigma_N(A) \vartheta(A, Q) + \sigma_L(A) (1 - \vartheta(A, Q)) \quad \text{and} \quad H_S = 1 - H_G. \tag{15}$$

Equation (15) shows that sectoral employment is the (expenditure share weighted) average of sectoral cost shares. As such, the sectoral allocation of labor is fully determined by A and Q.

Equation (15) highlights the distinct roles of Q and A in driving structural change

[8] See Appendix B.1 for a derivation of (12). The term $\frac{w}{\xi - 1}$ captures the profit generated by intermediate manufacturing and service firms, while the term $\frac{1}{\xi - 1} \left(\frac{\xi}{\xi - 1} w \right)$ captures the profit generated by final good firms.

[9] Since $\left(\frac{\xi - 1}{\xi} \right)^2 e = 1$, monopolies do not introduce distortions: market power only affects the distribution of income between wages and profits, leaving the overall allocation unaffected. This result arises because firms exhibit uniform market power across all sectors and labor supply is exogenous—cf. Epifani and Gancia (2011).

toward the service sector. It is easy to show (details in Appendix B.1) that the service employment share H_S is: (i) increasing in Q; (ii) increasing in A if $\rho < 1$; (iii) decreasing in A if $\rho > 1$ and ϕ is small, and (iv) constant in the case of Cobb-Douglas preferences and technologies (i.e. $\rho \to 1$ and $\phi = 0$). An increase in Q shifts demand toward luxury goods, which are more service-intensive, thereby raising the service employment share. An increase in A has two effects. Like Q, it shifts demand toward service-intensive luxury goods. In addition, it also alters the cost structure of final goods. If goods and services are complements ($\rho < 1$), higher A raises the service cost share, amplifying the income effect. If they are substitutes ($\rho > 1$), it raises the cost share of goods and can reduce service employment if income effects are small.

3.2 Directed Innovation and Equilibrium Dynamics

A key aspect of our analysis is the endogenous direction of technological progress between quantity and quality growth. We assume a total innovative research capacity of mass R, which can be directed toward improving either the productivity of manufacturing firms (A_i) or the quality of luxury goods (Q_i). Let R_Q and R_A denote the research effort allocated to Q and A, respectively, subject to the market-clearing condition $R_Q + R_A = R$.

The rate at which a unit of research effort in sector $s \in \{A, Q\}$ yields a successful innovation is given by $\eta_s R_s^{-\zeta}$, where R_s denotes the aggregate amount of research effort allocated to sector s and ζ captures the strengths of congestion in research. The parameter η_s denotes research efficiency. A successful innovation increases the quality or productivity of a randomly selected firm by a factor $\gamma > 1$.

From our earlier assumption, researchers reap profits only for a single period. Let V_s denote the expected value of directing research towards $s \in \{A, Q\}$. Then:

$$V_Q = \eta_Q R_Q^{-\zeta} \times \int_0^1 \pi_{iL}\, di \qquad \text{and} \qquad V_A = \eta_A R_A^{-\zeta} \times \int_0^1 \pi_{iG} di,$$

where $\int_0^1 \pi_{iL} di$ denotes the expected profits of replacing a randomly selected firm producing luxury final goods and $\int_0^1 \pi_{iG} di$ denotes the profits of a goods-producing intermediate input firm.

Standard analysis (see Appendix B.1) establishes that equilibrium profits are equal

to:

$$\int_0^1 \pi_{iL} di = \frac{1-\vartheta(A,Q)}{\xi} \quad \text{and} \quad \int_0^1 \pi_{iG} di = \frac{\vartheta(A,Q)\sigma_N(A) + (1-\vartheta(A,Q))\,\sigma_L(A)}{\xi} = \frac{H_G}{\xi}.$$

Note that the demand for quality improvements arises directly from the demand for luxury goods and hence only depends on $\vartheta(Q,A)$. By contrast, the demand for material productivity innovation is derived from the cost share of goods in producing both necessities and luxuries, and therefore ultimately tied to the employment share of goods.

In equilibrium, the value of the marginal product of researchers is equalized: $V_A = V_Q$. Substituting the expressions for V_Q and V_A, the resulting arbitrage condition is:

$$\frac{R_Q}{R_A} = \left(\frac{\eta_Q}{\eta_A}\right)^{\frac{1}{\zeta}} \left(\frac{1-\vartheta(A,Q)}{H_G(A,Q)}\right)^{\frac{1}{\zeta}}. \tag{16}$$

Equation (16) is central to our theory, as it determines the equilibrium allocation of researchers and, hence, the direction of technological progress. It shows that this allocation is driven by two forces, both linked to market size. First, a higher spending share on luxuries increases the incentive to improve quality, as demand for quality-sensitive goods is large. Second, holding the final spending share ϑ fixed, a higher employment share in goods production increases the incentive to enhance quantity rather than quality TFP. Intuitively, when ϑ is constant, low employment in goods reflects a high cost share of services. When services account for the lion's share of total costs, incentives to raise quantity TFP are weak. The direction of technological change is thus jointly determined with the structural transformation: the expansion of the service sector goes hand in hand with a shift toward quality-led growth.

Because the right-hand side of (16) is fully determined from A and Q, the dynamic evolution of both quality and quantity TFP is easy to characterize. The respective growth rates are given by

$$g_{At} \equiv \frac{\dot{A}_t}{A_t} = R_A^{1-\zeta}\eta_A(\gamma-1) - \delta \quad \text{and} \quad g_{Qt} \equiv \frac{\dot{Q}_t}{Q_t} = R_Q^{1-\zeta}\eta_Q(\gamma-1) - \delta. \tag{17}$$

Here, $\delta \geq 0$ captures knowledge depreciation, which may also be interpreted as a stand-in for overhead costs in idea production—a positive threshold mass of researchers is necessary to generate improvements in technology or quality. Equations (16) and (17), together with the condition $R = R_Q + R_A$, determine the entire path of $\{A_t, Q_t\}_t$ given an initial condition (A_0, Q_0).

3.3 Asymptotic Balanced Growth Equilibrium

In this section, we study the asymptotic equilibrium dynamics of the model.

Definition 1. *An Asymptotic Balanced Growth Equilibrium (ABGP) is a dynamic equilibrium path such that, as $t \to \infty$, expenditure shares, sectoral employment shares and growth rates converge to a constant: $\vartheta(A_t, Q_t) \to \omega$, $H_{jt} \to \overline{H}_j$ for $j \in \{G, S\}$, and $g_{st} \to g_s \geq 0$ for $s \in \{A, Q\}$.*

Our theory jointly determines the structural transformation toward services, the direction of technological change, and the equilibrium path of emissions. The following Proposition contains a characterization of these objects for the case of complementarities ($\rho < 1$). The analysis of the case $\rho \geq 1$, which is not the main focus of the paper, is deferred to Appendix B.2.

Proposition 1. *Given initial conditions $\{A_0, Q_0\} \gg 0$, there exists $\bar{\delta} > 0$ such that, for all $\delta \in [0, \bar{\delta})$, the economy converges to a unique ABGP. If goods and services are complements, i.e., $\rho < 1$, the ABGP has the following features:*

1) $g_A = 0$ and $g_Q > 0$; that is, growth is entirely quality-led.

2) If $\delta > 0$, the production of goods Y_{Gt} converges in finite time to a constant level $\overline{Y}_G = \overline{A}\,\overline{H}_G < \infty$, where $\overline{H}_G > 0$ and $\overline{A}$ denote the asymptotic employment share in goods and level of quantity TFP, respectively. These are given by:

$$\overline{H}_G \equiv (1-\omega)\frac{\eta_Q}{\eta_A}\left(\left(\frac{\eta_A(\gamma-1)}{\delta}\right)^{\frac{1}{1-\zeta}} R - 1\right)^{-\zeta}, \tag{18}$$

and $\overline{A} = \overline{A}(\omega, \lambda_N, \lambda_L, \rho)$ is implicitly defined by:

$$\overline{H}_G = \omega\sigma_N(\overline{A}) + (1-\omega)\sigma_L(\overline{A}). \tag{19}$$

If $\delta = 0$, then $A_{Gt} \to \infty$, $h_{Gt} \to 0$, and $Y_{Gt} \to \infty$—that is, the production of goods grows unboundedly, while its growth rate converges to zero.

3) Suppose that $\mathcal{E}(Y_{Gt}, z_t) = \kappa_{\mathcal{E}} z^{-t} Y_{Gt}$ where $\kappa_{\mathcal{E}} > 0$ is a constant and $z > 1$ denotes the (exogenous) rate of abatement. Then, emissions converge to zero in the ABGP.

Proof. We separately analyze two cases: $\delta > 0$ and $\delta = 0$. Assume $\delta > 0$. The R&D arbitrage condition (16) implies $\frac{R_{Qt}}{R_{At}} \to \Phi_C \equiv \left(\frac{\eta_Q}{\eta_A}\frac{1-\omega}{\overline{h}_G}\right)^{\frac{1}{\zeta}}$, where $\overline{h}_G$ denotes the asymptotic

employment share in the goods sector. Hence, $R_{At} \to R_A^- \equiv \frac{1}{1+\Phi_C} R$ and $R_{Qt} \to R_Q^+ \equiv \frac{\Phi_C}{1+\Phi_C} R$. This implies that, in the ABGP, $g_Q = \eta_Q(\gamma-1)(R_Q^+)^{1-\zeta} - \delta$, which is strictly positive for sufficiently small δ. To determine $\overline{h}_G$, note that the proposed ABGP features $g_A = 0$. Equation (17) thus implies that $(R_A^-)^{1-\zeta}\eta_A(\gamma-1) = \delta$. Substituting the expression for R_A^- and solving for $\overline{h}_G$ yields Eq. (18) in the proposition. The expression for $\overline{A}$ in Eq. (19) follows from the goods market clearing condition (Eq. (15)). Finally, material production Y_{Gt} satisfies $Y_{Gt} \to \bar{Y}_G = \overline{A}\,\overline{h}_G < \overline{A} < \infty$.

If instead $\delta = 0$, then $A_t \to \infty$ and $\overline{h} \to 0$ because $\rho < 1$. The fact that $\overline{h} \to 0$ implies that $R_{At} \to 0$ and hence $g_A = 0$. At the same time it can be shown that $Y_{Gt} = A_t H_{Gt} \to \infty$. This proves the existence of an ABGP with $g_Q > 0$ and $g_A = 0$.[10] □

Proposition 1 characterizes the long-run properties of the theory in terms of growth, structural change, and emissions. When goods and services are complements—the empirically relevant case—the goods-producing sector gradually recedes, and employment converges almost entirely toward services. As the relative size of the goods sector shrinks, so does the demand for innovations that raise quantity TFP, redirecting research toward quality. The economy thus becomes effectively "weightless," with long-run growth driven entirely by qualitative improvements.

If $\delta > 0$, material output stabilizes at a finite level in finite time. The goods sector becomes small relative to the rest of the economy, with both quantity TFP and the total material production converging to constant, finite levels; a form of "quantity degrowth". Crucially, economic growth does not stop: quality continues to grow, becoming the sole source of rising living standards. If $\delta = 0$, the same mechanism is active, except that employment in goods production vanishes only asymptotically, while quantity TFP and material output do not converge to a finite limit. Rather, they continue to grow at a declining rate that converges to zero.[11]

Part 3 of Proposition 1 summarizes the implications for emissions. Because in our model, emissions stem solely from material production and the growth rate of material production falls to zero (either asymptotically, if $\delta = 0$, or in finite time, if $\delta > 0$), any

[10] To establish convergence, note that if $\rho < 1$, then R_{Qt}/R_{At} increases over time. As a result, g_{At} decreases monotonically, while g_{Qt} weakly increases. Depending on initial conditions, g_{Qt} may initially be zero, but the continued growth of A_t, together with small δ, ensures that Q_t eventually begins to rise. Thereafter, $g_{Qt} > 0$, and the expenditure share $\vartheta(Q, A)$ converges to ω.

[11] This prediction hinges on Inada conditions in R&D: as research input falls, its marginal product rises without bound, sustaining productivity growth even as its contribution to GDP fades. Any overhead R&D cost would play a role similar to knowledge depreciation, implying that R&D is eventually fully directed toward quality, and material output converges to a finite level.

technological progress in abatement ($z > 1$) ensures that emissions eventually decline and the economy reaches net-zero emissions in the long-run. Note that this result does not hinge on our assumption that only material production matters for emissions. Even if the output of services, Y_S, were to contribute to emissions, Y_S is necessarily bounded in our model because service productivity is, by assumption, constant. What is crucial is that rising quality does not increase emissions.

Despite the promising long-run outlook, the transition may proceed too slowly to avert environmental disasters, particularly in the presence of tipping points. We will come back to this point in Section 5, when we discuss environmental policies.

The characterization in Proposition 1 focuses on the case of complementarity. If goods and services are gross substitutes ($\rho > 1$), our economy still admits an ABGP, but labor is increasingly drawn into goods production and away from services—an implication which runs against the empirical fact of a rising service employment share over the last century. Eventually, $\sigma_k \to 1$ and the service sector disappears because even luxury goods, whose spending share approaches $1-\omega$, are produced entirely with material inputs. As a consequence, research effort is split between A and Q (see (16)), and material output Y_G continues to grow at a positive rate in the long run. Net-zero emissions might therefore not be attainable even in the presence of continuing progress in green technologies.

