NBER WORKING PAPER SERIES

SKIPPING THE FACTORY:
SERVICE-LED GROWTH AND STRUCTURAL TRANSFORMATION
IN THE DEVELOPING WORLD

Michael Peters
Youdan Zhang
Fabrizio Zilibotti

Working Paper 34692
http://www.nber.org/papers/w34692

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2026

We are grateful to Penny Goldberg for her excellent comments, to Doug Gollin for his insightful discussion, and to Charles Gottlieb for sharing some of his data. Nicolas Fajardo provided outstanding research assistance. We thank the Tobin Center for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Skipping the Factory: Service-Led Growth and Structural Transformation in the Developing World
Michael Peters, Youdan Zhang, and Fabrizio Zilibotti
NBER Working Paper No. 34692
January 2026
JEL No. O14

ABSTRACT

In today's developing world, many economies appear to bypass industrialization and transition directly from agriculture to services. The largest rise in service employment has occurred in non-tradable consumer services, such as retail and hospitality, especially in urban areas, where many cities resemble consumer hubs built around local demand. These patterns of growth raise fundamental questions: Can service-led growth sustain a growth of living standards over time? Is service-led growth inherently biased toward affluent urban consumers? What role should policy play? We propose a parsimonious general equilibrium framework that incorporates non-homothetic preferences and locally non-tradable consumer services. Our framework bridges macro and micro perspectives, and enables counterfactual and welfare analysis that accounts for both individual and spatial heterogeneity. We relate our framework to the recent literature and suggest several extensions and directions of future research.

Michael Peters
Yale University
Department of Economics
and NBER
m.peters@yale.edu

Fabrizio Zilibotti
Yale University
Department of Economics
and NBER
fabrizio.zilibotti@yale.edu

Youdan Zhang
Yale University
Department of Economics
youdan.zhang@yale.edu

Skipping the Factory: Service-Led Growth and Structural Transformation in the Developing World *

Michael Peters, Youdan Zhang, and Fabrizio Zilibotti

January 2026

Abstract

In today's developing world, many economies appear to bypass industrialization and transition directly from agriculture to services. The largest rise in service employment has occurred in non-tradable consumer services, such as retail and hospitality, especially in urban areas, where many cities resemble consumer hubs built around local demand. These patterns of growth raise fundamental questions: Can service-led growth sustain a growth of living standards over time? Is service-led growth inherently biased toward affluent urban consumers? What role should policy play? We propose a parsimonious general equilibrium framework that incorporates non-homothetic preferences and locally non-tradable consumer services. Our framework bridges macro and micro perspectives, and enables counterfactual and welfare analysis that accounts for both individual and spatial heterogeneity. We relate our framework to the recent literature and suggest several extensions and directions of future research.

1 Introduction

In the canonical narrative of structural transformation, development unfolds in three stages: at an early stage, agriculture is the dominant economic activity; then, agriculture sheds labor into a growing manufacturing sector; eventually, at a late stage, the bulk of economic activity shifts to services. This pattern, which is in line with the trajectories of Western Europe, the United States, and East Asia, has long shaped our expectations of how growth should proceed. Yet in much of today's developing world, structural change appears to be taking a different path. Since the early

*We are grateful to Penny Goldberg for her excellent comments, to Doug Gollin for his insightful discussion, and to Charles Gottlieb for sharing some of his data. Nicolas Fajardo provided outstanding research assistance. We thank the Tobin Center for financial support.

2000s, many developing countries, including fast-growing economies like India, have experienced a direct transition from agriculture to services, largely bypassing industrialization.

Moreover, the fastest growing component of the service sector consists of non-tradable consumer-facing activities such as retail, hospitality, and transportation, especially in urban areas where cities operate as demand-driven hubs. This pattern, sometimes labeled "premature deindustrialization" (Rodrik, 2016), raises concerns about the sustainability of this growth process over time. If manufacturing has historically been the main source of productivity gains, what are the consequences when developing economies do not industrialize? These developments also have distributional implications. When economic activity expands primarily in consumer cities sustained by local demand (Gollin et al., 2016), the gains may accrue disproportionately to affluent urban households, potentially leaving rural populations behind.

In this paper, we develop a parsimonious model to study these questions. Since the literature on the role of consumer services in structural change and growth is, to date, very limited, our aim is not to provide a comprehensive review. Rather, we present a theoretical framework designed to clarify the mechanisms that we view as central and show how the framework can be taken to the data. We relate our analysis to the existing literature in the concluding section.

Building on our earlier work (Fan et al., 2023), the paper reflects the central empirical role of non-tradable, consumer-facing services. The theory emphasizes two key mechanisms: spatial frictions and non-homothetic preferences. On the one hand, many consumer services are non-tradable and therefore depend on local endowments, technologies, and demand. On the other hand, these services are luxury goods, primarily consumed by richer households, which makes the local income distribution a key determinant of the sector's expansion. In addition, our model accommodates some fundamental drivers of growth documented in the literature: (i) productivity improvements across sectors; (ii) gains in allocative efficiency as labor moves from low-productivity informal activities to more productive formal ones; and (iii) the accumulation of human capital through rising educational attainment. Finally, the model captures multiple aspects of inequality by distinguishing individuals along three dimensions: skills (proxied by education), idiosyncratic labor productivity, and geographic location (e.g., whether they reside in rural or urban areas.)

Our approach provides a unified framework for quantifying both positive and normative effects of structural change. By integrating a spatial general equilibrium model with granular data on labor allocations, earnings, educational attainment, and sectoral composition across regions and over time, the framework bridges a gap in the literature between macroeconomic analyses of aggregate trends and microeconomic studies that often abstract from general equilibrium forces.

Topics Our framework naturally connects a range of topics that have often been analyzed in isolation. We organize the discussion around four themes.

1. Service-Led vs. Service-Biased Growth. We draw a sharp distinction between *service-led* and *service-biased* growth. Growth is *service-biased* when the expansion of services reflects rising incomes generated by forces *outside* the service sector, such as productivity growth in agriculture or manufacturing or broad-based human-capital accumulation, which shift demand toward services under non-homothetic preferences. In contrast, growth is *service-led* when the main driver of structural change is productivity improvements *within* the service sector itself, which raise real incomes and expand both the demand and the supply of local services. Our empirical analysis recovers sectoral productivity growth at a geographically granular level, allowing us to distinguish these mechanisms and to quantify their relative importance.

2. The Rise of Consumer Cities. As shown by Jedwab et al. (2022), many urban centers in developing countries (especially in sub-Saharan Africa) have recently expanded not around tradable industrial production (and the service activities providing inputs to the manufacturing sector) but around domestic consumption. Unlike the manufacturing-driven agglomeration observed in East Asia, these *consumer cities* grow through traditional service activities such as food retail, small-scale hospitality, personal care, and transportation, which are inherently local and highly responsive to income effects and urban density. Our framework is designed to analyze these patterns by incorporating spatial heterogeneity, non-tradable service demand, and frictions in formal employment, all of which are central to the emergence and persistence of consumer cities.

3. Informality. Our theory also sheds light on the role of informality, long viewed through the lens of Harris and Todaro (1970) as a temporary stopgap for rural migrants searching for formal employment. Microdata from sub-Saharan Africa challenge this narrow view: households reliant on consumer service jobs exhibit living standards broadly comparable to those in manufacturing, as measured by durable goods ownership, housing quality, and access to basic services. These descriptive findings suggest that service jobs may be more than a mere reserve of disguised unemployment. This raises broader questions: Is the growth of consumer services merely a symptom of stalled industrialization, or can it spark long-run development? Do urban service jobs create broad opportunities for all sectors of the population, or do they benefit mainly affluent city dwellers?

4. Sources of Productivity Growth in Consumer Services. Although consumer services are not typically regarded as fertile ground for technological progress, our framework highlights several mechanisms for productivity gains in services. First, in urban settings, labor market frictions and limited formal job creation generate misallocation as workers crowd into low-productivity informal services. The reduction of misallocation can be a powerful engine of growth through a mechanism similar to that highlighted for the industrial sector by Song et al. (2011). Institutional reforms, rising education, or macroeconomic improvements can ease these frictions and speed up transitions into higher-productivity formal jobs. Second, demand externalities can be important. The endogenous entry of formal firms for which overhead costs are significant is limited by the extent of local demand: at low incomes, demand is too low to justify their entry, so services are supplied by informal self-employed workers. As demand expands, more formal firms enter, thereby improving allocative efficiency, expanding product variety and quality, and further contributing to an expansion of local demand. Third, productivity gains can also arise within the informal sector itself: advances in mobile connectivity, transportation (for example, cheaper motorbikes), and access to basic financial services enhance efficiency and incomes for informal providers.

These shifts may not immediately lead to formalization but can nonetheless improve livelihoods and reduce poverty. Our framework allows the separate identification of these distinct sources of productivity growth.

For pedagogical reasons, we introduce the framework in a stripped-down form that abstracts from several important forces. We then extend it along several dimensions. First, we enrich the tradable sector by distinguishing agriculture, manufacturing, and business services such as finance and IT. Next, we embed the model in a regional and international trade environment, allowing for specialization and terms-of-trade effects. We then allow sectors to differ in skill intensity and show how variation in skill supplies can drive structural change. Finally, we incorporate internal migration, linking the framework to the literature on rural–urban mobility, spatial wage gaps, and labor misallocation.

Application: Sub-Saharan Africa in the XXIst Century While our earlier work studied the recent economic development of India (Fan et al., 2023), this paper shifts the focus to a set of fast-growing sub-Saharan African economies, whose economic activity is not dominated by natural resources like oil. These economies are geographically clustered, part in East Africa and part in West Africa. Altogether, they have a population of about 500 million people in 2024, which represents one-third of the total African population. In these economies, similar to India and in contrast to many East Asian countries, urbanization and the decline in agriculture's employment share have not been accompanied by significant industrialization but by a rise of traditional consumer services such as retail, hospitality, and local transport.

As a showcase of our methodology, we calibrate the model to salient empirical moments from this sample of African countries and study the role of productivity growth in consumer services for rising living standards. We find that improvements in service productivity account for a substantial share of welfare gains, especially in urban areas. This pattern is reminiscent of India, as documented in our earlier work. We contrast this African experience with Vietnam, which, since the late 1990s, has followed a more conventional industrialization-led path. Using the same empirical methodology, we find that improvements in living standards in Vietnam have relied much more heavily on productivity growth in tradable manufacturing and producer services.

Summary Our goal for this paper is twofold. On the one hand, it provides a coherent framework to understand how different sources of growth shape the structure of employment and the distribution of welfare. On the other hand, it shows how the model can be disciplined using microdata moments on labor allocation, income, expenditure, and geography that are increasingly available for developing countries.

We view the application to sub-Saharan Africa as an illustrative exercise that yields informative insights into the role of productivity growth in the service sector. Determining whether, and under what conditions, services can become engines of growth is a central policy challenge for countries shaping their development trajectories in the twenty-first century. Our framework provides a foundation for future research aimed at informing policy design in this domain.

Road Map The paper proceeds as follows. Section 2 documents empirical patterns of structural transformation and service expansion in a set of fast-growing African economies using detailed household and labor market data. Section 3 presents the baseline version of our theoretical framework and contains the main analytical result, which derives the aggregate demand system over sectoral value added from the individual-level demand system defined over final goods. This section also provides a step-by-step guide to calibrating the framework using both micro and macro data from sub-Saharan Africa.

Section 4 incorporates labor market frictions and the process of formalization through the entry of formal service providers, and shows how these mechanisms can operate as sources of productivity growth in services. Section 5 discusses various extensions of the baseline model. Section 6 reviews the related literature and situates our contribution within it. Section 7 concludes with a discussion of promising avenues for future research. An Appendix and Supplemental Material available on our websites provide detailed derivations of all theoretical results as well as additional information on the empirical analysis.

2 Structural Change in Sub-Saharan Africa

In this section, we document that structural change in most of today's developing world takes the form of an expanding service sector, particularly consumer services (henceforth, CS) concentrated in urban centers. We mainly focus on two geographic clusters of high growth within Africa. Specifically, we consider countries located south of the Sahel that (i) recorded average annual GDP growth of at least 3.5% between 2000 and 2019, (ii) were not classified as oil exporters by the IMF, and (iii) did not experience prolonged conflict, such as Somalia or the Democratic Republic of Congo.[1]

This selection yields two fast-growing clusters in West and East Africa, shown in Figure 1.[2] The East African cluster comprises Ethiopia, Kenya, Malawi, Mozambique, Rwanda, Tanzania, and Uganda, with a combined population of approximately 350 million in 2024. The West African cluster includes Benin, Burkina Faso, Côte d'Ivoire, Ghana, Guinea, Guinea-Bissau, Senegal, Sierra Leone, and Togo, home to around 150 million people in 2024. We contrast these experiences with that of Vietnam, an export-oriented, rapidly industrializing economy with a population of about 100 million.

All East African countries in our sample are classified as non-resource-intensive by the IMF, except Tanzania. Even in Tanzania, natural resources play a modest role: mining contributes just 3.3% of GDP, though tourism is a significant sector. In West Africa, four of the nine selected countries are classified as resource-intensive, but only Guinea and Sierra Leone are heavily reliant on mineral exports.[3] We emphasize the limited role of natural resources because a common narrative attributes Africa's economic revival primarily to rising commodity prices. By contrast, growth in our sample is unlikely to be the mere effect of a persistent resource revenue boom.

We structure our analysis along three dimensions. First, we document the substantive rise of CS, particularly in urban centers, but not limited to them. We do so using both aggregate and micro data. Second, we provide suggestive evidence that CS jobs are not necessarily worse than available alternatives. Using microdata on the ownership of durable goods, we show that

[1] We use the International Monetary Fund (2024) classification, which divides countries into Oil Exporters, Other Resource-Intensive Countries, and Non-Resource-Intensive Countries. We exclude oil exporters such as Angola, Cameroon, Chad, and Nigeria, as well as small island economies including Cabo Verde, Mauritius, São Tomé e Príncipe, and Seychelles.

[2] Namibia also meets the growth threshold, having grown at an average annual rate of 3.7%. However, its economy is heavily resource-dependent, with mining accounting for 25% of GDP, and does not belong to a broader high-growth region. We therefore exclude it.

[3] Excluding these two countries, which together account for just 15% of the Western cluster's population, does not affect any of our results.

Figure 1: The Countries in Our Sample

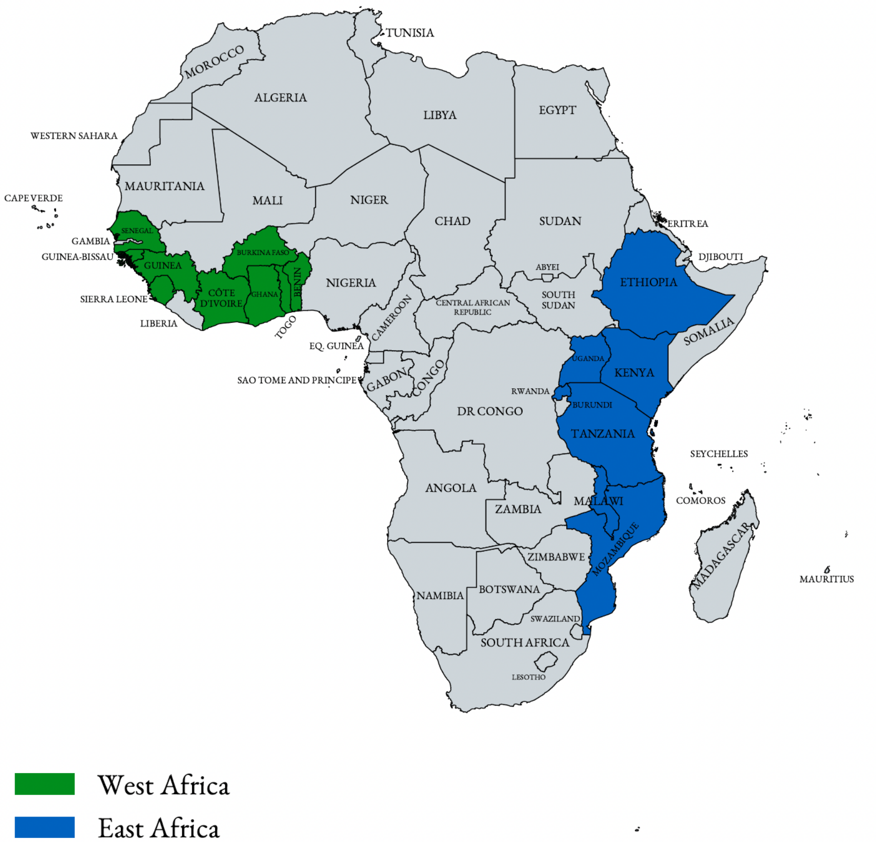

Notes: This figure displays the set of African countries included in our analysis. The green-shaded regions correspond to nine West African countries: Benin, Burkina Faso, Côte d'Ivoire, Ghana, Guinea, Guinea-Bissau, Senegal, Sierra Leone, and Togo. The blue-shaded regions indicate seven East African countries: Ethiopia, Kenya, Malawi, Mozambique, Rwanda, Tanzania, and Uganda. Shading indicates the geographical clustering used in our empirical design. Map lines delineate study areas and do not necessarily depict accepted national boundaries.

individuals employed in the CS sector report, if anything, higher ownership rates than individuals working in manufacturing. Finally, we contrast the experience of sub-Saharan Africa with that of Vietnam and show that this looks markedly different.

2.1 Data

We use both aggregate and micro data. At the aggregate level, we use the World Development Indicators (WDI) from the World Bank to measure GDP, population, urbanization, and sectoral employment shares. We augment these data with information on educational attainment from

Barro and Lee (2013). At the micro-level, we use microdata from the Integrated Public Use Microdata Series (IPUMS). IPUMS harmonizes nationally representative censuses and large-scale survey samples across countries and years, providing consistent individual- and household-level information on sectoral employment, occupation, education, and urban/rural residence. We extract the most recent census rounds available between 1997 and 2019. When we aggregate surveys of different countries, we use national populations as weights. In Appendix Section B-1 we describe the data and our measurement choices in detail.

2.2 Consumer Services, Structural Change, and Economic Development

We begin by documenting the growth performance and the central role of services, and in particular CS, in the process of structural change. Figure 2 reports average growth in total GDP and GDP per capita (left panel), as well as changes in sectoral employment shares (right panel), between 2000 and 2019. Total GDP grew at a stellar average annual rate of 7% in our East African sample, with especially strong performances in Ethiopia (9%), Rwanda (7.7%), and Mozambique (6.9%). Growth was also substantial, though more moderate, in West Africa, averaging 4.9%, with Ghana emerging as the top performer at 6.2%.

Although rapid population growth tempers the magnitude of GDP per capita gains relative to total GDP growth, the figure shows that countries in our African sample nonetheless experienced significant per capita growth during the first two decades of the twenty-first century. In Ethiopia, GDP per capita grew at an average annual rate of 6%, one of the highest rates observed worldwide over this period. This performance stands in sharp contrast to the lost decades of African growth in the preceding period, as documented by Artadi and Sala-i Martin (2003) or Easterly and Levine (1997).

In Appendix Section B-2, we show that this rise in income per capita has been accompanied by rising educational attainment and urbanization. We also show that subjective measures of life satisfaction from the Gallup World Poll increased substantially in all countries in our sample except Malawi and Rwanda.

Economic growth has been accompanied by a sharp shift in sectoral employment. The right panel of Figure 2 documents this reallocation using the standard agriculture, industry, and services classification. In every country, agricultural employment declined, and most of the released labor moved into services. Industry expanded only modestly. On average, agriculture fell by 18 percentage points, with services accounting for 15 of those points and industry for 3. The patterns are similar across East and West Africa.

Figure 2: Growth and Structural Change Towards Services in Sub-Saharan Africa

(a) GDP Growth (b) Growth of the service sector

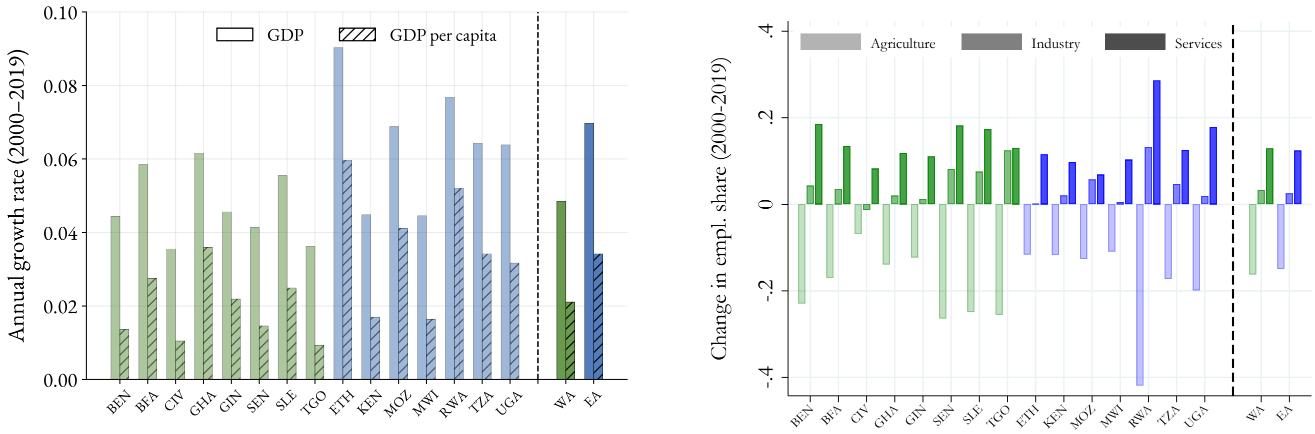

Notes: Panel (a) reports the annual growth rate of GDP and GDP per capita at constant 2021 International Dollars (PPP) over 2000–2019. Panel (b) reports the change in sectoral employment shares between 2000 and 2019. Regional averages are population-weighted using 2002 population shares. Country codes follow ISO-3 standards (e.g., KEN = Kenya, CIV = Côte d'Ivoire, ...), and regional abbreviations are WA = West Africa and EA = East Africa. Guinea-Bissau (GNB) is not shown by the individual bar but included in the average West Africa calculation. Data source: WDI.

Figure 3 shows that Africa is not exceptional. Rather, growth without industrialization is a common feature of recent structural change in the developing world. We consider all non OECD countries whose per capita income in 2019 was below that of China.[4] Countries in our African sample appear in blue (East Africa) and green (West Africa). The figure plots the cross country relationship between the change in the agricultural employment share and the corresponding change in the service share (left panel) and the manufacturing share (right panel) over 2000–2019. In the typical developing country, most of the decline in agricultural employment is offset by rising service

[4] The figure is an adapted version of panels (c) and (d) from Figure 1 in Fan et al. (2023), using a different sample period. Fan et al. (2023) also show that a low pace of industrialization is not associated with slower aggregate growth. In their sample, which also includes the 1990s, countries with average GDP growth above 4% experienced faster declines in agricultural employment (18 percentage points) and faster growth in service employment (13 percentage points).

Figure 3: Structural Change toward Services in the Developing World

(a) Reallocation Toward Services (b) Reallocation Toward Manufacturing

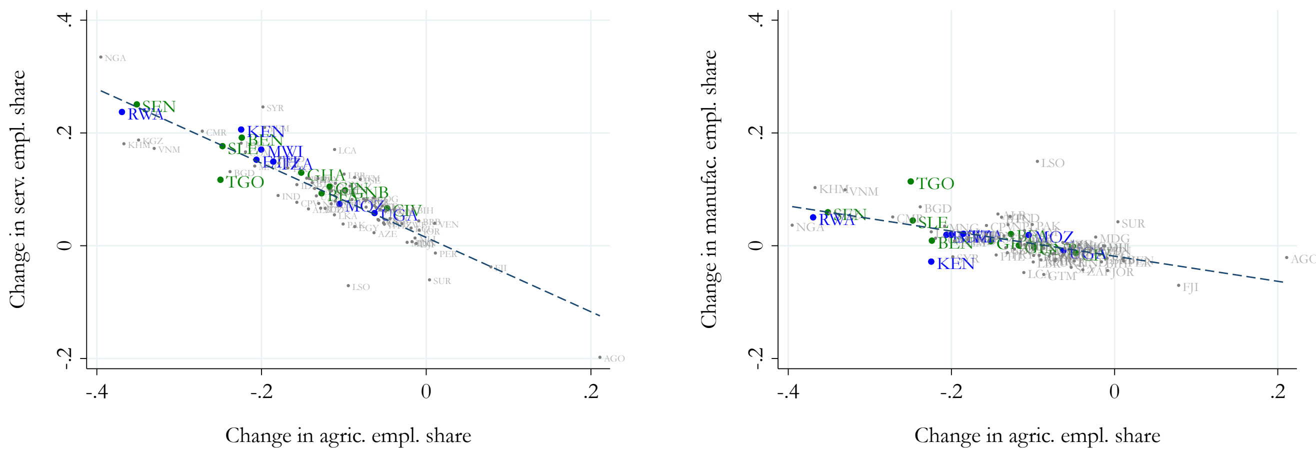

Notes: The figure shows the cross-country correlation between changes in agricultural and service (left panel) and manufacturing (right panel) employment shares over the same period (2000-2019) for all non-OECD countries poorer than China in 2019. The countries in our African sample are highlighted in blue for East Africa and green for West Africa. Data sources: World Bank, International Labor Organization, and Penn World Tables.

employment, whereas manufacturing expands only modestly. While the shift toward services is a common trend worldwide, the pattern is even more pronounced in Africa: most African countries lie above the regression line in the left panel.

To examine these patterns more closely, we turn to the IPUMS microdata and estimate employment shares across detailed sectors for rural and urban locations. We classify workers into four broad sectors: agriculture, manufacturing (including mining and construction), private services, and public services. Private services are further divided into three mutually exclusive groups: (i) retail, leisure, and health, which capture non-tradable CS; (ii) finance, business, and ICT, representing producer services; and (iii) transport and storage, which serves both households and firms.

We focus on countries for which at least two survey rounds are available, allowing us to compare structural changes over time. Specifically, we include six countries, two in West Africa (Benin (2002 and 2013) and Ghana (2000 and 2010)) and four in East Africa (Mozambique (1997 and 2017), Malawi (1998 and 2008), Rwanda (2002 and 2012), and Tanzania (2002 and 2012).)[5] Given the

[5] Data for the remaining countries were not available from IPUMS in a format that allows us to perform a consistent

limited number of countries with usable data, we pool the West and East African cases, weighting observations by country population in 2002.

Panels (a) and (b) in Figure 4 show the sectoral employment shares for rural and urban areas. Expectedly, employment in rural areas is predominantly agricultural. By contrast, cities are "service hubs" with private services being the largest sector of employment. Agricultural employment in both rural and urban Africa declined over time, with most of the reallocation absorbed by private services. The employment shares of manufacturing remained largely stable, while public sector employment declined slightly as a share of the total.

In panels (c) and (d), we zoom into private services and present the disaggregated composition of private service employment. The lion's share of employment growth is accounted for by CS, such as retail, health, or hospitality. Producer services are on the rise but still account for only a negligible share of total employment in African countries, especially in rural areas.

2.3 Are Jobs in Consumer Services Bad Jobs?

The patterns of structural change documented above raises some concerns. In particular, CS jobs in retail and hospitality are frequently informal, and may therefore be interpreted as disguised unemployment rather than activities with a high marginal product.

In this section, we provide suggestive evidence that grants a more optimistic view. While direct measures of household income are unavailable in the IPUMS data, these data report several indicators of living conditions and ownership of durable goods that can serve as proxies for household income. Specifically, we consider access to six basic amenities: electricity, piped water, radio, flush toilet, brick or metal walls, and good-quality roofing.[6] These indicators are commonly used measures of living standards in development economics (see Young (2012)).

We estimate separate household-level regressions for each of these durables of the form

$$y_{ilct} = \delta_{ct} + \sum_s \beta_s \; \mathbf{1}\{s\}_{ilct} + x'_{ilct}\gamma + u_{ilct}, \tag{1}$$

where y_{ilct} is a dummy variable indicating whether household i in country c and location l at time

pre–post analysis.

[6] Roofing is classified as good if made of cement, tile, or metal, and poor if made of grass, mud, or other low-quality materials.

Figure 4: The Rise of Consumer Services in sub-Saharan Africa

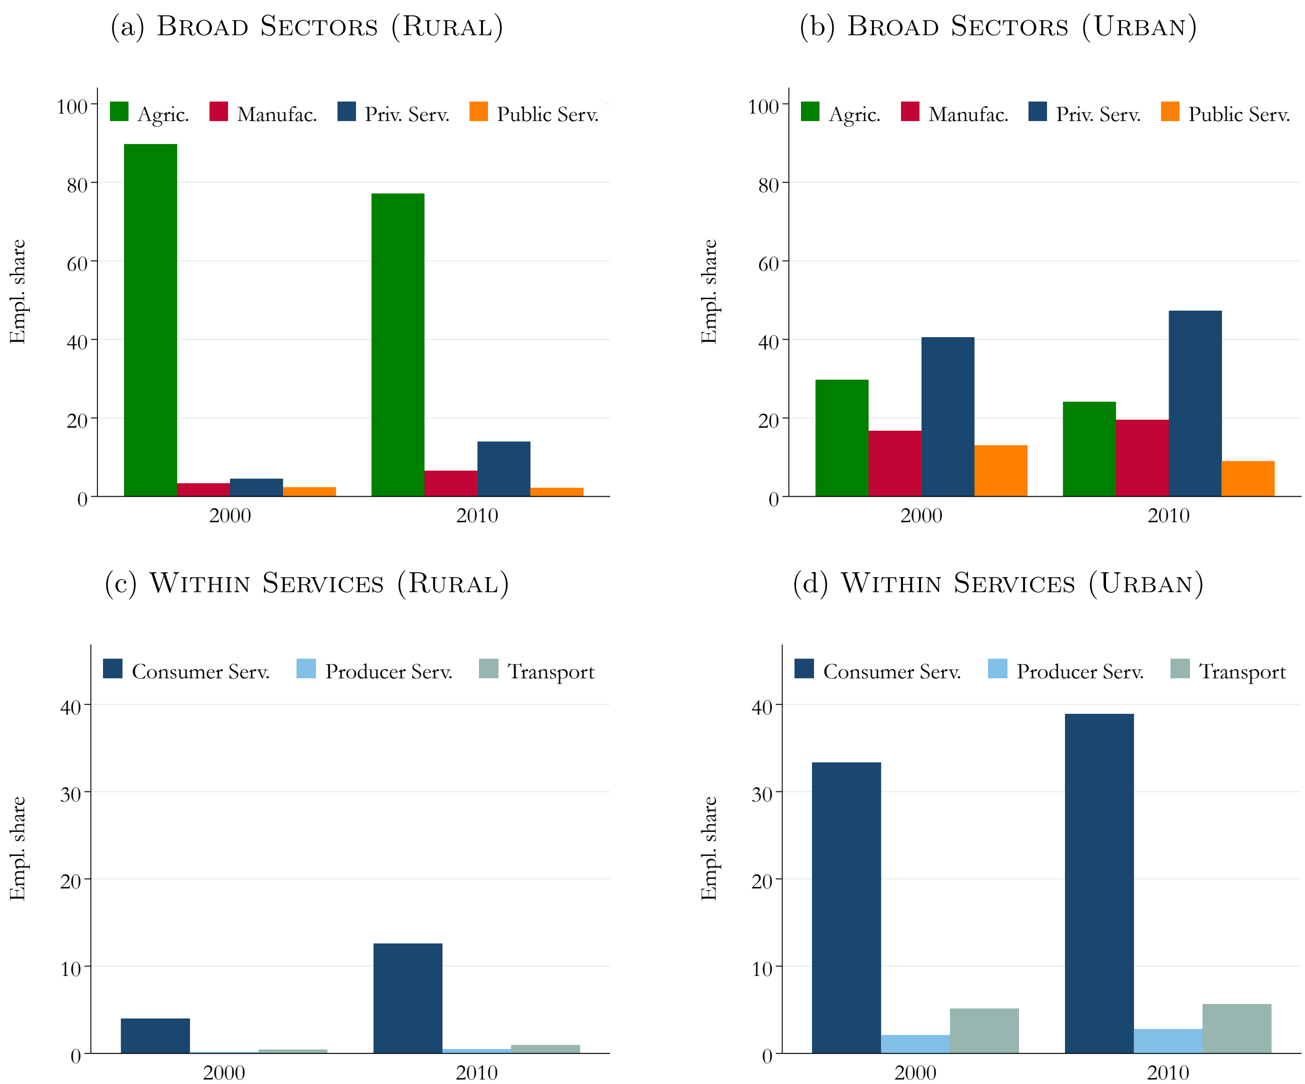

Notes: Panels (a) and (b) report broad sectoral employment shares. Panels (c) and (d) report employment shares within private services. We split private services into three mutually exclusive industries: (i) retail, leisure, and health (CS); (ii) finance, business, and ICT (Producer Services); and (iii) transport and storage (Transport). Surveys used for calculating these results include: Benin (2002 and 2013), Ghana (2000 and 2010), Mozambique (1997 and 2017), Malawi (1998 and 2008), Rwanda (2002 and 2012), and Tanzania (2002 and 2012). Regional averages are population-weighted using 2002 population shares. Data sources: IPUMS International.

t owns the respective amenity, $\mathbf{1}\{s\}$ is an indicator variable whether the head of household i works in sector s, x_{ilct} contains additional covariates, and δ_{ct} denotes country-year fixed effects.