GDP and Degrowth: The implications of our theory for the degrowth debate are best understood by examining its predictions for GDP—the metric around which much of the debate is framed. We focus on the asymptotic equilibrium in which expenditure shares are constant, thereby making the analysis more transparent.

At the heart of our theory is the insight that quality growth allows for long-run welfare growth even with a constant amount of physical production. Let $GDP_{\mathcal{M}}$ denote GDP as measured at market prices without quality adjustment and $GDP_{\mathcal{W}}$ the welfare-relevant metric that accounts for rising quality. Then

$$GDP_{\mathcal{M}} = \frac{e}{\tilde{p}_L^{1-\omega}\tilde{p}_N^{\omega}} = \frac{1}{\psi_L\left(A\right)^{1-\omega}\psi_N\left(A\right)^{\omega}} \quad \text{and} \quad GDP_{\mathcal{W}} = \frac{e}{p_L^{1-\omega}p_N^{\omega}} = Q^{1-\omega}GDP_{\mathcal{M}}. \tag{20}$$

If we continue to assume that $\rho < 1$, $\psi_j \to \lambda_j^{1/(1-\rho)}$ so that market GDP remains bounded. The service share of GDP approaches one, but productivity in services remains stagnant—a manifestation of Baumol's disease. However, quality-adjusted GDP, $GDP_{\mathcal{W}}$, and living standards continue to rise even if material output stagnates.

In practice, statistical offices attempt to incorporate quality adjustments, resulting in

some partial correction. In our quantitative analysis, we calibrate the extent to which these adjustments are reflected in official statistics. As long as some share of quality improvements is captured, long-run GDP growth remains positive.

3.4 Taking Stock

Our theory shows that degrowth in material production can arise endogenously, even in the absence of resource constraints, as part of the economy's structural transformation. As material production approaches its asymptotic limit, measured GDP growth may slow or plateau, but this does not imply economic stagnation. Structural change, shaped by preferences and production technologies, redirects innovation and enables sustained progress through quality improvements. As long as quality improvements are recognized—whether through measurement or conceptually—and quality growth is clean, economic activity continues to expand without compromising environmental sustainability.

4 Quantitative Analysis

In this section, we perform a quantitative analysis based on a calibrated model. Our results indicate that structural change and quality-led growth played an important role in the decline in the emission intensity of the U.S. economy. In Section 5 we use our calibrated model to quantify the impacts of environmental policies.

4.1 Data and Measurement

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

We use the CEX to measure the distribution of individual spending across final goods. Specifically, we use household-level data from the 2002, 2012, and 2017 waves, each covering approximately 12,000 households and reporting consumption expenditures across roughly 400 final good categories.[12] In addition, we use BLS estimates of average per-person expenditure in 2002 at the most detailed UCC level, based on the diary and survey

[12] The CEX includes 407 categories in 2002, 406 in 2012, and 352 in 2017.

samples of the CEX. These data cover 549 distinct UCC (Universal classification code) categories.

To map these data to our model, we aggregate final goods into two mutually exclusive groups defined by their service content. To compute the service content, we combine a concordance table linking CEX categories to BEA PCE categories with the 2002 Input-Output (IO) tables, which report intermediate input use by sector. As detailed in Appendix A.1, we use the IO table and BEA bridge tables to compute the total service content embodied in each industry's output up to final consumption. We then apply the industry-to-product crosswalk observed in the CEX to assign a service share to each final good.

We classify a final good as a necessity if its service share falls below the expenditure-weighted median, and as a luxury if it exceeds this threshold (see Appendix A.3). Our distinction is conceptually based on the service content of different expenditure categories. The terms "necessity" and "luxury" are mere labels, reflecting the empirical relationship between the service (or goods) intensity of a category and its income elasticity. Based on this classification, we compute each individual's spending shares on luxuries and necessities, ϑ_{iL} and ϑ_{iN}, as well as the aggregate cost shares σ_N and σ_L. To measure the prices of luxuries and necessities, we use disaggregated price indices from the CPI microdata, matched by Jaravel (2024) to the UCC item level in the CEX, to construct annual price deflators.

We use the NEI to calculate the environmental footprint of each final product. The NEI reports total CO_2 emissions by industry, which we convert into emissions intensity by dividing total emissions by total sales. We then use sectoral linkages from the IO tables to compute the emissions intensity of each final product. In addition to CO_2, the NEI reports emissions for five other pollutants: particulate matter smaller than 10 microns (PM10), volatile organic compounds (VOC), nitrogen oxides (NOx), sulfur dioxide (SO2), and carbon monoxide (CO). While these measures are not used to calibrate the model, we show below that the service share is also a key predictor of environmental damages across these pollutants.

4.2 Calibration Strategy

Our model is characterized by 15 structural parameters describing preferences and technology and an emission function $\mathcal{E}$. We also allow for potential mismeasurement in official price indices related to quality growth. We capture this with a single parameter, μ, which

we estimate from the data on relative inflation rates between luxuries and necessities. The full set of structural parameters is:

$$\mathcal{P} = \{\underbrace{\varepsilon, \omega, \phi, \beta, \xi}_{\text{Preferences}}, \underbrace{\rho, \lambda_N, \lambda_L}_{\text{Technology}}, \underbrace{R, \zeta, \eta_Q, \eta_A, \gamma, \delta}_{\text{Innovation}}, \underbrace{\mu}_{Q\text{-measurement}}, \underbrace{\mathcal{E}(.)}_{\text{Environment}} \} \tag{21}$$

We calibrate these parameters by targeting salient moments of the structural transformation of the U.S. economy over the past century. The calibration does not assume that the economy has reached its ABGP. Although all parameters are calibrated jointly, there is a clear mapping between specific moments and individual parameters.

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. We normalize the base-year levels of productivity and quality to unity without loss of generality, i.e., $A_{2002} = Q_{2002} = 1$. Equation (9) then directly implies that $\lambda_k = 1 - \sigma_{k,2002}$.

The parameter ρ governs the elasticity of substitution between goods and services. Following the literature, we set $\rho = 0.5$, implying that goods and services are complements.[13]

Household Preferences ω, ε, β, ϕ, and ξ: The Engel elasticity ε governs the strength of income effects, while $\beta = \omega - \varsigma/\varepsilon$ determines the weights assigned to p_N and p_L in determining Υ. Equations (11) and (14) imply that

$$\ln(\vartheta_N(e_i, p_{Lt}, p_{Nt}) - \omega) = \ln \phi - \underbrace{\varepsilon \ln e_i}_{\text{Cross-section}} + \underbrace{\varepsilon\beta \ln p_{N,t} + \varepsilon(1-\beta) \ln p_{L,t}}_{\text{Time-series}}. \tag{22}$$

Equation (22) shows how we can estimate ε from cross-sectional variation across households within periods, and β from the comovement of aggregate expenditure shares and prices in the time-series. In the cross-section, the income elasticity of the distance between the expenditure share on necessities and its asymptotic value ω is constant and equal to the Engel elasticity ε. Following Boppart (2014) and Fan et al. (2023), we estimate ε using the cross-sectional correlation between household income and ϑ_N.

To implement (22), we need to specify the value of ω. According to our theory, all

[13] Comin et al. (2021) calibrate the sectoral elasticity of substitution to 0.5 based on panel estimates. Acemoglu and Guerrieri (2008) use 0.76, while Buera and Kaboski (2009) set the asymptotic elasticity to 0.5. See Herrendorf et al. (2014) for a survey.

Table I: Nonhomothetic 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: The table reports coefficients from individual regressions of the log of expenditure shares on necessities on the log of total expenditure. Column 1 reports the OLS estimate, including year fixed effects. Column 2 presents the IV estimate using occupational fixed effects as instruments for household expenditure across all years, while columns 3, 4, and 5 report the IV estimates separately for each year. All specifications control for fixed effects related to household size, geographic location, education, race, and marital status.

individuals' expenditure shares on necessities should be bounded below by ω. Empirically, the 1st percentile of the observed distribution is 16%—already very low and still declining. We therefore set the asymptotic share ω equal to zero. Equation (22) then implies that ε can be estimated from the regression

$$\ln(\vartheta_{iNt}) = d_t - \varepsilon \ln(e_{it}) + x_{it}'\gamma + u_{it}, \tag{23}$$

where ϑ_{iNt} is the expenditure share on necessities of household i at time t, e_{it} is total household expenditure, and the time fixed effect d_t controls for the prices $p_{Nt}^{\beta} p_{Lt}^{1-\beta}$, which are common across households. In addition, in (23) 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 I. In column 1, we estimate (23) using OLS, pooling CEX data from all years. We control for household size, marital status, race, education, and geographic location. The estimated elasticity is 0.07. In column 2, we implement an IV strategy. We do so for two reasons. First, measurement error in individual expenditure likely biases the OLS coefficient. Second, we aim to capture variation in permanent income, which we view as more informative for identifying nonhomothetic demand. Accordingly, following Fan et al. (2023), we instrument individual log expenditure e_{it} with a full set of occupation fixed effects, implying that ε is identified from systematic differences

in spending across occupations with persistently high or low income.

Column 2 shows that the resulting estimate of ε is 0.318—substantially larger in magnitude. In columns 3–5, we run the same regression separately for each year and find estimates ranging from 0.25 to 0.37—generally close to the pooled IV estimate. For our quantitative analysis, we use the estimate from column 2 as our baseline.

Our estimate of ε allows us to control for income effects to obtain β from the time-series relationship between aggregate expenditure shares and the relative price of necessities and luxuries. Following an approach similar to Boppart (2014), we express the aggregate necessity share $\vartheta_{N,t}$, net of income effects, as a function of relative prices:

$$\ln(\vartheta_N(e_t, p_{N,t}, p_{L,t}) - \omega) + \varepsilon \ln\left(\frac{e_t}{p_{N,t}}\right) = \ln\phi - \varepsilon(1-\beta)\ln\left(\frac{p_{N,t}}{p_{L,t}}\right) \tag{24}$$

We take the time series of $\vartheta_{N,t}$ and e_t directly from the BEA's NIPA Table 2.4.5: Personal Consumption Expenditures by Type of Product, which provides aggregate expenditure broken down by PCE categories.[14]

To construct time series for $p_{N,t}$ and $p_{L,t}$, we use microdata from Jaravel (2024), which provides a mapping to CEX UCC categories. We build necessity and luxury baskets using 2002 CEX expenditure shares, then aggregate item-level prices to form price indices for necessities and luxuries from 1983 to 2019. For luxury prices, we allow for a potential mismeasurement of quality and we estimate the parameters β and the degree of mismeasurement jointly—see below. We use our estimate of ε and the time series of $\frac{e_t}{p_{N,t}}$ to adjust for income effects on the left-hand side of (24). We then estimate the coefficient $-\varepsilon(1-\beta)$ from (24) using OLS. We obtain -0.209, which implies $\beta = 0.343$. Further details are provided in Appendix A.4.

We identify the preference parameter ϕ directly from the observed employment share of services in 2002. Using the normalization that $A_{2002} = Q_{2002} = 1$, (13) implies that $\Upsilon_{2002} = 1$ and that the expenditure share of necessities, $\vartheta(A, Q)$, is simply equal to ϕ—see (14)). Labor market market clearing therefore requires that the employment share of services in 2002—see (15)—satisfies $H_{S,2002} = (1 - \sigma_{N,2002})\phi + (1 - \sigma_{L,2002})(1 - \phi)$. Given that $\sigma_{N,2002}$ and $\sigma_{L,2002}$ are the cost shares of goods for necessities and luxuries, respectively, this equation can be solved for ϕ. Finally, we set the elasticity of substitution ξ to 5, a consensus estimate in the literature.

The Measurement of Quality μ: We explicitly allow for the possibility that quality

[14] To calculate the expenditure share on necessities, we classify each PCE category as a necessity or luxury by its service share in the 2002 IO tables, using the same procedure described in Section 4.1.

growth is only partially captured in official statistics. Specifically, we assume that the BLS price index for luxury goods is given by

$$p_{Lt}^{BLS} = Q_t^{1-\mu} p_{Lt} = \left(\frac{\xi}{\xi - 1}\right)^2 \frac{\psi_L(A_t)}{Q_t^{\mu}} w. \tag{25}$$

When $\mu = 1$, quality is fully measured, and p_{Lt}^{BLS} coincides with the welfare-relevant, quality-adjusted price p_{Lt}. When $\mu = 0$, quality is not measured at all, and the BLS index understates the decline in luxury prices.

To estimate μ, we use item-level inflation data from Jaravel (2024) to target the measured inflation of luxuries relative to necessities. Using expenditure shares from the 2002 CEX and our classification of all items as luxuries or necessities, we construct luxury and necessity baskets to measure price inflation over time and track the evolution of the relative measured price $p_{L,t}^{BLS}/p_{N,t}$.

The Innovation Process R, ζ, γ, η_A, and η_Q: The number of researchers R, the step size γ, and the efficiency parameters η_k are not separately identified. We therefore set $R = 0.1$, implying that 10% of labor is devoted to research, and $\gamma = 1.5$, so each successful innovation increases productivity by 50%. Following Akcigit et al. (2021), we assume the innovation cost function has an elasticity of 2, setting $\zeta = 0.5$. We set $\delta = 0.001$, reflecting the view that knowledge depreciates slowly. This choice has negligible effects over the period of interest and only matters in the very long run. For our horizon, results are virtually unchanged if we set $\delta = 0$.

This leaves the two R&D efficiency parameters, η_A and η_Q, which we calibrate using three moments: the service employment share in 1950, GDP per capita growth from 1950 to 2000, and from 2000 to 2020. Intuitively, both quantity and quality growth raise GDP per capita, but affect service employment differently. Quality growth influences it only through income effects, while quantity growth also raises service employment via technological substitution—the Baumol channel. In addition, the shift in research toward quality impacts the relative growth contributions across the two periods. These three moments thus allow us to separately identify η_Q and η_A.

CO_2 Emissions $\mathcal{E}_t$: We parameterize the emissions function $\mathcal{E}$ as in Proposition 1. The term $z \geq 1$ captures the effect of green technology on emission reduction, which we treat as exogenous for simplicity. Intuitively, for a given level of material output, emissions decline over time due to abatement or improved fuel efficiency. We calibrate z and $\kappa_{\mathcal{E}}$ to match observed CO_2 emissions in 1980 and 2000. Also, to parametrize the stock of

pollution following (5), we set $\varphi = 1 - 0.5^{1/30} \approx 0.023$ to capture the fact that the half-life of emissions in the atmosphere is 30 years (see Solomon et al. (2007)). This parameter does not play a role for our calibration, but is required to evaluate the welfare impact of policies in Section 5.

4.3 Estimated Parameters and Model Fit

We summarize all parameters and their corresponding target moments in Table II. Consistent with observed cost shares, luxuries are substantially more service-intensive than necessities: $\lambda_L = 0.93 > \lambda_N = 0.41$. We also find that research productivity is higher for quantity growth than for quality growth ($\eta_A > \eta_Q$). We estimate a significant role for green technological progress: $z = 1.017$, implying that the environmental footprint of material output declines by 1.7% per year. Finally, we estimate $\mu = 0.67$, implying that 67% of quality improvements are captured in official statistics—in the same ballpark as Aghion et al. (2019) who estimate that about one third of total measured growth is missed in official statistics. In terms of the targeted moments, our model successfully replicates overall GDP per capita growth (both between 1950 and 2000 and 2000 and 2019) and the structural transformation toward services. In addition, our model is also consistent with the data on relative price inflation, whereby the relative price of service-intensive, luxury products increases by 1% a year.

In Figure 2 we display the overall fit of our model graphically. The upper left panel compares model predictions with historical GDP per worker since 1950; the upper right panel shows the corresponding growth rates. In both cases, the data appears in gray, the model in red. The model captures both the rise in GDP per worker over this period and the subsequent growth slowdown. To illustrate the role of quality measurement, we also present two polar scenarios, one assuming all quality improvements are fully captured (dotted lines), and the other assuming none are measured (dashed lines). Figure 2 shows that our model predicts a growth slowdown even if all quality gains were recorded. However, the decline in measured growth is more pronounced because quality growth accelerates over time. If quality improvements were entirely unmeasured, average measured growth between 2000 and 2019 would have been only 0.7%. By 2100, our model implies that GDP per capita would still rise by 1.5%, while measured growth would be close to zero if quality improvements were entirely unaccounted for.

The two middle panels of Figure 2 show the employment share in services (left) and the time series of emissions (right). The model closely tracks the rise in the service em-

Table II: Structural parameters

Parameter	Value	Target	Target value	Model
ε	0.318	Engel curve slope	0.318	-
β	0.343	Time series impact of $\ln \frac{p_{N,t}}{p_{L,t}}$ on $\ln \vartheta_{N,t}$	-0.209	-
λ_L	0.93	Cost share of services in L (IO table)	0.93	0.93
λ_N	0.41	Cost share of services in N (IO table)	0.41	0.41
η_A	0.331	$\text{GDP}_{1950}/\text{GDP}_{2000}$	0.413	0.413
η_Q	0.103	$\text{GDP}_{2019}/\text{GDP}_{2000}$	1.269	1.265
		1950 U.S. service share	54.5	54.6
ϕ	0.314	2002 U.S. service share	77	77
μ	0.67	2002-2017 Measured rel. price inflation $\pi_{L,t}^{BLS}/\pi_{N,t}$	0.01	0.010
z	1.017	1980 U.S. CO_2 emissions	4.722	4.722
$\kappa_{\mathcal{E}}$	145.503	2000 U.S. CO_2 emissions	5.724	5.724
ω	0	Set exogenously	-	-
ξ	5	Set exogenously	-	-
ρ	0.5	Set exogenously	-	-
ζ	0.5	Based on Akcigit et al. (2021)	-	-
δ	0.001	Set exogenously	-	
φ	0.023	Based on Solomon et al. (2007)	-	

Notes: The table reports the calibrated parameters and target moments. We normalize $R = 0.1$, $\gamma_A = \gamma_Q = 1.5$, and $A_{2002} = Q_{2002} = 1$.

ployment share, partly due to calibration targets for the U.S. in 1950 and 2002. Notably, it also approximates the 1900 value (not shown) well, despite this data point not being used in the estimation. Looking ahead, the model predicts that the shift toward services will continue, though at a slower pace, as most of the population is already employed in the sector.[15]

Regarding emissions, the model closely tracks the long-term trend since 1900, despite being calibrated only to match total emissions in 1980 and 2000. Notably, it also replicates the pronounced hump-shaped pattern of CO_2 emissions, even though green technological progress is assumed constant over time. This pattern arises from the relative decline in material production, driven by both a reallocation of labor toward services and a slowdown in quantity TFP growth, combined with continuing green technological progress. However, relative to the data, our model predicts a later peak in total emissions, suggesting that clean technology has advanced especially rapidly in recent decades. This acceleration likely reflects a combination of reduced fossil fuel use, policy-driven investment in renewables, and gains in green innovation efficiency. As a result, the observed decline

[15] In the model, the rise in service employment reflects both a shift toward services within each good's production and a reallocation of expenditure toward luxuries. This aligns with the data, where both the service cost share and luxury expenditure share have increased. Quantitatively, the observed shift toward luxuries between 2002 and 2017 is stronger than what the model predicts.

in emissions is steeper than what structural change and directed innovation alone would predict, underscoring the importance of targeted policy and technological breakthroughs.

Panel (e) highlights how our model is able to generate rising living standards and falling emissions simultaneously: research effort is redirected toward product quality in conjunction with the structural transformation toward services. In 1900, quantity TFP A (dark red) grew at nearly 5% per year, while quality growth (light red) was minimal. Over time, research increasingly prioritizes quality. However, the rise in quality growth is smaller than the decline in productivity growth, since research is more efficient at advancing A than enhancing Q ($\eta_A > \eta_Q$).

This change in the direction of technical change is quantitatively consistent with the data on relative price inflation. In panel (f) of Figure 2 we show that our model fits well the observed change in measured relative prices between service-intensive luxuries and good-intensive necessities between 2002 and 2017. In the data, measured prices of service-intensive products grew, on average, 1 percentage point per year faster than the prices of goods-intensive products (dashed gray line). In our simulated model, we reproduce the same time series of measured relative prices $p_{L,t}^{BLS}/p_{N,t}$ using the assumption on quality measurement in (25). Our model (red bar) exactly matches this gap with $\mu = 0.67$. For comparison, if no quality improvements were measured ($\mu = 0$), relative inflation would appear 1.7 percentage points higher (left-hand pink bar). If all quality improvements were measured ($\mu = 1$), the gap would be only 0.6 percentage points.

Pollution and Service Intensity: Our analysis assumes that material production is the sole driver of emissions and that neither product quality nor the value added of service workers has an environmental footprint. Figure 3 shows that this assumption is, in fact, quantitatively consistent with the cross-sectional correlation between service cost shares and pollution intensity at the industry level shown in panel (c) of Figure 1.

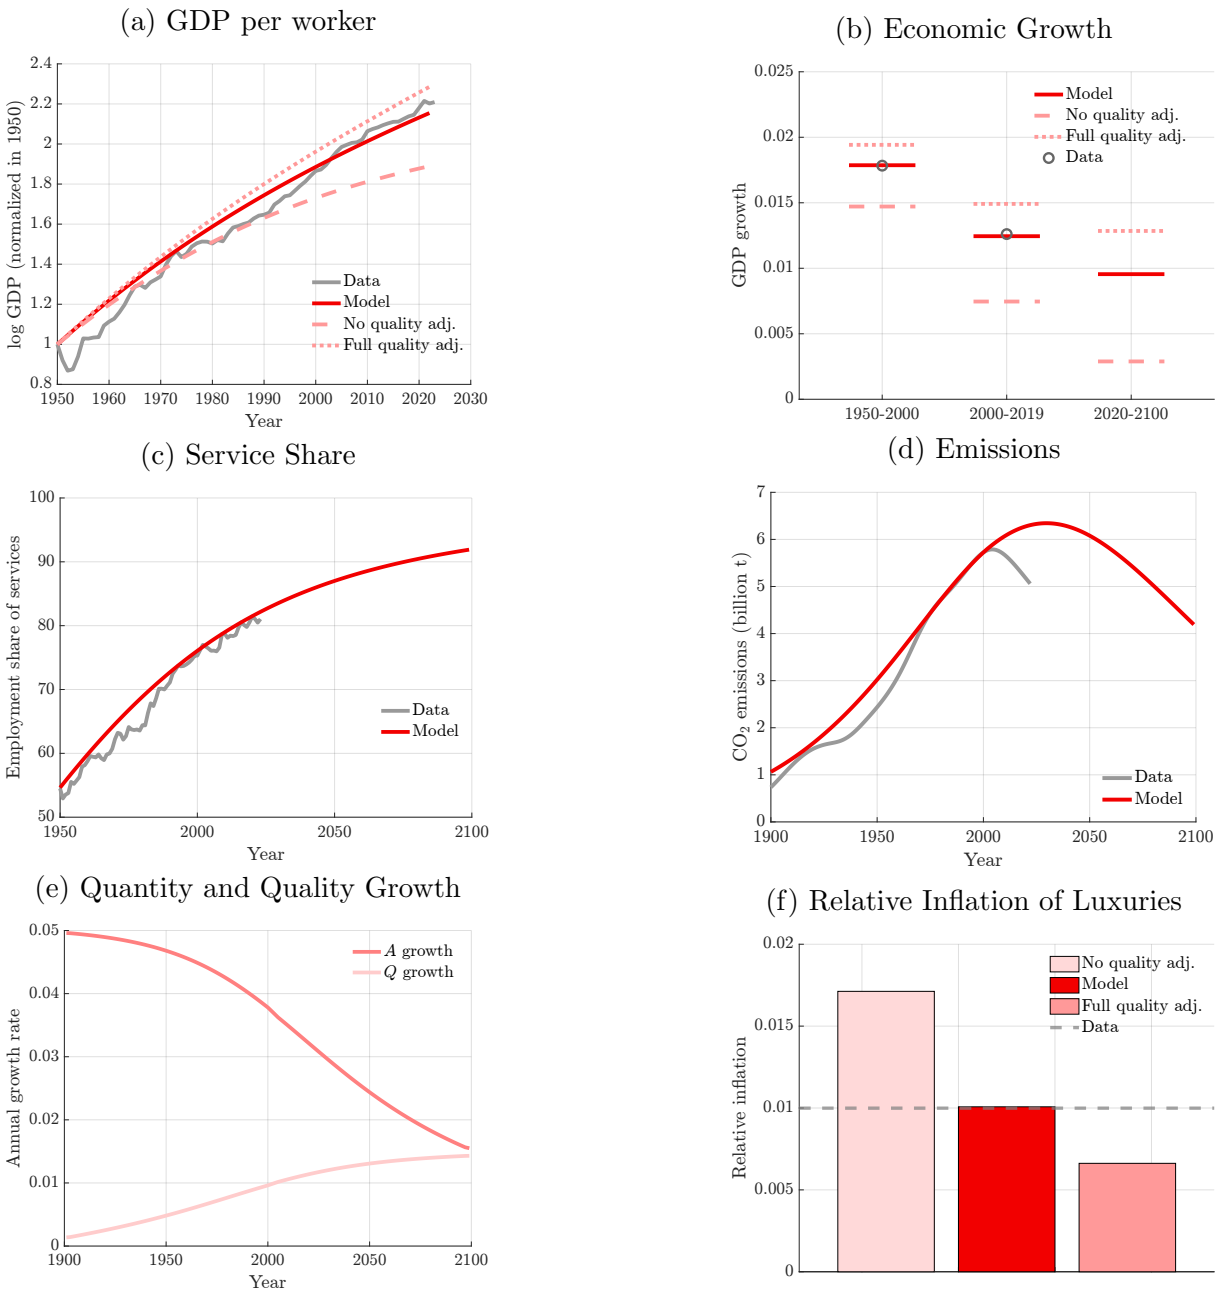
More specifically, as detailed in Appendix A.5, our model yields a precise functional relationship between industrial emissions and service cost shares: (log) emissions per dollar of output in industry k, $\bar{\mathcal{E}}_{kt}$, are given by

$$\ln \bar{\mathcal{E}}_{kt} = d_t + \ln \sigma_k = d_t + \ln \left(1 - \text{service share}_k\right), \tag{26}$$

where d_t is a time-varying effect that reflects equilibrium prices and quantities, but is constant across industries within a given year. As a result, the cross-sectional relationship between emission intensity and service shares should display a unit elasticity.