Table 1: Service Employment and Durable Ownership

	Electricity		Piped Water		Radio		Flush Toilet		Wall		Roof	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Consumer Serv.	0.041**	0.023**	0.033**	0.014*	-0.033***	-0.033***	0.010	0.002	0.027*	0.012*	-0.003	-0.014
	(0.016)	(0.009)	(0.013)	(0.006)	(0.006)	(0.006)	(0.021)	(0.019)	(0.012)	(0.006)	(0.016)	(0.013)
Producer Serv.	0.171***	0.116***	0.130***	0.064**	0.023	0.018	0.115*	0.093*	0.118***	0.072***	0.061	0.022
	(0.024)	(0.020)	(0.032)	(0.024)	(0.022)	(0.018)	(0.052)	(0.049)	(0.024)	(0.020)	(0.034)	(0.026)
Transport	0.121***	0.088***	0.072**	0.035*	0.033**	0.031***	0.028	0.014	0.053*	0.026	0.034	0.013
	(0.028)	(0.020)	(0.027)	(0.018)	(0.010)	(0.007)	(0.019)	(0.012)	(0.027)	(0.022)	(0.022)	(0.012)
Agricul.	-0.294***	-0.164***	-0.275***	-0.124***	-0.073**	-0.066***	-0.104*	-0.048*	-0.258***	-0.153***	-0.199***	-0.116***
	(0.028)	(0.021)	(0.019)	(0.017)	(0.022)	(0.008)	(0.046)	(0.025)	(0.039)	(0.030)	(0.040)	(0.020)
Public Serv.	0.187***	0.171***	0.082**	0.065**	0.093***	0.091***	0.110***	0.104***	0.105***	0.092***	0.090**	0.080**
	(0.030)	(0.023)	(0.031)	(0.020)	(0.023)	(0.022)	(0.032)	(0.029)	(0.019)	(0.011)	(0.032)	(0.024)
Urban		0.294***		0.334***		0.008		0.127**		0.238***		0.183**
		(0.025)		(0.021)		(0.034)		(0.051)		(0.022)		(0.058)
High Ed.		0.040***		0.058***		0.025***		0.007		0.002		-0.003
		(0.011)		(0.016)		(0.007)		(0.014)		(0.034)		(0.021)
Obs.	3,360,990	3,294,771	3,359,615	3,293,408	2,893,900	2,827,675	3,356,296	3,290,103	3,361,271	3,295,053	3,361,647	3,295,423
R-squared	0.361	0.413	0.132	0.207	0.048	0.049	0.166	0.198	0.261	0.298	0.270	0.280

Notes: This table reports the results of regression (1). The dependent variables are different categories of home amenities: electricity, piped water, radio, flush toilet, brick or metal walls, and good-quality roof (cement, tile, or metal). The omitted employment category is manufacturing, which also includes construction and utilities, and mining. All specifications control for country-year fixed effects. Standard errors, reported in parentheses, are clustered at the country-year level. The surveys used are: Benin (2013), Burkina Faso (2006), Ethiopia (2007), Ghana (2010), Kenya (2019), Malawi (2008), Mozambique (2017), Rwanda (2012), Senegal (2013), Sierra Leone (2015), Tanzania (2012), Togo (2010), and Uganda (2014).

The results of estimating (1) are reported in Table 1. The findings paint a consistent picture. Workers in service sectors enjoy a sizable premium relative to manufacturing workers (the excluded group) and an even larger one relative to agricultural workers. While such differences might be expected for employees in the public sector or in producer services, the comparison between CS and manufacturing is particularly informative. Contrary to the notion that CS consist predominantly of low-quality jobs, our results do not indicate substantial disadvantages for workers in this sector.[7]

The even-numbered columns show that these patterns do not merely reflect spatial sorting (for example, CS workers being disproportionately located in urban areas) nor differences in educational attainments. As expected, the likelihood of having access to basic amenities is significantly higher for households living in urban locations or headed by more educated individuals. However, controlling for these differences does not alter the main conclusions.

[7] Radio ownership may convey an ambiguous signal in an era of digital devices. For instance, urban households are not more likely to own a radio than their rural counterparts.

In Figure B-3 in Appendix Section B-2, we use the information on durable ownership to document changes in ownership over time. The results are qualitatively consistent with the cross-sectional patterns reported in Table 1. In both rural and urban areas, households have experienced rising access to electricity, piped water, flush toilets, and improved wall and roof materials, and are increasingly likely to own a radio.[8]

2.4 Growth with Industrialization: Vietnam

This recent pattern of structural change differs markedly from the path followed by many earlier industrializers including China as documented by Chen et al. (2023). Vietnam offers an interesting contrast. During the period 2000–19, Vietnam experienced very rapid growth in both total GDP (6.6%) and GDP per capita (5.3%), rates that are comparable to (and in per capita terms even higher than) those of the average East African economy in our sample. However, the process of structural transformation in Vietnam differs sharply from that observed in African economies. Figure 5, which is the analogue of Figure 4, highlights these differences. The differences are especially pronounced in rural areas, where the rapid decline in agricultural employment is absorbed by gains in both manufacturing and services.[9] Such rapid rural industrialization remains largely absent in Africa. Another important difference is that Vietnam exhibits stronger growth in business-related services, particularly in urban areas. This indicates that in Vietnam the expansion of the service sector is to a large extent associated with the production of tradable goods.

These patterns are consistent with recent empirical research showing that manufacturing plays an important role in Vietnam's economic development (McCaig and Pavcnik, 2017, Nguyen, 2018). Vietnam also differs sharply from the African economies we study in its strong export orientation.

[8] These findings are consistent with Gollin et al. (2016), who show that resource-rich countries, who satisfy the narrative of consumption cities built around local consumption spending, have higher urbanization rates, higher employment in non-tradable services, and lower measures of poverty. They also show that conditional on overall urbanization, income per capita, and the non-agricultural employment share, resource-rich countries do not feature higher poverty rates.

[9] Ethiopia, the fastest-growing African economy in our sample, has been seeking to emulate the experience of China and Vietnam by promoting industrialization through dedicated industrial parks and special economic zones. While this strategy has achieved some success in attracting foreign investment and generating manufacturing employment, its spillovers to the rest of the local economy remained limited, reflecting weak backward and forward linkages with domestic suppliers and service providers, in contrast to the experience of East Asian economies. Evidence on place-based policy in China and Vietnam is provided, among others, by Alder et al. (2016) and Huynh (2025), respectively; for reviews of industrial parks in Ethiopia, see Butler et al. (2022), Zhang et al. (2018).

As documented in Appendix Section B-2, its export to GDP ratio ranged from 60 to 80 percent between 2000 and 2020, compared with about 30 percent in our African sample. This difference reflects Vietnam's stronger specialization in tradable production, whereas African economies shifted mainly toward local CS.

Figure 5: Structural Change in Vietnam

(a) Broad Sectors (Rural)

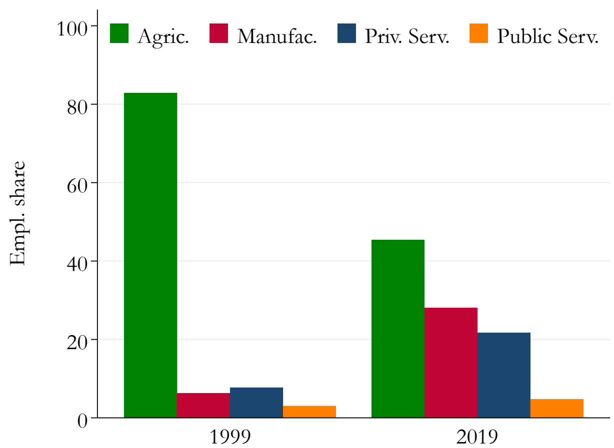

(b) Broad Sectors (Urban)

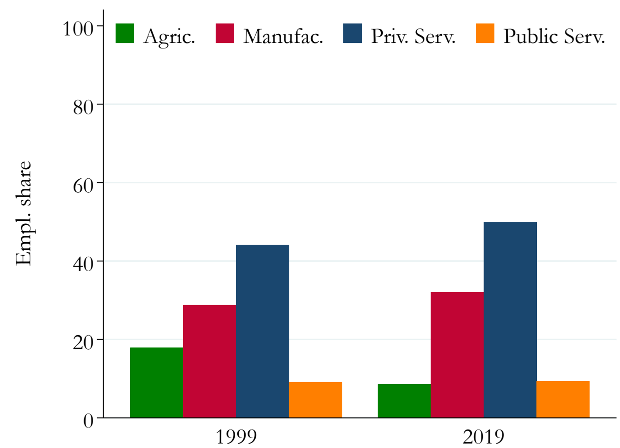

(c) Within Services (Rural)

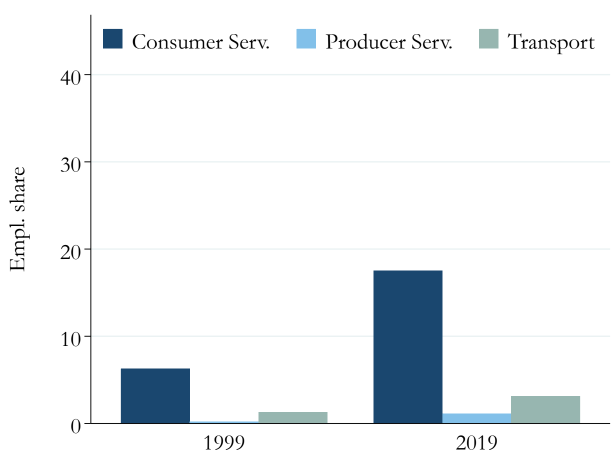

(d) Within Services (Urban)

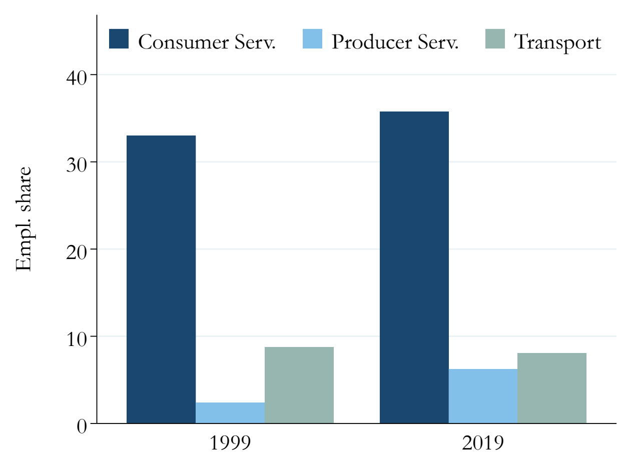

Notes: See notes of Figure 4. Data sources: IPUMS.

3 A Baseline Model of Service-Led Growth

In this section, we develop a model that links the emergence of CS to the broader process of economic development. Building on our earlier work (Fan et al., 2023), the model is analytically tractable and yet flexible enough to capture heterogeneity across both space and individuals. Moreover, it is amenable to estimation using macro and micro data and to conduct counterfactual exercises. For clarity, and reflecting the spatial patterns documented in Section 2, we focus on two regions: cities and rural areas.[10]

We present the framework in a modular fashion. We begin with a baseline setup that highlights two central forces behind the spatially unbalanced expansion of the service sector and the unequal welfare effects of structural change. First, we assume that *CS are non-tradable* and therefore depend on local factor endowments, technologies, and demand.[11] Second, *CS are luxury goods*: their expenditure share rises with income. As a result, they are primarily consumed by richer households, making the local income distribution a key determinant of the sector's expansion. Moreover, rising productivity in CS will mostly benefit richer households.

We show that even a model with these minimal ingredients can be taken to the data to interpret the empirical patterns documented in Section 2. Then, in the following sections, we extend this framework in several directions. In particular, we discuss: *agglomeration economies*, whereby service productivity increases with market size; *labor market frictions*, to capture the idea that service employment may reflect not only high marginal productivity but also the rationing of industrial jobs. We also show to capture *sectoral heterogeneity*, to distinguish between agricultural and manufacturing employment; *international trade*, to examine whether shifts in global competition such as the rise of China as the "producer for the world" have shaped recent growth patterns in developing countries; and *internal migration*, to highlight how unbalanced structural transformation influences location choices.

[10] This restriction is not essential. In Fan et al. (2023), we carry out the analysis at a much finer spatial scale and study the role of service-led growth across hundreds of Indian districts.

[11] This does not imply that such services can only be consumed by local residents. Tourism, for instance, illustrates how locations can effectively "export" non-tradable services such as restaurant meals or hospitality to visitors from other regions or countries (Faber and Gaubert, 2019). In this paper we abstract from this possibility, but it can easily be incorporated. Blanchard et al. (2025) study a model related to ours, which explicitly incorporates the possibility of rural residents traveling to cities to consume non-tradable services.

3.1 The Basic Environment

Geography and Skills The model economy comprises two regions, indexed by $r \in \{R, C\}$, representing rural and urban (*city*) areas, respectively. Labor is the sole factor of production. Individual heterogeneity in productivity and income is captured by differences in endowments of efficiency units of labor, denoted by z, which are distributed according to a region-specific probability distribution function $F_r(z)$. This regional heterogeneity can be interpreted as reflecting differences in human capital.

Preferences We postulate that consumers have preferences belonging to the PIGL class. These preferences feature two important properties. First, they admit aggregation, which enables us to bring a spatial demand system to the data and to perform welfare analysis. Second, they provide a tractable mapping from final demand to value added, facilitating the integration of data on final consumption and sectoral production.

More formally, consumers have preferences over a unit continuum of non-tradable final goods $n \in [0, 1]$. Their indirect utility takes the PIGL form[12]

$$\mathcal{V}^{FE}(e, \mathbf{p}_r) = \frac{1}{\varepsilon}\left(\frac{e}{B(\mathbf{p}_r)}\right)^{\varepsilon} - D(\mathbf{p}_r). \tag{2}$$

where e is total expenditure, $\mathbf{p}_r = \{p_{rn}\}_{n\in[0,1]}$ denotes the local price vector, and $\varepsilon \in (0, 1)$ is the Engel elasticity. The functions $B(\mathbf{p}_r)$ and $D(\mathbf{p}_r)$ are homogeneous of degree one and zero, respectively. We parameterize these functions as

$$B(\mathbf{p}_r) = \exp\left(\int_0^1 \beta_n \ln p_{rn}\, dn\right) \qquad \text{and} \qquad D(\mathbf{p}_r) = \int_0^1 \kappa_n \ln p_{rn}\, dn, \tag{3}$$

where $\int_0^1 \beta_n\, dn = 1$ and $\int_0^1 \kappa_n\, dn = 0$.

By Roy's identity, the expenditure share on good n by a consumer in region r with total spending e is given by:

$$\vartheta_n^{FE}(e, \mathbf{p}_r) = \beta_n + \kappa_n \Upsilon(e, \mathbf{p}_r)^{-\varepsilon} \qquad \text{where} \qquad \Upsilon(e, \mathbf{p}_r) \equiv \left(\frac{e}{\exp\left(\int_0^1 \beta_n \ln p_{rn}\, dn\right)}\right). \tag{4}$$

[12] PIGL preferences do not admit a closed-form utility representation but are characterized by the indirect utility function in (2).

Figure 6: Engel Curves

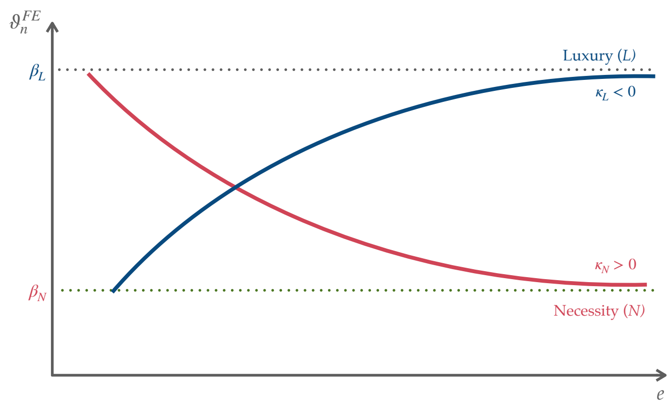

Notes: The figure shows expenditure shares for necessities (red) and luxuries (blue) as a function of expenditure; see (4).

Equation (4) shows that our model features a non-homothetic demand system, consisting of a Cobb-Douglas core and an income-dependent adjustment term. The scalar $\Upsilon(e, \mathbf{p}_r)$ parameterizes "real income," as it increases with nominal spending e and decreases with final good prices $\mathbf{p}_r$.[13]

Demand for each product n is governed by two product-specific parameters, β_n and κ_n. β_n parameterizes product n's asymptotic expenditure share: As income rises, the expenditure share on good n converges to β_n. κ_n determines the product's income elasticity. A good is classified as a luxury if $\kappa_n < 0$, and as a necessity if $\kappa_n > 0$. Figure 6 provides an illustration of how expenditure shares vary with income for a necessity and a luxury good. Cobb-Douglas preferences arise as the special case $\kappa_n = 0$. The slope of the Engel curves, and hence the strength of income effects, is governed by the parameter ε.

Technology Each final good $n \in [0, 1]$ is produced using a combination of tradable inputs x_n and local CS inputs $y_{CS,n}$:

$$y_n = \overline{\lambda}_n \, x_n^{1-\lambda_n} \, y_{CS,n}^{\lambda_n}. \tag{5}$$

[13] We put "real income" in quotes because preferences are non-homothetic and $\Upsilon(e, \mathbf{p}_r)$ is therefore not a welfare metric.

The parameter $\lambda_n \in (0,1)$ denotes the service intensity of product n.[14] Equations (4) and (5) together allow for a flexible modeling of product-level heterogeneity. Goods such as vegetables sold in markets rely primarily on tradable inputs (i.e., small λ_n), whereas goods like haircuts depend more heavily on local labor (i.e., large λ_n). The extent to which richer consumers demand more service-intensive goods is governed by the correlation between λ_n and κ_n.

Turning to the production of the two inputs, we assume that the tradable input x combines the output of the agricultural sector, the manufacturing sector, and producer services, such as finance or consulting. Specifically,

$$x = f_r(A_r^{AG}H_r^{AG}, A_r^{MAN}H_r^{MAN}, A_r^{PS}H_r^{PS}), \tag{6}$$

where H_r^j denotes total employment in sector j in region r, A_r^j is the associated productivity term, and f_r is an aggregate production function that can vary across space. As long as f_r has constant returns to scale, the production of x can be represented as

$$x = \mathcal{A}_r^T \left(H_r^{AG} + H_r^{MAN} + H_r^{PS}\right) \equiv \mathcal{A}_r^T H_r^T,$$

where $\mathcal{A}_r^T$ denotes tradable productivity in region r and H_r^T denotes efficiency units of labor employed in the three tradable sectors. Crucially, $\mathcal{A}_r^T$ only depends on the production function f_r and the associated productivity terms A_r^j and is independent of prices and employment allocations.[15] Hence, $\mathcal{A}_r^T$ might be large in cities because of high productivity in manufacturing or producer services, while rural locations might have high tradable productivity because of their efficiency in agriculture. In Section 5.2, we put more structure on f_r and, in doing so, also on the allocation of employment between agriculture, manufacturing, and producer services.

For the CS sector, we assume that production is linear in labor and given by

$$y_r^{CS} = \mathcal{A}_r^{CS} H_r^{CS}.$$

A resource constraint requires that $H_r^T + H_r^{CS} \leq H_r$, where H_r denotes the regional endowment of effective units of labor.

[14] We define $\overline{\lambda}_n \equiv \lambda_n^{-\lambda_n}(1-\lambda_n)^{-(1-\lambda_n)}$ as a constant that simplifies the algebra.

[15] If, for example, f_r was a CES production function with elasticity ρ, then $\mathcal{A}_r^T = \left(\sum_j \left(A_r^j\right)^{\rho-1}\right)^{1/(\rho-1)}$.

3.2 Two Results: Aggregation and the Demand over Value Added

Aggregation An appealing feature of PIGL preferences is that they are amenable to aggregation: despite being non-homothetic at the individual level and falling outside the Gorman polar form, they nonetheless admit an aggregate demand representation via a fictitious representative household. More precisely, although the associated demand system depends on the distribution of income, this dependence can be summarized by a set of distributional statistics, as shown in Proposition 1 below. This property justifies the use of a representative consumer, provided appropriate adjustments are made for inequality.[16] This tractability is particularly useful for quantitative applications and counterfactual analysis.

To derive the *aggregate* spending share at the regional level from the individual-level expenditure shares from equation (4), let w_r denote the prevailing wage in location r. Total income and expenditure by an individual with efficiency z are then zw_r. We can then compute a closed-form expression for the aggregate spending share on good n in region r.

Proposition 1. *Let* $\vartheta_{rn}^{FE} \equiv \frac{\int \vartheta_n^{FE}(zw_r,\mathbf{p}_r)zw_r\, dF_r(z)}{\int zw_r\, dF_r(z)}$ *denote the aggregate spending share on good* n *in region* r*. Then,*

$$\vartheta_{rn}^{FE} = \beta_n + \kappa_n \mathcal{I}_r\, \Upsilon\left(\mathbb{E}_r[z]w_r, \mathbf{p}_r\right)^{-\varepsilon} \tag{7}$$

where $\Upsilon(\cdot)$ *is defined in* (4)*, and*

$$\mathcal{I}_r = \frac{\int z^{1-\varepsilon}\, dF_r(z)}{\left(\int z\, dF_r(z)\right)^{1-\varepsilon}}. \tag{8}$$

Proof. See Appendix Section A-1.1. □

Proposition 1 implies that aggregate demand behaves *as if* generated by a representative consumer with average spending $\mathbb{E}_r[z]w_r$, but with the non-homotheticity parameter κ_n scaled by the inequality term $\mathcal{I}_r$. This statistic summarizes how inequality affects demand and allows the entire distribution $F_r(z)$ to be captured by a single number. Interestingly, it is a nonlinear monotonic transformation of the Atkinson index (Atkinson, 1970).[17] Hence, even though preferences are non-homothetic, one does not need the full income distribution to compute aggregate demand;

[16] See Chapter 6 of Deaton and Muellbauer (1980b) for a detailed discussion.

[17] Note that $\mathcal{I}_r$ as defined in (8) equals $(1 - AI_r)^{1-\varepsilon}$, where AI_r is the standard Atkinson index. With some abuse of terminology, we refer to $\mathcal{I}_r$ as the Atkinson index in the rest of the paper.

the scalar $\mathcal{I}_r$ is sufficient. This aggregation property is particularly useful for quantitative work. Moreover, in some particular cases, $\mathcal{I}_r$ has a closed-form expression. For example, suppose that z is Fréchet distributed with common tail parameter ζ:

$$F_r(z) = \exp\left(-\left(z/\overline{z}_r\right)^{-\zeta}\right). \tag{9}$$

Then $\mathcal{I}_r = \Gamma_{\frac{\zeta}{1-\varepsilon}}/\Gamma_{\zeta}^{1-\varepsilon} = \mathcal{I}$, where Γ_x denotes the gamma function.[18] In this case the inequality adjustment is constant across regions.

Proposition 1 also highlights how inequality affects aggregate demand for different goods. Because $0 < \varepsilon < 1$, Jensen's inequality implies that a higher inequality (in the sense of a mean-preserving spread) reduces $\mathcal{I}_r$. Moreover, equation (7) implies that

$$\frac{\partial\,\vartheta_{rn}^{FE}}{\partial\,\mathcal{I}_r} = \kappa_n\,\Upsilon\left(\mathbb{E}_r[z]w_r, \mathbf{p}_r\right)^{-\varepsilon}.$$

Hence, higher inequality increases the aggregate expenditure share of luxuries (where $\kappa_n < 0$) and reduces the spending share on necessities (where $\kappa_n > 0$).

From Final Goods to Value Added The second important property of PIGL preferences is that they permit a seamless mapping from final goods to sectoral value added. This mapping is essential for empirical applications: sectoral employment data refer to value added, whereas household expenditure data concern final goods. In general, preferences defined over final goods need not translate cleanly into preferences over value added—see Herrendorf et al. (2013). Within the PIGL class, together with Cobb–Douglas technologies as in (5), this mapping takes a particularly tractable form.

To derive the demand system over sectoral value added, let p_x denote the price of the tradable good and q_r the local price of CS inputs. Because services are local in nature, q_r varies across regions. The price of final good n is then

$$p_{rn} = p_x^{1-\lambda_n} q_r^{\lambda_n}. \tag{10}$$

Let $\mathcal{S}_r^{CS}$ denote total spending on CS *inputs* in region r. Then,

$$\mathcal{S}_r^{CS} = \int_n \underbrace{\lambda_n}_{\text{CS share}} \times \underbrace{\vartheta_{rn}^{FE} L_r w_r \mathbb{E}_r[z]}_{\text{Total spending on } n}\, dn, \tag{11}$$

[18] Γ_x is the gamma function with parameter $1 - 1/x$, that is, $\Gamma_x = \Gamma(1 - 1/x)$.

since the service component of good n accounts for a share λ_n of total expenditure on that good. Substituting the expression for ϑ_{rn}^{FE} from Proposition 1 leads to the following proposition.

Proposition 2. *Let $\vartheta_r^{CS} \equiv \frac{\mathcal{S}_r^{CS}}{L_r w_r \mathbb{E}_r[z]}$ denote the aggregate spending share on value added in the CS sector. Then,*

$$\vartheta_r^{CS} = \omega_{CS} + \nu_{CS}\mathcal{I}_r \Upsilon\left(\mathbb{E}_r[z]w_r; p_x, q_r\right)^{-\varepsilon}, \tag{12}$$

where

$$\omega_{CS} = \int \lambda_n \beta_n \, dn, \qquad \nu_{CS} = \int \lambda_n \kappa_n \, dn, \qquad \Upsilon\left(\mathbb{E}_r[z]w_r; p_x, q_r\right) = \frac{\mathbb{E}_r[z]w_r}{p_x^{1-\omega_{CS}} q_r^{\omega_{CS}}}.$$

Proof. See Appendix Section A-1.1. □

Proposition 2 shows that the demand system over sectoral value added again belongs to the PIGL class. The preference parameters of this value-added demand system, ω_{CS} and $\nu_{CS}\mathcal{I}_r$, reflect preferences (β_n and κ_n), inequality ($\mathcal{I}_r$), and technology (λ_n). The parameter ω_{CS} is a weighted average of the consumer-service content λ_n, with weights given by the asymptotic expenditure shares β_n. Likewise, ν_{CS} is a weighted average of λ_n, weighted by the non-homotheticity parameters κ_n.

The proposition also clarifies how the technology shapes sectoral demand. If all final goods rely on CS with the same intensity, that is if $\lambda_n = \lambda$ for all n, then $\nu_{CS} = 0$ because $\int_0^1 \kappa_n \, dn = 0$. In this case, factor demand remains homothetic even when the underlying final-good demand system is non-homothetic. In contrast, service value added behaves as a luxury good (that is, $\nu_{CS} < 0$) when service intensity and income elasticity are negatively correlated, meaning that goods with high λ_n tend to have low κ_n.

Finally, Proposition 2 establishes a tight link between inequality and demand for service value added because the overall shifter of the factor demand system, $\nu_{CS}\mathcal{I}_r$, is given by

$$\nu_{CS}\mathcal{I}_r = \underbrace{\left(\int \kappa_n \lambda_n \, dn\right)}_{\text{Technology + Preferences}} \times \underbrace{\mathcal{I}_r}_{\text{Inequality}}.$$

Higher inequality, namely, lower $\mathcal{I}_r$ raises the aggregate value-added share of services when luxury goods are service-intensive. Using CEX data, Aghion et al. (2025) document that this is indeed the case for U.S. households.[19]

3.3 Equilibrium: Service-Led or Service-Biased Growth?

To characterize the equilibrium allocations, recall that the tradable good is freely traded between cities and rural areas. Letting the tradable good serve as the numeraire, $p_x \equiv 1$, nominal wages are determined by local productivity in the tradable sector, and the price of CS inputs in region r reflects the relative productivity of the tradable and non-tradable sectors:

$$w_r = \mathcal{A}_r^T \qquad \text{and} \qquad q_r = \frac{w_r}{\mathcal{A}_r^{CS}} = \frac{\mathcal{A}_r^T}{\mathcal{A}_r^{CS}}. \tag{13}$$

To solve for the equilibrium, note that the *employment share* in CS in region r, denoted ℓ_r^{CS}, equals the aggregate *spending share* on CS inputs given in (12) (see Appendix A-1.2). Substituting for the equilibrium prices $p_x = 1$ and q_r in (13) yields

$$\ell_r^{CS} = \vartheta_r^{CS} = \omega_{CS} + \nu_{CS}\,\mathcal{I}_r\left(\mathbb{E}_r[z](\mathcal{A}_r^T)^{1-\omega_{CS}}(\mathcal{A}_r^{CS})^{\omega_{CS}}\right)^{-\varepsilon}. \tag{14}$$

Equation (14) expresses the CS employment share as a function of structural parameters, sectoral productivities, local human capital, and inequality. This relationship is central for our analysis because it separates two channels of service expansion. The first is service-biased growth: higher income, driven by human capital $E_r[z]$, tradable productivity $\mathcal{A}_r^T$, or greater inequality, raises demand for income elastic services. Since these services are non-tradable, stronger demand translates directly into higher local employment. The second channel is service led growth: higher CS productivity raises local real incomes, which in turn increases expenditure on services and expands CS employment.

3.4 Unequal Incidence: Why Service-Led Growth is Pro-Rich

A key implication of our theory is that the welfare gains from sectoral productivity growth differ systematically across individuals. Substituting final good prices p_{rn} in (10) into the indirect utility

[19] The view that goods are necessities and services are luxury goods is also consistent with traditional Engel curve analysis and with the existence of hierarchical preferences—see, e.g., Foellmi and Zweimüller (2008) and Chai et al. (2023). Banerjee and Duflo (2007) document that extremely poor households in developing countries allocate between 56 and 78 percent of their budgets to food, whereas food accounts for less than 10 percent of spending among rich households in advanced economies. Although the model in this paper abstracts from differences in the income elasticity of agricultural and manufacturing goods, the framework readily accommodates this refinement, as shown in Fan et al. (2023).

function ((2) and (3)) allows us to express indirect utility in terms of the input prices p_x and q_r:

$$\mathcal{V}(zw_r, p_x, q_r) = \frac{1}{\varepsilon}\left(\frac{zw_r}{p_x^{1-\omega_{CS}} q_r^{\omega_{CS}}}\right)^{\varepsilon} - \nu_{CS} \ln\left(\frac{q_r}{p_x}\right). \tag{15}$$

Let $e(\mathcal{V}, p_x, q_r)$ denote the expenditure a consumer, facing input prices p_x and q_r, requires to attain a level of utility $\mathcal{V}$. The welfare effect of a change in input prices can be measured using the equivalent variation, that is, the change in expenditure needed to maintain a fixed utility level $\mathcal{V}$. Based on the properties of the expenditure function (see Appendix Section A-1.3), one can show that:

$$\frac{\partial \ln e(\mathcal{V}, p_x, q_r)}{\partial \ln q_r} = \vartheta_r^{CS}(zw_r, p_x, q_r) \quad \text{and} \quad \frac{\partial \ln e(\mathcal{V}, p_x, q_r)}{\partial \ln p_x} = 1 - \vartheta_r^{CS}(zw_r, p_x, q_r), \tag{16}$$

where $\vartheta_r^{CS}(zw_r, p_x, q_r)$ denotes the share of expenditure by a consumer with skill level z in region r on the value-added component of the service sector, which (in analogy to (12)) is given by

$$\vartheta_r^{CS}(zw_r, p_x, q_r) = \omega_{CS} + \nu_{CS}\left(\frac{zw_r}{p_x^{1-\omega_{CS}} q_r^{\omega_{CS}}}\right)^{-\varepsilon}. \tag{17}$$

Note that the inequality term $\mathcal{I}_r$ is absent because (17) describes the sectoral spending share of an *individual* rather than the aggregate share.

Equations (16) and (17) lead to a sharp distributional conclusion: *service-led growth is pro-rich*, whereas growth driven by productivity increases in the tradable sector is pro-poor. Equation (16) implies that the welfare effect of a price change is proportional to the value-added share in an individual's expenditure. A one percent fall in the service price q_r increases the utility of a consumer with skill z by $\vartheta_r^{CS}(zw_r, p_x, q_r)$. Because richer households allocate a larger share of their spending to services, improvements in $\mathcal{A}_r^{CS}$ benefit them more than poorer households.

3.5 Quantifying the Importance of Service-Led Growth

Methodology In this section, we illustrate how one can take the model to the data. A key benefit of our approach is an empirical strategy that tackles the well-known difficulty of measuring productivity in service sectors by recovering it from the system of equilibrium conditions implied by the theory.

Directly measuring CS productivity $\mathcal{A}_r^{CS}$ is challenging. Because services are non-tradable, one would need local price indices, which are rarely available. More importantly, differences in *quality* are especially salient in services: expensive restaurants, hairdressers or hotels typically provide a better experience rather than operating less efficiently. Without proper quality adjustment, such differences remain unmeasured. Our methodology circumvents these issues by inferring $\mathcal{A}_r^{CS}$ from consumers' *willingness-to-pay*, as revealed by regions' sectoral employment patterns given consumers' income.

Suppose we observe average local earnings $\{y_r\}$, CS employment shares $\{\ell_r^{CS}\}$, and summary measures of human capital such as years of schooling $\{h_r\}$.[20] Equations (13) and (14) imply that

$$\ell_r^{CS} = \omega_{CS} + \nu_{CS}\,\mathcal{I}\left(\mathbb{E}_r[z](\mathcal{A}_r^T)^{1-\omega_{CS}}(\mathcal{A}_r^{CS})^{\omega_{CS}}\right)^{-\varepsilon} \qquad \text{and} \qquad y_r = \mathbb{E}_r[z]w_r = \mathbb{E}_r[z]\mathcal{A}_r^T. \tag{18}$$

Take the structural parameters ω_{CS} and ε as given (we return to their estimation below). Because productivity $\mathcal{A}_r^{CS}$ has no natural units, we can, without loss of generality, normalize $\nu_{CS} = -1/\mathcal{I}$. The negative sign implies that CS value added is a luxury, which is the empirically relevant case. Under the assumption that local earnings are measured in units of the tradable good, and that local human capital $\mathbb{E}_r[z]$ can be linked to observed educational attainment (e.g. via $\mathbb{E}_r[z] = e^{\rho h_r}$, where ρ is the Mincerian return to education), one can infer local productivities $\{\mathcal{A}_r^T, \mathcal{A}_r^{CS}\}$ directly from (18).

The intuition is straightforward. Local earnings y_r reflect both tradable sector productivity $\mathcal{A}_r^T$ and local human capital $\mathbb{E}_r[z]$. As in the development accounting literature (Caselli, 2005, Gancia et al., 2013), observed earnings allow us to recover $\mathcal{A}_r^T$ once human capital is measured. Given ω_{CS} and ε, the CS employment share ℓ_r^{CS} then identifies $\mathcal{A}_r^{CS}$. Because the CS spending share is sensitive to CS productivity through income effects, this approach separates service-biased from service-led growth: a high CS employment share may reflect high incomes or high CS productivity. The structure of the model delivers the required identification.

The discussion above presupposes knowledge of the demand parameters ω_{CS} and ε. An advantage of our approach is that they can be estimated using micro-level data on consumption expenditure. For this purpose, the result that the Engel elasticity ε governs *both* the demand system for final goods (ϑ_{rn}^{FE}) and the demand for value added (ϑ_r^s) is crucial. Equation (4) implies

[20] If one has data on the local income distribution, one can also directly measure $\mathcal{I}_r$. In our application we lack such data and assume that inequality is constant across locations, i.e. $\mathcal{I}_r = \mathcal{I}$.

that the expenditure share of individual i for final good n is given by

$$\vartheta_n^{FE}(e_i, \mathbf{p}_r) = \beta_n + \kappa_n \ \exp\left(\int_0^1 \beta_n \ln p_{rn}\, dn\right)^{\varepsilon} \ e_i^{-\varepsilon} \tag{19}$$

To take (19) to the data, suppose that a particular category of consumption goods observed in standard consumer expenditure surveys corresponds to a collection of final goods, $\mathcal{J} \subset [0,1]$. For example, the observed spending share on "basic food" comprises expenditures on different varieties of edible products (rice, fruit, etc.). Recall that these are final products, meaning that their prices p_{rn} reflect the full cost of providing each product, including the CS content embodied in local retail services. Integrating (19) over the set $\mathcal{J}$ and assuming that this set consists, on average, of necessities, i.e. $\kappa_{\mathcal{J}} \equiv \int_{j\in\mathcal{J}} \kappa_n > 0$, we obtain

$$\ln\left(\vartheta_{\mathcal{J}}^{FE}(e_i, \mathbf{p}_r) - \beta_{\mathcal{J}}\right) = \ln \kappa_{\mathcal{J}} + \left(\int_0^1 \beta_n \ln p_{rn}\, dn\right)^{\varepsilon} - \varepsilon \ln e_i = \delta_r - \varepsilon \ln e_i \tag{20}$$

where $\beta_{\mathcal{J}} \equiv \int_{j\in\mathcal{J}} \beta_n\, dn$. Equation (20) shows that one can estimate the Engel elasticity ε from a *cross-sectional* regression of spending shares on individual income. In particular, the unobserved final-good prices $\mathbf{p}_r$ are absorbed by the regional fixed effect δ_r. The additional parameter $\beta_{\mathcal{J}}$ can either be estimated jointly with ε or handled via an additional restriction, such as $\beta_{\mathcal{J}} \approx 0$ when the choice set $\mathcal{J}$ comprises goods whose asymptotic expenditure share is small. For an empirical implementation of this methodology, we refer to Fan et al. (2023), who apply it using consumption expenditure data from India.[21] To discipline the parameter ω_{CS}, one can exploit information from the time-series variation of employment shares in rich countries. Recall that ω_{CS} and $\omega_T = 1 - \omega_{CS}$ parameterize the asymptotic spending and employment shares on services and tradable value added as incomes grow large, so they can be identified from a long-run restriction. For example, if $\omega_{CS} = 0.7$, this would imply that, in the long run, 70% of individuals are employed in the CS sector. Under the assumption that preferences are stable across countries, data on

[21] The discussion above assumes that $\kappa_{\mathcal{J}} \equiv \int_{j\in\mathcal{J}} \kappa_n > 0$. The only reason for this assumption is to ensure that $\ln \kappa_{\mathcal{J}}$ is well defined. If the set $\mathcal{J}$ comprises luxuries, $\kappa_{\mathcal{J}}$, (20) can implemented as

$$\ln\left(\beta_{\mathcal{J}} - \vartheta_{\mathcal{J}}^{FE}(e, \mathbf{p}_r)\right) = \ln(-\kappa_{\mathcal{J}}) + \left(\int_0^1 \beta_n \ln p_{rn}\, dn\right)^{\varepsilon} - \varepsilon \ln e_i = \delta_r - \varepsilon \ln e_i,$$

reflecting the fact that the expenditure share on luxuries is increasing in income and approaches $\beta_{\mathcal{J}}$ from below.

sectoral employment shares in rich economies such as the US can therefore be used to discipline ω_{CS}.

Application: Structural Change Towards Services in Sub-Saharan Africa We now apply this methodology to our sample of African countries analyzed in Section 2. The exercise is intended to be illustrative rather than a comprehensive estimation.

To classify sectors into tradable goods and CS, we rely on the data reported in Figure 4. Our analysis focuses exclusively on private employment, excluding public sector jobs. We define CS as encompassing retail, leisure, and health. Consequently, all other sectors are classified as part of the tradable sector.[22] In Section 5 we illustrate how our framework can be extended to incorporate a more disaggregated sectoral classification, distinguishing among manufacturing, agriculture, and producer services within the tradable sector.

In Table 2 we report the empirical moments used to estimate our model. We restrict the sample to countries for which two surveys are available in the period of interest in order to construct a balanced panel. The moments include the relative population shares of rural and urban areas, the allocation of employment between tradable goods and CS, measures of educational attainment, and aggregate GDP per capita. In addition, we need information on the urban productivity premium in the tradable sector, $\gamma \equiv \mathcal{A}_C^T/\mathcal{A}_R^T$. We estimate this premium from microdata on relative earnings from the "Harmonized World Labor Force Survey (HWLFS)".[23] As we show in detail in Appendix Section A-2, we can infer the urban productivity premium γ directly from the microdata on earnings. Consistent with our results in Section 2, urbanization, educational attainment, and income per capita have increased over time, and employment has shifted toward CS. The data also

[22] In this classification, the tradable sector includes industries whose products are non-tradable in themselves but serve as inputs in the production of tradable manufacturing goods, such as business services. Some industries are ambiguous. For instance, transport services and construction are used by both firms and consumers. In Fan et al. (2023), we use Indian survey data to measure the share of each industry that serves consumers and firms and split value added accordingly. Because comparable data are not available for the African sample, we classify Construction and Utilities, Transport and Storage, and Finance, Business and ICT as producer services that ultimately belong to the tradable goods sector. Since we emphasize the rapid expansion of CS, this choice is conservative.

[23] We are immensely grateful to Charles Gottlieb, who developed the Harmonized World Labor Force Survey (HWLFS), and ran these regressions for us. The HWLFS is a unique global micro-dataset harmonizing employment and job information from thousands of surveys in more than 100 countries and has also been used in Goldberg et al. (2025), Gottlieb et al. (2024), and Gottlieb et al. (2025).

reveal a pronounced rural–urban gradient: urban locations are more productive, more educated, and more specialized in the service sector. Interestingly, the rural-urban productivity gap has risen between 2000 and 2010 from about 5% to 15%.

Table 2: Spatial Structural Change in sub-Saharan Africa

Area	Year	Pop. Share	Sectoral Emp. Share		Yrs of	Urban-rural	GDP per
			Tradable	CS	schooling	prod. gap γ	capita
Rural	2000	0.76	0.96	0.04	3.7	1.05	1,902
Urban		0.24	0.62	0.38	6.9		
Rural	2010	0.63	0.87	0.13	4.4	1.15	2,693
Urban		0.37	0.57	0.43	7.8		

Notes: The table reports empirical moments for the sample of African countries with at least two surveys in the period 1997–2017. The sample includes Benin (2002, 2013), Ghana (2000, 2010), Malawi (1998, 2008), Mozambique (1997, 2017), Rwanda (2002, 2012), and Tanzania (2002, 2012). The data are aggregated using 2002 population weights. CS comprise Retail, Leisure, and Health. The goods sector comprises Agriculture, Manufacturing, Transport and Storage, and Finance, Business, and ICT. Public services are excluded from total employment. GDP per capita is measured in constant 2021 international dollars.

As discussed above, to estimate sectoral productivity $\{\mathcal{A}_{rt}^T, \mathcal{A}_{rt}^{CS}\}$ from the data reported in Table 2, we need to take a stand on the structural parameters ε and ω_{CS} and the Mincerian returns to education ρ. For the former, we rely on the estimates of Fan et al. (2023), who estimate that $\varepsilon = 0.395$ and $\omega_{CS} = 0.692$. To map observed years of schooling into human capital, we assume an annual Mincerian return of $\rho = 0.1$. As we describe in detail in Appendix Section A-2, we can then uniquely solve for the productivity terms $\{\mathcal{A}_{rt}^T, \mathcal{A}_{rt}^{CS}\}$.

The results are reported in Table 3. For ease of comparison, we normalize sectoral productivity in rural locations in 2000 to unity, and we report both the estimated productivity levels and the corresponding annual growth rates. Table 3 reveals three important findings. First, there is a substantial rural–urban productivity gap in both the tradable and non-tradable sectors, with the disparity especially pronounced in CS. Second, productivity increased in both sectors, and if anything, CS productivity increased faster. Quantitatively, annual productivity growth in the tradable sector was between 1% and 3%; productivity growth in CS was about 3–4%.

Counterfactuals Using the estimated productivity terms in Table 3, we conduct counterfactual analyses to quantify the welfare consequences of the inferred patterns of productivity growth.

Table 3: ESTIMATED SECTORAL PRODUCTIVITY IN SUB-SAHARAN AFRICA

	Tradable ($\mathcal{A}_r^T$)			CS ($\mathcal{A}_r^{CS}$)		
Region	2000	2010	Growth	2000	2010	Growth
Rural	1.00	1.19	0.017	1.00	1.40	0.035
Urban	1.05	1.37	0.027	9.37	13.03	0.033

Notes: The table reports the estimated productivity levels for each region (urban and rural). The growth columns report the corresponding annual growth rates. Productivity in both the tradable and CS sectors in rural areas in 2000 is normalized to 1.

Specifically, we compute the equivalent variation for consumers at different points of the income distribution and residing in different locations, separately for service-biased growth (that is, rising tradable productivity) and service-led growth (that is, rising CS productivity).[24] Intuitively, this statistic measures the share of income that an individual would be willing to give up for the respective change in productivity to occur. We define a "poor" consumer as one at the 5th percentile of the local income distribution and a "rich" consumer as one at the 95th percentile.

Panel (a) of Figure 7 shows two main patterns. First, rising CS productivity generates large welfare gains. Second, the incidence of these gains varies across locations and income groups. Across space, the welfare effects of service-led growth are larger in urban areas where CS account for a larger share of both employment and expenditure—see Table 2. Rich urban households enjoy welfare gains larger than those accruing from tradable-sector productivity growth. Across income groups, richer households benefit more than poorer households from service-led growth because they devote a larger share of their income to CS. In contrast, poorer households benefit mainly from productivity growth in agriculture and manufacturing. The contrast is most dramatic for poor rural households, who experience no welfare gains from service-led growth because their expenditure share on services is zero at low income levels. In summary, richer urban households are the main beneficiaries of productivity improvements in the provision of CS, while the living standards of poor households, especially in rural areas, hinge primarily on productivity growth in agriculture and other tradable goods.

We can also assess how much of the structural change toward services is driven by service-biased

[24] Appendix Section A-1.3 provides the formal derivation of the equivalent variation in our model.

growth versus service-led growth. Panel (b) of Figure 7 shows the results. The gray bars report the empirically observed change in the CS employment share. Note that, although the CS employment share is larger in urban Africa, the change over time is larger in rural areas, reflecting the rapid decline of agriculture, which employed almost all rural workers in 2000; see again Table 2.

The blue bars plot the counterfactual change if only CS productivity, $\mathcal{A}_{rt}^{CS}$, had changed, holding $\mathcal{A}^{T}$ and human capital fixed at their 2000 levels. They therefore measure the structural change attributable solely to service-sector productivity growth. The height of the blue relative to the gray bars reflects the importance of service-led growth relative to the total change, which is similar in rural and urban areas. The red bars plot the counterfactual change under service-biased growth, that is, the change that would occur if only $\mathcal{A}_{rt}^{T}$ had increased, with all other primitives fixed at their 2000 levels. The importance of service-led growth for African economies is apparent: to rationalize the observed structural transformation toward services, the model points to a central role for rising CS productivity in both rural and urban areas.

A Different Development Path: Vietnam In Section 2.4, we have shown that Vietnam followed a different development trajectory. Although income growth has been faster than in our sample of African countries, Vietnam experienced a slower increase in the CS employment share, with much of the decline in agricultural employment absorbed by manufacturing growth. As a result, our framework points to very different conclusions about the sources of growth and their welfare incidence.[25] As shown in panel (d) of Figure 7, the lion's share of structural transformation toward services is driven by productivity growth in the tradable sector. Our methodology infers very little productivity growth in the service sector. In fact, in urban Vietnam we even estimate a slight decline in CS productivity. As seen in panel (c), this implies that Vietnam's welfare gains are driven almost entirely by tradable productivity growth, with virtually no contribution from service-led growth.

[25] We use the same data sources for Vietnam as we use for the countries in sub-Saharan Africa. Because we did not have access to microdata on earnings, we took an urban-rural productivity premium γ of 1.15 as in Table 2 for ease of comparison.

Figure 7: Service-Led Growth: Structural Change and Welfare Effects

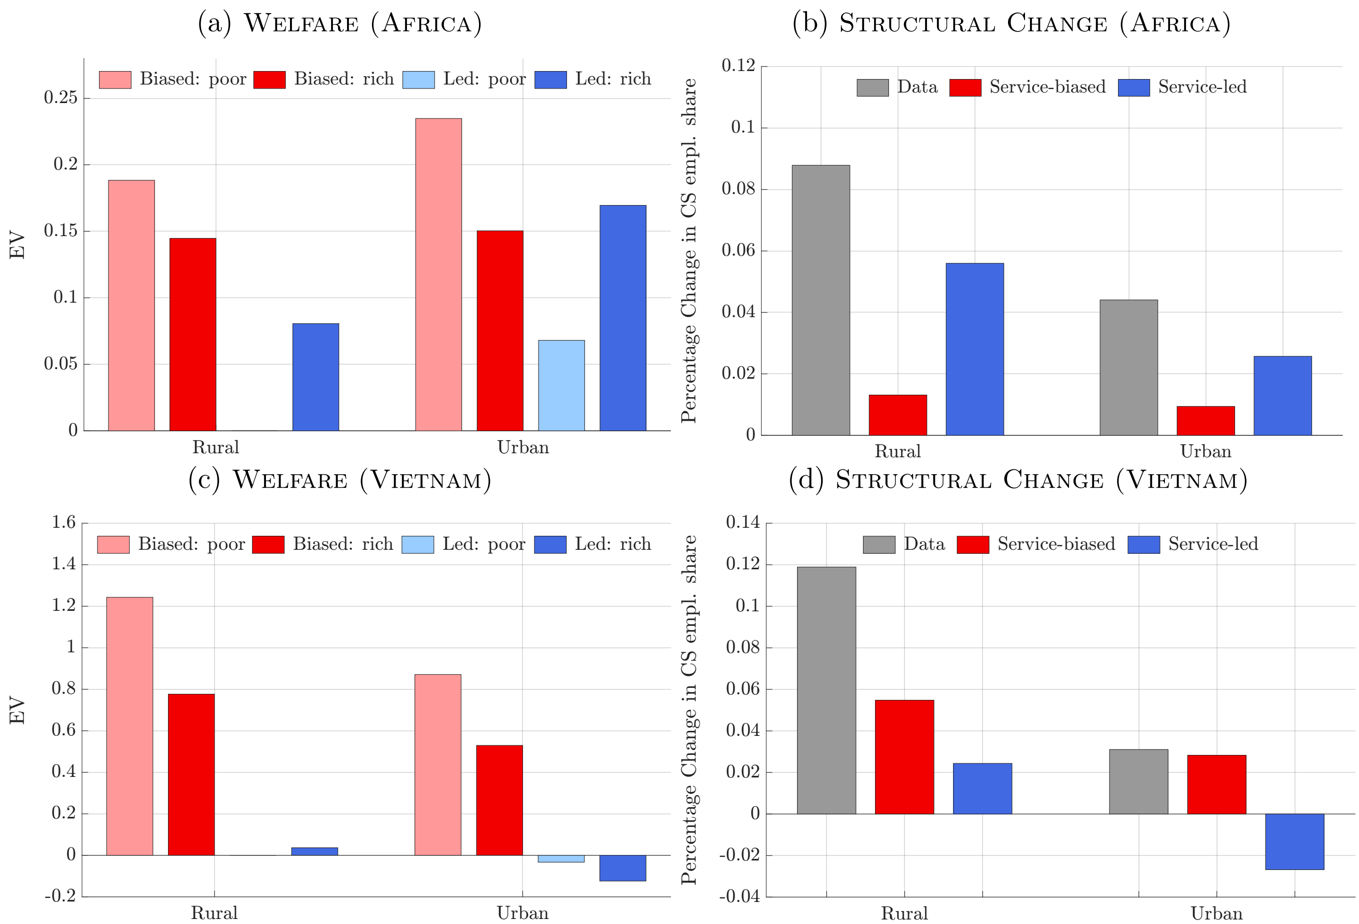

Notes: The left panels report the associated welfare changes, measured as percent changes in income. The red bars correspond to service-biased growth (allowing only $\mathcal{A}_{rt}^{T}$ to change) and the blue bars to service-led growth (allowing only $\mathcal{A}_{rt}^{CS}$ to change). Light bars show the effects for the poor (the 5% of the local income distribution) and dark bars show the effects for the rich (the 95%). The right panels show the change in the CS employment share in the data (gray) and the counterfactual changes under service-biased growth (red) and service-led growth (blue). Panels (a) and (b) correspond to our calibration for our sample of African countries. Panels (c) and (d) correspond to Vietnam.

Taking Stock. In this section, we developed a framework that captures the key forces behind structural transformation in developing countries. It also highlights the distributional implications of service-led growth: richer households devote a larger budget share to services and therefore benefit more from falling service prices (or rising quality).

The model can be taken to the data in a transparent way. The PIGL demand system allows a consistent mapping between final expenditures and sectoral value added and permits the construction of a representative household whose demand system aggregates correctly. We illustrated the approach using representative features of fast-growing African economies and contrasted them with Vietnam's industrialization path. While the quantitative conclusions are necessarily suggestive given the simplicity of the baseline framework, they align with the richer empirical analysis in Fan et al. (2023). Future work with alternative parameterizations and microdata can assess the robustness of these findings.

An important question left open for future research concerns the mechanisms and extent of formalization across developing economies. In many countries, including fast-growing ones such as India, service-led growth has so far not been accompanied by a substantial decline in informality. By contrast, economies such as Vietnam, where growth has been more industry-led, appear to have experienced a faster reduction in informality, particularly among younger and more skilled workers (see McCaig and Pavcnik (2017)). Whether this contrast reflects differences in income levels or instead the nature of jobs created in services versus manufacturing remains unclear. Relatedly, recent work by Dix-Carneiro et al. (2026) shows that trade liberalization reduces informality in the tradable manufacturing sector, as exposure to foreign competition and reallocation toward formal exporters lower the share of informal firms producing tradable goods, while informality in the non-tradable sector remains largely unchanged.

4 Sources of Service-Led Growth: Agglomeration and Labor Market Frictions

Our results above suggest that rising productivity in CS can play an important role in overall growth and structural change. This conclusion naturally raises the question of where such productivity gains originate. After all, retail or hospitality are usually not viewed as activities in which innovation or other productivity-enhancing investments play a major role.

In this section, we discuss three potential mechanisms that could drive productivity growth in local services. First, in Section 4.1, we extend our baseline framework above with a simple model of *service agglomeration*. Specifically, the range of services that can be offered in equilibrium is limited by the size of the local market. Production often requires a fixed cost to open a store or a restaurant, but is otherwise subject to constant returns. Therefore, in a world where consumers love variety and are willing to pay for it, the effective productivity of services (which in our model reflects consumers' willingness to pay) depends on the scale of local demand. As such, this mechanism can explain the emergence of *consumer cities*, that is, why service productivity is particularly high in urban areas.[26] In addition, this extension also allows to quantitatively assess the relative importance of this endogenous mechanism relative to other forces of service-biased productivity growth such as digitization.

Second, in Section 4.2, we explicitly introduce labor market frictions to capture the possibility that some service jobs such as informal retail represent a form of disguised unemployment. This implies that labor reallocation within the service sector can be a source of productivity growth if it reduces misallocation. If, for example, a large share of service output is produced by informal, self-employed workers, falling labor market frictions can lead to formalization and productivity gains.

Third, in Section 4.3, we discuss the implications of actual productivity growth in the informal provision of service output. These improvements can, for example, reflect access to digital technologies through the rising penetration of cell-phones or improvements in the financial system such as small-scale loans that allow for larger inventories.

[26] A recent illustration of this mechanism is offered by Graff (2025), who studies the provision of repair services for equipment in Uganda. He documents that the limited market size restricts the availability of prompt repair services, thereby reducing effective productivity and increasing the depreciation rate of capital in poor economies. Different from his paper, we emphasize the role of love for variety in consumption. However, the mechanism is similar: market size is a crucial constraint on specialization and on the effective productivity of service activities.

4.1 Agglomeration and Consumer Cities

Our empirical results above and the findings of Fan et al. (2023) indicate that the effective productivity of services is particularly high in urban areas. Given their non-tradable nature, the fact that cities are dense and the urban population is relatively rich suggests the presence of market size effects. To capture this possibility, we extend our baseline framework with a model of agglomeration along the lines of Krugman (1980), as in Blanchard et al. (2025).

Specifically, rather than assuming that services are produced with exogenous productivity $\mathcal{A}_r^{CS}$ as before, assume that CS are produced by a CES composite of different service providers, i.e.

$$y_r^{CS} = \left(\int_0^{N_r} s_r^{CS} \left(\nu \right)^{\frac{\sigma-1}{\sigma}} d\nu \right)^{\frac{\sigma}{\sigma-1}}. \tag{21}$$

Here N_r is the number of service stores in region r and $s_r^{CS}(\nu)$ is the output of variety ν. Intuitively, to produce a final good n one needs physical goods and CS *content*. The CS content for each good consists of a variety of different activities: transportation, retail workers, wholesale workers, clerks, etc. The CES structure (21) parsimoniously summarizes the breadth of these CS activities. And because such activities are local in nature, they depend on the local mass of varieties N_r.

Each service variety ν is produced with productivity α_r, and opening a new establishment requires f_r workers. Note that both the entry costs and productivity can vary across space, reflecting, for example, differences in regulation or access to other forms of infrastructure such digitization of electricity. Each firm charges the standard CES markup, so its price is $(\sigma/(\sigma-1))\, w_r/\alpha_r$. Aggregating across varieties (see Section SM-2.1 for details) yields

$$q_r = \frac{\sigma}{\sigma-1} \frac{w_r}{\alpha_r N_r^{\frac{1}{\sigma-1}}} \equiv \frac{\sigma}{\sigma-1} \frac{\mathcal{A}_r^T}{\mathcal{A}^{CS}(\alpha_r, N_r)}, \tag{22}$$

Thus, CS productivity A_r^{CS} combines exogenous productivity α_r with the endogenous effect of market scale, N_r. In particular, a larger number of establishments increases CS productivity through the provision of more service varieties.

To solve the equilibrium number of service establishments, note that each store's profits are a fraction $1/\sigma$ of overall revenue. And because free entry requires that $\pi_r^{CS} = w_r f_r$, the equilibrium mass of establishments in region r is given by

$$N_r = \frac{1}{\sigma} \frac{\mathcal{S}_r^{CS}}{w_r f_r} = \frac{1}{\sigma f_r} \vartheta_r^{CS} L_r \mathbb{E}_r \left[z \right], \tag{23}$$

where $\mathcal{S}_r^{CS} = L_r w_r \mathbb{E}_r[z]\,\vartheta_r^{CS}$ is aggregate spending on CS inputs (see (11)). Hence, a higher service spending share, higher population density, and more human capital allow a larger number of CS establishments to break even.

To compute the equilibrium, note that ϑ_r^{CS} still satisfies the expression in Proposition 2:

$$\vartheta_r^{CS} = \omega_{CS} + \nu_{CS}\mathcal{I}_r \left(\frac{\mathbb{E}_r[z]\,\mathcal{A}_r^T}{q_r^{\omega_{CS}}}\right)^{-\varepsilon}. \tag{24}$$

Equations (22), (23), and (24) jointly determine the number of establishments N_r, the local service price q_r, and the service spending share ϑ_r^{CS}. Together they generate a self-reinforcing agglomeration mechanism. Higher income or population density raises service demand, which attracts additional service providers. Entry expands the range of available varieties, increases effective productivity, lowers the price q_r, and further stimulates service spending. This feedback can be strong enough to generate multiple equilibria, as in the classic "Big Push" literature (Buera et al., 2021, Murphy et al., 1989). When entry is limited, service productivity and demand remain low; when entry is extensive, a dense service sector emerges and spending on services rises.

To illustrate this dynamics graphically, consider Figure 8 (see Section SM-2.1 for details). The blue curve depicts the free entry condition, that is, the supply of service establishments, given by (23). The red curve shows the equilibrium CS expenditure share as a function of the number of establishments.[27] For very low levels of N_r, the CS expenditure share is equal to zero because service prices are prohibitively high. As N_r increases, ϑ_r^{CS} rises, since falling service prices make consumers richer in real terms. When N_r becomes large, ϑ_r^{CS} approaches the asymptotic spending share ω_{CS}.

The equilibrium is described by the intersection between the supply and demand schedules. In Figure 8, we indicate the relevant stable equilibrium with a grey circle.[28] In the left panel, we depict how growth in tradable productivity endogenously *increases* service productivity: higher incomes increase overall service spending, which allows for more entry.[29] In the right panel, we

[27] Formally, this is equation (24) after substituting for q_r in (22).

[28] As we describe in more detail in Section SM-2.1, we consider the other two intersections as less relevant. The interaction at zero means that if there is no demand, no one enters. The intersection in between is an unstable equilibrium.