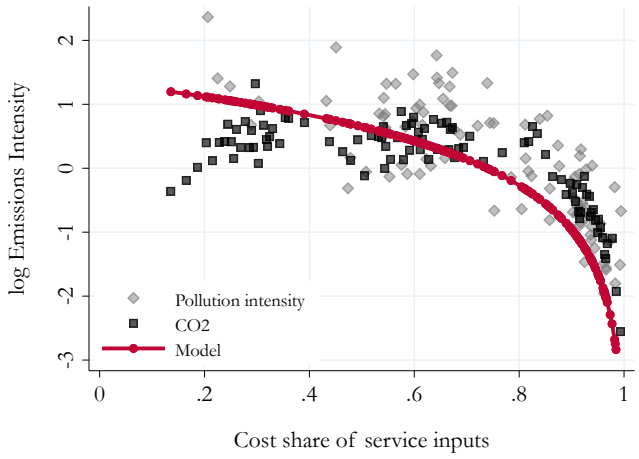
In Figure 3, we plot measured CO_2 intensity (black squares) against service shares

Figure 2: Model Fit



Notes: Panel (a) plots log GDP per worker in the data (gray line) and in the calibrated model (red line), along with two counterfactuals: full quality measurement ($\mu = 1$, dotted line) and no quality measurement ($\mu = 0$, dashed line). Panel (b) displays average GDP per worker growth from 1950–2000 and 2000–2019 in the data (gray bars) and the model (red bars), as well as under the full (light red) and zero (light pink) quality measurement scenarios. Growth rates are computed using chain-weighted price indices with Tornqvist weights. Panels (c) and (d) show the service share and the flow of emissions in the data (gray line) and in the calibrated model (red line). Panel (e) displays the growth rates of A and Q as implied by our model. Panel (f) displays measured price inflation of luxuries relative to necessities, averaged over the period 2002–2017 (dashed gray line). Setting $\mu = 0.67$ exactly matches this relative inflation in the model (red bar). We also report the model's implications if no or all quality growth was measured.

Figure 3: Non-Targeted Moment: Emissions and Service Intensity



Notes: The figure shows the non-targeted relationship between the cost share of services and the log of emissions for each industry, as predicted by the model (red connected dots, see (26)) and as observed in the data. We show the log of CO_2 emissions per dollar (black squares) from the EPA, and the log of pollution intensity (gray diamonds) from Levinson and O'Brien (2019). To construct model-predicted emissions, we consider a set J of final goods, where each $j \in J$ corresponds to an industry indexed by its service cost share $1 - \sigma_j$.

from the data (cf. Figure 1), and superimpose the model's predicted relationship from (26). We calibrate the scale of d_t so that the CO_2 intensity of the median industry in the model matches that of the median industry in the data. 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 goods production.

Figure 3 also includes an alternative measure of industrial pollution intensity from Levinson and O'Brien (2019)—see Appendix A.1. This composite index aggregates five major air pollutants. The figure shows that these pollutants (gray diamonds) are also strongly negatively correlated with the service share, and that our functional form fits this relationship equally well.

4.4 Quantity Degrowth and Falling Emissions

A central prediction of our theory is the endogenous decline of quantity TFP growth. This form of "quantity degrowth," together with the simultaneous rise of quality, reconciles rising GDP and welfare with falling emissions.

To highlight the importance of this channel, in the left panel of Figure 4 we offer a structural decomposition of the change in overall emissions. The literature identifies three ways how economic growth shapes carbon emissions (Brock and Taylor, 2005, Shapiro and

Walker, 2018). The *scale effect*—scale-up of material production with economic growth—increases emissions. The *technique effect*—improvements in the environmental efficiency of production—and the *composition effect*—the shift of economic activity towards environmentally friendly products—counteract the scale effect to reduce emissions.

All three of these channels are present in our theory. To see this, note that emissions are $\mathcal{E}_t = \frac{\kappa_\varepsilon}{z^t} A_t H_{G,t}$, where $H_{G,t} = \vartheta_{N,t}\sigma_{N,t} + \vartheta_{L,t}\sigma_{L,t}$. The growth rate of emissions can therefore be decomposed into the scale effect, stemming directly from growth of quantity TFP, the technique effect, summarized by the combination of green technological progress and the declining cost share of material inputs, and the composition effect, that is, the shift of demand towards service-intensive luxury goods:

$$g_{\mathcal{E}} = \underbrace{g_A}_{\text{Scale}} \underbrace{- \ln z + \frac{\sum_j \vartheta_{jt} \frac{d\sigma_{jt}}{d \ln A_t}}{H_{G,t}} g_A}_{\text{Technique}} + \underbrace{\frac{(\sigma_{Nt} - \sigma_{Lt}) \left(\frac{d\vartheta_{Nt}}{d \ln A_t} g_A + \frac{d\vartheta_{Nt}}{d \ln Q_t} g_Q \right)}{H_{G,t}}}_{\text{Composition}}. \tag{27}$$

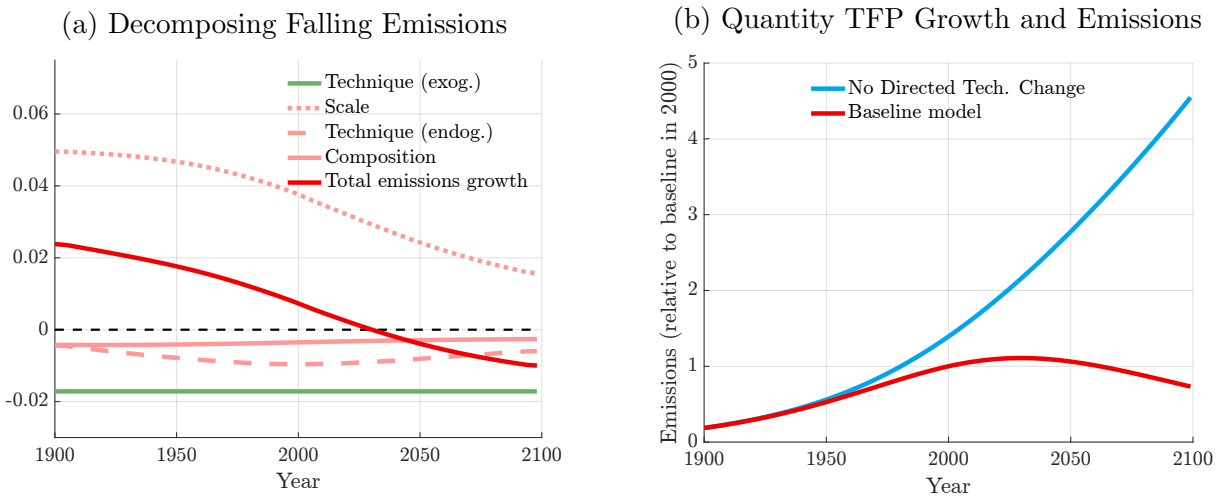
While the scale effect increases emissions, the other two effects are negative as long as goods and services are complements ($\frac{d\sigma_{jt}}{d \ln A_t} < 0$), necessities are more goods-intensive than luxuries ($\sigma_{Nt} > \sigma_{Lt}$), and demand is nonhomothetic ($\frac{d\vartheta_{Nt}}{d \ln A_t} < 0$ and $\frac{d\vartheta_{Nt}}{d \ln Q_t} < 0$).

Our model makes tight predictions about the changing importance of these different channels. First, as highlighted in Proposition 1, our theory predicts a falling scale effect. Because g_A falls over time and eventually converges to zero, in the long run economic growth becomes totally decoupled from material production and is entirely quality-led. Second, our model highlights that technological substitution toward services acts as a technique effect: because rising quantity TFP makes material inputs relatively cheap, the production of final goods becomes less material-intensive over time. Third, our model predicts a pro-green composition effect driven by nonhomothetic preferences. As consumers get richer, demand shifts towards service-intensive, luxury products that have a lower environmental footprint. Finally, the term z represents the transition to cleaner production, which is calibrated but left exogenous in our theory given its prominent role in the existing literature (see, e.g., Acemoglu et al. (2012a)).

In the left panel of Figure 4, we quantify the importance of the different terms in (27) in our calibrated model. The scale effect (dotted line) initially dominates the technique and composition effects, so that emissions (shown in solid red) grow alongside GDP. Over time, the scale effect weakens and is overtaken by the other pro-green effects. As a consequence, the growth of emissions slows down and, eventually, turns negative. Quantitatively, we find that technological substitution toward services (dashed line) and the reallocation of

spending toward service-intensive goods (solid light red line) combined are roughly as important as the effect of green technologies (green line).

Figure 4: The Cleansing Effect of Quantity Degrowth



Notes: Panel (a) implements the decomposition contained in (27). Panel (b) compares emissions under the benchmark model (red) to an alternative simulation where technical change is exogenous and fixed at the levels in 1900.

The right panel of Figure 4 illustrates the critical role of declining quantity TFP growth in a different way. It compares emissions under our benchmark model (red line) to an alternative scenario with exogenous technical change (blue line), in which the research allocation remains fixed at its 1900 level. As a result, quantity and quality growth stay constant at approximately 5% and 0.2%, respectively (see Panel (e) of Figure 2). Unlike the benchmark, continued growth in quantity TFP leads to explosive emissions growth. Interestingly, this increase is driven by faster growth in quantity TFP, *not* by slower reallocation of employment into services. In fact, as we show in Appendix A.6, the service employment share is slightly *larger* in the scenario without directed technological change because faster quantity TFP growth pushes employment out of manufacturing. In summary, while the rise of service employment contributes to lower emissions, the main driver of falling emissions is the slowdown in material productivity growth.

5 Environmental Policy

The coexistence of two growth engines (productivity- and quality-led innovation) and two input types (goods and services) with different environmental footprints creates scope for

policy to influence environmental outcomes by shaping the pace of structural change. In this section, we examine the effects of policies that introduce static or dynamic sectoral wedges, thereby speeding up or slowing the phaseout of polluting activities. While such policies can improve aggregate welfare by internalizing a negative externality, they also alter the relative prices of necessities and luxuries, with redistributive effects that shape political support across income groups.

Goods-Services Wedge: We first consider a policy experiment in which, starting in 2000, the government imposes a constant 20% tax on the purchase of material inputs.[16] The proceeds finance a time-varying subsidy on service inputs, with the government budget balanced in every period. See Appendix B.3 for implementation details.

Figure 5 shows the effects of the policy by contrasting the equilibrium trajectory under the policy (green) with the laissez-faire benchmark (red). All panels also include a counterfactual (blue), where the policy does not influence the direction of technical change, which remains exogenous and identical to the benchmark. The difference between the green and blue lines thus reflects the role of endogenous redirection of innovation in response to the policy: relative to the baseline, A-growth falls and Q-growth rises.

Panel (a) shows that the policy reduces the employment share in goods. The initial drop is sharp (3 percentage points), after which the decline tapers off. In the long run, $H_G \to \overline{H}_G$, as given by (19), which is independent of environmental policy. Notably, total employment in goods is *higher* with directed technical change. The reason is that the tax *reduces* quantity TFP growth, which keeps more workers in manufacturing.

Panel (b) displays the path of pollution. The policy leads to a significant and sustained decline in emissions, with pollution peaking twenty years earlier than in the benchmark. The endogenous shift from A to Q accounts for roughly one-third of the total reduction. This highlights that in the short run, material production declines due to labor reallocation out of the manufacturing sector; in the long run, the drop reflects the cumulative effect of reduced innovation in material productivity.