[29] Note that this link relies crucially on the presence of non-homotheticities. If $\nu^{CS} = 0$, then $N_r = \frac{\omega_{CS} L_r \mathbb{E}_r[z]}{\sigma f_r}$. In that case, N_r *only* depends on human capital and population density and entry costs but is independent of $\mathcal{A}_r^T$.

Figure 8: Service Agglomeration and Endogenous Productivity

(a) Higher tradable productivity

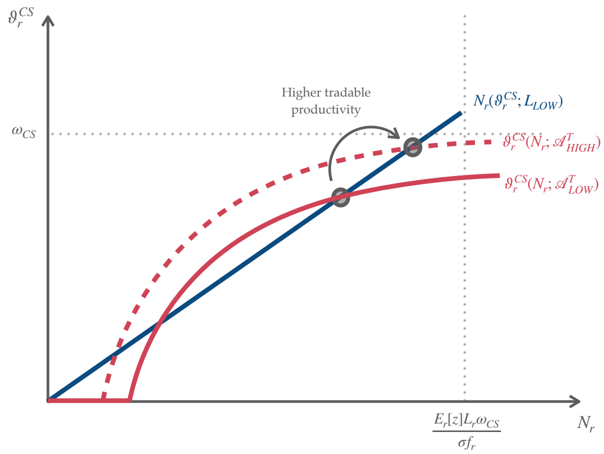

(b) Higher Population Density

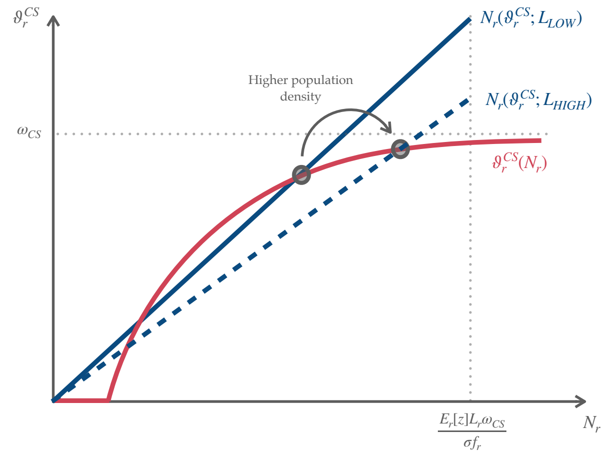

Notes: The figure shows the equilibrium of the model with agglomeration. In the left (right) panel, we depict the consequences of an increase in tradable productivity (population density).

depict the effect of population growth (in the time-series) or higher population density (across locations). While service demand does not depend on population size directly, the supply of service establishments increases. The accompanying fall in service prices in turn increases service spending, which amplifies the initial effect of higher population density. Hence, our theory implies that more service establishments will open in places with high density, lots of human capital, low entry costs, and high incomes. As such, it can at least qualitatively rationalize the emergence of consumption cities with a vibrant service sector.[30]

An appealing feature of this formulation of local service agglomeration is that we can directly link it to the productivities estimated in our baseline model. As we show in detail in Section SM-2.1, this model allows for an *exact* decomposition of the productivity estimates from the baseline model. More specifically, let $d\ln\mathcal{A}^{CS}_{rt}$ denote the estimated growth rate of CS productivity in the baseline

[30] Interestingly, as can be seen from (22), whether services are cheaper in cities is theoretically ambiguous. High tradable productivity makes restaurants expensive. But variety is also higher in cities, both because of higher $\mathcal{A}^T_r$ and because of higher density L_r and human capital $\mathbb{E}_r[z]$. The overall effect on prices, therefore, depends on whether the "Krugman force" of higher variety dominates the neoclassical Balassa-Samuelson type effect. The tug of war between these two effects is discussed in detail in Blanchard et al. (2025).

model reported in Table 3. One can then write

$$d\ln\mathcal{A}_{rt}^{CS} = \underbrace{d\ln\left(\alpha_{rt} f_{rt}^{-\frac{1}{\sigma-1}}\right)}_{\text{Exogenous CS productivity}} \quad \times \quad \underbrace{\frac{1}{\sigma-1} d\ln\left(\vartheta_r^{CS} L_r \mathbb{E}_r\left[z\right]\right)}_{\text{Endogenous agglomeration}}. \tag{25}$$

Hence, given the parameter σ, CS productivity growth from our baseline model can be exactly split into an *exogenous* component $\alpha_r f_r^{-\frac{1}{\sigma-1}}$ and an *endogenous* component driven by overall service spending $\vartheta_r^{CS} L_r \mathbb{E}_r\left[z\right]$.[31] As a concrete example of actual productivity growth in consumer services, $d\ln\alpha_{rt}$, consider the impact of improvements in digital technologies and rising access to digital means of communication for the functioning of retail, which we discuss in more detail in Section 4.3 below.

For a given value of σ, we can perform the decomposition in (25) with the same data that we used to calibrate the baseline model. As we show in detail in Section SM-2.1 in the Supplementary Material, this mechanism is quite powerful. Even for an elasticity of 8, which is consistent with many service varieties being highly substitutable, the expansion in market size driven by human capital accumulation, population growth, and rising CS employment can account for a large share of CS productivity growth.

4.2 Labor Market Frictions and the Rise of the Formal Sector

The agglomerative nature of services discussed above provides one plausible reason for why service productivity is larger in cities and why it has been growing over time. In this section, we discuss a different mechanism that can shed light on the sources of service-led growth: labor reallocation in the presence of labor market frictions. The analysis builds on the recent work of Schwartzman (2025).

Informal employment is widespread in developing economies, particularly in retail and transportation, where self-employment and informal arrangements are common (Ulyssea, 2020). This dual structure has two implications. First, misallocation may shape measured sectoral productivity: if informal jobs are less productive, formalization can raise aggregate efficiency. Second, a

[31] Note also that average service productivity α_r is not separately identified from f_r, which implies that fundamental spatial differences in service productivity could be due either to differences in labor productivity α_r or to differences in entry costs f_r.

large employment share in CS need not signal high productivity or strong demand. It may instead reflect frictions that keep workers in low productivity informal activities, which must be taken into account when estimating service sector productivity.

In this section, we show that our framework can be easily extended to incorporate such consideration. Specifically, we consider an environment in which formal jobs are scarce, and frictions constrain access to the formal sector. We show that improvements in labor market fluidity are sometimes equivalent to productivity gains in the CS sector and are therefore part of service-led growth.

To capture the possibilities of labor market frictions within our baseline framework, we modify the labor market to feature a dual structure comprising both *formal* and *informal* jobs. For simplicity, we assume that output from the two modes of production is perfectly substitutable. The formal sector operates as in the baseline model: each efficiency unit of labor produces $\mathcal{A}_r^s$ units of output in sector s. In contrast, informal workers produce only $\phi_r^s < \mathcal{A}_r^s$ units of output, reflecting the productivity penalty associated with informality.[32] Hence, absent frictions, all output would be produced in formal establishments.

However, access to formal jobs is rationed: only a fraction μ_r of individuals can obtain formal employment, while the remaining $1 - \mu_r$ are restricted to informal work. We interpret μ_r as a reduced-form measure of both policy-induced and institutional labor market frictions. On the labor supply side, we retain the same structure as in the baseline model: each individual is endowed with z efficiency units, drawn from the Fréchet distribution $F_r(z)$ specified in (9). For simplicity, we assume that individual ability z and access to the formal labor market are independent, that is, there is no selection on ability.[33]

Letting q_r again denote the price of CS output and w_r the equilibrium formal wage, the earnings of informal (self-employed) and formal workers are given by:

$$y_r^{CS|I}(z) = q_r\, z\, \phi_r^{CS}, \qquad y_r^{T|I}(z) = z\, \phi_r^{T}, \qquad y_r^{F}(z) = z w_r.$$

Because informal workers chose the sector that maximizes their earnings, let the return per informal

[32] The productivity of informal production can still vary across localities because it could depend on mobile connectivity, basic transportation infrastructure, or access to basic financial services.

[33] In Section SM-2.3 of the Supplementary Material, we show how to extend the theory to allow for selection.

efficiency unit be defined as $\phi_r \equiv \max\{q_r \phi_r^{CS}, \phi_r^T\}$. Note that, in equilibrium, it must hold that $w_r \geq \phi_r$; otherwise, workers with access to formal jobs would have no incentive to supply labor in the formal market.

Per capita earnings and spending in this economy are given by:

$$\mathcal{Y}_r / L_r = \mathbb{E}_r[z] \overline{w}_r \quad \text{where} \quad \overline{w}_r = \mu_r w_r + (1 - \mu_r)\phi_r, \tag{26}$$

where $\overline{w}_r$ denotes the average returns per efficient unit. Equation (26) shows that labor market frictions depress per capita income. Since $w_r \geq \phi_r$, the marginal revenue product of labor is higher in the formal sector than in the informal sector. Hence, improvements in labor market institutions, captured by an increase in μ_r, or a rise in the service price q_r, raise overall income.

To compute the overall demand for CS output, we must account for the fact that informal workers are poorer and therefore spend less on CS. Leveraging our earlier results on the aggregative properties of PIGL contained in Proposition 2, we show in Section SM-2.2 in the Supplemental Materials that the aggregate service spending share is given by

$$\vartheta_r^{CS} = \omega_{CS} + \nu_{CS}\, \mathcal{I}_r \left(\frac{\mathbb{E}_r[z]\, \overline{w}_r}{q_r^{\omega_{CS}}} \right)^{-\varepsilon} \tag{27}$$

where the inequality index, $\mathcal{I}_r$, is now given by

$$\mathcal{I}_r = \underbrace{\frac{\int z^{1-\varepsilon} dF_r(z)}{\left(\int z dF_r(z)\right)^{1-\varepsilon}}}_{\text{Within group inequality}} \underbrace{\frac{(1-\mu_r)\left(\frac{\phi_r}{w_r}\right)^{1-\varepsilon} + \mu_r}{\left((1-\mu_r)\left(\frac{\phi_r}{w_r}\right) + \mu_r\right)^{1-\varepsilon}}}_{\text{Across group inequality}}.$$

Equation (27) shows that CS demand has the same structure as in our baseline model except that the demand shifter $\mathcal{I}_r$ is now endogenous and reflects the combination of *within-group inequality* (i.e., inequality within the group of formal and informal workers) and *across-group inequality* (i.e., inequality between formal and informal workers). As such, overall demand depends on labor market frictions (μ_r) and the relative sectoral earnings wedge (ϕ_r / w_r). As $\mu_r \to 1$ or $\phi_r \to w_r$, labor misallocation and inequality across groups disappear, and the model converges to the benchmark case.

This expression also shows how the presence of labor market frictions interferes with the identification of CS productivity. First, holding average earnings $\mathbb{E}_r[z]\, \overline{w}_r$ fixed, the aggregate CS expenditure share ϑ_r^{CS} depends both on CS productivity (via CS prices q_r) *and* on labor market frictions (via μ_r and ϕ_r/w_r). Second, in the presence of frictions, the spending share no longer coincides with the employment share, precisely because some workers might be "underpaid". Letting $s_r^{CS|I}$ denote the share of informal workers employed in the CS sector, we show in Section SM-2.2 in the Supplemental Material that the CS employment share satisfies

$$\ell_r^{CS} = \vartheta_r^{CS} + (1-\mu_r)\left(1-\frac{\phi_r}{w_r}\right)\left(s_r^{CS|I} - \vartheta_r^{CS}\right). \tag{28}$$

Employment therefore exceeds expenditure when formal jobs are rationed ($\mu_r < 1$), formal wages exceed informal earnings ($w_r > \phi_r$), and informal workers concentrate more heavily in CS than the population as a whole ($s_r^{CS|I} > \vartheta_r^{CS}$).

Equilibrium To determine the equilibrium, note that generally one sector will fully specialize in formal work, while the other will exhibit a dual structure of formal and informal employment.[34] To highlight the prevalence of informal workers in CS, we assume that informal workers have a technological comparative advantage in the CS sector, i.e. $\phi_r^{CS}/\mathcal{A}_r^{CS} > \phi_r^T/\mathcal{A}_r^T$. Economically, informal food stalls may be closer substitutes to formal supermarkets than self-employed craftsmen are to formal industrial plants.[35] This implies that the equilibrium can take three forms:

1. All formal workers are employed in the tradable sector, and informal workers are employed in both CS and tradables;

2. Tradable production is entirely formalized and CS production is purely informal;

3. Both sectors feature formal workers and all remaining informal workers are employed in CS.

[34] If formal workers are employed in both sectors, then $w_r = \mathcal{A}_r^T = q_r\mathcal{A}_r^{CS}$. Similarly, if informal workers are present in both sectors, $\phi_r^T = q_r\phi_r^{CS}$. These two equations can only hold under the knife-edge condition $\frac{\mathcal{A}_r^T}{\mathcal{A}_r^{CS}} = \frac{\phi_r^T}{\phi_r^{CS}}$.

[35] This assumption is well suited to analyzing the contrast between manufacturing and service activities. The CS sector is often the primary source of informal employment in developing countries, as shown by Schwartzman (2025). In the case of agriculture, the distinction is more blurred. One could, in principle, distinguish between subsistence farming and mechanized farms, as in Storesletten et al. (2025). However, in very poor countries, most farming activity is conducted in traditional labor-intensive agriculture. We do not pursue this distinction in this section and defer a discussion of the separate role of agriculture to Section 5.2 below.

Figure 9: The Consumer Service Sector with Labor Market Frictions

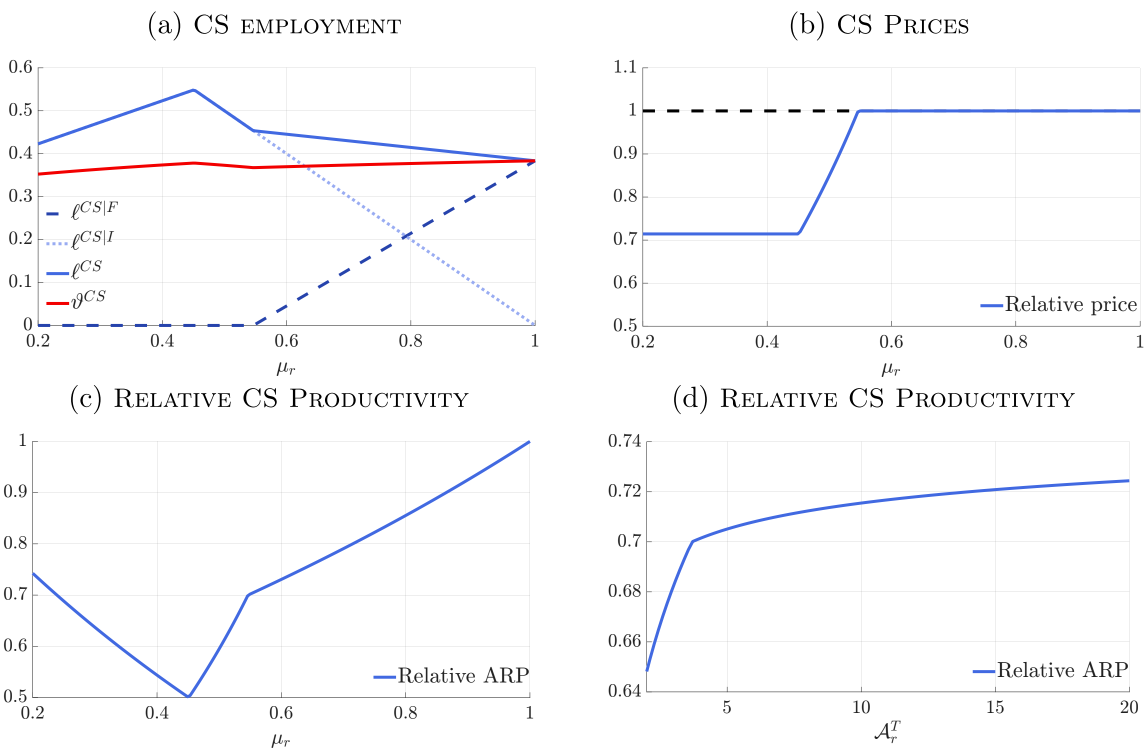

Notes: Panel (a) displays the aggregate CS employment share ℓ_r^{CS}, the CS expenditure share ϑ_r^{CS}, and total CS employment among formal, $\ell_r^{CS|F} = \mu_r s_r^{CS|F}$, and informal, $\ell_r^{CS|I} = (1 - \mu_r) s_r^{CS|I}$, workers as a function of μ_r. Panel (b) depicts the CS price q_r as a function of μ_r. Panels (c) and (d) depict relative CS revenue productivity, ARP_r^{CS}/ARP_r^T, as a function of μ_r and $\mathcal{A}_r^T$, respectively.

In Figure 9 we display the equilibrium properties graphically.[36] Panels (a) and (b) show, respectively, the allocation of CS employment and CS prices as functions of labor market frictions μ. A striking feature is that total CS employment (shown in blue) first increases with μ and then declines. The initial increase is intuitive. Given our assumptions on comparative advantage, informal workers are present in both sectors, but workers who formalize do so first in the tradable sector rather than in CS. This reallocation raises aggregate income, which in turn increases the demand for CS output. The additional demand is met by a reallocation of remaining informal workers from the tradable to the CS sector. In this sense, a higher μ reallocates informal tradable workers both into the formal tradable sector and into the informal CS sector.

Panel (b) also shows that CS prices rise as μ rises. Labor market frictions keep workers poor and depress aggregate service demand, which lowers CS prices. At some point, all remaining informal workers are employed in the service sector. With further formalization, total CS employment then declines, because the marginal informal worker enters formal firms in the tradable sector. Hence, total service output falls, and prices rise in response.

Finally, as formalization progresses, service prices become high enough for formal jobs to be viable in the service sector, and further formalization increases formal employment in services at constant prices (shown in dashed blue). Total CS employment continues to decline because each formal worker produces more output than an informal worker, reducing the number of workers needed to meet demand. Note also the disconnect between the employment share and the expenditure share (shown in red). As long as frictions remain, employment exceeds expenditure. As $\mu \to 1$, the difference disappears, and the employment and expenditure shares coincide.

In the bottom panels, we depict the implications for *revenue* productivity in the CS sector. Specifically, letting ARP_r^{CS} and ARP_r^T denote the average revenue product in the CS and tradable sector, we plot *relative* service productivity, ARP_r^{CS}/ARP_r^T, as a function of labor market frictions μ (left panel) and tradable productivity $\mathcal{A}_r^T$ (right panel).[37] The relationship between

[36] See Section SM-2.2 for a formal characterization.

[37] As we show in Section SM-2.2, these are given by

$$ARP_r^{CS} = q_r \mathcal{A}_r^{CS}\left(1 - \frac{\left(1 - \frac{\phi_r^{CS}}{\mathcal{A}_r^{CS}}\right)(1-\mu_r)s_r^{CS|I}}{\ell_r^{CS}}\right) \quad \text{and} \quad ARP_r^T = \mathcal{A}_r^T\left(1 - \frac{\left(1 - \frac{\phi_r^T}{\mathcal{A}_r^T}\right)(1-\mu_r)s_r^{T|I}}{\ell_r^T}\right),$$

formalization and relative revenue productivity in services is U-shaped. The reason is the asymmetric reallocation process shown in panel (a). At low levels of formalization, it is the tradable sector that formalizes first. Hence, the gains from reallocation are first realized in tradables, and relative revenue productivity in services, which is already lower, declines further. Once all informal workers are employed in services, rising service prices, as shown in panel (b), *raise* the revenue productivity of services. Finally, as informal self-employment gives way to more productive formal jobs, aggregate CS productivity increases even as total CS employment continues to fall. This decoupling arises from changes in the composition of employment and from the productivity gap between formal and informal work. Relative to our baseline model, labor market frictions therefore break the tight relationship between CS employment shares and CS productivity.

Finally, in panel (d), we show the link between physical productivity in the tradable sector, $\mathcal{A}_r^T$, and the sectoral revenue productivity gap. Interestingly, in the presence of labor market frictions, higher physical efficiency in tradable goods increases relative service productivity. As consumers become richer, the demand for services rises, which increases prices and therefore revenue productivity.

The patterns shown in Figure 9 paint a plausible picture of the development process. In very poor economies, most workers are excluded from formal labor markets, and the few with access to formal jobs tend to work in the manufacturing sector (and possibly in the public sector, which we abstract from in this model). Productivity and earnings in informal service jobs remain low, and there is a large reservoir of informal workers in the tradable sector (which in this simple two-sector version of our framework also contains agriculture). Improvements in labor market institutions allow informal tradable workers to obtain formal jobs, incomes rise, and informal work shifts from agricultural activities to the informal provision of CS. At later stages of development, however, the service sector acquires a dual structure comprising both formal and informal employment.

and therefore reflect service prices (q_r), the relative employment shares across sectors (ℓ_r^s), and the allocation of informal workers ($s_r^{s|I}$). Note that if there are no frictions ($\mu = 1$) or if there is no productivity gap ($\phi^s = \mathcal{A}^s$), revenue productivity is equalized, because $w_r = q_r \mathcal{A}_r^{CS} = \mathcal{A}_r^T$.

4.3 Productivity Growth in the Informal Sector

The discussion in Section 4.2 highlighted the link between formalization and service productivity. While this might be relevant in countries such as Brazil that are already further along the development process (Schwartzman, 2025), the reality of many countries in Africa looks quite different: informal means of production dominate the service sector and broad-based formalization has yet to take place. This, however, does not mean that service productivity is necessarily stagnant, because many new technologies have the potential to make the provision of services more efficient, even if they are performed informally. As mentioned above, examples are improvements in digital technologies, that enable firms to improve inventories, facilitate the matching between buyers and seller, or reduce price dispersion through lower search costs (see, e.g., Jensen (2007) or Bergquist et al. (2024)).

To see the implications of this form of technological progress, we build on the labor market friction model of Section 4.2 and consider an increase of informal CS productivity, ϕ_r^{CS}, relative to the formal provision of CS. The results of this exercise are summarized in Figure 10. In the initial equilibrium, where $\phi_r^{CS}/\mathcal{A}_r^{CS} = 0.5$, all formal workers are employed in the tradable sector and informal workers are employed both in tradables and in consumer services. As we show in Panel (a), an increase of informal CS productivity raises both CS employment and CS spending. This is due to the fact that overall service productivity rises (panel (c)) and that the price of CS inputs, q_r, falls (panel (b)). Interestingly, the impact on *revenue* productivity is non-monotone. First, relative revenue productivity in the CS sector declines. The reason is that informal workers leave the tradable sector and that the price of CS output falls. Once, the tradable sector is solely populated by formal firms and no labor reallocation takes place, the decline in prices continues at a slower pace and revenue productivity in services increases. As such, Figure 9 highlights how rising productivity of informal CS activities can play a role for service-led growth, that is rising living standards and rising service employment.

Figure 10: Productivity Growth of Informal CS Providers

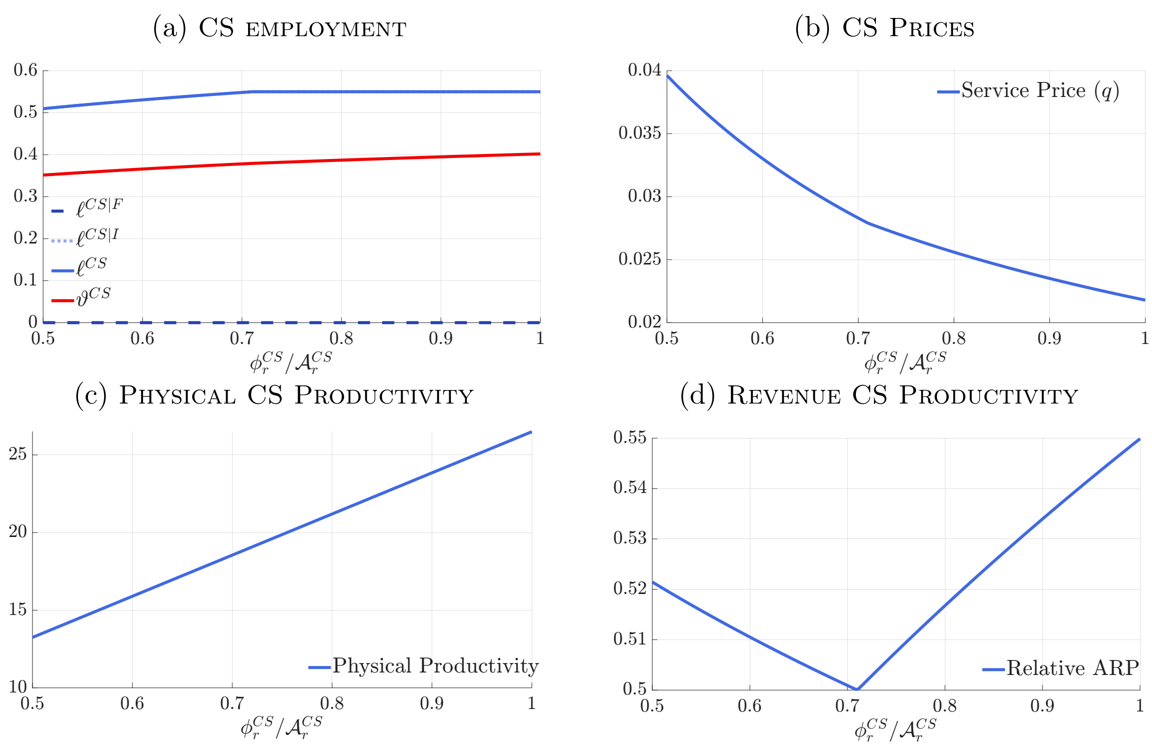

Notes: Panel (a) displays the aggregate CS employment share ℓ_r^{CS}, the CS expenditure share ϑ_r^{CS}, and total CS employment among formal, $\ell_r^{CS|F} = \mu_r s_r^{CS|F}$, and informal, $\ell_r^{CS|I} = (1-\mu_r)s_r^{CS|I}$, workers as a function of $\phi_r^{CS}/\mathcal{A}_r^{CS}$. Panel (b) depicts the CS price q_r as a function of $\phi_r^{CS}/\mathcal{A}_r^{CS}$. Panels (c) and (d) depict physical and relative CS revenue productivity, ARP_r^{CS}/ARP_r^{T}, as a function of $\phi_r^{CS}/\mathcal{A}_r^{CS}$, respectively.

4.4 Taking Stock

In this section, we have discussed three perspectives on service productivity. Fist, economies of scale in the provision of services can generate a form of agglomeration, whereby rising demand triggers the entry of new service providers, and productivity rises as a consequence. This mechanism can, at least qualitatively, explain the presence of consumption cities and our finding that service productivity increases over time. Second, labor market frictions can shape both service employment and service productivity. Formalization can raise service productivity if informal workers are channeled into formal employment (even though the relationship is non-monotone), and rising tradable productivity can do so as well. Third, new technologies might have the potential to increase the productivity of the informal provision and can, at least temporary, be a form of service-led growth.

For expositional clarity, we have discussed these aspects separately. However, especially between the first two there is an obvious interconnection, which is explored in Schwartzman (2025). If few formal service establishments enter, overall productivity in the service sector is low because informal jobs dominate. This depresses income and aggregate demand, which in turn can make limited entry optimal. Conversely, widespread entry fosters a vibrant service sector, which encourages consumers to shift their spending toward services. Improved labor market institutions can therefore support formal job creation, greater firm entry, and higher aggregate income and service consumption.

We view the mechanisms discussed here as plausible, theoretical possibilities that we hope future, empirically-oriented work could analyze in more detail. The link between economic development and formalization (or the lack thereof) seems particularly relevant to us.[38]

[38] The existing literature contains some evidence that rising exports leads to formalization within the manufacturing sector (see, e.g., McCaig and Pavcnik (2018)). This is consistent with the theory proposed in Dix-Carneiro et al. (2026), that also implies a shift toward formal employment within the manufacturing sector in response to trade opening. At the same time, Schwartzman (2025) shows that the time-series decline in informality across countries is to a large extent driven by formalization within non-tradable services.

5 Extending the Baseline Framework

A strength of our framework is its portability. In this section, we show how it can be extended to different contexts. In Section 5.1, we allow different regions to produce regional varieties of tradable goods that are imperfect substitutes. In Section 5.2, we disaggregate the tradable sector into industries such as agriculture and manufacturing, as well as into goods production and producer services. These extensions make it possible to study natural patterns of specialization, such as rural regions specializing in agriculture and urban regions specializing in producer services. Next, we consider international trade and the role of natural resources in shaping industrialization patterns (Section 5.3). In Section 5.4, we extend the framework to incorporate heterogeneous skills, where human capital is not perfectly substitutable, and labor supply choices can influence structural change. Finally, we allow for regional migration between urban and rural areas to illustrate how service-led growth in cities can lead to spatial sorting (Section 5.5).

5.1 Regional Varieties, Specialization, and Spatial Inequality

In the baseline model the tradable composite x was homogeneous across regions, so no trade occurred in equilibrium and the urban-rural wage gap was fully determined from productivity. We now allow x to be a CES aggregate of regional varieties,

$$x = \left(x_U^{(\varrho-1)/\varrho} + x_R^{(\varrho-1)/\varrho} \right)^{\varrho/(\varrho-1)},$$

where $\varrho > 1$ is the elasticity of substitution.[39] The baseline framework is nested as the special case of $\varrho \to \infty$.

When regional varieties are imperfect substitutes, our model features *endogenous regional wage differences*. To derive the urban-rural wage gap, note that trade balance requires that total spending by rural households on the urban variety be equal to total spending by urban households on

[39] The EOS ϱ plays the role of the trade elasticity, which has consistently been estimated to be larger than one (see, e.g., Caliendo and Parro (2015)).

the rural variety.[40] This condition implies that

$$\frac{w_U}{w_R} = \underbrace{\left(\frac{\mathbb{E}_U[z] L_U}{\mathbb{E}_R[z] L_R}\right)^{-\frac{1}{\varrho}} \times \left(\frac{\mathcal{A}_U^T}{\mathcal{A}_R^T}\right)^{\frac{\varrho-1}{\varrho}}}_{\text{Standard CES Properties}} \times \underbrace{\left(\frac{1-\vartheta_U^{CS}}{1-\vartheta_R^{CS}}\right)^{-\frac{1}{\varrho}}}_{\text{Novel PIGL Effect}}. \tag{29}$$

In general, under non-homothetic preferences the urban wage premium is not only driven by tradable technologies and labor endowment, but also by the local service spending in the two regions, which in turn depends on the relative income. Specifically, if the urban region is richer and allocates a larger share of income to non-tradable CS goods (as in the consumer city parable), this reduces the supply of the urban tradable good, thereby increasing its price. The resulting terms-of-trade effect hurts rural households and increases the urban wage premium.

To fully solve for the equilibrium, observe that the CS expenditure share is still given by the expression in Proposition 2,

$$\vartheta_r^{CS} = \omega_{CS} + \nu_{CS}\, \mathcal{I}_r \, \left(\left(\mathcal{A}_r^{CS}\right)^{\omega_{CS}} \mathbb{E}_r[z]\, w_r^{1-\omega_{CS}}\right)^{-\varepsilon}.$$

Together with equation (29) and the choice of the tradable good as the numeraire, this expression fully determines the equilibrium.