(Unequal) Welfare Effects: Our model also allows us to compute the welfare effects of said policy. Because the government's budget is balanced each period, the policy affects consumers' welfare through two channels. First, it changes the path of final good prices. Second, it lowers pollution. Letting $[p_{Nt}, p_{Lt}, \mathcal{P}_t]$ and $[p^\tau_{Nt}, p^\tau_{Lt}, \mathcal{P}^\tau_t]$ denote the path of prices and pollution in the laissez-faire economy and under the policy τ, we compute the

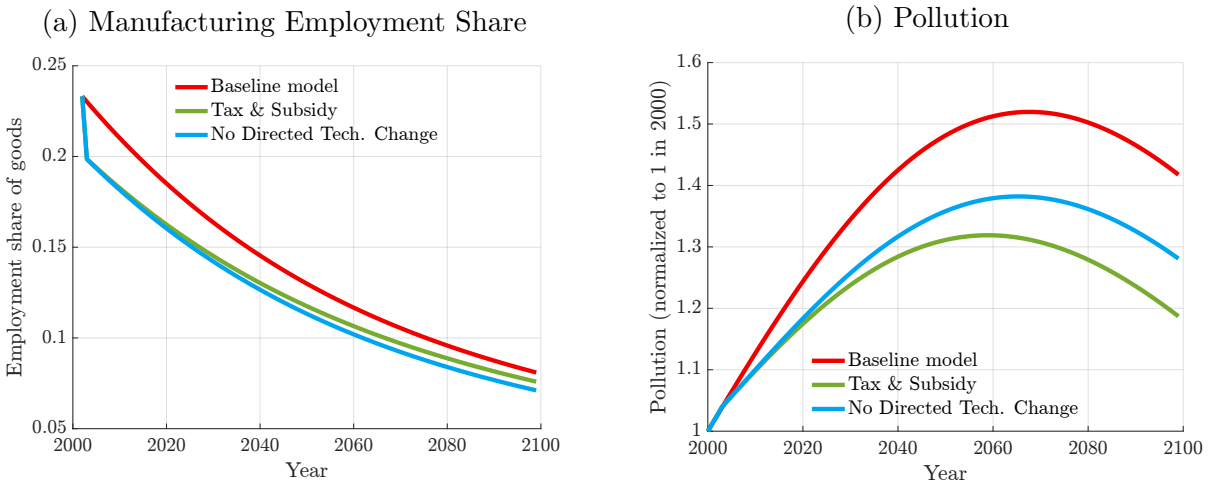
[16] An excise tax on the purchase of Y_G would have the same qualitative effect and is, in fact, identical to a carbon tax.

time t welfare impact of the policy, Λ_t as the solution to

$$\mathcal{V}\left(e, p_{Nt}^{\tau}, p_{Lt}^{\tau}, \mathcal{P}_t^{\tau}\right) \equiv \mathcal{V}\left((1+\Lambda_t)e, p_{Nt}, p_{Lt}, \mathcal{P}_t\right). \tag{28}$$

Hence, Λ_t measures the percentage change (equivalent variation) in income that the representative household would require to achieve the same utility in the status quo as under the counterfactual allocation implemented by the policy.[17]

Figure 5: The Impact of Environmental Policies



Notes: The figure shows the effect of the environmental policy described in the text. The red trajectories represent the laissez-faire benchmark. The green trajectories represent the policy intervention. The blue trajectories represent the policy intervention under the counterfactual scenario in which innovation remains at the benchmark. Panel (a) displays the employment share of services. Panel (b) displays pollution.

Calculating the welfare effect requires taking a stand on the disutility of pollution, $v(\mathcal{P})$. As we describe in Appendix B.3, we parametrize $v(\mathcal{P})$ as a convex polynomial and calibrate the parameters to match a value of the social cost of carbon of \$190 in 2020, which is in line with the estimates of the EPA.[18]

[17] Our welfare measure Λ_t captures the utility flow accruing to the representative agent. This approach avoids taking a stance on the representative agent's discount rate, which is not explicitly defined in the model, as agents make no intertemporal choices. However, note that our empirical strategy implicitly places some weight on the future, because we calibrate $v(\mathcal{P})$ to the social cost of carbon (which reflects the net present value of utility losses of an additional ton of carbon today). If, for instance, a more forward-looking decision-maker would assign a higher social cost of carbon, our methodology would infer a larger willingness to pay in terms of current consumption.

[18] Specifically, let Λ_t^{SCC} denote the social cost of carbon relative to GDP per capita. Using (28), we then calibrate $v(\mathcal{P})$ to ensure that $\mathcal{V}\left(e, p_{Nt}, p_{Lt}, \mathcal{P}_t + 1\right) \equiv \mathcal{V}\left((1-\Lambda_t^{SCC})e, p_{Nt}, p_{Lt}, \mathcal{P}_t\right)$. Hence, holding prices fixed, the representative consumer would be willing to give up a fraction Λ_t^{SCC} of average income to prevent overall CO_2 from rising by one unit.

As we show in Appendix B.3, we find that the welfare impact of the policy is time-varying. In the very short run, welfare declines because the misallocative cost of the policy outweighs the benefits of lower pollution. Over time, Λ_t rises due to two forces. First, since $v(\mathcal{P})$ is convex, the utility cost of pollution increases with its level, raising the value of being on the green path rather than the red one. Second, as the marginal utility of consumption declines, economic growth increases the dollar value of pollution reduction, further raising the consumption-equivalent gain. Quantitatively, we estimate a welfare gain of up to 9% by 2060, with about half attributable to the endogenous shift of innovation toward quality. This underscores the importance of quality-led growth for the overall effectiveness of environmental policy.

An important implication of our theory is that the welfare impact and hence the political support for environmental policies varies across the income distribution. First, due to diminishing marginal utility of consumption, richer households are more willing to accept lower market consumption in exchange for reduced emissions. Second, because of nonhomothetic preferences, the tax raises the relative price of goods-intensive necessities, placing a larger burden on poorer households, who devote a greater income share to such goods. Third, the shift of innovation from material productivity to quality disproportionately benefits the rich. For all these reasons, *ceteris paribus*, poorer households are less likely to support the policy than richer ones.

Our framework—in particular, the aggregation properties of PIGL preferences—is well-suited to assess such distributional effects and, therefore, informs how political support for environmental policy may vary across socioeconomic groups. As detailed in Appendix B.3, we can use use (28) to compute Λ_t across the empirical U.S. income distribution.[19] We find substantial heterogeneity: the poorest 5% *lose* 5.6% of income at the time of implementation, while the richest 1% *gain* 1.6% (see Table B-1 in Appendix B.3). Over time, the policy yields positive welfare effects for all individuals, but they remain "pro-rich". This underscores that decarbonization policies may generate significant political conflict, precisely because they shift resources toward high-quality goods consumed disproportionately by affluent households.

Research Wedge: An alternative policy directly subsidizes research directed toward quality. We consider a constant tax on A-directed research, rebated through time-varying subsidies to Q-directed research. For comparability, the tax is set to match the 2020 emissions level achieved under the other policy (goods-services wedge). As discussed in

[19] Formally, for a consumer with income x, the welfare change $\Lambda_t(x)$ solves $\mathcal{V}\left(x, p_{Nt}^{\tau}, p_{Lt}^{\tau}, \mathcal{P}_t^{\tau}\right) \equiv \mathcal{V}\left((1+\Lambda_t(x))x, p_{Nt}, p_{Lt}, \mathcal{P}_t\right)$.

Appendix B.3.4, the effects of this policy are similar to those of the previous one, though more back-loaded. By taxing quantity R&D, it reduces the growth rate of A, which *increases* the manufacturing employment share due to complementarity. As a result, pollution falls less in the short run, and welfare gains are smaller. In the long run, however, the decline in pollution is larger because A is lower.

Tipping Points: Research in climate science emphasizes the risk of crossing environmental tipping points—thresholds beyond which degradation triggers irreversible damage. Examples include the melting of polar ice sheets, thawing permafrost, and ecosystem collapse. These risks can be modeled as discontinuous utility losses that capture nonlinear damages (see, e.g., Lenton et al. (2008), Wagner and Weitzman (2015).[20] The existence of tipping points makes the *timing* of pollution reduction critical. Although emissions endogenously decline in the long run, the drop may come too late to avoid crossing $\bar{\mathcal{P}}$. In such cases, environmental policies such as the ones discussed above, might be essential to avert a disaster.

6 International Trade

So far we have considered the U.S. as a closed economy and abstracted from international trade. This might overstate the "cleansing effect" of the structural transformation if the decline in U.S. manufacturing is due to international outsourcing. In an important extension available in the working paper version of this paper—that could not be included in this submission because of the journal's article length constraint—we extend our model to an open economy setting that allows us to speak to this question of carbon leakage whereby manufacturing production and emissions are shipped abroad. The details are available in our accompanying working paper (Aghion et al., 2025).

We consider an economy with two countries that produce and trade differentiated goods. One country (China) is poorer and has a comparative advantage in manufacturing, while the other (the U.S.) is richer and specializes in services. We allow for trade imbalances to calibrate the model to the data. Our theory predicts that trade and economic integration have opposite effects on emissions in the two countries. China deepens its specialization in manufacturing, causing innovation to shift toward that sector. At the same time, the gains from trade make Chinese consumers more affluent, shifting their demand toward service-intensive luxury goods. The net effect on emissions is ambiguous.

[20] More formally, $\lim_{\mathcal{P}\to\bar{\mathcal{P}}} v'(\mathcal{P}) = \infty$, where $\bar{\mathcal{P}} < \infty$ is an environmental disaster threshold.

In contrast, the US moves further into services, leading to cleaner production and lower emissions.

Crucially, over a multi-decade horizon—including the critical phase when global emissions peak—our calibrated model predicts that free trade reduces overall emissions by hastening the global shift toward cleaner, service-oriented economies. This quantitative result remains robust even when the poorer country is assumed to use more polluting technologies for the same type of industrial production.

7 Conclusion

In this paper, we develop and quantify a growth model where: (i) consumers have nonhomothetic preferences with respect to quality; (ii) the direction of technological progress—between material productivity versus improving quality—is endogenous; (iii) quality is service-intensive and the production of services has a lower environmental footprint than that of material goods.

The model delivers several novel insights. First, as the economy develops, consumers progressively shift their demand toward quality-sensitive goods, which in turn redirects innovation away from increasing material productivity and toward enhancing quality. Second, the transition to quality-driven growth leads to a decline in measured GDP growth and, ultimately, to convergence to a finite level of material production—an endogenous *limit to growth*—even as quality-adjusted GDP and welfare continue to rise. Third, environmental policy can speed up the ongoing structural transformations. However, these policies tend to have adverse distributional effects.

Future research can extend our analysis in several directions. First, one could consider additional factors that affect the environmental footprint of goods at different quality levels. In the paper, we link environmental impact to the service share (which is observable), but this does not capture the full variation in emissions across goods. Moreover, one could consider how the emergence of energy-intensive technologies to produce services like AI might change the picture. Second, we abstract from the fact that rich and poor consumers typically consume different quality versions of the same good category—for example, inexpensive restaurants versus gourmet alternatives. Third, while we emphasize weightless luxury goods, degrowth proponents have long advocated a move toward relational goods and non-consumption-oriented activities. Such a transition, whether spontaneous or policy-induced, could have implications for environmental sustainability.

Finally, future research could aim to quantify the extent to which the recent slowdown in measured growth reflects an accelerated shift toward quality-based innovation. These avenues could deepen our understanding of how weightless growth can support both prosperity and environmental goals.

References

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

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

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi: 2012b, 'The network origins of aggregate fluctuations'. *Econometrica* **80**(5), 1977–2016.

Acemoglu, D. and V. Guerrieri: 2008, 'Capital deepening and nonbalanced economic growth'. *Journal of Political Economy* **116**(3), 467–498.

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.

Aghion, P., T. Boppart, M. Peters, M. Schwartzman, and F. Zilibotti: 2025, 'A Theory of Endogenous Degrowth and Environmental Sustainability'. NBER WP 33634.

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.

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.

Brock, W. A. and M. S. Taylor: 2005, 'Economic growth and the environment: a review of theory and empirics'. *Handbook of Economic Growth* **1**, 1749–1821.

Buera, F. J. and J. P. Kaboski: 2009, 'Can traditional theories of structural change fit the data?'. *Journal of the European Economic Association* **7**(2-3), 469–477.

Comin, D., D. Lashkari, and M. Mestieri: 2021, 'Structural change with long-run income and price effects'. *Econometrica* **89**(1), 311–374.

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'. Discussion Paper CFM-DP2025-09, 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.

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 international 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.

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

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

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: 2014, 'Growth and structural transformation'. In: *Handbook of Economic Growth*, Vol. 2. Elsevier, pp. 855–941.

Hickel, J., G. Kallis, T. Jackson, D. W. O'Neill, J. B. Schor, J. K. Steinberger, P. A. Victor, and D. Ürge-Vorsatz: 2022, 'Degrowth can work — here's how science can

help'. *Nature* **612**, 400–403.

Jaravel, X.: 2024, 'Distributional consumer price indices'. CEPR Discussion Paper DP19802.

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

Klenow, P. J., E. Pastén, and C. Ruane: 2024, 'Carbon taxes and misallocation in Chile'. Technical report, Mimeo, Stanford.

Lenton, T. M., H. Held, E. Kriegler, J. W. Hall, W. Lucht, S. Rahmstorf, and H. J. Schellnhuber: 2008, 'Tipping elements in the earth's climate system'. *Proceedings of the National Academy of Sciences* **105**(6), 1786–1793.

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. Behrens 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'. *Economic Journal* **101**(407), 920–937.

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.

Shapiro, J. S. and R. Walker: 2018, 'Why is pollution from US manufacturing declining? The roles of environmental regulation, productivity, and trade'. *American Economic Review* **108**(12), 3814–3854.

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.

Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K. Averyt, M. Tignor, and H. Miller: 2007, *Climate change 2007: the physical science basis. Contribution of working group I to the fourth assessment report of the intergovernmental panel on climate*

change. Cambridge University Press.

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.