A particular transparent expression for spatial inequality arises if the asymptotic expenditure share on services, ω_{CS}, is close to unity. In that case, the urban-rural *earnings* gap, accounting for human capital differences, is given by

$$\Delta_{RU} \equiv \frac{w_U\ \mathbb{E}_U[z]}{w_R\ \mathbb{E}_R[z]} \approx \underbrace{\left(\frac{L_U}{L_R}\right)^{-\frac{1}{\varrho}}}_{\text{Rel. size}} \times \underbrace{\left(\frac{\mathcal{A}_U^T}{\mathcal{A}_R^T}\right)^{\frac{\varrho-1}{\varrho}}}_{\text{Technology}} \times \underbrace{\left(\frac{\mathcal{I}_U}{\mathcal{I}_R}\right)^{-\frac{1}{\varrho}}}_{\text{Inequality}} \times \underbrace{\left(\frac{\mathbb{E}_U[z]}{\mathbb{E}_R[z]}\right)^{\frac{\varrho-1+\varepsilon}{\varrho}}}_{\text{Human Capital}} \times \underbrace{\left(\frac{\mathcal{A}_U^{CS}}{\mathcal{A}_R^{CS}}\right)^{\frac{\varepsilon}{\varrho}}}_{\text{CS Productivity}}. \tag{30}$$

Equation (30) shows that the urban earning premium depends on five terms. The first two follow from standard properties of the CES production function: An increase in the relative productivity of urban tradable goods raises the urban-rural wage gap, and an increase in the supply of urban

[40] Letting the tradeable good be the numeraire, trade balance requires $\mathcal{S}_{RU}^T = \mathcal{S}_{UR}^T$, where

$$\mathcal{S}_{rk}^T = \underbrace{\left(1-\vartheta_r^{CS}\right) w_r \mathbb{E}_r[z] L_r}_{\text{Total spending on tradables in } r} \cdot \underbrace{\left(w_k/\mathcal{A}_k^T\right)^{1-\varrho}}_{\text{Share on variety from } k}.$$

varieties reduces the terms of trade and hence urban wages. The more novel effects can be broken down into three terms:

1. **[Inequality]** As we already discussed above, higher inequality increases the aggregate demand for services, holding average income constant. Rising inequality in cities (relative to rural areas) therefore reduces the relative supply of tradable goods and *increases* the urban-rural earnings gap.

2. **[Human Capital]** Human capital differences also lead to disparities in spatial earnings. First, higher urban human capital increases the supply of the urban tradable variety and reduces skill prices. The elasticity of human-capital adjusted earnings with respect to human capital is therefore not equal to 1 but only equal to $(\varrho - 1)/\varrho < 1$. Second, higher human capital increases individuals' income and shifts demand towards CS. This second mechanism operating via non-homothetic demand introduces an additional elasticity $\varepsilon/\varrho > 0$.

3. **[CS Productivity]** The third effect is associated with CS productivity. A TFP increase in CS in the urban region (holding rural CS TFP constant) induces a shift in demand toward non-tradable CS, and as a consequence, raises average urban wages (in terms of the tradable numeraire good!).

Expression (30) highlights that the elasticity of substitution ϱ is the key parameter governing these terms-of-trade effects. When $\varrho \to \infty$, as in the baseline framework, rural and urban varieties become perfect substitutes, wages reflect tradable productivity and human capital, thus, $\Delta_{RU} = \frac{\mathcal{A}_U^T}{\mathcal{A}_R^T}\frac{\mathbb{E}_U[z]}{\mathbb{E}_R[z]}$. When $\varrho \to 1$, corresponding to Cobb–Douglas aggregation of regional varieties, local tradable prices fall with local output with unitary elasticity. Spatial wage gaps, therefore, no longer depend on tradable productivity but instead reflect population size, local inequality, CS productivity, and human capital.

This discussion highlights that the welfare consequences of productivity growth differ sharply depending on whether the improvement originates in the tradable or the non-tradable sector. For example, productivity growth in the production of the rural tradable variety, such as that associated with the Green Revolution, lowers the price of tradable goods and thereby raises *real* incomes for both rural and urban residents. By contrast, productivity gains in non-tradable services accrue exclusively to residents of the location in which the improvement occurs, and can adversely affect residents of other areas through the terms-of-trade effect we highlighted. Thus, an expansion of consumer cities driven by urban productivity growth in non-tradable CS reduces the welfare of rural populations that remain behind.

5.2 Inside the Tradable Sector

So far, we have treated the tradable sector as a single aggregate. Tradable goods are produced with some constant returns to scale production function that combines agricultural, manufacturing, and "white collar" service workers; see equation (6). In this section, we introduce additional structure into the aggregate production function f_r to distinguish explicitly between these industries. Doing so allows us to decompose productivity, $\mathcal{A}_r^T$, into its sectoral components and to capture the idea that poor locations are often specialized in agricultural production, whereas business services are more prevalent in urban regions.

To this aim, we assume that the composite tradable good x_r is produced using a CES aggregate of agricultural and industrial goods, while the latter are produced with manufacturing workers and producer services. Unlike CS, producer services, which comprise industries such as IT, business services, and finance, become embedded in tradable goods and thus generate tradable value added

$$x_r = \left[\left(A_r^{AG} H_r^{AG}\right)^{\frac{\eta-1}{\eta}} + \left(y_r^{IND}\right)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \quad \text{and} \quad y_r^{IND} = \left[\left(A_r^{PS} H_r^{PS}\right)^{\frac{\beta-1}{\beta}} + \left(A_r^{MAN} H_r^{MAN}\right)^{\frac{\beta-1}{\beta}}\right]^{\frac{\beta}{\beta-1}}.$$

Here η denotes the EOS between agricultural and industrial inputs, β governs the substitutability between manufacturing inputs and producer services, and A_r^s are region-sector-specific labor productivities that reflect the quality and quantity of land (in the case of agriculture), differences in production technologies (in the case of manufacturing), or agglomeration externalities (for producer services).

Profit maximization implies that $x_r = \mathcal{A}_r^T \left(H_r^{AG} + H_r^{MAN} + H_r^{PS}\right)$, where

$$\mathcal{A}_r^T = \left(\left(A_r^{AG}\right)^{\eta-1} + \left(\left(A_r^{PS}\right)^{\beta-1} + \left(A_r^{MAN}\right)^{\beta-1}\right)^{\frac{\eta-1}{\beta-1}}\right)^{\frac{1}{\eta-1}}. \tag{31}$$

Equation (31) highlights that tradable productivity, $\mathcal{A}_r^T$, can ultimately be viewed as a CES aggregate of local agricultural, manufacturing, and business service productivity. This implies that the underlying reasons for high tradable productivity across space may differ: rural locations might benefit mostly from high agricultural productivity (e.g., due to the abundance of fertile land), while urban centers might leverage their high productivity in business services.

Differences in sectoral specialization can be inferred directly from relative employment shares. Profit maximization yields

$$\frac{H_r^{AG}}{H_r^{MAN} + H_r^{PS}} = \left(\frac{A_r^{AG}}{\left((A_r^{PS})^{\beta-1} + (A_r^{MAN})^{\beta-1} \right)^{\frac{1}{\beta-1}}} \right)^{\eta-1} \quad \text{and} \quad \frac{H_r^{PS}}{H_r^{MAN}} = \left(\frac{A_r^{PS}}{A_r^{MAN}} \right)^{\beta-1}.$$

Given η and β, these ratios identify A_r^{AG}, A_r^{MAN}, and A_r^{PS} and thus allow a decomposition of tradable productivity A_r^T. For example, as long as $\eta > 1$, the large agricultural employment share in rural labor markets indicates a high level of agricultural productivity, relative to manufacturing and producer services, reflecting the abundance of agricultural land.[41]

Producer services play, so far, a marginal role in sub-Saharan Africa. However, they are important in economies like Vietnam or China, where industrial productivity growth has been the main driver of growth (see Chen et al. (2023) for the case of China). Note that, in contrast to CS, fast productivity growth in producer services has beneficial general equilibrium effects on tradable sectors and spreads welfare gains across regions by reducing the price of goods. The relationship between service-led growth and inequality, therefore, crucially depends on whether productivity rises in non-tradable services such as CS or in tradable business services.

[41] The elasticity of substitution between agricultural and non-agricultural goods is a key parameter in many models of structural change. Given the structure of our model, assuming that the two inputs are complements ($\eta < 1$) would generate counterfactual predictions across space. Specifically, land abundance in rural areas implies that relative labor productivity in agriculture is higher in rural than in urban areas. Assuming complementarity would then require that the employment ratio between agriculture and manufacturing be *higher* in urban than in rural areas, which would clearly be at odds with the data. Kelly et al. (2023) provide evidence in line with agricultural and non-agricultural output being substitutes. In Section 6 we discuss the existing evidence in detail. Note that any assumption about η affects only the estimates of relative TFP in agriculture and manufacturing; it does not influence the estimation of TFP in CS.

5.3 International Trade

Our framework can be further extended to be incorporated into an open-economy model that is suitable to study the link between structural change and international trade. Consider a small open developing economy and suppose that final tradable goods are produced as a CES Armington aggregate of domestic and foreign goods:

$$x = \left(x_h^{\frac{\varsigma-1}{\varsigma}} + \varkappa_f^{\frac{1}{\varsigma}} x_f^{\frac{\varsigma-1}{\varsigma}} \right)^{\frac{\varsigma}{\varsigma-1}}, \qquad \varsigma > 1,$$

where $\varkappa_f$ captures the importance of foreign goods and ς is the Armington elasticity of substitution.

The domestic good x_h should be interpreted as the same national tradable composite introduced in the closed-economy setting of Section 5.1, where it is defined as a CES aggregator over the regional tradable varieties x_r. The Armington structure introduced here simply augments that national composite with an imported variety x_f, while preserving the internal spatial aggregation of the closed economy. In the context of Section 5.2, and for the case in which the home economy is a poor developing country, it is natural to interpret x_h as being intensive in agricultural inputs.[42]

Service-Biased Growth in the Open Economy To understand how international trade can induce structural change toward services, note that natural resources play a central role for many low income economies. To capture this, suppose the country is endowed with a resource that requires no labor to extract and is not consumed domestically, so it is fully exported. Let $\mathfrak{b}$ denote the export value of the endowment, which reflects the quantity of the resource, its world price, and the trade costs associated with exporting it.

Consider first the limiting case in which home and foreign goods are perfect substitutes, $\varsigma \to \infty$. In this case, if $\mathfrak{b} = 0$, although goods are tradable, no trade occurs in equilibrium, and the allocation coincides with that of the closed economy. When $\mathfrak{b} > 0$ (while maintaining that $\varsigma \to \infty$), there is a positive trade flow whereby the domestic economy exports the natural resource and imports foreign goods. Sectorally, the availability of natural resources reduces the domestic production of

[42] This characterization refers to equilibrium factor use rather than to a change in the underlying technology. The nested CES structure still combines agricultural and industrial inputs, but in an economy with a comparative advantage in agriculture, the endogenous value share of agricultural inputs is high. This is the sense in which x_h is agriculture-intensive.

tradable goods and shifts employment toward services. Changes in the value of $\mathfrak{b}$ have similar effects. For instance, suppose that $\mathfrak{b}$ increases due to a boom in natural resource prices. Since the domestic economy can now finance more imports of foreign goods, domestic employment shifts away from agriculture and toward services.

When domestic and foreign goods are imperfect substitutes ($\varsigma < \infty$), the domestic economy also exports its own variety of tradable goods. If we think of the domestic economy as a poor sub-Saharan country, it is natural to assume that it has a comparative advantage in agriculture relative to industry. Consequently, it exports agricultural-intensive tradable goods (alongside natural resources) and imports industrial tradable goods. Suppose again that the value of the endowment $\mathfrak{b}$ increases, for example, due to a boom in natural resource prices. Since the domestic economy can now finance more imports of foreign industrial goods, domestic employment shifts away from agriculture and toward services, while exports of primary resources fund higher imports of foreign tradable goods.

The export of natural resources increases domestic service employment through two reinforcing channels: production specialization and non-homothetic demand. First, raw materials provide a cost-effective means of financing imports, reducing the need for domestic goods production. Second, as domestic residents become richer, they shift expenditure from goods to services, further lowering employment in the domestic goods sector. Both trade liberalization and increases in the world price of tradable resources can, therefore, trigger a structural transformation that resembles premature industrialization.

The Global Economy So far, we have focused on a small open economy. The framework, however, can be readily extended to large countries representing different world regions. Consider, for instance, a three-region world comprising the *North*, the *South*, and the *East*. The North is altogether richer and characterized by high wages. Since most of its labor force is employed in services, the goods produced in the North are relatively expensive. The East has a comparative advantage in the production of goods, especially owing to its high productivity in the industrial sector.

In a two-region North-South world, the South might specialize in the production and export of goods owing to its low wages. However, two factors limit these opportunities in the three-region scenario. First, the existence of the East that has a strong technological advantage in the production of goods constrains the South's competitiveness in this sector. Second, the South possesses raw materials that are highly valued in international markets, either for consumption or as inputs in production. In this world, the North naturally specializes in services, exporting a few high-priced varieties to the rest of the world. The East becomes the world supplier of goods, while the South exports primary resources and only a limited range of (agriculturally intensive) goods.[43]

This picture is consistent with the structure of international trade in the contemporary world. Further economic development in the East, accompanied by rising local wages, together with productivity gains in the goods sector of the South, could, in principle, alter this pattern. However, reversing the current configuration of comparative advantages is likely to remain an uphill battle. An alternative path may lie in achieving further productivity gains in the service sector. Whether this path has a growth potential comparable to export-led industrialization and growth remains an open question.

5.4 Imperfect Substitution and Skill Bias in Technology

In our baseline model, workers with different efficiency units are perfect substitutes. As a result, changes in the skill endowment, either across regions or over time, affect structural transformation only through an income effect, since a more skilled workforce raises productivity and income. In reality, changes in the skill distribution can also influence technology adoption and the sectoral allocation of labor. For example, an increase in the supply of educated workers reduces the skill premium and encourages the expansion of sectors that use skilled labor more intensively. Our framework can be extended to incorporate this mechanism by allowing imperfect substitution between skills, so that increases in the skill endowment operate as a technological driver of structural change, similar to Porzio et al. (2022) and Hendricks and Schoellman (2023).

Specifically, suppose that workers with different educational attainments are imperfect substitutes, and tradable goods and CS vary in their skill intensity. If, for example, there are two skill

[43] In Aghion et al. (2025), we postulate that the North is endowed with financial services (such as the provision of safe assets) that are valued by other countries. We interpret this as capturing the empirical notion of trade imbalances. In this scenario, the North can also exchange financial services for goods.

groups, production could be modeled as

$$x_r = \mathcal{A}_r^T \left(\left(H_{rT}^- \right)^{\frac{\rho-1}{\rho}} + \left(\mathcal{B}_r^T H_{rT}^+ \right)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}} \quad \text{and} \quad y_r^{CS} = \mathcal{A}_r^{CS} \left(\left(H_{rCS}^- \right)^{\frac{\rho-1}{\rho}} + \left(\mathcal{B}_r^{CS} H_{rCS}^+ \right)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}},$$

where H^+ and H^- denote high- and low-skilled workers, respectively, and $\mathcal{B}_r^s$ parameterizes the technological skill-bias, which can differ between tradables and CS and between rural and urban locations.

As before, we allow for heterogeneous productivity among workers within the same educational group and assume that individuals of skill type $j \in \{-, +\}$ draw their efficiency z from a distribution $F_z^j(z)$. Total income of an individual i of skill type j in region r is thus given by $y_r^{i,j} = w_r^j z_i^j$.

The aggregative properties of PIGL preferences still yield a characterization of the equilibrium that is similar to our baseline model. In particular, the overall spending share on CS is given by the same functional form as in Proposition 2

$$\vartheta_r^{CS} = \omega_{CS} + \nu_{CS} \, \mathcal{I}_r \left(\frac{\overline{\mathbb{E}_r[z] w_r}}{q_r^{\omega_{CS}}} \right)^{-\varepsilon},$$

where $\overline{\mathbb{E}_r[z] w_r} = \sum_j s_r^j w_r^j \mathbb{E}_r \left[z^j \right]$ denotes average income, $s_r^j = L_r^j / L_r$ is the share of people of type j, and the inequality index $\mathcal{I}_r$ is given by

$$\mathcal{I}_r = \frac{\sum_j s_r^j \left(w_r^j \right)^{1-\varepsilon} \int z^{1-\varepsilon} dF_r^j(z)}{\left(\sum_j s_r^j w_r^j \int z dF_r^j(z) \right)^{1-\varepsilon}}.$$

Hence, as in our baseline model, the overall demand for CS output is driven by average local income and the inequality shifter $\mathcal{I}_r$, which now depends both on the distribution of z within skill-groups (*"within-inequality"*) and skill-specific factor prices w_r^j (*"across-inequality"*). Moreover, given the data on sectoral earnings shares and average earnings by skill group in each region, the model still uniquely identifies the productivity parameters $\mathcal{A}_r^s$ and $\mathcal{B}_r^s$.[44]

5.5 Regional Migration

Our baseline model treats households as immobile. However, it is straightforward to allow for regional productivity differences to generate migration flows. Following the spatial economics

[44] See Section SM-3.4 in the Supplementary Materials for details.

literature (e.g., Redding and Rossi-Hansberg (2017) or Allen and Arkolakis (2025)), we model location choice as a discrete choice problem with idiosyncratic preferences and region specific amenities. Let the utility of living in region r be

$$\mathcal{U}_r(z^i, u_r^i) = \mathcal{V}(z, w_r, q_r)\ \mathcal{Q}_r\, u_r^i,$$

where $\mathcal{V}(z, w_r, q_r) > 0$ is indirect utility, $\mathcal{Q}_r$ is an amenity term, and u_r^i is an idiosyncratic preference draw.[45]

To fully describe individuals' migration choices, one needs to take a stand on whether individuals know the realization of the skill index z, or whether they only know the distribution in each region, but observe their individual productivity draw z after settling in region r. In the latter case, which is implemented in Fan et al. (2023), the optimal location of residence of individual i, $r^\star([u_r^i]_r)$ is given by

$$r^\star(\left[u_r^i\right]_r) = \left\{ r \middle| \int_z \mathcal{U}_r(z, u_r^i) dF_r(z) \geq \int_z \mathcal{U}_l(z, u_l^i) dF_l(z) \quad \forall\, l \right\}. \tag{32}$$

It is convenient to assume that the vector $[u_r^i]_r$ is drawn from a Fréchet distribution with shape parameter θ. Equation (32) then implies that the probability to choose location r, ϖ_r, is given by

$$\varpi_r = \frac{\left(\int_z \mathcal{V}(z, w_r, q_r)\, dF_r(z)\right)^\theta\ \mathcal{Q}_r^\theta}{\sum_l \left(\int_z \mathcal{V}(z, w_l, q_l)\, dF_l(z)\right)^\theta\ \mathcal{Q}_l^\theta}.$$

With a continuum of consumers L in the aggregate economy, the total population in location r is then given by $L_r = \varpi_r L$. Note that the distribution of efficiencies in each location is still given by F_r, as in the baseline model. Hence, the estimation of regional productivities would proceed exactly as in Section 3.5, using the *observed* population distribution. This observed population distribution can always be rationalized with an appropriate choice of local amenities $\left[\mathcal{Q}_r\right]_r$. However, allowing for migration affects the counterfactual results: higher real wages in cities can draw workers from rural areas, and vice versa.

An alternative, and arguably more interesting, assumption would be that individuals know their skill realization z *before* making their moving choices. In that case, the moving probabilities ϖ_r would be a function of z, because they would depend on $\mathcal{V}(z, w_r, q_r)$ as opposed to *expected*

[45] Note that we must require that $\mathcal{V}(z, w_r, q_r) > 0$; otherwise $\mathcal{U}_r(z^i, u_r^i)$ might be decreasing in local amenities $\mathcal{Q}_r$.

utility. Using the indirect utility function in (15), and that $p_x = 1$, the *relative* probability of individuals of type z to settle in r as opposed to k is given by

$$\frac{\varpi_r(z)}{\varpi_k(z)} = \frac{\mathcal{V}(z, w_r, q_r)^\theta \ \mathcal{Q}_r^\theta}{\mathcal{V}(z, w_k, q_k)^\theta \ \mathcal{Q}_k^\theta} = \left(\frac{\mathcal{Q}_r}{\mathcal{Q}_k}\right)^\theta \left(\frac{\frac{1}{\varepsilon} z^\varepsilon \left(\frac{w_r}{q_r^{\omega CS}}\right)^\varepsilon - \nu_{CS} \ln q_r}{\frac{1}{\varepsilon} z^\varepsilon \left(\frac{w_k}{q_k^{\omega CS}}\right)^\varepsilon - \nu_{CS} \ln q_k}\right)^\theta . \tag{33}$$

Equation (33) highlights that there can be *spatial sorting* because non-homothetic preferences for consumption goods also generate a non-homotheticity in individuals' spatial choices: as long as $\nu_{CS} \neq 0$, the relative moving probabilities are a function of z, that is individuals will sort across space in a selected way. In particular, if $\nu_{CS} < 0$ and services are luxuries, high z individuals will have an incentive to move towards "urban service hubs" where wages are high and services are, quality-adjusted, cheap; see Gaubert and Robert-Nicoud (2025) for a recent theoretical treatment of spatial sorting, which also covers the case of PIGL preferences.

6 Literature Review

In this section, we discuss and interpret the results of several earlier papers from diverse literatures (development, macroeconomics, urban economics, and trade) through the lens of our framework.

We build on a long-standing macroeconomic literature on structural transformation. The dual economy model of Lewis (1954) and the empirical work of Chenery and Syrquin (1975) emphasize the reallocation of labor from agriculture to manufacturing and, eventually, to services. Seminal models by Kongsamut et al. (2001) and Ngai and Pissarides (2007) highlight the roles of income and substitution effects. A growing literature explores the drivers of these transitions, including capital deepening (Acemoglu and Guerrieri, 2008, Storesletten et al., 2025), reductions in misallocation over time (Cheremukhin et al., 2016, Song et al., 2011), productivity growth in consumption and investment sectors (García-Santana et al., 2021, Herrendorf et al., 2020), and the accumulation of human capital (Caselli and Coleman II, 2001, Porzio et al., 2022). See Herrendorf et al. (2014) for a survey of this literature.

The remainder of the literature review is organized into seven thematic areas: (1) the decline of agriculture; (2) demand-driven structural change and non-homothetic preferences; (3) premature industrialization; (4) human capital accumulation and structural change toward services; (5)

spatial structure, urbanization and consumer cities; (6) informality, labor market frictions, and marketization; (7) development accounting. In each area, we highlight how our framework relates to and extends existing contributions.

Decline of Agriculture The development process is typically rooted in a shift out of agriculture, which is the dominant employer in low-income economies. In line with our empirical analysis in Section 2, McMillan and Harttgen (2014) documents that structural change, and more specifically the reallocation of labor from agriculture to services, has been a salient feature of the recent growth experience in sub-Saharan Africa.

Lagakos and Waugh (2013) and, more recently, Storesletten et al. (2025) document that differences in labor productivity between poor and rich countries are especially large in agriculture. They emphasize, respectively, the productivity gains arising from worker sorting across sectors and from the mechanization of agricultural production.[46] In a similar vein, Herrendorf and Schoellman (2018) show that differences in sectoral wages across agriculture, manufacturing, and services largely reflect differences in workers' human capital.

Alvarez-Cuadrado and Poschke (2011) distinguish between "push" and "pull" factors behind the decline of agriculture. The main push factor is productivity growth in agriculture, which allows the same amount of food to be produced with less labor, thereby releasing excess workers from agriculture. This mechanism is consistent with the narrative in the classical work of Ashton (1948) and Rostow (1960). More recently, Bustos et al. (2016) exploit the introduction of genetically engineered soybeans in Brazil as an exogenous source of technological change that raised agricultural productivity, particularly in regions suited to soybean cultivation. They find that regions with higher exposure to this technological shock experienced labor reallocation out of agriculture toward manufacturing and services.

In contrast, pull factors are forces that increase the attractiveness of nonagricultural sectors, thereby drawing workers out of agriculture. Specifically, faster productivity growth in industry and services raises wages outside agriculture and attracts workers toward these expanding sectors. In

[46] Part of these productivity differences likely reflects measurement error. For instance, Gollin et al. (2014) point to differences in hours worked per worker and underreporting of self-employment or subsistence output.

line with this view, Kelly et al. (2023) show that during the British Industrial Revolution, industrial development took root in regions where agricultural productivity and wages were relatively low. A similar pattern is observed in contemporary rural India: Foster and Rosenzweig (2004) document that nonfarm employment expanded more rapidly in regions where agricultural productivity growth was slower, as surplus labor was reallocated into rural industry. This result is echoed in the recent findings of Fan et al. (2023). Michaels et al. (2012) provide complementary evidence from a macro and spatial perspective, showing that a high initial reliance on agriculture can slow the transition toward urban, high-productivity sectors, particularly when geographic frictions and infrastructure constraints are binding. See also Goldberg and Reed (2023) who, using cross-country data, also find the link between agricultural productivity growth and poverty reduction to be weak.

Storesletten et al. (2025) propose a synthesis of push and pull factors. In their model, capital accumulation triggers both a shift of workers out of agriculture and a reallocation within the agricultural sector itself. Specifically, they distinguish between two types of agricultural activity: traditional, labor-intensive farming and modern, capital-intensive agriculture. As capital accumulates, total agricultural employment declines, and the sector becomes increasingly dominated by its modern, capital-intensive segment.

In our framework, agriculture is part of the tradable sector. While productivity growth in agriculture contributes to overall economic development, its benefits are transmitted across regions through trade, primarily via a decline in the relative price of agricultural goods. This contrasts with the productivity gains in the regional non-tradable CS sector that, as noted in the extension presented in Section 5.2, generate negative effects on other regions through a terms of trade channel.

Non-homothetic Preferences Non-homothetic preferences play an important role in our analysis.[47] We adopt PIGL preferences, a class of non-Gorman form preferences first introduced by Muellbauer (1976) and further discussed by Deaton and Muellbauer (1980a). Boppart (2014) revamped the interest in this class of preferences by proposing a flexible formulation of PIGL that has significantly shaped subsequent research. Alder et al. (2022) show that this class of preferences can

[47] As such, we also relate to a growing body of research that studies the role of income effects in growth theory and trade more generally (see, e.g., Fieler (2011), Foellmi and Zweimüller (2008), Matsuyama (2000), or Oberfield (2023)).

successfully match the aggregate dynamics of structural change across countries. León-Ledesma and Moro (2020) develop a two-sector model under PIGL preferences in which structural transformation toward services helps account for observed long-run trends in investment, output growth, and capital returns. Eckert and Peters (2022) incorporate PIGL preferences into a spatial general equilibrium model with trade and migration and develop a quantitative theory of spatial structure change that can rationalize the first structural transformation of the US between 1880 and 1920 across both time and space.

An alternative class of non-homothetic preferences based on nested CES functions originally proposed by Sato (1975) has also gained traction. These formulations, recently applied in Matsuyama (2019) and Comin et al. (2021), are particularly well-suited to capturing demand complementarities and long-run compositional shifts.[48]

In this paper we stress two features of PIGL preferences that are particularly useful for our analysis. First, they permit tractable aggregation: inequality enters the aggregate demand system only through a single statistic closely related to the Atkinson index. Second, they provide a transparent mapping between demand for final goods and sectoral value added, a key distinction in studies of structural change (Herrendorf et al., 2013). These properties make it possible to analyze structural transformation in a spatial general equilibrium setting with regional inequality and labor market segmentation using both macro and micro data.

Premature Industrialization The expansion of services during the development process has often been described as a mere consequence of rising income rather than an independent source of growth. A common argument is that traditional services offer limited scope for innovation and technological progress, and that their expansion primarily reflects disguised unemployment or underemployment. Rodrik (2016) defines *premature deindustrialization* as the tendency for developing countries to experience a decline in manufacturing employment and value-added shares at much lower levels of income than today's advanced economies did historically.[49] Rodrik argues that this pattern implies that developing economies face a narrower window for industrialization

[48] See Matsuyama (2023) for a recent survey of Non-CES preferences.
[49] See also Huneeus and Rogerson (2023) for a quantitative assessment of such different paths of industrialization.

as a path to rapid productivity growth. In a similar vein, Duarte and Restuccia (2019) show that cross-country productivity gaps in services are large and persistent. In their view, sluggish convergence in services helps explain why structural change in many developing countries has not translated into faster aggregate growth. Lagakos (2016) argues that retail productivity in developing countries because stores rationally use traditional, small-scale technologies, given the absence of complementary investments such as cars of fridges on the household side.

A critical issue in this debate is the measurement of service productivity. Reliance on published price indexes is problematic insofar as these indexes do not accurately reflect changes in service quality. Moreover, given their non-tradable nature, the correct price indices should be local in nature. Our methodology addresses this challenge by retrieving productivity changes from expenditure and employment data without relying on any information on prices.

Another important dimension is the heterogeneity of service industries. Hallward-Driemeier and Nayyar (2018) and Nayyar et al. (2021) argue that the direction of global technological change, especially automation, together with contemporary trade patterns, has reduced the scope for developing countries to replicate East Asian–style manufacturing-led industrialization. At the same time, they emphasize that modern tradable services such as ICT, professional, and business services can generate productivity gains, spillovers, and export opportunities analogous to those traditionally associated with manufacturing. The role of these services for the future of growth in the US is discussed in Duernecker et al. (2024).

The distinction between non-tradable CS and tradable producer services can be incorporated into our analysis, as discussed in Section 5.2. Empirically, producer services have played and are playing an important role in East Asian countries such as Vietnam and China (see, e.g., Chen et al. (2023)).

Urbanization, Spatial Structure and Consumer Cities Urbanization plays a pivotal role in the process of structural change. While producer services are important in many industrializing economies, the experience of much of sub-Saharan Africa is quite different: the expansion of services has been dominated by consumer-oriented, non-tradable activities. This distinction has important spatial implications. Because CS depend heavily on local demand and face limited scope for longdistance trade, their growth is closely tied to patterns of urbanization and regional labor-market dynamics. Understanding how these spatial forces interact with service-led structural change is an important step in our analysis.

The idea that urban growth can be driven by CS amenities originates in the "consumer city" hypothesis of Glaeser et al. (2001). They document that cities offering richer amenity bundles—restaurants, entertainment, and other consumption-oriented services—tend to experience faster population growth. This insight has been influential in recent work examining why many cities in developing countries have expanded in the absence of large-scale manufacturing.

Empirically, several studies characterize these patterns as the rise of consumer (or consumption) cities in the developing world. Gollin et al. (2016) and Jedwab et al. (2022) show that urbanization in Africa and parts of South Asia has been fueled primarily by domestic demand for CS rather than by tradable production. These cities are dominated by comparatively low-productivity and often informal activities such as local retail, hospitality, and transport. Reflecting this, much of the literature offers a pessimistic interpretation: for example, Henderson and Turner (2020) argue that some countries urbanize "too early," before the institutional and technological preconditions for productivity-enhancing agglomeration are in place.

Our analysis offers a different perspective. Although urban growth in poor countries may appear disorderly, the expansion of CS can still be associated with substantial improvements in productivity and living standards. By explicitly modeling non-homothetic demand, the composition of service industries, and spatial labor market frictions, our framework quantifies how the distribution of population across regions interacts with sectoral specialization and informality to shape the welfare consequences of service-led urbanization. Consistent with this emphasis on cities as centers for non-tradable consumption, Blanchard et al. (2025) use cell phone data from Tanzania, Nigeria, and Kenya to show that rural residents travel frequently to cities to visit service establishments and consume services that are unavailable, or only weakly available, in the countryside.

Our approach is also more broadly related to the literature in spatial economics that studies how geography, trade, and migration determine the regional allocation of economic activity—see, e.g., Alder (2025), Allen and Arkolakis (2014), Donaldson and Storeygard (2016), Redding and Rossi-Hansberg (2017), and Bryan et al. (2020) and Bryan et al. (2025) for an emphasis on the developing world. Within this tradition, our paper highlights the central role of non-tradable CS in shaping spatial inequality and structural change.

Finally, our discussion (especially, the extension discussing internal migration) relates to Budí-Ors and Pijoan-Mas (2022), who develop a spatial general equilibrium model with rural–urban migration and regional labor-market frictions. They show how mobility costs and local labor-market conditions interact to determine regional specialization patterns.

Skill-Biased Service Growth A recent literature also provides a more optimistic view of services in the development process by linking structural transformation to the interaction between rising skills and the diffusion of skill-biased technologies. Buera and Kaboski (2012) develop a model in which market services require fixed costs and specialized skills, generating scale economies and productivity advantages relative to home production. Higher demand for market services thereby increases the returns to skill acquisition, reinforcing the transition from home to market services and generating skill-biased structural change. Building on this mechanism, Buera et al. (2022) propose a multisector general equilibrium model in which skill-biased technological change and the rise of services are jointly determined. The evidence reviewed in Gollin and Kaboski (2024) further emphasizes that modern structural transformation is tightly linked to the diffusion of skill-biased technologies, including in low- and middle-income settings where services are increasingly absorbing skilled labor.

Recent empirical work provides complementary causal evidence connecting human-capital accumulation to the exit from agriculture. Porzio et al. (2022) show that rising cohort-level schooling explains a large share of the global decline in agricultural employment, suggesting the rising human capital shifts labor supply out of agriculture. This is consistent with Cox Lescano (2024), who studies a large-scale expansion of higher education in Brazil. Exploiting regional variation in educational supply constraints, he shows that rising education leads to a shrinking agricultural and expanding service sector. Engbom et al. (2025) provide a quantitative analysis of the link between rising skill supplies and structural change.

Our baseline framework incorporates human capital by allowing for regional heterogeneity in the distribution of efficiency units and by modeling how skills affect the demand for CS because human capital accumulation raises average incomes. The extension discussed in Section 5.4 allows one to incorporate the supply effect: if sectors differ in their skill-intensive, spatial differences in educational attainment emerge as a determinant of sectoral specialization. Using this methodology, Fan et al. (2023) document substantial variation in skill bias across Indian states and over time. This relates to the theory of Acemoglu and Zilibotti (2001), who argue that poor countries and regions adopt low-skill technologies in industries where advanced economies employ high-skill ones. It also connects to the recent findings of Rossi (2022), who shows that cross-country heterogeneity in the skill bias of technology helps explain why high-skill workers are relatively more productive in rich economies.

Informality Informal labor markets are a pervasive characteristic of developing economies, particularly in the service sector. The transition from informal to formal provision of services plays an important role in our analysis.

A rich literature investigates the origins and consequences of informality—see Ulyssea (2020) for a survey. La Porta and Shleifer (2014) distinguish between "dualist" and "survivalist" forms of informality, underscoring that informal activity can reflect both structural segmentation and low-productivity self-employment. Ulyssea (2018) develops a structural model in which imperfect enforcement and entry barriers shape firms' formalization decisions, linking informality to policy and institutional frictions. On the workers' side, Meghir et al. (2015) emphasize wage segmentation and limited mobility across formal and informal jobs. Taken together, these contributions highlight how enforcement, entry costs, and labor mobility shape the size and productivity of the informal service sector. These mechanisms play a central role in our analysis.

Recent work has embedded informality into general equilibrium environments as we do. Dix-Carneiro and Kovak (2017) and Dix-Carneiro et al. (2026) examine how trade liberalization and labor market regulation affect the allocation of workers across formal and informal sectors within an economy. These models highlight how exposure to global competition, combined with institutional frictions, can generate persistent declines in formal employment and long-lasting spatial and sectoral adjustment patterns.

Our analysis of informality in the service sector specifically relates to the recent paper of

Schwartzman (2025), who analyzes the transformation of services in Brazil. He shows that formalization has involved the replacement of traditional own-account workers by modern service firms, particularly in sectors such as retail and hospitality. Formalization is associated with large productivity and earnings gains, especially in poorer regions. These dynamics closely align with the mechanisms emphasized in our paper: formalization can boost productivity, both through improved worker-firm matching and through the entry of more productive firms.

Development Accounting Finally, our paper connects more broadly with the development accounting literature, which aims to quantify the sources of cross-country income differences (Caselli, 2005, Hall and Jones, 1999, Klenow and Rodríguez-Clare, 1997) A growing body of research extends these insights to sectoral and spatial comparisons (see, e.g., Gollin et al. (2014), that we discussed above). Gancia et al. (2013) develop a structural approach to development accounting that, like ours, leverages general equilibrium conditions to infer sectoral productivities from observed wages and employment shares.

Our work builds on Fan et al. (2023), who apply a similar strategy to the Indian context, combining microdata on income and employment with general equilibrium modeling to identify productivity growth in the non-tradable service sector. This methodology clarifies the distinction between service-biased and service-led growth. A rising service employment share may reflect either high productivity in CS or demand shifts driven by income effects. Our framework provides a tractable way to separate these forces using restrictions imposed by preferences and equilibrium behavior. In doing so, we offer a new lens on structural transformation that connects micro-level labor allocation to macro-level productivity gaps.

7 Conclusion and Future Work

This paper has developed a spatial general equilibrium framework to study recent trends in the structural transformation of developing economies, with a particular focus on fast-growing African countries where urbanization and the decline of agricultural employment have taken place largely in the absence of large-scale industrialization. By integrating non-homothetic preferences, the non-tradable nature of CS, spatial heterogeneity, and labor market frictions, our framework distinguishes between *service-led* growth, where the service sector acts as an engine of productivity growth, and *service-biased* growth, where the expansion of services arises as a corollary of productivity growth in other sectors and of human capital accumulation. It also allows us to quantify their distinct implications for productivity, welfare, and inequality.

A key methodological contribution of our framework is that it recovers sectoral productivity, particularly in non-tradable CS, from equilibrium allocations rather than from price data, which are often unavailable or poorly measured. This feature makes the framework well-suited for studying service-led growth in settings where data limitations are severe. In addition, our analysis of agglomeration mechanisms highlights how growing urban markets can endogenously raise service productivity, reinforcing the emergence of consumer cities and amplifying the gains from service-led growth.

Our application to sub-Saharan Africa shows that, in line with the recent findings of Fan et al. (2023) for India, much of the observed shift toward services reflects productivity improvements in locally oriented CS rather than the mere byproduct of rising incomes. These gains have been particularly important in urban areas, where formalization, agglomeration, and specialization effects reinforce productivity growth. At the same time, the welfare effects are uneven: productivity growth in tradables tends to benefit poorer households across space, whereas service-led growth delivers larger gains to richer and more urban consumers. This distributional asymmetry highlights the importance of accounting for non-homotheticity and spatial structure when evaluating development trajectories.

Our framework is modular and can accommodate extensions that incorporate regional specialization, regional and international trade, skill heterogeneity and skill-biased technological change, and endogenous migration. These extensions make it possible to study, for example, how educational expansion interacts with sectoral demand patterns, how internal migration mediates the spatial incidence of productivity shocks, or how natural resources and the presence of globally competitive goods producers abroad shape the industrial opportunities available to latecomer economies.

We conclude that the path to development need not replicate the factory-first model of earlier industrializers. Under the right institutional and market conditions, CS can contribute meaningfully to aggregate productivity growth. Yet this is not automatic, and the sustainability of a service-led path depends on complementary transformations such as formalization, infrastructure investments that enable agglomeration externalities, and local demand growth.

While service-led growth can play a significant role in reducing poverty over the long run, its short-term distributional consequences are more problematic. In particular, our analysis shows that the rise of consumer cities driven by productivity growth in non-tradable services can generate adverse welfare effects for those who remain in rural areas. Non-homothetic preferences imply that low-income households benefit less from improved service availability, while terms-of-trade effects can erode their real incomes. As a result, richer households, and especially urban residents, tend to capture the largest gains. Understanding these dynamics is critical for policy design, as complementary measures such as improving rural access to markets, investing in human capital, and facilitating mobility may be needed to ensure that service-led growth translates into broad-based welfare improvements.

Several promising directions for future research remain. First, we did not discuss the gender dimension, which could shed light on how changes in female labor force participation affect the composition and productivity of services—see, for example, Rendall (2013), Ngai and Petrongolo (2017), and Peters et al. (2024). Second, future work along the lines of Acemoglu and Zilibotti (2001) could more explicitly model the direction of technological change to analyze whether growth can indeed be service-led in the long run. Third, a fuller integration of producer services and global value chains would help assess whether service-led growth can generate spillovers comparable to those historically associated with manufacturing—see Antràs and Chor (2022), Baldwin (2016), Caliendo and Parro (2015), Costinot et al. (2013) and the excellent overview provided by Goldberg and Ruta (2026) in this volume. Fourth, a dynamic extension with capital accumulation along the lines of Chen et al. (2023) would enable the analysis of transitional dynamics and long-run growth in settings where savings, investment, and sectoral productivity evolve jointly.

Last but not least, the interaction between service-led growth and environmental sustainability deserves particular attention. More specifically, could economic development be "greener" if today's developing economies skip the factory stage? In recent research, we develop a framework in which long-run growth features a sustained slowdown in the expansion of material production while living standards and service employment continue to rise (Aghion et al., 2025). This analysis complements the themes of this paper, embedding non-homothetic PIGL preferences, sectoral heterogeneity, and directed technical change into a unified model of growth, structural transformation, and environmental sustainability. We view this as a particularly promising direction for future work.

By combining a tractable theoretical structure with micro-founded empirical measurement, the approach developed here offers a template for studying the positive and normative implications of structural transformation across a broad range of development paths, whether growth is driven primarily by services or by manufacturing. We hope it will provide a foundation for future work on the diverse pathways through which developing economies can achieve sustained and inclusive growth.

References

Acemoglu, D. and V. Guerrieri (2008): "Capital Deepening and Nonbalanced Economic Growth," *Journal of Political Economy*, 116, 467–498.

Acemoglu, D. and F. Zilibotti (2001): "Productivity Differences," *Quarterly Journal of Economics*, 116, 563–606.

Aghion, P., T. Boppart, M. Peters, M. Schwartzman, and F. Zilibotti (2025): "A Theory of Endogenous Degrowth and Environmental Sustainability," NBER Working Paper 33634, National Bureau of Economic Research.

Alder, S. (2025): "Chinese Roads in India: The Effect of Transport Infrastructure on Economic Development," *Journal of International Economics*, 157, 104140.

Alder, S., T. Boppart, and A. Müller (2022): "A Theory of Structural Change That Can Fit the Data," *American Economic Journal: Macroeconomics*, 14, 160–206.

Alder, S., L. Shao, and F. Zilibotti (2016): "Economic Reforms and Industrial Policy in a Panel of Chinese Cities," *Journal of Economic Growth*, 21, 305–349.

Allen, T. and C. Arkolakis (2014): " Trade and the Topography of the Spatial Economy," *Quarterly Journal of Economics*, 129, 1085–1140.

——— (2025): "Chapter 1 - Quantitative Regional Economics," in *Handbook of Regional and Urban Economics*, ed. by D. Donaldson and S. J. Redding, Elsevier, vol. 6, 1–72.

Alvarez-Cuadrado, F. and M. Poschke (2011): "Structural Change Out of Agriculture: Labor Push versus Labor Pull," *American Economic Journal: Macroeconomics*, 3, 127–58.

Antràs, P. and D. Chor (2022): "Global Value Chains," in *Handbook of International Economics*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, Amsterdam: Elsevier, vol. 5, 285–378.

Artadi, E. V. and X. Sala-i Martin (2003): "The Economic Tragedy of the XXth Century: Growth in Africa," NBER Working Paper 9865, National Bureau of Economic Research.

Ashton, T. (1948): *The Industrial Revolution, 1760–1830*, Oxford: Oxford University Press.

Atkinson, A. B. (1970): "On the Measurement of Inequality," *Journal of Economic Theory*, 2, 244–263.

Baldwin, R. (2016): *The Great Convergence: Information Technology and the New Globalization*, Cambridge, MA: Harvard University Press.

Banerjee, A. V. and E. Duflo (2007): "The Economic Lives of the Poor," *Journal of Economic Perspectives*, 21, 141–167.

Barro, R. J. and J. W. Lee (2013): "A New Data Set of Educational Attainment in the World, 1950–2010," *Journal of Development Economics*, 104, 184–198.

Bergquist, L. F., C. McIntosh, and M. Startz (2024): "Search Costs, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets," NBER Working Paper 33221, National Bureau of Economic Research.

Blanchard, P., D. Gollin, M. Kirchberger, and M. Peters (2025): "Making Non-Tradables Tradable: Human Mobility and the Spatial Reach of African Cities," CEPR Discussion Paper, Centre for Economic Policy Research.

Boppart, T. (2014): "Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences," *Econometrica*, 82, 2167–2196.

Bryan, G., E. Glaeser, and N. Tsivanidis (2020): "Cities in the Developing World," *Annual Review of Economics*, 12, 273–297.

Bryan, G. T., K. Frye, and M. Morten (2025): "Spatial Economics for Low-and Middle-income Countries," NBER Working Paper 33453, National Bureau of Economic Research.

Budí-Ors, T. and J. Pijoan-Mas (2022): "Macroeconomic Development, Rural Exodus, and Uneven Industrialization," CEPR Discussion Paper No. DP17086.

Buera, F. J., H. Hopenhayn, Y. Shin, and N. Trachter (2021): "Big Push in Distorted Economies," NBER Working Paper 28561, National Bureau of Economic Research.

Buera, F. J. and J. P. Kaboski (2012): "The Rise of the Service Economy," *American Economic Review*, 102, 2540–2569.

Buera, F. J., J. P. Kaboski, R. Rogerson, and E. Vizcaíno (2022): "Skill-Biased Structural Change," *Review of Economic Studies*, 89, 889–927.

Bustos, P., B. Caprettini, and J. Ponticelli (2016): "Agricultural Productivity and Structural Transformation: Evidence from Brazil," *American Economic Review*, 106, 1320–1365.

Butler, M., S. Fanuel, and P. Grinsted (2022): "On the Path to Industrialization: A Review of Industrial Parks in Ethiopia," Policy Report, World Bank, Washington, DC.

Caliendo, L. and F. Parro (2015): "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 82, 1–44.

Caselli, F. (2005): "Chapter 9 Accounting for Cross-Country Income Differences," in *Handbook of Economic Growth*, ed. by P. Aghion and S. N. Durlauf, Elsevier, vol. 1, 679–741.

Caselli, F. and W. J. Coleman II (2001): "The US Structural Transformation and Regional Convergence: A Reinterpretation," *Journal of Political Economy*, 109, 584–616.

Chai, A., E. Stepanova, and A. Moneta (2023): "Quantifying Expenditure Hierarchies and the Expansion of Global Consumption Diversity," *Journal of Economic Behavior & Organization*, 214, 860–886.

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

Chenery, H. B. and M. Syrquin (1975): *Patterns of Development, 1950–1970*, Oxford University Press.

Cheremukhin, A., M. Golosov, S. Guriev, and A. Tsyvinski (2016): "The Industrialization and Economic Development of Russia through the Lens of a Neoclassical Growth Model," *Review of Economic Studies*, 84, 613–649.

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

Costinot, A., J. Vogel, and S. Wang (2013): "An Elementary Theory of Global Supply Chains," *The Review of Economic Studies*, 80, 109–144.

Cox Lescano, A. (2024): "Fostering Development through Higher Education: College Attainment, Firms and Economic Growth," PhD Thesis, Yale University.

Deaton, A. and J. Muellbauer (1980a): "An Almost Ideal Demand System," *American Economic Review*, 70, 312–326.

——— (1980b): *Economics and Consumer Behavior*, Cambridge: Cambridge University Press.

Dix-Carneiro, R. and B. K. Kovak (2017): "Margins of Labor Market Adjustment to Trade," *Journal of International Economics*, 104, 112–129.

Dix-Carneiro, R., P. K. Goldberg, C. Meghir, and G. Ulyssea (2026): "Trade and Domestic Distortions: The Case of Informality," *Econometrica*, forthcoming.

Donaldson, D. and A. Storeygard (2016): "The View from Above: Applications of Satellite Data in Economics," *Journal of Economic Perspectives*, 30, 171–198.

Duarte, M. and D. Restuccia (2019): "Relative Prices and Sectoral Productivity," *Journal of the European Economic Association*, 18, 1400–1443.

Duernecker, G., B. Herrendorf, and A. Valentinyi (2024): "Structural Change within the Services Sector and the Future of Cost Disease," *Journal of the European Economic Association*, 22, 428–473.

Easterly, W. and R. Levine (1997): "Africa's Growth Tragedy: Policies and Ethnic Divisions," *Quarterly Journal of Economics*, 112, 1203–1250.

Eckert, F. and M. Peters (2022): "Spatial Structural Change," NBER Working Paper 30489, National Bureau of Economic Research, revision date: July 2025.

Engbom, N., H. Malmberg, T. Porzio, F. Rossi, and T. Schoellman (2025): "Economic Development According to Chandler," NBER Working Paper 34483, National Bureau of Economic Research.

Faber, B. and C. Gaubert (2019): "Tourism and Economic Development: Evidence from Mexico's Coastline," *American Economic Review*, 109, 2245–2293.

Fan, T., M. Peters, and F. Zilibotti (2023): "Growing like India–the Unequal Effects of Service-Led Growth," *Econometrica*, 91, 1457–1494.

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

Foellmi, R. and J. Zweimüller (2008): "Structural Change, Engel's Consumption Cycles and Kaldor's Facts of Economic Growth," *Journal of Monetary Economics*, 55, 1317–1328.

Foster, A. D. and M. R. Rosenzweig (2004): "Agricultural Productivity Growth, Rural Economic Diversity, and Economic Reforms: India, 1970–2000," *Economic Development and Cultural Change*, 52, 509–542.

Gancia, G., A. Müller, and F. Zilibotti (2013): "Structural Development Accounting," in *Advances in Economics and Econometrics: Theory and Applications (10th World Congress of the Econometric Society)*, ed. by D. Acemoglu, M. Arellano, and E. Dekel, Cambridge Univ. Press, 373–418.

García-Santana, M., J. Pijoan-Mas, and L. Villacorta (2021): "Investment Demand and Structural Change," *Econometrica*, 89, 2751–2785.

Gaubert, C. and F. Robert-Nicoud (2025): "Sorting to Expensive Cities," NBER Working Paper 33652, National Bureau of Economic Research.

Glaeser, E. L., J. Kolko, and A. Saiz (2001): "Consumer City," *Journal of Economic Geography*, 1, 27–50.

Goldberg, P. and M. Ruta (2026): "The Changing Nature of International Trade and its Implications for Development," in *Handbook of Development Economics*, ed. by P. Dupas, P. Goldberg, and R. Pande, Elsevier, vol. 6.

Goldberg, P. K., C. T. Gottlieb, S. Lall, M. Mehta, M. Peters, and A. L. Ratan (2025): "The Global Gender Distortions Index (GGDI)," NBER Working Paper 34142, National Bureau of Economic Research.

Goldberg, P. K. and T. Reed (2023): "Presidential Address: Demand-Side Constraints in Development. The Role of Market Size, Trade, and (In) Equality," *Econometrica*, 91, 1915–1950.

Gollin, D., R. Jedwab, and D. Vollrath (2016): "Urbanization with and without Industrialization," *Journal of Economic Growth*, 21, 35–70.

Gollin, D. and J. P. Kaboski (2024): "New Views of Structural Transformation: Insights from Recent Literature," STEG Synthesis Paper 8, Structural Transformation and Economic Growth (STEG) Programme, London.

Gollin, D., D. Lagakos, and M. E. Waugh (2014): "The Agricultural Productivity Gap," *American Economic Review*, 104, 1651–1699.

Gottlieb, C., C. Doss, D. Gollin, and M. Poschke (2024): "The Gender Division of Work across Countries," IZA Discussion Paper 16896, Institute of Labor Economics (IZA), Bonn.

Gottlieb, C., J. Grobovsek, and A. Monge-Naranjo (2025): "From Skills to Occupations: Comparative Advantage and Cross-Country Income Differences," CEPR Working Paper 20630, Centre for Economic Policy Research.

Graff, T. (2025): "Depreciation and Growth: Evidence from Machine Repair in Uganda," Unpublished manuscript, PhD dissertation, Harvard University.

Hall, R. E. and C. I. Jones (1999): "Why Do Some Countries Produce So Much More Output per Worker than Others?" *Quarterly Journal of Economics*, 114, 83–116.

Hallward-Driemeier, M. and G. Nayyar (2018): *Trouble in the Making? The Future of Manufacturing-Led Development*, Washington, DC: World Bank Group.

Harris, J. R. and M. P. Todaro (1970): "Migration, Unemployment and Development: A Two-Sector Analysis," *Quarterly Journal of Economics*, 84, 488–521.

Henderson, J. V. and M. A. Turner (2020): "Urbanization in the Developing World: Too Early or Too Slow?" *Journal of Economic Perspectives*, 34, 150–173.

Hendricks, L. and T. Schoellman (2023): "Skilled Labor Productivity and Cross-Country Income Differences," *American Economic Journal: Macroeconomics*, 15, 240–268.

Herrendorf, B., R. Rogerson, and Ákos Valentinyi (2014): "Growth and Structural Transformation," in *Handbook of Economic Growth*, ed. by P. Aghion and S. N. Durlauf, Elsevier, vol. 2, 855–941.

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

——— (2020): "Structural Change in Investment and Consumption: A Unified Analysis," *Review of Economic Studies*, 88, 1311–1346.

Herrendorf, B. and T. Schoellman (2018): "Wages, Human Capital, and Barriers to Structural Transformation," *American Economic Journal: Macroeconomics*, 10, 1–23.

Huneeus, F. and R. Rogerson (2023): "Heterogeneous Paths of Industrialization," *Review of Economic Studies*, 91, 1746–1774.

Huynh, N. Q. (2025): "Place-based Policy, Migration Barriers, and Spatial Inequality," SSRN Working Paper 5281304.

International Monetary Fund (2024): "Regional Economic Outlook: Sub-Saharan Africa—A Tepid and Pricey Recovery," Report, International Monetary Fund, Washington, DC, prepared under the supervision of Wenjie Chen, Luc Eyraud, and Catherine Pattillo.

Jedwab, R., E. Ianchovichina, and F. Haslop (2022): "Consumption Cities vs. Production Cities: New Considerations and Evidence," SSRN Working Paper 4166684.

Jensen, R. (2007): "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector," *The Quarterly Journal of Economics*, 122, 879–924.

Kelly, M., J. Mokyr, and C. Ó Gráda (2023): "The Mechanics of the Industrial Revolution," *Journal of Political Economy*, 131, 59–94.

Klenow, P. J. and A. Rodríguez-Clare (1997): "The Neoclassical Revival in Growth Economics: Has It Gone Too Far?" *NBER Macroeconomics Annual*, 12, 73–103.

Kongsamut, P., S. Rebelo, and D. Xie (2001): "Beyond Balanced Growth," *Review of Economic Studies*, 68, 869–882.

Krugman, P. (1980): "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70, 950–959.

La Porta, R. and A. Shleifer (2014): "Informality and Development," *Journal of Economic Perspectives*, 28, 109–126.

Lagakos, D. (2016): "Explaining Cross-country Productivity Differences in Retail Trade," *Journal of Political Economy*, 124, 579–620.

Lagakos, D. and M. E. Waugh (2013): "Selection, Agriculture, and Cross-Country Productivity Differences," *American Economic Review*, 103, 948–980.

Lewis, W. A. (1954): "Economic Development with Unlimited Supplies of Labour," *The Manchester School*, 22, 139–191.

León-Ledesma, M. and A. Moro (2020): "The Rise of Services and Balanced Growth in Theory and Data," *American Economic Journal: Macroeconomics*, 12, 109–146.

Matsuyama, K. (2000): "A Ricardian Model with a Continuum of Goods Under Nonhomothetic Preferences: Demand Complementarities, Income Distribution, and North-South Trade," *Journal of Political Economy*, 108, 1093–1120.

——— (2019): "Engel's Law in the Global Economy: Demand-Induced Patterns of Structural Change, Innovation, and Trade," *Econometrica*, 87, 497–528.

——— (2023): "Non-CES Aggregators: a Guided Tour," *Annual Review of Economics*, 15, 235–265.

McCaig, B. and N. Pavcnik (2017): "Moving Out of Agriculture: Structural Change in Vietnam," in *Structural Change, Fundamentals, and Growth: A Framework and Case Studies*, ed. by M. McMillan, D. Rodrik, and C. P. Sepúlveda, International Food Policy Research Institute, 91–118.

——— (2018): "Export Markets and Labor Allocation in a Low-income Country," *American Economic Review*, 108, 1899–1941.

McMillan, M. S. and K. Harttgen (2014): "What is Driving the African Growth Miracle?" NBER Working Paper 20077, National Bureau of Economic Research.

Meghir, C., R. Narita, and J.-M. Robin (2015): "Wages and Informality in Developing Countries," *American Economic Review*, 105, 1509–1546.

Michaels, G., F. Rauch, and S. J. Redding (2012): "Urbanization and Structural Transformation," *Quarterly Journal of Economics*, 127, 535–586.

Muellbauer, J. (1976): "Community Preferences and the Representative Consumer," *Econometrica*, 44, 979–999.

Murphy, K. M., A. Shleifer, and R. W. Vishny (1989): "Industrialization and the Big Push," *Journal of Political Economy*, 97, 1003–1026.

Nayyar, G., M. Hallward-Driemeier, and E. Davies (2021): *At Your Service? The Promise of Services-led Development*, World Bank Publications.

Ngai, L. R. and B. Petrongolo (2017): "Gender Gaps and the Rise of the Service Economy," *American Economic Journal: Macroeconomics*, 9, 1–44.

Ngai, L. R. and C. A. Pissarides (2007): "Structural Change in a Multisector Model of Growth," *American Economic Review*, 97, 429–443.

Nguyen, H. C. (2018): "Empirical Evidence of Structural Change: The Case of Vietnam's Economic Growth," *Asian Economic and Financial Review*, 8, 1341–1355.

Oberfield, E. (2023): "Inequality and Measured Growth," NBER Working Paper 31096, National Bureau of Economic Research.

Peters, M., P. Torola, L. Uniat, and F. Zilibotti (2024): "From Rural Fields to Urban Kitchens: Structural Change and the Decline of Women's Work in India," EGC Discussion Paper 1111, Economic Growth Center, Yale University.

Porzio, T., F. Rossi, and G. Santangelo (2022): "The Human Side of Structural Transformation," *American Economic Review*, 112, 2774–2814.

Redding, S. J. and E. Rossi-Hansberg (2017): "Quantitative Spatial Economics," *Annual Review of Economics*, 9, 21–58.

Rendall, M. (2013): "Structural Change in Developing Countries: Has It Decreased Gender Inequality?" *World Development*, 45, 1–16.

Rodrik, D. (2016): "Premature Deindustrialization," *Journal of Economic Growth*, 21, 1–33.

Rossi, F. (2022): "The Relative Efficiency of Skilled Labor across Countries: Measurement and Interpretation," *American Economic Review*, 112, 235–266.

Rostow, W. W. (1960): *The Stages of Economic Growth: A Non-Communist Manifesto*, Cambridge: Cambridge University Press.

Sato, R. (1975): "The Most General Class of CES Functions," *Econometrica*, 43, 999–1003.

Schwartzman, M. (2025): "From Street Markets to Shopping Malls: The Transformation of the Service Sector," STEG Working Paper WP116.

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

Storesletten, K., B. Zhao, and F. Zilibotti (2025): "Business Cycle During Structural Change: Arthur Lewis' Theory from a Neoclassical Perspective," *The Economic Journal*, forthcoming.

Ulyssea, G. (2018): "Firms, Informality, and Development: Theory and Evidence from Brazil," *American Economic Review*, 108, 2015–2047.

——— (2020): "Informality: Causes and Consequences for Development," *Annual Review of Economics*, 12, 525–546.

Young, A. (2012): "The African Growth Miracle," *Journal of Political Economy*, 120, 696–739.

Zhang, X., D. Tezera, C. Zou, Z. Wang, J. Zhao, E. A. Gebremenfas, and J. Dhavle (2018): "Industrial Parks Development in Ethiopia," Working Paper 21/2018, United Nations Industrial Development Organization (UNIDO), Department of Policy, Research and Statistics.

APPENDIX A: THEORETICAL AND QUANTITATIVE RESULTS

A-1 The Baseline Model (Section 3)

In this section, we present detailed derivations for the baseline model presented in Section 3.

A-1.1 Aggregation (Section 3.2)

Proof of Proposition 1 In Proposition 1 we derived the aggregate expenditure share on product n. Using the expression for the expenditure share at the individual level, equation (4), we get that

$$\begin{aligned}\vartheta_{rn}^{FE} \equiv \frac{\int \vartheta_{n}^{FE}\left(zw_r,\mathbf{p}_r\right)zw_r dF_r\left(z\right)}{\int zw_r dF_r\left(z\right)} &= \beta_n + \kappa_n\left(\exp\left(\int_0^1 \beta_n \ln p_{rn}\, dn\right)\right)^{\varepsilon}\frac{\int \left(zw_r\right)^{1-\varepsilon}dF_r\left(z\right)}{\int zw_r dF_r\left(z\right)}\\ &= \beta_n + \kappa_n \mathcal{I}_r\left(\frac{\mathbb{E}_r\left[z\right]w_r}{\exp\left(\int_0^1 \beta_n \ln p_{rn}\, dn\right)}\right)^{-\varepsilon}\\ &\equiv \beta_n + \kappa_n \mathcal{I}_r \varUpsilon\left(\mathbb{E}_r\left[z\right]w_r,\mathbf{p}_r\right)^{-\varepsilon} \qquad \text{(A-1)}\end{aligned}$$

where $\mathcal{I}_r = \frac{\int z^{1-\varepsilon}\, dF_r(z)}{\left(\int z\, dF_r(z)\right)^{1-\varepsilon}}$ and $\varUpsilon(.)$ is defined in (4). Note also that $\mathcal{I}_r$ is a transformation of the *Atkinson index*. The Atkinson index is given by $AI_r = 1 - \frac{\left(\int z^{1-\varepsilon}\, dF_r(z)\right)^{\frac{1}{1-\epsilon}}}{\int z\, dF_r(z)}$ (see Atkinson (1970)). Hence, $\mathcal{I}_r = (1-AI_r)^{1-\varepsilon}$.