Wagner, G. and M. L. Weitzman: 2015, *Climate shock: the economic consequences of a hotter planet*. Princeton, NJ: Princeton University Press.

A Empirical Analysis

A.1 Data

Consumer Expenditure Survey (CEX) The CEX is a nationwide household survey conducted by the U.S. Bureau of Labor Statistics on consumer spending. It consists of two components: the Interview Survey, which collects data on major and recurring expenses, and the Diary Survey, which covers smaller and frequently purchased items.

Our analysis relies primarily on the Interview Survey, which accounts for 80% to 95% of total household spending. We restrict the sample to households with heads aged 25–64, excluding students, retirees, and military personnel. Using all quarters from the 2002 release, we obtain a sample of approximately 12,000 households.

Consumption and income data are organized by Universal Classification Codes (UCC). We exclude UCCs related to assets and gifts, as well as observations with negative expenditures (e.g., due to reimbursements from government programs).

Input–Output Tables The Input–Output (I–O) Table, published every five years by the Bureau of Economic Analysis (BEA), provides an account of industry-level production and intermediate input use. We focus on the 2002 Use Table, which includes roughly 400 industries and offers a breakdown of value-added components and intermediate inputs.

We emphasize the 2002 dataset due to its concordance with the CEX categories, as established by Levinson and O'Brien (2019). To ensure consistency, we exclude scrap and non-comparable imports, which are difficult to classify as goods or services. We initially group I–O codes 1–3 as goods and codes 4–8 (plus government spending) as services. This classification allows us to compute the share of services used by each industry.

Table A-1: 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), compiled by the

U.S. Environmental Protection Agency every three years, is a comprehensive source of air emissions data. 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 this paper, we primarily rely on total emissions coefficients by industry from Levinson and O'Brien (2019).

We focus on five major air pollutants: PM10 (particulates smaller than 10 microns), VOCs (volatile organic compounds), NOx, SO_2, and CO. Because these pollutants are measured in different units, we use pollutant fixed effects when aggregating. Each pollutant's industry-level emission coefficient reflects pollution per dollar of final output, including both direct and input-related emissions. Combining these with CEX data allows us to estimate household-level emissions from consumption.

We also incorporate CO_2 emissions per dollar by industry from the EPA's Supply Chain Greenhouse Gas dataset. This source reports CO_2-equivalent (CO_2e) emissions in kilograms per dollar, capturing both direct emissions from production and upstream emissions from input acquisition, distribution, and storage.

Inflation between luxuries and necessities To construct relative inflation between luxuries and necessities, we use BLS price data processed by Jaravel (2024). As shown in that study, the most granular price classification is the entry-level item (ELI), for which we have data on 205 items from 1983 to 2023. We match each ELI to its corresponding expenditure in the CEX using the BLS concordance between UCCs and ELIs.

The CEX is available in two formats: (i) quarterly interviews, which report expenditures for broad bundles over the previous quarter, and (ii) daily interviews, which record weekly expenditures. To fully utilize both the BLS price and CEX expenditure data, we use the published CEX tables at the UCC level for 2002, which aggregate data across both interview types and report total annual expenditures.

We then link each UCC to its corresponding industry in the I-O tables and assign a service share. Luxuries and necessities are defined as UCCs with service shares above and below the median, respectively, weighted by 2002 expenditure. We compute annual average inflation for both groups from 2002 to 2017, weighting by UCC-level expenditure, and use this to calculate relative inflation. We also compute inflation paths from 1983 to 2023, again weighting by 2002 expenditures, to obtain the long-run relative inflation path between luxuries and necessities.

Service Share Construction To construct a measure of the service share for each industry in the Input–Output (I–O) tables, we use the Total Requirements Table (TRT). This table reports the value of intermediate inputs used along each industry's supply chain,

allowing us to measure the share of value attributed to goods and services (Medeiros and Howels III, 2017).

However, the TRT is based on gross output, so intermediate input values may include embedded inputs from earlier production stages, leading to double counting. In addition, it excludes costs incurred to deliver goods to final consumers, such as wholesale, retail, and transportation margins. We address both issues in a two-step procedure.

Step 1: Correcting for double counting Following Levinson and O'Brien (2019), let the linear production structure be defined by $X = CX + Y$, where $Y_{1\times n}$ is the vector of household consumption, $X_{1\times n}$ is total output, and $C_{n\times n}$ is the Direct Requirements Table with entries c_{ij} indicating the dollar input from industry i required to produce one dollar of output in industry j. Solving yields $X = [I - C]^{-1} \times Y$, or $X = TY$, where $T := [I - C]^{-1}$ is the Total Requirements Table. Each column of T shows the gross output required across industries to deliver one dollar of final demand.

To correct for embedded intermediate inputs, we compute $T^{-1} = I - C$ and sum each column of T^{-1} to approximate value added per dollar of output. Let $V_{n\times n}$ be a diagonal matrix with these sums on the diagonal. We then define the adjusted Total Requirements Table as $T_{adj} := VT$, where each entry $t_{ij}^{adj} = v_i t_{ij}$ reflects the value-added share embodied in gross output. This adjustment removes double-counting across the supply chain.

Step 2: Incorporating final delivery costs To account for delivery costs (transportation, wholesale, and retail), we use the Bridge PCE Tables, which compare producer and purchaser prices. These tables do not disaggregate trade and transport costs by industry (e.g., air vs. truck transport). We allocate these margins across industries proportionally to their input usage shares from the TRT.

We compile this information into a matrix $B_{n\times n}$, with each column corresponding to an industry. Diagonal entries indicate the share of purchase value attributed to production; off-diagonal entries for transportation, retail, and wholesale are allocated proportionally, while all other entries are zero. We then compute:

$$T_{Badj} := T_{adj} B,$$

which incorporates both value-added and final delivery costs into each industry's effective production structure.

Given the matrix T_{Badj} we compute the final service share at the industry level. We normalize each column of T_{Badj} by its total (i.e., the column sum), yielding $\hat{T}_{Badj}$. We

define a vector $S_{n\times 1}$, where $S_i = 1$ if the first digit of industry i's NAICS code is greater than 3 (indicating a service), and 0 otherwise. The service share vector $\hat{S}_{1\times n}$ is then given by:

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

Each entry of $\hat{S}$ gives the share of value added and delivery cost attributable to services in the total purchase price of a final good.

A.2 Services and Emission Intensity: Cross-country and Cross-county Variation

This section contains more details on estimating equation (1). The results are reported in Table A-2. Column 1 shows a significant negative correlation between emissions intensity and the service employment share: Columns 2 and 3 add controls for GDP per capita, total population, and land area; the relationship remains robust. Column 3 refers to the specification shown in Panel (d) of Figure 1. Column 4 includes country fixed effects, so that β is identified from within-country variation in service employment and emissions over time. Although the coefficient declines in magnitude, a substantial effect remains: a one-point increase in service employment is associated with a 2.2% decline in emissions intensity. Columns 5 to 8 replicate the analysis using cross-sectional data on counties *within* the U.S. As in the cross-country results, we find a significant negative relationship between service employment and emissions intensity.

A.3 Classification of Products: Luxuries versus Necessities

In our model, the parameter λ_k maps directly to the service share of final goods observed in the Input–Output table. To incorporate this structure consistently, we simplify the empirical classification to match the model's two-good framework: luxuries and necessities. Specifically, we classify final goods into two mutually exclusive groups—service-intensive ("luxuries") and goods-intensive ("necessities")—based on their service cost share.

We rank final products by their service cost share (which corresponds one-to-one with λ_k) and classify those above the median, weighted by expenditure, as service-intensive. This yields 189 "luxury" products and 218 "necessity" products. The average service cost share is 0.93 for luxuries and 0.41 for necessities (see Table A-3). In the model, we calibrate λ_L and λ_N to match these empirical moments.

Table A-2: Services and Pollution Intensity: Cross-regional evidence

	Across countries				Across counties within US			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Service Empl. Sh.	-5.381***	-5.380***	-4.905***	-2.181***	-2.799***	-1.901***	-1.586***	-1.406***
	(1.005)	(0.993)	(0.780)	(0.183)	(0.534)	(0.554)	(0.403)	(0.267)
log GDP pc		-0.002	-0.015	-0.503***		-0.991***	-0.643***	-0.381***
		(0.078)	(0.079)	(0.023)		(0.176)	(0.174)	(0.094)
Year FE	Yes	Yes	Yes	Yes				
Agriculture 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	5004	5004	4940	4940	3139	3139	3139	3139
R^2	.23	.23	.293	.913	.119	.158	.25	.315

Notes: The table reports coefficients from regressions of the log of emissions over GDP on the service employment share and log GDP per capita. Columns 1–4 are annual cross-country panel regressions for 1990–2019. Columns 5–8 are cross-sectional regressions across U.S. counties, using 2017 emissions data and 2010 Census characteristics. All specifications control for the employment share in agriculture. Columns 4 and 8 additionally include country and state fixed effects, respectively.

Table A-3: 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 CEX products classified as Necessities and Luxuries in 2002, along with their expenditure-weighted service cost shares.

A.4 Joint Identification of β and μ

While we describe their identification separately in the main text, estimating the relative price effect β depends on the overall price index and hence on the quality adjustment governed by μ. Conversely, calibrating μ to match measured relative inflation requires a simulation of the economy's path of quality and quantity TFP that, in turn, depends on the value of β. We therefore solve for a fixed point (β, μ) that jointly matches both the relative price effect and observed relative inflation.

Building a price series to estimate β Estimating β from Equation (24) requires the true relative price $\frac{p_{N,t}}{p_{L,t}}$, not the observed one $\frac{p_{N,t}}{p_{L,t}^{BLS}}$. Recovering this true relative price requires estimates of A_t and Q_t, since $\frac{p_{N,t}}{p_{L,t}} = Q_t \cdot \frac{\psi_N(A_t)}{\psi_L(A_t)}$.

We estimate A_t using nominal expenditures relative to necessity prices, $\frac{e_t}{p_{N,t}}$, which, in our framework, are unaffected by the measurement of quality. Specifically, $\frac{e_t}{p_{N,t}} = \psi_N(A_t)^{-1}$, where $\psi_N(\cdot)$ is calibrated from IO-based estimates. We then infer Q_t using the measured relative price, since

$$\frac{p_{N,t}}{p_{L,t}^{BLS}} = Q_t^{\mu} \cdot \frac{\psi_N(A_t)}{\psi_L(A_t)}.$$

This yields the following expression for the true, quality-adjusted relative price:

$$\frac{p_{N,t}}{p_{L,t}} = \left(\frac{p_{N,t}}{p_{L,t}^{BLS}}\right)^{1/\mu} \left(\frac{\psi_L(A_t)}{\psi_N(A_t)}\right)^{(1-\mu)/\mu}. \tag{A-1}$$

We use this quality-adjusted price series to estimate the coefficient $-\varepsilon(1-\beta)$ in Equation (24). Since this adjustment depends on μ, the estimate of β is contingent on a value of μ.

Estimating μ To identify μ, we utilize an iterative procedure. For each μ, we apply the quality adjustment in Equation (A-1), estimate β using Equation (24) and simulate the economy's path, including the implied series of measured relative inflation $\left(\frac{p_{L,t}^{BLS}}{p_{N,t}}\right)_{\text{Sim}}$. The correct value of μ is the one for which the simulated relative inflation matches the observed series $\left(\frac{p_{L,t}^{BLS}}{p_{N,t}}\right)_{\text{Data}}$. Since we have reliable data from 2002 to 2017, we target the average relative inflation over this period.

We report the estimates of Equation (24)—and the corresponding value of β under our calibrated μ—in Table A-4.

Table A-4: Estimating β from Equation (24)

	$\ln(\vartheta_{N,t}-\omega)+\varepsilon\ln\left(\frac{e_t}{p_{N,t}}\right)$
$\ln\left(\frac{p_{N,t}}{p_{L,t}}\right)$	-0.209***
	(0.054)
R^2	.226
μ	0.67
β	.343
N	37

Notes: The table reports the coefficient from regressing $\ln(\vartheta_{N,t}-\omega)+\varepsilon\ln\left(\frac{e_t}{p_{N,t}}\right)$ against the log of the ratio between the price of necessities and luxuries under $\mu = 0.67$. We use the Personal Consumption Expenditure from the BEA and our classification between necessities and luxuries to build the time series of the expenditure and the expenditure share in necessities. We additionally use the data from Jaravel (2024) to build the series for the price of necessities and luxuries between 1983 and 2019. The underlying coefficient implies $\beta = 0.343$.