Proof of Proposition 2 Substituting the expression from Proposition 1 into $\mathcal{S}_r^{CS}$ (see (11)) yields the result:

$$\begin{aligned}\vartheta_r^{CS} = \frac{L_r w_r \mathbb{E}_r\left[z\right]\int_n \lambda_n \vartheta_{rn}^{FE}\, dn}{L_r w_r \mathbb{E}_r\left[z\right]} = \int_n \lambda_n \vartheta_{rn}^{FE}\, dn &= \int_n \lambda_n\left(\beta_n + \kappa_n \mathcal{I}_r \varUpsilon\left(\mathbb{E}_r\left[z\right]w_r,\mathbf{p}_r\right)^{-\varepsilon}\right)\, dn\\ &= \int_n \lambda_n\beta_n\, dn + \int_n \lambda_n\kappa_n\, dn\ \mathcal{I}_r\varUpsilon\left(\mathbb{E}_r\left[z\right]w_r,\mathbf{p}_r\right)^{-\varepsilon}\\ &= \omega_{CS} + \nu_{CS}\mathcal{I}_r\varUpsilon\left(\mathbb{E}_r\left[z\right]w_r;p_x,q_r\right)^{-\varepsilon},\end{aligned}$$

where $\omega_{CS} = \int_n \lambda_n\beta_n\, dn$ and $\nu_{CS} = \int_n \lambda_n\kappa_n\, dn$. Using $p_{rn} = p_x^{1-\lambda_n}q_r^{\lambda_n}$ (see (10)), it follows that

$$\exp\left(\int_{n=0}^1 \beta_n \ln p_n dn\right) \equiv p_x^{1-\omega_{CS}}q_r^{\omega_{CS}}. \qquad \text{(A-2)}$$

Thus, we can rewrite

$$\varUpsilon\left(\mathbb{E}_r\left[z\right]w_r,\mathbf{p}_r\right) = \frac{\mathbb{E}_r\left[z\right]w_r}{p_x^{1-\omega_{CS}}q_r^{\omega_{CS}}} = \varUpsilon\left(\mathbb{E}_r\left[z\right]w_r;p_x,q_r\right) \qquad \text{(A-3)}$$

A-1.2 Equilibrium

Equation (13) determines the local wage and the price of CS inputs as $w_r = \mathcal{A}_r^T$ and $q_r = \mathcal{A}_r^T/\mathcal{A}_r^{CS}$. Hence, we merely need to determine the the sectoral labor allocations $\{H_r^T, H_r^{CS}\}$. These solve the following equilibrium conditions:

1. Market clearing for local CS:
$$w_r H_r^{CS} = \vartheta_r^{CS} w_r H_r$$

2. Market clearing for tradable goods:
$$w_r H_r^T = \left(1 - \vartheta_r^{CS}\right) w_r H_r$$

3. Labor market clearing: $H_r = H_r^{CS} + H_r^T$

With $H_r = \mathbb{E}_r[z] L_r$, the market clearing for CS can be written as

$$\ell_r^{CS} = \frac{L_r^{CS}}{L_r} = \frac{H_r^{CS}}{H_r} = \vartheta_r^{CS}$$

A-1.3 Welfare Effects

In Section 3.4 we argued that the welfare effects of sectoral productivity growth are summarized by the expenditure shares. To see this, note that the expenditure function to achieve utility $\mathcal{V}$ at prices (p_x, q_r) is given by

$$e(\mathcal{V}, p_x, q_r) = p_x^{1-\omega_{CS}}\, q_r^{\omega_{CS}} \left\{\varepsilon\big[\mathcal{V} + \nu_{CS}\ln(q_r/p_x)\big]\right\}^{1/\varepsilon}. \tag{A-4}$$

Taking logs of (A-4) and differentiating w.r.t. $\ln q_r$ (holding $\mathcal{V}$ constant) yields

$$\frac{\partial \ln e(\mathcal{V}, p_x, q_r)}{\partial \ln q_r} = \omega_{CS} + \frac{1}{\varepsilon}\cdot\frac{\nu_{CS}}{\mathcal{V} + \nu_{CS}\ln(q_r/p_x)}. \tag{A-5}$$

By Shephard's Lemma,

$$\frac{\partial \ln e}{\partial \ln q_r} = \frac{q_r}{e}\frac{\partial e}{\partial q_r} = \frac{q_r\, x_r^{CS}(\mathcal{V}, p_x, q_r)}{e(\mathcal{V}, p_x, q_r)} \equiv \vartheta_r^{CS}(\mathcal{V}; p_x, q_r),$$

and similarly $\partial \ln e/\partial \ln p_x = 1 - \vartheta_r^{CS}$. Also note that (A-4) implies that

$$\varepsilon\big[\mathcal{V} + \nu_{CS}\ln(q_r/p_x)\big] = \left(\frac{e(\mathcal{V}, p_x, q_r)}{p_x^{1-\omega_{CS}} q_r^{\omega_{CS}}}\right)^{\varepsilon}. \tag{A-6}$$

Using (A-6), we can derive

$$\frac{1}{\varepsilon}\cdot\frac{\nu_{CS}}{\mathcal{V} + \nu_{CS}\ln(q_r/p_x)} = \nu_{CS}\left(\frac{e(\mathcal{V}, p_x, q_r)}{p_x^{1-\omega_{CS}} q_r^{\omega_{CS}}}\right)^{-\varepsilon} = \vartheta_r^{CS}(\mathcal{V}; p_x, q_r) - \omega_{CS},$$

which together with (A-5) also yields the result.

A-1.4 Equivalent variation (EV)

In Figure 7 we reported the welfare effects of different sources of productivity growth. To compute the effects on welfare, we rely on individuals' *equivalent variation*. Formally, consider the we summarize the welfare effect of moving from the *status-quo* allocation (w, p_x, q) to the *counterfactual* allocation $(\hat{w}, \hat{p}_x, \hat{q})$ for an individual with skills z with the equivalent variation

$$EV(z) \equiv e(\hat{\mathcal{V}}(z), p_x, q) - e(\mathcal{V}(z), p_x, q) = e(\hat{\mathcal{V}}(z), p_x, q) - zw, \tag{A-7}$$

where $e(.)$ is the expenditure function and

$$\mathcal{V}(z) \equiv \mathcal{V}(e(z); w, p_x, q) = \frac{1}{\varepsilon}\left(\frac{zw}{p_x^{1-\omega_{CS}} q^{\omega_{CS}}}\right)^{\varepsilon} - \nu_{CS} \ln\left(\frac{q}{p_x}\right) \tag{A-8}$$

is the indirect utility function. Indirect utility in the counterfactual allocation is $\hat{\mathcal{V}}(z) \equiv \mathcal{V}(z\hat{w}; \hat{w}, \hat{p}_x, \hat{q})$ and $\mathcal{V}(z) \equiv \mathcal{V}(zw; w, p_x, q)$ denotes the indirect utility in the status quo. The second equality uses that $e(\mathcal{V}(z), p_x, q) = zw$, because zw is total spending in the status quo. In words: EV is the amount of money one would need to give consumer z in excess of their income zw, to reach utility $\hat{\mathcal{V}}(z)$ at the old prices (p_x, q). Note that the EV can be negative (in case the new allocation reduces welfare) or positive (in case the new allocation increases welfare.), because $e(\mathcal{V}, p_x, q)$ is increasing in $\mathcal{V}$.

To compute EV, we need to take into account the restriction $\vartheta_r^{CS} \geq 0$. In particular, because CS value added is luxury, the consumer's expenditure share on CS value added might be zero, in case the consumer is too poor. In that case, the indirect utility function in (A-8) and the associated expenditure function

$$e(\mathcal{V}, p, q) = \left[\varepsilon\left(\mathcal{V} + \nu_{CS} \ln\left(\frac{q}{p_x}\right)\right)\right]^{1/\varepsilon} p_x^{1-\omega_{CS}} q^{\omega_{CS}} \tag{A-9}$$

do not apply. To perform our welfare analysis while taking into account the non-negativity constraints in the CS expenditure share, we propose an extension of $\mathcal{V}(.)$ and $e(.)$ over the full domain of household expenditure that we can use to compute EV in (A-7) for all z.

Extending $\mathcal{V}(.)$ and $e(.)$ To extend $\mathcal{V}(.)$ and $e(.)$ over the full domain of household expenditure, it is useful to define

$$B(p_x, q) \equiv p_x^{1-\omega_{CS}} q^{\omega_{CS}} \quad \text{and} \quad \lambda \equiv -\frac{\nu^{CS}}{\omega^{CS}} > 0. \tag{A-10}$$

For the expenditure share on CS to be weakly positive, $\vartheta_i^{CS} \geq 0$, equation (17) implies that total expenditure has to satisfy

$$e \geq \overline{e}(p_x, q) = \left(\frac{-\nu_{CS}}{\omega_{CS}}\right)^{1/\varepsilon} p_x^{1-\omega_{CS}} q^{\omega_{CS}} = \lambda^{1/\varepsilon} B(p_x, q). \tag{A-11}$$

Equation (A-8) implies that indirect utility at $e = \overline{e}(p_x, q)$ is given by

$$\overline{\mathcal{V}}(p_x, q) \equiv \mathcal{V}^I(\overline{e}(p_x, q), p_x, q) = \frac{1}{\varepsilon}\left(\frac{\overline{e}(p_x, q)}{B(p_x, q)}\right)^{\varepsilon} - \nu_{CS}\ln\left(\frac{q}{p_x}\right) = \frac{1}{\varepsilon}\lambda + \lambda\omega_{CS}\ln\left(\frac{q}{p_x}\right), \quad \text{(A-12)}$$

where we use the superscript I to indicate that this is the branch of the indirect utility function that applies in the interior range of ϑ^{CS}. Note that (A-9) implies that the associated expenditure at $\overline{\mathcal{V}}(p_x, q)$ is given by $e^I\left(\overline{\mathcal{V}}(p_x, q), p_x, q\right) = \lambda^{1/\varepsilon}B(p_x, q) = \overline{e}(p_x, q)$ as required.

To extend $\mathcal{V}$ and e to the case of negative CS shares, we want to find $\mathcal{V}^C(e, p_x, q)$, "C" for corner, which satisfies three properties:

1. $\mathcal{V}$ is continuous, including at $\overline{e}(p_x, q)$, i.e. $\mathcal{V}^C(\overline{e}(p_x, q), p_x, q) = \mathcal{V}^I(\overline{e}(p_x, q), p_x, q)$,
2. $\mathcal{V}$ is differentiable, including at $\overline{e}(p_x, q)$, i.e. $\frac{\partial \mathcal{V}^C(\overline{e}(p_x,q),p_x,q)}{\partial e} = \frac{\partial \mathcal{V}^I(\overline{e}(p_x,q),p_x,q)}{\partial e}$,
3. Applying Roy's identity to $\mathcal{V}^C(e, p_x, q)$ indeed yields $\vartheta^{CS} = 0$ and $\vartheta^T = 1$.

As we show in Section SM-1.1 in the Supplementary Material, the following function form for $\mathcal{V}^C(e, p_x, q)$ satisfies all three properties:

$$\mathcal{V}^C(e, p_x, q) = \mathcal{V}^I(\overline{e}(p_x, q), p_x, q) + \lambda\ln\left(\frac{e}{\overline{e}(p_x, q)}\right) = \lambda\ln\left(\frac{e}{p_x}\right) + \frac{1}{\varepsilon}\lambda(1 - \ln\lambda), \quad \text{(A-13)}$$

where the second equality follows upon substituting $\mathcal{V}^I(\overline{e}(p_x, q), p_x, q)$ and $\overline{e}(p_x, q)$. The associated expenditure function is given by

$$e^C(\mathcal{V}, p_x, q) = \exp\left(\frac{\mathcal{V}}{\lambda} - \frac{1}{\varepsilon}(1 - \ln\lambda)\right)p_x. \quad \text{(A-14)}$$

Note that neither $\mathcal{V}^C(e, p_x, q)$ nor $e^C(\mathcal{V}, p_x, q)$ depend on q, which is intuitive because spending on CS is zero. The overall indirect utility and expenditure functions are thus given by

$$\mathcal{V}(e, p_x, q) = \begin{cases} \lambda\ln\left(\frac{e}{p_x}\right) + \frac{1}{\varepsilon}\lambda(1 - \ln\lambda) & \text{if } e < \overline{e}(p_x, q) = \lambda^{1/\varepsilon}B(p_x, q) \\ \frac{1}{\varepsilon}\left(\frac{e}{B(p_x,q)}\right)^{\varepsilon} + \lambda\omega_{CS}\ln\left(\frac{q}{p_x}\right) & \text{if } e \geq \overline{e}(p_x, q) = \lambda^{1/\varepsilon}B(p_x, q) \end{cases} \quad \text{(A-15)}$$

and

$$e(\mathcal{V}, p_x, q) = \begin{cases} \exp\left(\frac{\mathcal{V}}{\lambda} - \frac{1}{\varepsilon}(1 - \ln\lambda)\right)p_x & \text{if } \mathcal{V} < \overline{\mathcal{V}}(p, x) = \lambda\left(\frac{1}{\varepsilon} + \omega_{CS}\ln\left(\frac{q}{p_x}\right)\right) \\ \left[\varepsilon\left(\mathcal{V} - \lambda\omega_{CS}\ln\left(\frac{q}{p_x}\right)\right)\right]^{1/\varepsilon}B(p_x, q) & \text{if } \mathcal{V} \geq \overline{\mathcal{V}}(p, x) = \lambda\left(\frac{1}{\varepsilon} + \omega_{CS}\ln\left(\frac{q}{p_x}\right)\right) \end{cases}, \quad \text{(A-16)}$$

where $\overline{e}(p_x, q)$ and $\overline{\mathcal{V}}(p, x)$ are given in (A-11) and (A-12). In Section SM-1.1 in the Supplementary Material we also show that $e(\mathcal{V}, p_x, q)$ as defined in (A-16) is continuous, differentiable, and applying Shepard's lemma yields $\vartheta^{CS} = 0$ and $\vartheta^T = 1$.

Computing EV With $e(\mathcal{V}, p_x, q)$ and $\mathcal{V}(e, p_x, q)$ as defined in (A-15) and (A-16) at hand, we can compute the equivalent variation in (A-7). Because $\mathcal{V}$ and e have different branches depending on z, the equivalent variation is also defined as a function of z. As we show in detail in Section SM-1.2 in the Supplementary Material, depending on whether relative price of services are increasing or falling in the counterfactual. the following situations can occur.

Proposition 3. *Define*

$$\widetilde{z} \equiv \left(\frac{z\hat{w}}{B(\hat{p}_x, \hat{q})}\right)^{\varepsilon} \frac{1}{\lambda} \quad \text{and} \quad \kappa \equiv \left(\frac{q/p_x}{\hat{q}/\hat{p}_x}\right)^{\varepsilon \omega_{CS}}.$$

1. *Suppose the $\kappa > 1$, i.e. the relative price of services is higher in the status quo, compared to the counterfactual. Then:*
   (a) $\widetilde{z} \leq 1$ ***("Full corner"):*** *The consumer consumes only tradable goods both in the counterfactual and when trying to reach $\hat{\mathcal{V}}(z)$ at the old prices.*
   (b) $1 < \widetilde{z} < 1 + \ln \kappa$ ***("Interior / Corner")****: The consumer consumes positive quantities of CS value added in the counterfactual but only tradable goods when trying to reach $\hat{\mathcal{V}}(z)$ at the old prices.*
   (c) $1 + \ln \kappa \leq \widetilde{z}$ ***("Full Interior")****: The consumer consumes positive quantities of CS value added in both the counterfactual and when trying to reach $\hat{\mathcal{V}}(z)$ at the old prices.*
2. *Suppose the $\kappa < 1$, i.e. the relative price of services is lower in the status quo, compared to the counterfactual. Then:*
   (a) $\widetilde{z} \leq \kappa$ ***("Full corner"):*** *The consumer consumes only tradable goods both in the counterfactual and when trying to reach $\hat{\mathcal{V}}(z)$ at the old prices.*
   (b) $\kappa < \widetilde{z} < 1$ ***("Corner / Interior")****: The consumer consumes only tradable goods in the counterfactual but positive quantities of CS value added when trying to reach $\hat{\mathcal{V}}(z)$ at the old prices.*
   (c) $1 \leq \widetilde{z}$ ***("Full Interior")****: The consumer consumes positive quantities of CS value added in both the counterfactual and when trying to reach $\hat{\mathcal{V}}(z)$ at the old prices.*

Proof. See Section SM-1.2 in the Supplementary Material. □

These configurations are intuitive. Consider the case of $\kappa > 1$, i.e. services are relatively expensive at the old prices. If the consumer only consumes tradable goods in the new allocation ($\widetilde{z} < 1$), she will also only consume tradable goods when trying to achieve the same utility level at the old prices, i.e. when services are relatively expensive. By contrast, if the consumer choses an interior allocation in the new allocation ($\widetilde{z} > 1$), it is possible that the consumer would only consume tradables in case services are are sufficiently expensive ($1 < \widetilde{z} < \ln \kappa$). A similar intuition holds for the case of $\kappa < 1$, i.e. the case where services relatively cheap at old prices. There, if the consumer ends up buying services in the new allocation ($\widetilde{z} > 1$), she will also do so at old prices where services are cheap. However, it might be that the consumer only buys tradables in the new allocation ($\widetilde{z} < 1$) but ends up consuming services when facing old prices where services are cheap ($\kappa < \widetilde{z} < 1$).

Given this characterization, we can now formally state the implied equivalent variation. Specifically, we compute

$$\chi(z) \equiv \frac{EV(z)}{wz} = \frac{e(\hat{\mathcal{V}}(z), p_x, q)}{wz} - 1, \tag{A-17}$$

that is the percentage change in current income the individual has to receive to achieve the new utility level at current prices. Using the results in Proposition 3, $\chi(z)$ is given in the following Proposition.

Proposition 4. *Consider $\chi(z)$ as defined in (A-17). Let $\tilde{z}$ and κ be as defined in Proposition 3. Then*

1. *If $\kappa \geq 1$, $\chi(z)$ is given by*

$$\chi(z) = \begin{cases} \frac{\hat{w}/w}{B(\hat{p}_x,\hat{q})/B(p_x,q)}\left(1-\frac{\ln\kappa}{\tilde{z}}\right)^{1/\varepsilon} - 1 & \text{if } 1+\ln\kappa \leq \tilde{z} \\ \frac{exp\left(\frac{1}{\varepsilon}(\tilde{z}-\ln\kappa-(1-\ln\lambda))\right)}{\frac{zw}{B(p_x,q)}} - 1 & \text{if } 1 \leq \tilde{z} < 1+\ln\kappa \\ \frac{\hat{w}/\hat{p}_x}{w/p_x} - 1 & \text{if } \tilde{z} < 1 \end{cases}.$$

2. *If $\kappa < 1$, $\chi(z)$ is given by*

$$\chi(z) = \begin{cases} \frac{\hat{w}/w}{B(\hat{p}_x,\hat{q})/B(p_x,q)}\left(1-\frac{\ln\kappa}{\tilde{z}}\right)^{1/\varepsilon} - 1 & \text{if } 1 \leq \tilde{z} \\ \frac{\lambda^{1/\varepsilon}[\ln\tilde{z}+1-\ln\kappa]^{1/\varepsilon}}{\frac{zw}{B(p_x,q)}} - 1 & \text{if } \kappa \leq \tilde{z} < 1 \\ \frac{\hat{w}/\hat{p}_x}{w/p_x} - 1 & \text{if } \tilde{z} < \kappa \end{cases}.$$

With our choice of numeraire $(p_x = \hat{p}_x = 1)$, we have $w = \mathcal{A}^T$, $\hat{w} = \hat{\mathcal{A}}^T$, $q = \mathcal{A}^T/\mathcal{A}^{CS}$, and $\hat{q} = \hat{\mathcal{A}}^T/\hat{\mathcal{A}}^{CS}$.

Proof. See Section SM-1.2 in the Supplementary Material. □

Proposition 4 contains a full characterization of the welfare consequences of sectoral productivity growth. To build intuition, focus on the first and third branches, which are less knife-edge and which are the empirically relevant cases for our quantitative analysis. For poor consumers, i.e. for individuals with $\tilde{z}$ sufficiently low, we have the expected result that welfare is fully determined by $\frac{\hat{w}/\hat{p}_x}{w/p_x} = \frac{\hat{w}}{w}$ because they only consume tradable goods. For rich consumers, who consume CS in positive quantities at both prices vectors, the welfare effect is determined by the product of two terms

$$\frac{\hat{w}/w}{B(\hat{p}_x,\hat{q})/B(p_x,q)} \times \left(1-\frac{\ln\kappa}{\tilde{z}}\right)^{1/\varepsilon} = \underbrace{\frac{\hat{w}/w}{B(\hat{p}_x,\hat{q})/B(p_x,q)}}_{\text{Real wage}} \times \underbrace{\left(1-\lambda\varepsilon\omega_{CS}\frac{\ln(q/\hat{q})}{\left(\frac{z\hat{w}}{B(\hat{p}_x,\hat{q})}\right)^{\varepsilon}}\right)^{1/\varepsilon}}_{\text{Non-homo. adjustement}}.$$

The first term is akin to the change in the real wage as determined from the price index, $B(p_x, q)$, which is the price index evaluated at the asymptotic expenditure shares ω_{CS}. The second term captures the non-homothetic adjustment and hence the unequal gains from sectoral growth. Note first that if preferences are homothetic, $\nu_{CS} = \lambda = 0$, the adjustment disappears and the welfare gains are equal across consumers. Furthermore, note that

$$\frac{\partial \left(-\lambda \varepsilon \omega_{CS} \frac{\ln(q/\hat{q})}{\left(\frac{z\hat{w}}{B(\hat{p}_x,\hat{q})}\right)^{\varepsilon}} \right)}{\partial z} = \frac{\varepsilon^2 \omega_{CS}}{\left(\frac{z\hat{w}}{B(\hat{p}_x,\hat{q})}\right)^{\varepsilon} z} \times \lambda \ln(q/\hat{q}) .$$

This expression again highlights the unequal welfare effects of service-led growth. If CS output is a luxury ($\lambda > 0$), the adjustment is increasing in z (higher welfare gains for the rich) if relative service prices fall ($\ln(q/\hat{q}) = \ln\left(\frac{\hat{\mathcal{A}}^{CS}}{\mathcal{A}^{CS}}\frac{\mathcal{A}^T}{\hat{\mathcal{A}}^T}\right) > 0$), i.e. service productivity grows faster that tradable productivity. Vice versa, the adjustment is decreasing in z (higher welfare gains for the poor) if productivity in the tradable sector grows faster than productivity in services ($\ln(q/\hat{q}) = \ln\left(\frac{\hat{\mathcal{A}}^{CS}}{\mathcal{A}^{CS}}\frac{\mathcal{A}^T}{\hat{\mathcal{A}}^T}\right) < 0$).

A-2 Quantification of Baseline Model: Details

This section provides detailed steps for the quantitative analysis discussed in Section 3.5. Specifically, we first discuss the data we use and the moments we target to calibrate the structural parameters. We then discuss how we implement our counterfactual analysis.

A-2.1 Estimation

Data We calibrate our model using three data sources: (1) the micro data on sectoral employment as measured from IPUMS; (2) data on aggregate GDP from the WDI; (3) microdata on individual earnings from the "Harmonized World Labor Force Survey (HWLFS)" (see, e.g., Gottlieb et al. (2024)).

Parameters and Moments To calibrate our model, we need to estimate 8 productivity parameters: $\mathcal{A}^s_{rt}$ for two sectors (T vs CS), two regions (R and C), and two time periods ("pre" and "post"). Furthermore, we need three preference parameters (ε, ω_{CS}, and ν_{CS}), the average amount of human capital, $\mathbb{E}_{rt}[z]$, and the inequality index, $\mathcal{I}_{rt}$.

The Engel elasticity, ε, and the asymptotic CS value-added share, ω_{CS}, are taken from Fan et al. (2023), who follow the procedure laid out in the main text using micro data on household expenditure from India. In terms of the income distribution, $F_{rt}(z)$, we for now only assume that $\mathcal{I}$, defined in (8), is independent of r and t. This would, for example, be satisfied for the Fréchet distribution and the Pareto distribution if the shape parameter is common across space and time but the scale differs across both space and time. Because we can always choose the units of CS productivity, $\mathcal{A}_r^{CS}$, without loss of generality, we normalize the inequality adjusted non-homotheticity term, $\nu_{CS}\mathcal{I}$ to be -1. To map observed schooling into human capital, we assume an annual Mincerian return of $\rho = 0.1$, and hence set $\mathbb{E}_{rt}[z] = e^{0.1\times\overline{h}_{rt}}$, where $\overline{h}_{rt}$ denotes the average years of schooling in region r at time t. We summarize these parameters in the first four rows of Table A-I.

Table A-I: Structural Parameters

Parameters	Targets	Values
ε	Engel elasticity	0.395
ω_{CS}	From Fan et al. (2023)	0.692
$\nu_{CS}\mathcal{I}$	Normalization	-1
ρ	Mincerian schooling returns	0.1
ζ	Earning distribution within regions	2

Notes: This table summarizes the structural parameters we used for quantitative exercises.

To implement our estimation over both time and space, we take our "pre period" to be 2000 and "post period" to be 2010 and matching the timing of the respective IPUMS surveys (pre $\approx$ 2000, post $\approx$ 2010). We partition the population into two mutually exclusive locations, rural (R) and urban (C), with $\ell_R + \ell_C = 1$. Average years of schooling $\overline{h}_{rt}$ and sectoral employment shares ℓ_{rt}^s (by $r \in \{R, C\}$ and $s \in \{T, CS\}$) are computed as the location-specific averages of educational attainment and sectoral employment shares from IPUMS (see Section B-1.2). To compute such regional averages from the microdata, we employ the sampling weights that are included IPUMS (see also the discussion in Section B-1).

In our model, we considered the case of two sectors: tradables and CS. To map the data to these two industries, we associate CS directly with non-tradable services (i.e. retail, leisure, health) as in Figure 4 and combine the remaining sectors (agriculture; manufacturing; producer services including finance, business, ICT; and transport & storage) into tradables. We exclude public services because our theory concerns the provision of private services.

GDP per capita (PPP, constant 2021 int'l \$) is taken at the country-year level from WDI and matched to the exact IPUMS country and survey year. We then construct aggregate GDP pc capita at time t as the population-weighted average.

Identification. Given the structural parameters in Table A-I and overall human capital $\mathbb{E}_r[z]$, we need to identify the eight productivity terms $\{\mathcal{A}_{rt}^s\}$. We do so from the following restrictions.

1. **Rural tradable productivity in 2000 ($\mathcal{A}_{R2000}^T$):** By choice of units we can normalize rural tradeable productivity in the early period to unity; $\mathcal{A}_{R2000}^T = 1$.

2. **The urban productivity premium in tradables ($\mathcal{A}_{Ct}^T/\mathcal{A}_{Rt}^T$):** To estimate the relative productivity of urban locations in the tradable sector, i.e. $\gamma_t \equiv \mathcal{A}_{Ct}^T/\mathcal{A}_{Rt}^T$, we exploit micro-data on relative earnings form the "Harmonized World Labor Force Survey (HWLFS)".[1] Our

[1] We are very grateful to Charles Gottlieb for his help with this empirical anlysis.

model implies that total earnings for individual i in location r are given by $y_r^i = w_r z^i = \mathcal{A}_r^T z^i$. We therefore estimate γ from a regression of the form

$$\ln y_{rcti} = \delta_{tc} + \gamma_t \text{Urban}_{rcti} + x'_{rcti}\eta_t + u_{rcti}, \tag{A-18}$$

where r denotes region (i.e., urban versus not-urban), c denotes country, t denotes whether the survey belongs to the early period (i.e.. before 2000) or the late period (i.e., after 2000), and i denote the individual. The covariates x_{rcti} control for z^i, which is correlated with the urban dummy Urban_{rcti} via the distribution F_r.

Table A-II: Estimating the Tradable Urban-Rural Productivity Premium γ

	Log wage			
	(1)	(2)	(3)	(4)
Urban	0.154***	0.061*	0.263***	0.142***
	(0.032)	(0.022)	(0.050)	(0.027)
Time Period	2000-2010	2000-2010	2011-2025	2011-2025
Education		Yes		Yes
Demographics		Yes		Yes
Country-Year FE	Yes	Yes	Yes	Yes
Number of Country-Year Surveys	16	16	24	24
Observations	71,706	71,266	847,508	834,597
R^2 (within)	0.010	0.222	0.022	0.238

Notes: The table reports the results of estimating the regression in equation (A-18)

The results are reported in Table A-II. The first two columns refer to the early period (2000-2010), the last two columns refer to the late period (following 2010). Columns 1 and 3 report the unconditional estimates, columns 2 and 4 control for individual-level education and other demographics such as age and gender. The productivity premium is therefore around 6% in the early period and 14% in the late period. Given the sampling variation in these estimates, we, for simplicity, take use 5% and 15% in our quantitative analysis.[2]

3. **Rural tradable productivity in 2010 ($\mathcal{A}_{R2010}^T$):** We determine tradeable productivity in rural areas in 2010, $\mathcal{A}_{R2010}^T$, from the growth rate of aggregate GDP. Specifically, aggregate GDP is given by

$$GDP_t = w_{Rt}\mathbb{E}_{Rt}[z]\ell_{Rt} + w_{Ct}\mathbb{E}_{Ct}[z]\ell_{Ct} = \left(\mathbb{E}_{Rt}[z]\ell_{Rt} + \gamma_t \mathbb{E}_{Ct}[z]\ell_{Ct}\right)\ \mathcal{A}_{Rt}^T, \tag{A-19}$$

[2] These estimates are lower than the average results for urban-rural premia reported in Gollin et al. (2014). In the HWLFS we only observe earnings for individuals who receive wages. Given the large size of the informal sector in many countries in sub-Saharan Africa, this is a highly selected sample. Moreover, if informal and self-employed individuals have lower earnings and are more likely to live in rural areas, the overall urban-rural premium would be larger.

where ℓ_{rt} are the rural/urban population shares (see the column "Pop. Share" in Table 2) and γ_t is estimated in Table A-II. Using (A-19) and the normalization $\mathcal{A}^T_{R2000} = 1$, it follows that

$$\frac{GDP_{2010}}{GDP_{2000}} = \frac{\mathbb{E}_{R,2010}[z]\ell_{R,2010} + \gamma_{2010}\mathbb{E}_{C,2010}[z]\ell_{C,2010}}{\mathbb{E}_{R,2000}[z]\ell_{R,2000} + \gamma_{2000}\mathbb{E}_{C,2000}[z]\ell_{C,2000}}\mathcal{A}^T_{R,2010}. \tag{A-20}$$

This equation can be directly solved for $\mathcal{A}^T_{R,2010}$.

4. **CS productivity ($\mathcal{A}^{CS}_{rt}$):** To estimate productivity in the CS sector, $\mathcal{A}^{CS}_{rt}$, we use the labor market clearing condition (14):

$$\begin{aligned}\ell^{CS}_{rt} \;=\; \vartheta^{CS}_{rt} \;&=\; \omega_{CS} \;+\; \nu_{CS}\mathcal{I}\left(\mathbb{E}_{rt}[z]\,(\mathcal{A}^T_{rt})^{1-\omega_{CS}}\,(\mathcal{A}^{CS}_{rt})^{\omega_{CS}}\right)^{-\varepsilon}\\ &= \omega_{CS} \;-\; \left(\mathbb{E}_{rt}[z]\,(\mathcal{A}^T_{rt})^{1-\omega_{CS}}\,(\mathcal{A}^{CS}_{rt})^{\omega_{CS}}\right)^{-\varepsilon},\end{aligned} \tag{A-21}$$

where the second line uses $\nu_{CS}\mathcal{I} = -1$. Given $\mathbb{E}_{rt}[z]$, ε, ω_{CS}, and the identified $\mathcal{A}^T_{rt}$, equation (A-21) can be used to uniquely determine $\mathcal{A}^{CS}_{rt}$ in each period for both rural and urban areas.