A.5 Construction of Figure 3

This section details the construction of Figure 3. Let σ_k denote the goods cost share for product k. Since final goods are priced at a constant markup $\frac{\xi}{\xi-1}$, total spending on goods relative to revenue when producing x_k units of product k is:

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

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

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

Taking logs, the log emission intensity becomes:

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

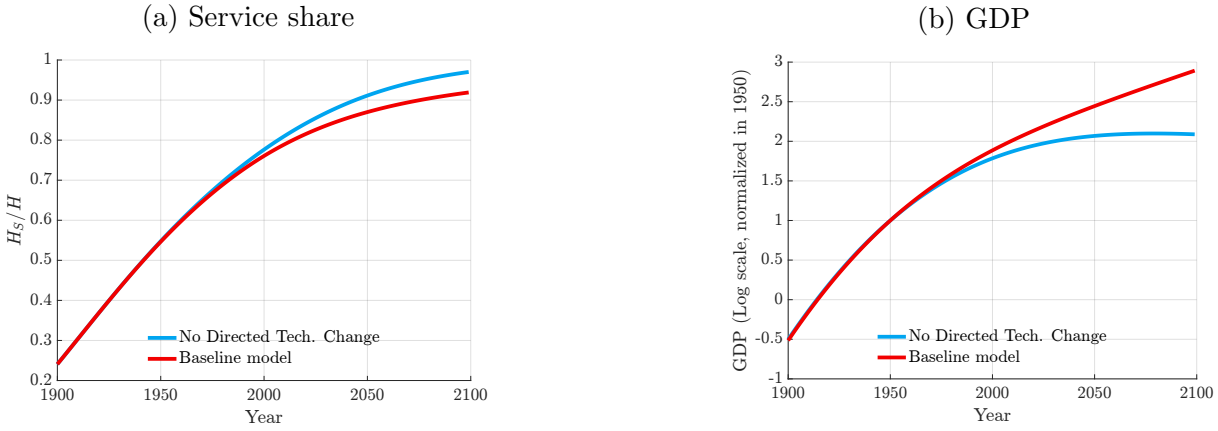
where d_t is an aggregate term that does not vary with k. Hence, the model predicts a log-linear relationship between product-level emission intensity and the service share. To construct Figure 3, which focuses on the cross-sectional variation, we chose d_t to match the median service intensity observed in the data. To construct model-predicted emissions, we consider a set J of final goods, where each $j \in J$ corresponds to an industry indexed by its service cost share $1-\sigma_j$.

A.6 Additional Quantitative Results

The Importance of Quality-Led Growth (Section 4.4) In Section 4.4, we examined the role of directed technical change by comparing our baseline model with a counterfactual in which the allocation of researchers remains fixed. As shown in Figure 4, this leads to substantially higher emissions. Figure A-1 reports the implications for the service employment share (left panel) and GDP per worker (right panel).

Two features stand out. First, faster quality growth is not necessary for the service employment share to rise. In fact, service employment also increases in the absence of directed technical change—and somewhat more rapidly. This is because faster quantity growth raises incomes and induces technological substitution (the "Baumol effect"). Second, GDP per worker grows slightly faster under directed technical change, as the direction of innovation responds to shifts in market size, boosting welfare-relevant production.

Figure A-1: The Impact of Directed Technical Change



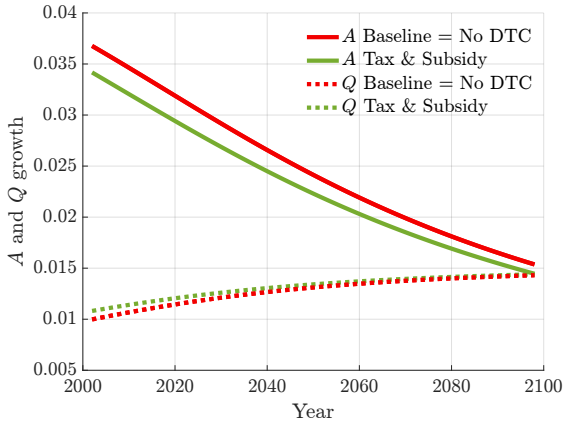
Notes: Panel (a) shows the employment share in services in the baseline model (red line) and in a counterfactual without directed technical change (blue line). Panel (b) shows GDP per worker in the baseline economy (red) and the counterfactual (blue).

The Effect of Environmental Policies on the Direction of Innovation (Section 5) One important implication of the environmental policy discussed in Section 5 is the redirection of innovation. By taxing material inputs, it discourages quantity-innovation and increases the incentives to increase quality.

In Figure A-2 we plot the resulting paths for quality and quantity growth in our baseline, laissez-faire economy and under the policy. Quantity growth is uniformly lower under the policy and in the long-run, the *level* of quantity TFP converges to a lower level. By contrast, the rate of quality growth is higher.

Figure A-2: Environmental Policies and Directed Technological Change

(a) Quantity and Quality Growth



Notes: The figure shows the effect of the environmental policy described in the text on quality and quantify growth. The red trajectories represent the laissez-faire benchmark. The green trajectories represent the policy intervention.

B Theory

B.1 Derivations of Theoretical Results

This section contains detailed derivations of some expressions we use in the main text.

Derivation of Eq. (4) [**Expenditure Shares**] Consider the indirect utility function in Eq. (3). The expenditure share for good k follows from Roy's identity. Standard algebra yields:

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

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 $\beta_j\equiv\varsigma_j/\varepsilon+\omega_j$, and letting $\phi_k=\tilde{\varsigma}\varsigma_k$ yields equation (4) in the text.

Microstructure I: Prices and Profits of Final Good Firms We derive the expression for prices and profits earned by firms producing varieties of final goods. Consider the firm producing variety i in the final good sector $j\in\{1,2,...,J\}$. Profit maximization implies that the monopolist will set a markup over its marginal cost

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

where $\mu=\frac{\xi}{\xi-1}$ is the optimal mark-up. The (quality-adjusted) price index for good j is therefore

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

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 imply that the profits ac-

cruing to the monopolist are

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

where $p_j Y_j$ is aggregate spending on good j. The aggregate profits generated by firms operating in the final good sector $j \in \{1,2,...J\}$ are therefore equal to

$$\Pi_j = \int_0^1 \pi_{ij}\; di = \frac{1}{\xi} p_j Y_j. \tag{B-2}$$

Microstructure II: Prices and Profits of Intermediate Good Firms 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

$$p_{iG} = \mu\; \frac{w}{A_i} \qquad \text{and} \qquad p_{iS} = \mu\; w.$$

The price indices in sector $k \in \{S,G\}$ are $p_S = \mu\, w$ 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. Total sectoral profits are given by

$$\Pi_G = \int_i \pi_{iG} di = \int_i \frac{1}{\xi}\mathcal{D}_G \left(\frac{A_i}{A}\right)^{\xi-1} di = \frac{1}{\xi}\mathcal{D}_G, \quad \Pi_S = \int_i \pi_{iS} di = \int_i \frac{1}{\xi}\mathcal{D}_S di = \frac{1}{\xi}\mathcal{D}_S. \tag{B-3}$$

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

$$\mathcal{D}_G \equiv \sum_{j=1}^{J} \sigma_j\; \frac{\xi-1}{\xi}\vartheta_j e \quad \text{and} \quad \mathcal{D}_S \equiv \sum_{j=1}^{J} \left(1-\sigma_j\right)\; \frac{\xi-1}{\xi}\vartheta_j e. \tag{B-4}$$

Here, $\vartheta_j e$ is total spending on product j, a fraction $\frac{\xi-1}{\xi}$ of which is paid to the variable factors G and S, and σ_j is the cost share of goods for product j. Summing over all final goods j yields total spending on goods, $\mathcal{D}_G$. The intuition for $\mathcal{D}_S$ is similar.

Total Income and Spending. We derive Eq. (12) and the expression $e = \left(\frac{\xi}{\xi-1}\right)^2$ used in the text. The representative household is the recipient of both labor income and overall profits.

Total profits equal $\Pi = \Pi_G + \Pi_S + \sum_j \Pi_j$, 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 accrue to innovative and parasitic researchers.

Using (B-2) and (B-3), we obtain

$$\Pi = \frac{1}{\xi}\left(\sum_j p_j Y_j + \mathcal{D}_G + \mathcal{D}_S\right) = \frac{1}{\xi}\left(e + \frac{\xi-1}{\xi}e\right) = \frac{1}{\xi}\left(1 + \frac{\xi-1}{\xi}\right)e. \tag{B-5}$$

Labor income is given by

$$w = \frac{\xi-1}{\xi}\left(\mathcal{D}_G + \mathcal{D}_S\right) = \left(\frac{\xi-1}{\xi}\right)^2 e, \tag{B-6}$$

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, $\Pi = \frac{2\xi-1}{\xi-1}w$. Hence, as required total income is equal to total spending

$$\text{Income} = w + \Pi = \left(\frac{\xi-1}{\xi}\right)^2 e + \frac{1}{\xi}\left(1 + \frac{\xi-1}{\xi}\right)e = e. \tag{B-7}$$

Note also that (B-6) implies that $e = \left(\frac{\xi}{\xi-1}\right)^2 w = \left(\frac{\xi}{\xi-1}\right)^2$, since $w = 1$.

Labor Market Clearing Now, consider the demand for labor. Labor market clearing for manufacturing workers requires that (see (B-4))

$$wH_G = \frac{\xi-1}{\xi}\mathcal{D}_G = \left(\frac{\xi-1}{\xi}\right)^2 \sum_j \sigma_j \vartheta_j = \left(\frac{\xi-1}{\xi}\right)^2 e \sum_j \sigma_j \vartheta_j, \tag{B-8}$$

because workers receive a share $\frac{\xi-1}{\xi}$ of overall spending on goods. This implies that $H_G = \sum_j \sigma_j \vartheta_j$. Similarly, $H_S = \sum_j (1-\sigma_j)\vartheta_j$.

Real Income $\Upsilon(A, Q)$. Equation (4) implies that expenditure shares can be written as

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

We now express Υ directly as a function of the state variables A and Q. Using the

expression for prices p_j in (B-1), 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

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

The Optimal Allocation of Research The value of directing research towards improving quality in final good j is given by

$$V_{Qj} = \eta_{Qj} R_{Qj}^{-\zeta} \int \pi_{ij} (\gamma Q_{ij}) \, di,$$

where $\pi_{ij} (\gamma Q_{ij})$ denotes the profits of providing variety i for good j at quality γQ_{ij}. Using the expression for equilibrium profits in (B-1), we obtain:

$$\int \pi_{ij} (\gamma Q_i) \, di = \frac{1}{\xi} p_j Y_j \gamma^{\alpha_j (\xi - 1)}.$$

Hence,

$$V_{Qj} = \eta_Q R_Q^{-\zeta} \frac{1}{\xi} p_j Y_j \gamma^{\alpha_j (\xi - 1)}. \tag{B-9}$$

Similarly, one obtains:

$$V_A = \eta_A R_A^{-\zeta} \int \pi_{iG} (\gamma A_i) \, di = \eta_A R_A^{-\zeta} \frac{1}{\xi} \gamma^{\xi - 1} \mathcal{D}_G, \tag{B-10}$$

where the last equation uses (B-3). Arbitrage 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-9) and (B-10) imply

that

$$\frac{R_Q}{R_A} = \left(\frac{\eta_Q}{\eta_A}\frac{p_L Y_L}{\mathcal{D}_G}\right)^{1/\zeta} = \left(\frac{\eta_Q}{\eta_A}\frac{\xi}{\xi-1}\frac{\vartheta_L}{\sigma_N\vartheta_N+\sigma_L\vartheta_L}\right)^{1/\zeta},$$

where the last equality uses that

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

Comparative Statics in Section 3.1 [Labor Market Equilibrium] Equation (15) implies that the service employment share is given by

$$H_S = 1 - \left[\sigma_N(A)\vartheta\left(A,Q\right) + \sigma_L(A)\left(1-\vartheta\left(A,Q\right)\right)\right]. \tag{B-11}$$

Because $\lambda_L > \lambda_N$, $\sigma_N(A) > \sigma_L(A)$. If $\phi > 0$, $\vartheta\left(A,Q\right)$ is decreasing in both Q and A. This implies that H_S is increasing in Q. Moreover, if $\rho < 1$, $\sigma_k(A)$ is decreasing in A. Hence, H_S is increasing in A. If $\rho > 1$, $\sigma_k(A)$ is increasing in A. Moreover, $\lim_{\phi\to 0}\vartheta_A\left(A,Q\right) = 0$, where the subscript denotes a partial derivative. This implies that H_S is decreasing in A if $\rho > 1$ and ϕ is sufficiently small. Finally, if $\phi = 0$ and $\rho = 1$, σ_k and ϑ are constant.

B.2 Proposition 1: The case of $\rho \leq 1$

In the main text, we proved Proposition 1 only for the case of complementarity. In this Section we prove Proposition 2 for the case of $\rho \leq 1$.

Proposition B-1. *Given initial conditions $\{A_0, Q_0\} \gg 0$, there exists $\bar{\delta} > 0$ such that, for all $\delta \in [0, \bar{\delta})$, the economy converges to a unique ABGP. If goods and services are (weak) substitutes, i.e., $\rho \geq 1$, the ABGP has the following features:*

1) *$g_A > 0$ and $g_Q > 0$*

2) *If $\rho > 1$, then $H_G \to 1$. If $\rho = 1$, $H_G \to (1-\lambda_N)\omega + (1-\lambda_L)(1-\omega)$*

3) *Suppose that $\mathcal{E}(Y_{Gt}, z_t) = \kappa_{\mathcal{E}} z^{-t} Y_{Gt}$ where $\kappa_{\mathcal{E}} > 0$ is a constant and $z > 1$ denotes the (exogenous) rate of abatement. As long as $z < 1 + g_A$, emissions grow exponentially in the ABGP.*

Proof. Suppose that $A_t \to \infty$. If $\rho > 1$, then $\sigma_k(A_t) \to 1$. If $\rho = 1$, $\sigma_k(A_t) = 1 - \lambda_k$. In either case, the R&D arbitrage condition implies $\frac{R_Q}{R_A} \to \Phi_{WS}(\rho)$, where

$$\Phi_{WS}(\rho) = \begin{cases} \left((1-\omega)\frac{\eta_Q}{\eta_A}\right)^{\frac{1}{\zeta}}, & \text{if } \rho > 1, \\ \left(\frac{1-\omega}{\omega(1-\lambda_N)+(1-\omega)(1-\lambda_L)}\frac{\eta_Q}{\eta_A}\right)^{\frac{1}{\zeta}}, & \text{if } \rho = 1. \end{cases}$$

Therefore, in the ABGP:

$$g_A = \left(\frac{1}{1+\Phi_{WS}(\rho)}R\right)^{1-\zeta}\eta_A(\gamma-1)-\delta, \quad g_Q = \left(\frac{\Phi_{WS}(\rho)}{1+\Phi_{WS}(\rho)}R\right)^{1-\zeta}\eta_Q(\gamma-1)-\delta.$$

For small δ, both $g_A > 0$ and $g_Q > 0$, establishing the result. The behavior of H_G in part (2) follows directly from the labor market clearing condition, $H_G = \sum_k \sigma_k \vartheta_k$. Part (3) follows from the fact that $H_{Gt} \to \overline{H}$ so that Y_{Gt} grows at the same rate as A_t. Overall emissions therefore satisfy $\mathcal{E}_t \propto (\frac{1+g_A}{z})^t$. □

B.3 Details for Section 5 (Environmental Policy)

In this section we provide more details for the results reported in Section 5.

B.3.1 Computing the Equilibrium with Environmental Policy

We compute these as the equilibrium prices of our economy where material inputs, Y_G are subject to a tax τ, i.e. the price of the G input is given by $(1+\tau)\mu w/A$. The tax receipts are used to subsidize the purchase of service inputs, i.e. the price of the S input is given by $(1-s)\mu w$. Given a fixed tax rate τ, we chose the time-varying subsidy rate s_t to ensure that the government budget clears. Specifically, set s_t according to

$$s_t = \tau \frac{\sum_{j\in L,N} \sigma_{jt}\vartheta_{jt}}{\sum_{j\in L,N}(1-\sigma_{jt})\vartheta_{jt}}. \tag{B-12}$$

In terms of implementation, we fix $\tau = 0.2$ and guess a value of $\{s_t\}$. We then compute the equilibrium path and check whether (B-12) is satisfied. We iterate over $\{s_t\}$ until (B-12) holds at each point in time.

Because the policy is budget neutral, nominal household spending is still equal to e. We denote the path of counterfactual prices $\widehat{p}_L$ and $\widehat{p}_N$ induced by the environmental policy τ. The paths of sectoral employment shares and overall pollution are displayed in

Panels (a) and (b) of Figure 5.

B.3.2 Measuring Welfare Changes

The indirect utility of a consumer with spending e facing prices p_N and p_L is given by

$$\mathcal{V}\left(e, p_N, p_L, \mathcal{P}\right) = \frac{1}{\varepsilon}\left(\frac{x}{p_L^{1-\omega} p_N^{\omega}}\right)^{\varepsilon} + \frac{\phi}{\varsigma}\left(\frac{p_L}{p_N}\right)^{\varsigma} - v(\mathcal{P}). \tag{B-13}$$

Now consider another path $\left[\hat{p}_N, \hat{p}_L, \hat{\mathcal{P}}\right]$. The welfare difference between these paths as measured by the relative consumption equivalent, Λ, relative to the status quo is implicitly defined by

$$\mathcal{V}\left(e, \hat{p}_N, \hat{p}_L, \hat{\mathcal{P}}\right) \equiv \mathcal{V}\left((1+\Lambda)e, p_N, p_L, \mathcal{P}\right).$$

Hence, Λ measures the percentage change in income e that the representative household would require to achieve the same utility in the status quo as under the counterfactual allocation. Using (B-13), we can solve for Λ as

$$1+\Lambda = \left[\left(\frac{p_L^{1-\omega} p_N^{\omega}}{\hat{p}_L^{1-\omega} \hat{p}_N^{\omega}}\right)^{\varepsilon} + \left(\frac{e}{p_L^{1-\omega} p_N^{\omega}}\right)^{-\varepsilon} \varepsilon\left[\frac{\phi}{\varsigma}\left(\left(\frac{\hat{p}_L}{\hat{p}_N}\right)^{\varsigma} - \left(\frac{p_L}{p_N}\right)^{\varsigma}\right) + v(\mathcal{P}) - v(\hat{\mathcal{P}})\right]\right]^{1/\varepsilon}. \tag{B-14}$$

Given estimates for the structural parameters $(\omega, \varepsilon, \phi, \varsigma)$ and the pollution cost function $v(.)$, we can use (B-14) to solve for Λ given the paths $[p_N, p_L, \mathcal{P}]$ and $\left[\hat{p}_N, \hat{p}_L, \hat{\mathcal{P}}\right]$.

To compute (B-14), we need to take a stand on the welfare costs of pollution encapsulated in the utility function $v(.)$. To do so, we link $v(.)$ to the social cost of carbon (SCC), that is the dollar value of the economic damages that result from emitting an additional ton of CO2.

To compute the SSC in our model, recall that $\mathcal{P}_t = (1-\varphi)\mathcal{P}_{t-1} + \mathcal{E}_t$. Because, in our model, preferences are static, consider a counterfactual allocation with the same prices but pollution is higher by 1 unit, i.e. $\widehat{\mathcal{P}} = \mathcal{P} + 1$. Eq. (B-14) then implies that

$$1+\Lambda = \left\{1 - \left(\frac{e}{p_L^{1-\omega} p_N^{\omega}}\right)^{-\varepsilon} \varepsilon\left[\left(v(\mathcal{P}+1) - v(\mathcal{P})\right)\right]\right\}^{1/\varepsilon}. \tag{B-15}$$

Hence, if pollution is higher, $\Lambda < 0$ because higher pollution corresponds to a utility and income *loss*.

For a given utility pollution disutility function $v(.)$, we can compute Λ according to

(B-15). The SCC are denominated in real dollars. Let $\overline{SCC}_t$ denote the social cost of carbon *relative to GDP per capita*. We then set

$$\overline{SCC}_t \equiv -\Lambda, \tag{B-16}$$

Hence, (B-16) implies that we are measuring the SCC, as a fraction of GDP per capita, by the percentage of spending, the representative agent would be willing to give up to not emit a marginal unit of pollution. We have "-" in (B-16) because the SCC is a cost.

We use this relationship to parametrize the disutility of pollution. Suppose that $v(\mathcal{P}) = v_0 \frac{1}{\iota+1}\mathcal{P}^{\iota+1}$. We then approximate

$$v(\mathcal{P}+1) - v(\mathcal{P}) \approx v'(\mathcal{P}) = v_0 \mathcal{P}^{\iota}. \tag{B-17}$$

Equation (B-16) then implies that

$$\overline{SSC}_t = 1 - \left\{1 - \left(\frac{e}{p_L^{1-\omega} p_N^{\omega}}\right)^{-\varepsilon} \varepsilon \left[(v(\mathcal{P}+1) - v(\mathcal{P}))\right]\right\}^{1/\varepsilon}. \tag{B-18}$$

For given ι, we can solve for v_0 as

$$v_0 = \frac{\left(1 - \left(1 - \overline{SSC}_t\right)^{\varepsilon}\right)\left(\frac{e}{p_{Lt}^{1-\omega} p_{Nt}^{\omega}}\right)^{\varepsilon}\frac{1}{\varepsilon}}{\mathcal{P}_t^{\iota}}. \tag{B-19}$$

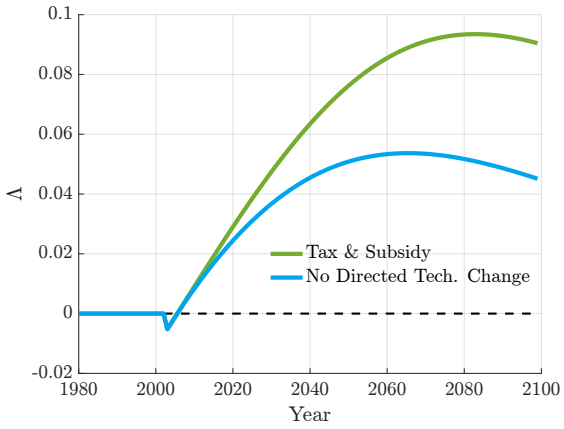
For our quantitative analysis we assume that $\iota = 0.3$. Given the value of v_0 in (B-19), we then compute the welfare consequences of the environmental policy, Λ, according to (B-14) and display them in Figure B-1.

B.3.3 Distributional Effects

Equation (B-14) highlights that Λ is a function of overall spending e due to the nonhomotheticity of preferences. Hence, individuals of different income differ in the welfare assessments between the paths $[p_N, p_L, \mathcal{P}]$ and $\left[\hat{p}_N, \hat{p}_L, \hat{\mathcal{P}}\right]$.

To capture this heterogeneity, we use (B-14) to compute Λ for different quantiles of the spending distribution. To do so, consider an individual on the qth quantile of the income distribution and let $x^q \equiv \Delta^q e$, i.e. Δ^q is q's income, *relative* to average income e. We can then define the welfare effects of a particular counterfactual on individuals at the

Figure B-1: The Welfare Impact of Environmental Policies



Notes: The figure shows the effect of the environmental policy described in the text on the welfare of the representative agent. The green line represents the policy intervention. The blue line represents the policy intervention under the counterfactual scenario in which innovation remains at the benchmark.

qth quantile of the income distribution as

$$\Lambda^q = \left[\left(\frac{p_L^{1-\omega} p_N^{\omega}}{\hat{p}_L^{1-\omega} \hat{p}_N^{\omega}} \right)^{\varepsilon} + \left(\frac{\Delta^q e}{p_L^{1-\omega} p_N^{\omega}} \right)^{-\varepsilon} \varepsilon \left[\frac{\phi}{\varsigma} \left(\left(\frac{\hat{p}_L}{\hat{p}_N} \right)^{\varsigma} - \left(\frac{p_L}{p_N} \right)^{\varsigma} \right) + v(\mathcal{P}) - v(\hat{\mathcal{P}}) \right] \right]^{1/\varepsilon} - 1. \tag{B-20}$$

In Table B-1 we report the results for Λ different quantiles of the US income distribution in 2003 (that is, immediately after the policy is put in place), in 2005, and in 2020. Table B-1 shows that the policy is pro-rich. People at the lower end of the income distribution lose from the policy (at least in the short run) while people at the upper end benefit. Overall, the cross-sectional dispersion in welfare assessments is between 3% and 5%.

Table B-1: The Unequal Effects of Environmental Policies

	Quantile of income distribution						
Year	5%	10%	25%	50%	75%	90%	99%
2003	−0.056	−0.035	−0.023	−0.010	−0.002	0.003	0.016
2005	−0.048	−0.028	−0.018	−0.006	0.002	0.007	0.018
2020	0.011	0.019	0.023	0.027	0.030	0.032	0.036

Notes: The table shows the welfare effects of the environmental policy as a function of incomes.

B.3.4 Taxing R&D

To put the overall impact of taxing the usage of goods into perspective, we also consider an alternative policy where the government taxes R&D directed to quantity TFP A and subsidizes quality-led growth. As before, we consider a constant tax $\tau^{R\&D}$ on research directed toward A and assume that the government uses these tax receipts to subsidize Q growth at rate $s_t^{R\&D}$, where $s_t^{R\&D}$ is chosen to balance the budget. To make the exercise comparable, we chose $\tau^{R\&D}$ so that the fall in pollution by 2020 under the R&D policy is the same as under the production tax τ of 20% considered above.

In Figure B-2 we plot the outcome of the R&D policy and compare it with both laissez-faire allocation and the production tax analyzed above. The structure of Figure B-2 is exactly the same as in Figure 5 in the main text. Panels (a) and (b) show that a tax on R&D directed towards quantity TFP reduces the growth rate of A and *increases* the employment share of goods. The latter effect is due to the complementarity between services and goods: less quantity TFP growth increases the price of goods and hence raises the cost and employment share. This discrepancy also explains the impact on overall pollution and welfare. In the short run, the effect of the reallocation of employment dominates and pollution is higher under the R&D tax. In the medium run, lower quantity TFP growth reduces the quantity of goods production, and pollution is lower. As a consequence, the welfare impact of the R&D tax is back-loaded. For the first four decades, the production tax yields a higher welfare gain, in the long run, welfare is higher under the R&D tax.

Figure B-2: R&D versus Production Taxes



Notes: The figure shows the effect of the R&D policy and the production tax on the employment share of goods (Panel (a)), the growth rates of A and Q (Panel (b)), overall pollution (Panel (c)) and the representative consumer's willingness to pay for the policy (Panel (d)).