The results for $\mathcal{A}^s_{rt}$ are reported in Table (3). To compute annual growth rates, we use for $g_{rs} = \left(\frac{\mathcal{A}^s_{r2010}}{\mathcal{A}^s_{r2000}}\right)^{1/10} - 1$.

A-2.2 Counterfactual analysis

Given the estimated model, we then conduct our counterfactual analysis; reported in Figure 7. Specifically, we refer to the *service-led* scenario as the equilibrium, where tradable productivity and human capital are held fixed ($\mathcal{A}^T_{r2010} = \mathcal{A}^T_{r2000}$ and $\mathbb{E}_{r,2010}[z] = \mathbb{E}_{r,2000}[z]$) and only CS productivity changes. In the same spirit, the *service-biased* scenario refers to the equilibrium where CS productivity and human capital is unchanged ($\mathcal{A}^{CS}_{r2010} = \mathcal{A}^{CS}_{r2000}$ and $\mathbb{E}_{r,2010}[z] = \mathbb{E}_{r,2000}[z]$) and tradable productivity changes.

Structural Change To quantity the consequences for structural change, we simply use the expression for the sectoral employment share in (A-21). The counterfactual *service-led* share is then given by

$$\widehat{\ell}^{CS,LED}_{r2010} = \omega_{CS} \;+\; \nu_{CS}\mathcal{I}\Big(\mathbb{E}_{r2000}[z]\;(\mathcal{A}^T_{r2000})^{1-\omega_{CS}}\;(\mathcal{A}^{CS}_{r2010})^{\omega_{CS}}\Big)^{-\varepsilon}.$$

The corresponding change due to service-led growth (blue bars in Figure 7), is then computed as $\widehat{\ell}^{CS,LED}_{r2010} - \ell^{CS}_{r,2000}$. Similarly, we compute CS employment in the *service-biased* scenario as

$$\widehat{\ell}^{CS,BIASED}_{r2010} = \omega_{CS} \;+\; \nu_{CS}\mathcal{I}\Big(\mathbb{E}_{r2000}[z]\;(\mathcal{A}^T_{r2010})^{1-\omega_{CS}}\;(\mathcal{A}^{CS}_{r2000})^{\omega_{CS}}\Big)^{-\varepsilon}.$$

Service-biased structural change (red bars in Figure 7) is then given by $\widehat{\ell}^{CS,BIASED}_{r2010} - \ell^{CS}_{r,2000}$.

Welfare To compute the associated changes in welfare, we compute equivalent variations (see Section A-1.4). To do so, we require the preference parameter ν_{CS}. Hence, it is at this point where we need to separately identify ν_{CS} and $\mathcal{I}$ as opposed to $\nu_{CS}\mathcal{I} = -1$. To make progress, we assume that z is Fréchet distributed as in (9): $F_{rt}(z) = \exp\left(-\left(z/\overline{z}_{rt}\right)^{-\zeta}\right)$. Hence, averge human capital is given by $\mathbb{E}_{r2000}[z] = \overline{z}_{rt}\,\Gamma_{\zeta}$, where $\Gamma_x \equiv \Gamma(1-\frac{1}{x})$ is the gamma function. Moreover, the scalar $\mathcal{I}$ si given by

$$\mathcal{I}_{rt} = \frac{\Gamma_{\frac{\zeta}{1-\varepsilon}}}{\Gamma_{\zeta}^{1-\varepsilon}} = \mathcal{I},$$

which shows that $\mathcal{I}_{rt}$ is indeed constant across time and space as long as ζ and ε are constant.

To calibrate ζ, recall that individual earnings in location r are given by $y_{rit} = w_{rt} z_i$. Hence, the distribution of earnings is given by

$$P(y_{rti} < \bar{y}) = F_{rt}\left(\frac{\bar{y}}{w_{rt}}\right) = e^{-\left(\frac{\bar{y}}{w_{rt}\bar{z}_{rt}}\right)^{-\zeta}},$$

i.e. earnings are also Fréchet distributed with shape ζ and scale $w_{rt}\bar{z}_{rt}$. This implies that the Gini coefficient of earnings is given by

$$GINI = 2^{1/\zeta} - 1.$$

Using data from Tanzania, Blanchard et al. (2025) report a Gini coefficient of earnings in both rural and urban locations of around 0.4 (see (Blanchard et al., 2025, Table 6)). This implies that $\zeta = \ln(2)/\ln(1+0.4) \approx 2.06$. We therefore use $\zeta = 2$ for our quantitative analysis (see last line of Table A-I). This implies that $\nu_{CS} = -\mathcal{I}^{-1}$, where $\mathcal{I} = \Gamma_{\frac{\zeta}{1-\varepsilon}}/\Gamma_{\zeta}^{1-\varepsilon}$ evaluated at $\zeta = 2$ and $\varepsilon = 0.395$.

We then compute the welfare gains from service-led growth, $\chi^{CS}(z)$, and the welfare gains from service-biased growth, $\chi^{T}(z)$, by applying Proposition 4. For service-led growth we set $\hat{\mathcal{A}}_r^T = \mathcal{A}_{r2000}^T$ and $\hat{\mathcal{A}}_r^{CS} = \mathcal{A}_{r2010}^{CS}$. For service-biased growth we set $\hat{\mathcal{A}}_r^T = \mathcal{A}_{r2010}^T$ and $\hat{\mathcal{A}}_r^{CS} = \mathcal{A}_{r2000}^{CS}$.

We evaluate these terms at different quantiles of the local income distribution. Recall that we assumed that z is Fréchet distributed according to (9). Define the τth quantile of the distribution in region r as z,

$$\tau = P\left(z_r \leq z^{\tau}\right) = F_r\left(z^{\tau}\right) = \exp\left[-\left(\frac{z^{\tau}}{\overline{z}_r}\right)^{-\zeta}\right].$$

This implies that $z^{\tau} = \left(-\ln\tau\right)^{-1/\zeta}\frac{1}{\Gamma_{\zeta}}\,\mathbb{E}_r[z]$, with τ the income quantile and Γ_{ζ} the normalization constant for the parametric income distribution. We set $\tau = 0.05$ for the poor and $\tau = 0.95$ for the rich, and report EVs for rural and urban groups.

APPENDIX B: EMPIRICAL RESULTS

B-1 Empirical Data Sources

This section documents the data sources and our measurement choices for the analysis in Section 2.

B-1.1 Aggregate Data

World Development Indicators (WDI). WDI is the World Bank's primary collection of development statistics, compiled from official sources. In this paper, we use the following series: (i) aggregate GDP and GDP per capita (PPP, constant 2021 international $); (ii) population; (iii) sectoral employment shares; (iv) urban population share (% of total population); and (v) exports of goods and services (% of GDP).

Table B-I lists all sub-Saharan African countries with average annual real GDP growth of at least 3.5% over 2000–2019, as reported in WDI, and highlights the countries included in our analysis—blue for East Africa and green for West Africa. We restrict attention to large, non-oil-exporting, and politically stable countries over this period. Unless otherwise noted, population refers to the year 2002, which we also use to weigh counties when computing aggregates.

The sectoral employment data used in panel (b) of Figure 2 follow the WDI's standard three-sector decomposition: (i) agriculture, including agriculture, hunting, forestry, and fishing; (ii) industry, including mining and quarrying, manufacturing, construction, and utilities; and (iii) services, including wholesale and retail trade; restaurants and hotels; transport, storage, and communications; finance, insurance, real estate, and business services; and community, social, and personal services. This taxonomy aligns with the micro-level classification we implement below using IPUMS.

Educational Attainment. We measure educational attainment at the country-level from the Barro–Lee data (Barro and Lee, 2013). These data are an internationally comparable panel reporting schooling distributions and average years of schooling by sex and 5-year age group for about 146 countries (1950–2015). We focus on average years of schooling for ages 25–34 to capture education improvements among new labor-market entrants, comparing 2000 and 2015 where available.

Gallup World Poll. We measure economic well-being using the Life Evaluation Index from the Gallup World Poll, based on a Cantril ladder. Respondents are asked to imagine a ladder with steps numbered from 0 (worst possible life) to 10 (best possible life) and answer:

- "On which step of the ladder would you say you personally feel you stand at this time?"
- "Just your best guess, on which step do you think you will stand in the future, say about five years from now?"

Table B-I: Descriptive Statistics of Sub–Saharan Countries

Country	Code	GDP	GDP pc	Population	Type	Small
Ethiopia	ETH	0.090	0.060	71,635,517	NRI	
Rwanda	RWA	0.077	0.052	8,400,180	NRI	
Mozambique	MOZ	0.069	0.041	18,958,338	NRI	
Somalia	SOM	0.074	0.040	9,550,063	N/A	
Cabo Verde	CPV	0.045	0.038	467,046	NRI	Yes
Ghana	GHA	0.062	0.036	20,729,388	RI	
Mauritius	MUS	0.039	0.035	1,204,621	NRI	Yes
Tanzania	TZA	0.064	0.034	36,121,304	RI	
Uganda	UGA	0.064	0.032	25,532,870	NRI	
Chad	TCD	0.066	0.029	9,116,192	OIL	
Nigeria	NGA	0.056	0.029	133,471,989	OIL	
Burkina Faso	BFA	0.058	0.027	12,680,897	RI	
Seychelles	SYC	0.037	0.027	83,723	NRI	Yes
Eswatini	SWZ	0.032	0.026	1,059,441	NRI	Yes
Sierra Leone	SLE	0.055	0.025	4,971,435	RI	
Equatorial Guinea	GNQ	0.072	0.024	797,957	OIL	Yes
Guinea	GIN	0.046	0.022	8,674,066	RI	
Sao Tome and Principe	STP	0.042	0.021	149,922	NRI	Yes
Congo, Dem. Rep.	COD	0.053	0.020	53,750,524	RI	
Lesotho	LSO	0.024	0.019	1,994,679	NRI	Yes
Mali	MLI	0.051	0.018	12,288,050	OIL	
Angola	AGO	0.055	0.017	17,327,699	OIL	
Kenya	KEN	0.045	0.017	32,629,808	NRI	
Namibia	NAM	0.037	0.017	1,889,970	RI	Yes
Botswana	BWA	0.035	0.017	1,744,763	RI	Yes
Malawi	MWI	0.045	0.016	11,880,440	NRI	
Niger	NER	0.054	0.016	12,349,144	RI	
Senegal	SEN	0.041	0.015	10,466,214	NRI	
Benin	BEN	0.044	0.014	7,675,426	NRI	
South Africa	ZAF	0.026	0.013	48,019,415	RI	
Cameroon	CMR	0.042	0.013	15,725,528	OIL	
Guinea-Bissau	GNB	0.036	0.011	1,289,307	NRI	Yes
Mauritania	MRT	0.040	0.011	2,738,655	N/A	
Côte d'Ivoire	CIV	0.036	0.010	18,654,771	NRI	
Togo	TGO	0.036	0.009	5,413,274	NRI	

Notes: This table reports all sub-Saharan African countries with average annual real GDP growth $\geq 3.5\%$ over 2000–2019 and indicates our analysis samples — blue for East Africa and green for West Africa. GDP and GDP per capita are average annual growth rates over 2000–2019 from the World Development Indicators (PPP, constant 2021 international $). Population is measured in 2002 and used as fixed weights throughout the paper. *Type* follows International Monetary Fund (2024): NRI = Non-Resource-Intensive; OIL = Oil Exporters; RI = Other Resource-Intensive; N/A = not classified. *Small* indicates population < 2 million.

Respondents are classified as *thriving* if their current life evaluation is between 7 and 10 and their expected future evaluation is between 8 and 10; as *suffering* if both current and expected future evaluations are 4 or below; and as *struggling* otherwise. We compute changes in the country-level shares in these three categories between 2008 and 2019, replacing 2008 with 2009 for Côte d'Ivoire and Malawi due to data availability. Guinea, Guinea-Bissau, and Ethiopia are excluded because no data are available during this period. By construction, changes in the three shares sum to zero.

B-1.2 Microdata: IPUMS

Source and Coverage. We use harmonized census microdata from IPUMS International. We select census rounds between 1997 and 2020 that consistently report information on industry/sector, occupation, education, and urban–rural residence. Vietnam is included as a comparator, using the 1999 and 2019 rounds. Responses coded as "Unknown" or "NIU" (not in universe) are treated as missing for the corresponding variables.

Table B-II summarizes the census rounds and how each country enters the analysis. Column "Emp. Analysis" indicates whether a country is included in the employment analyses in the main text. Column "Amenity Change" indicates the countries used to construct Figure B-3. The last column reports whether the country is included in the cross-sectional amenity regression in Table 1, where we use only the most recent census round available in each country.

Table B-II: IPUMS Coverage

Country	Census Rounds	Emp. Analysis	Amenity Change	Amenity Reg.
Benin	2002, 2013	Yes		Yes
Burkina Faso	2006			Yes
Ethiopia	2007			Yes
Ghana	2000, 2010	Yes		Yes
Kenya	2019			Yes
Malawi	1998, 2008	Yes		Yes
Mozambique	1997, 2017	Yes		Yes
Rwanda	2002, 2012	Yes	Yes	Yes
Senegal	2002, 2013		Yes	Yes
Sierra Leone	2004, 2015		Yes	Yes
Tanzania	2002, 2012	Yes	Yes	Yes
Togo	2010			Yes
Uganda	2002, 2014		Yes	Yes
Vietnam	1999, 2019	Yes		

Notes: This table lists all IPUMS census samples used in the analysis. Vietnam is used only as a comparator in the structural change exercises and is not included in the amenity analysis.

Unit of Analysis and Weights. To measure changes in sectoral employment, we focus on individuals. We restrict to working age (25–60) employed individuals (IPUMS `EMPSTAT` $= 1$) and use IPUMS person weights (PERWT) to compute aggregates. The availability of durable goods and other amenities is measured at the household level. We therefore aggregate the information at the household level using household weights (HHWT).

Classification of amenities and durable goods. We construct binary indicators for household amenities and availablity of durable goods using harmonized IPUMS variables. Table B-III summarizes the coding rules; all responses coded as "Unknown" or "NIU" are set to missing prior to recoding.

Table B-III: Classification of IPUMS Amenities

Amenity	Variable	Coding rule (indicator = 1)
Electricity	ELECTRIC	Access to electricity
Radio	RADIO	Access to a radio
Water supply	WATSUP	Piped water (codes 10-18)
Toilet	TOILET	Flush toilet (code 21)
Wall material	WALL	Good walls (codes 500-526, 550-557)
Roof material	ROOF	Good roof (codes 10-38)

Notes: This table lists all IPUMS amenity variables used in the analysis. Water sources piped inside or outside the dwelling are classified as piped water. Walls made of brick, stone, or metal are treated as "good", as are roofs made of cement, tile, or metal.

Sector Classification. We aggregate IPUMS industry codes into six broad sectors: agriculture, mining, manufacturing, construction & utilities, private services, and public services. Within private services, we distinguish three subgroups used in the main text: CS (retail, leisure, health), producer services (finance, business, ICT), and transport and storage. Table B-IV reports the mapping from the original IPUMS codes to our detailed and broad sector classifications.

Other Characteristics. We also recode several demographic variables to ensure consistency across countries:

- Years of schooling: we recode "some primary" (code 91) to 4 years; "some technical after primary" (code 92) and "some secondary" (code 93) to 10 years; "some tertiary" (code 94) to 14 years. "Not specified" (code 90), "Adult literacy" (code 95), and "Special education" (code 96) are set to missing.
- Educational attainment: we define "high-educated" as completed at least secondary education; all others are "low-educated".
- Class of worker (CLASSWK): we distinguish self-employed, wage/salary workers, and a residual group combining unpaid workers and "other" categories, since the latter are small.

All recodes are applied consistently across countries and waves; the replication files provide the exact maps from harmonized codes to analysis variables.

Table B-IV: Reclassification of IPUMS Industry Codes

Code	Original IPUMS Label	Detailed Sector	Broad Sector
010	Agriculture, fishing, and forestry	Agriculture	Agriculture
020	Mining and extraction	Mining	Manufacturing
030	Manufacturing	Manufacturing	Manufacturing
040	Electricity, gas, water and waste management	Construction & Utilities	Manufacturing
050	Construction	Construction & Utilities	Manufacturing
060	Wholesale and retail trade	Retail, Leisure, and Health	Private Services
070	Hotels and restaurants	Retail, Leisure, and Health	Private Services
080	Transportation, storage, and communications	Transport and Storage	Private Services
090	Financial services and insurance	Finance, Business, and ICT	Private Services
100	Public administration and defense	Education & Public Administration	Public Services
110	Services, not specified	Retail, Leisure, and Health	Private Services
111	Business services and real estate	Finance, Business, and ICT	Private Services
112	Education	Education & Public Administration	Public Services
113	Health and social work	Retail, Leisure, and Health	Private Services
114	Other services	Retail, Leisure, and Health	Private Services
120	Private household services	Retail, Leisure, and Health	Private Services
000	NIU (not in universe)	*Set to missing / excluded from analysis*	
130	Other industry, n.e.c.	*Set to missing / excluded from analysis*	
998	Response suppressed	*Set to missing / excluded from analysis*	
999	Unknown	*Set to missing / excluded from analysis*	

Notes: CS comprises retail, leisure, and health; producer services comprise finance, business, and ICT; transport and storage is treated as a mixed private service. Full code–sector mappings by country and year are provided in the replication files.

B-2 Additional Empirical Results

In this subsection, we provide additional results for our empirical analysis contained in Section 2.

Urbanization and Educational Attainment

Figure B-1 complements Figure 2 by documenting urbanization and educational attainment for the same set of African countries, using the WDI and Barro–Lee measures described in Section B-1.1.

Panel A reports trends in urbanization. The urban share of the population rose by more than 8 percentage points on average, with faster increases in West than in East Africa. Ghana, Tanzania, and Burkina Faso each experienced increases exceeding 12 percentage points, whereas Malawi and Rwanda remain low-urbanization outliers, with some of the lowest urban shares in the world.

Panel B reports changes in educational attainment for the 25–34 age group in countries where data are available. These data indicate substantial gains in human capital: on average, individuals aged 25–34 in 2015 had completed about 1.5 more years of schooling than the corresponding cohort in 2000. The gains are heterogeneous, ranging from almost no change in Kenya to an increase of roughly 4.5 years in Mozambique.

Happiness

We complement the income-based evidence with self-reported well-being, using the Gallup World Poll Life Evaluation Index defined in Section B-1.1. To maximize sample size, we use the 2008–

Figure B-1: Education and Urbanization

(a) Urbanization

(b) Rise in Educational Attainment

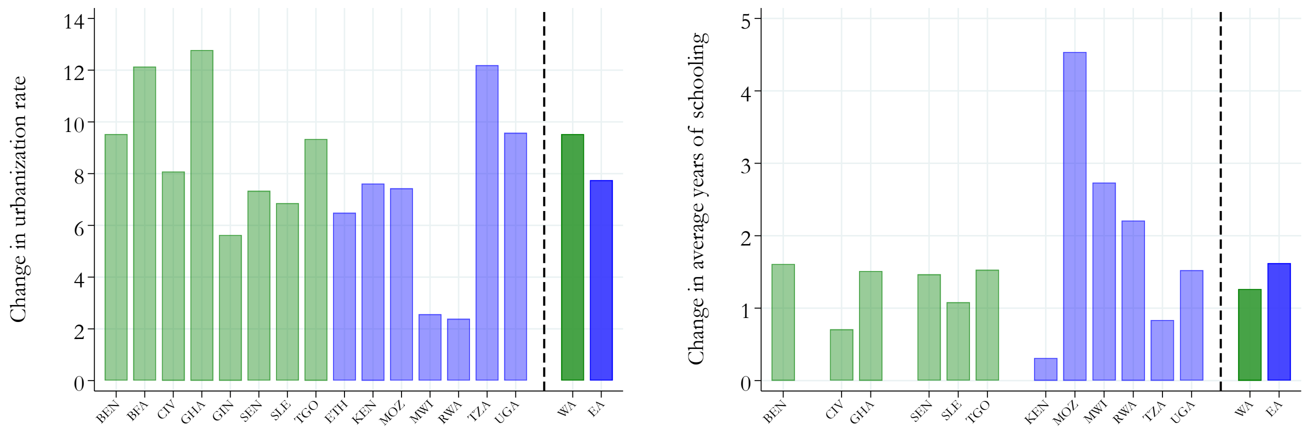

Notes: Panel A illustrates changes in the urbanization rate between 2000 and 2019. Panel B shows the rise in average years of schooling among the 25–34 age group between 2000 and 2015. BFA, GIN, GNB, and ETH are missing in the Barro–Lee data. Regional averages are population-weighted using 2002 population shares. Country codes follow ISO-3 standards (e.g., KEN = Kenya, CIV = Côte d'Ivoire), and regional abbreviations are WA = West Africa (green) and EA = East Africa (blue). Guinea-Bissau (GNB) is not shown as an individual bar but is included in the West Africa average. Data Source: WDI, Barro–Lee.

2019 window (replacing 2008 with 2009 for Côte d'Ivoire and Malawi due to data availability); Guinea, Guinea-Bissau, and Ethiopia are excluded because no data are available in this period.

Figure B-2 shows that, despite considerable heterogeneity, the share of people classified as thriving increased between 2008 and 2019 in all countries except Malawi and Rwanda, with an average increase of around 13 percentage points. The increase in the share of thriving respondents is primarily offset by a decline in the share of those classified as struggling, while the proportion of those suffering remained relatively stable. Overall, subjective well-being improves in this set of fast-growing poor economies, mainly through an expansion of the group reporting the highest life satisfaction.

Changes in Household Amenities Over Time

In Table 1 we analyzed cross-sectional patterns of durable ownership across individuals. Here, we examine how economic growth has affected access to these amenities over time by comparing multiple survey waves from the same countries, using the amenity definitions and coding described in Section B-1.2. For each country, we compute changes between the initial and final survey rounds, and then calculate a weighted average using population size as weights, disaggregated by rural and urban areas. Figure B-3 reports the results.

Given the heterogeneity across surveys, some degree of measurement error is to be expected. Nonetheless, the results broadly align with expectations: African households have experienced rising access to electricity, piped water, flush toilets, and improved wall and roof materials, and

Figure B-2: Changes in Happiness

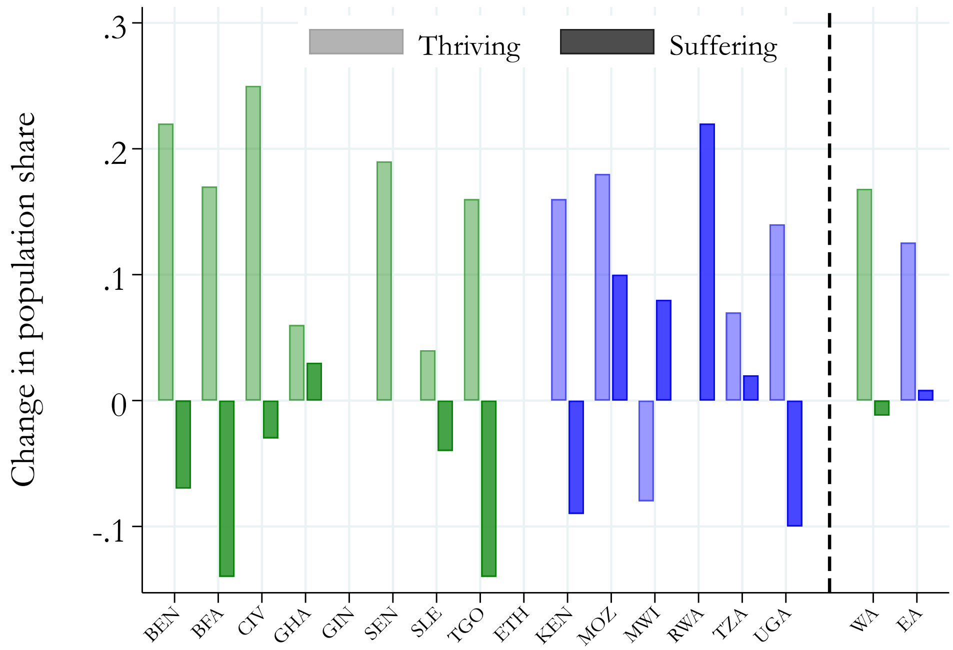

Notes: This figure shows the percentage-point change in the share of individuals classified as thriving (light shaded bars) or suffering (dark shaded bars) from the Gallup Life Evaluation Index between 2008 and 2019. Country codes follow ISO-3 standards (e.g., KEN = Kenya, CIV = Côte d'Ivoire), and regional abbreviations are WA = West Africa (green) and EA = East Africa (blue). For CIV and MWI, 2009 is used as the baseline year because 2008 data are unavailable; GIN, GNB, and ETH are excluded due to missing baseline data. The absence of a bar indicates no change between the two years. Regional averages are population-weighted using 2002 population shares. Data source: Gallup World Poll.

are increasingly likely to own a radio. The slight decline in the urban share of households with piped water likely reflects the challenge many governments face in expanding public utility provision amid rapid urbanization.

International Trade

Figure B-4 compares exports of goods and services as a share of GDP in our East and West African country groups with Vietnam, using WDI data. In East Africa, exports constitute a relatively modest share of GDP and show no clear upward trend over time. In West Africa, the export-to-GDP ratio is somewhat higher (around 25%), reflecting the presence of mineral-rich economies such as Ghana and Guinea, but export performance is again fairly flat. By contrast, Vietnam's export-to-GDP ratio rises sharply from about 57% to 82% between 2000 and 2019.

This divergence reflects deeper differences in development models. Vietnam exemplifies the canonical export-led path to structural transformation: rapid integration into global value chains has fueled industrialization, urban wage growth, and productivity gains. In contrast, most African economies have remained only weakly integrated into global markets, and the tradable sector has played a relatively modest role in employment and output growth. The limited dynamism of exports, combined with a sectoral composition tilted toward non-tradable CS, suggests that domestic demand has been the primary engine of structural change.

Figure B-3: Growth in Household Durables

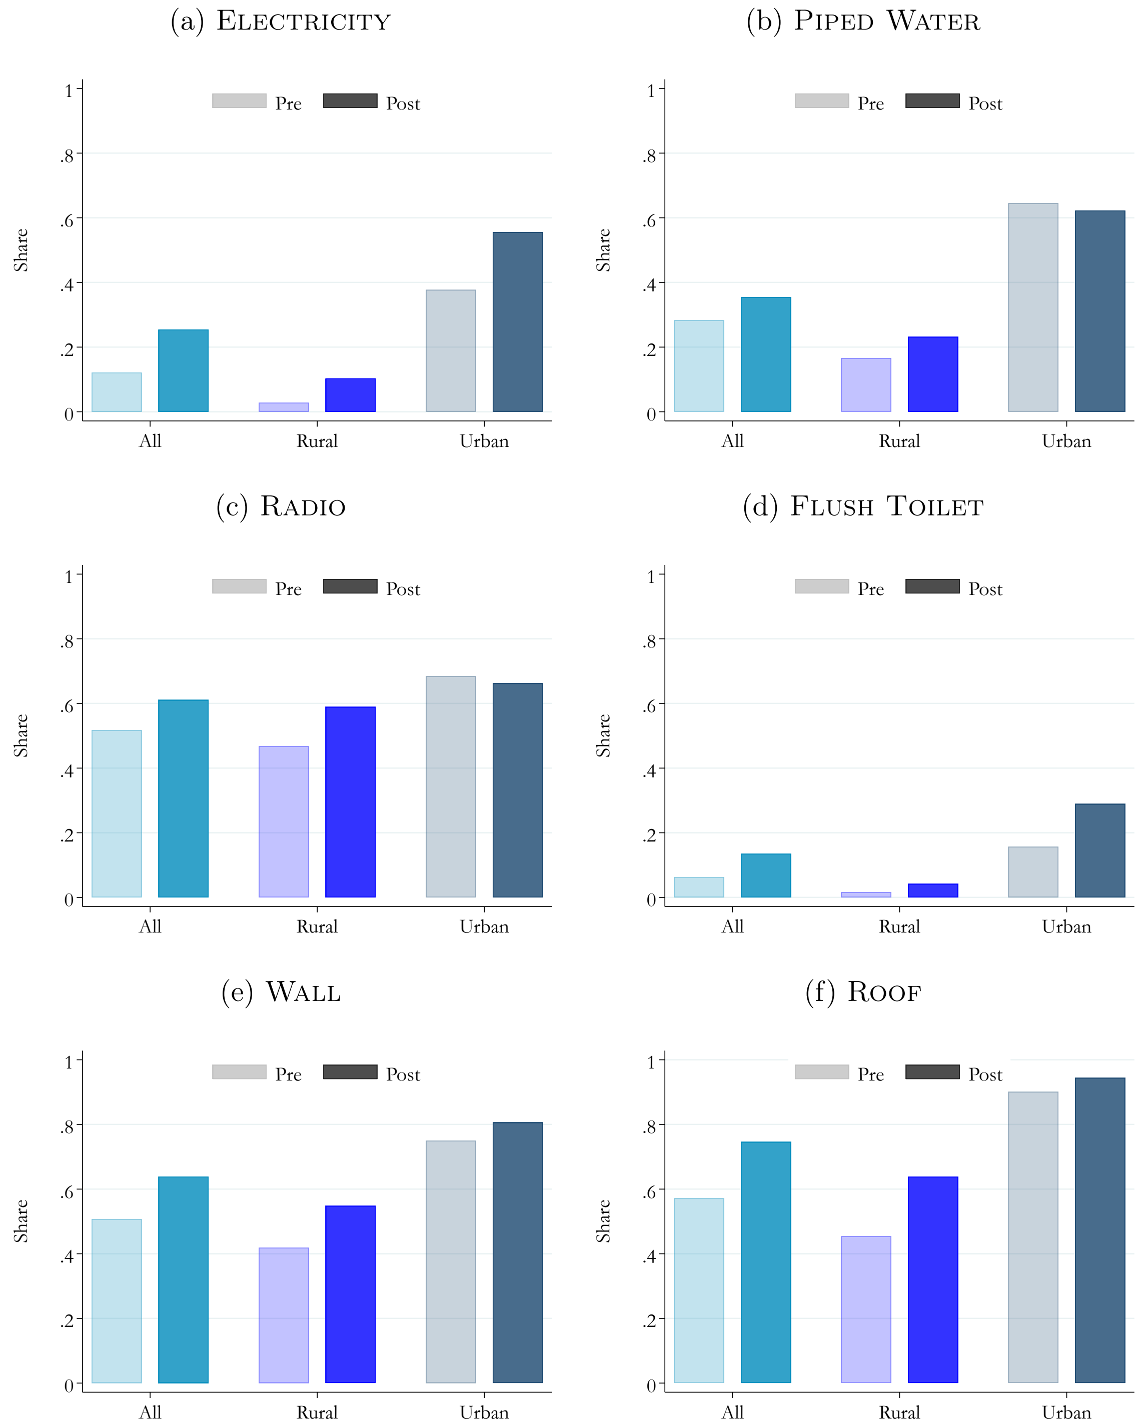

Notes: The figure reports the share of households that have access to electricity, piped water, radio, flush toilet, good-quality wall, and good-quality roof (as defined in Section B-1.2) in rural and urban districts between pre- and post-period. Surveys used for calculating averages include: Rwanda (2002, 2012), Tanzania (2002, 2012), Uganda (2002, 2014), Senegal (2002, 2013), Sierra Leone (2004, 2015). The pre-period is before 2005, while the post-period is after 2005. The shares are a population-weighted average of country-level surveys using 2002 population shares. Data source: IPUMS International.

Figure B-4: International Trade

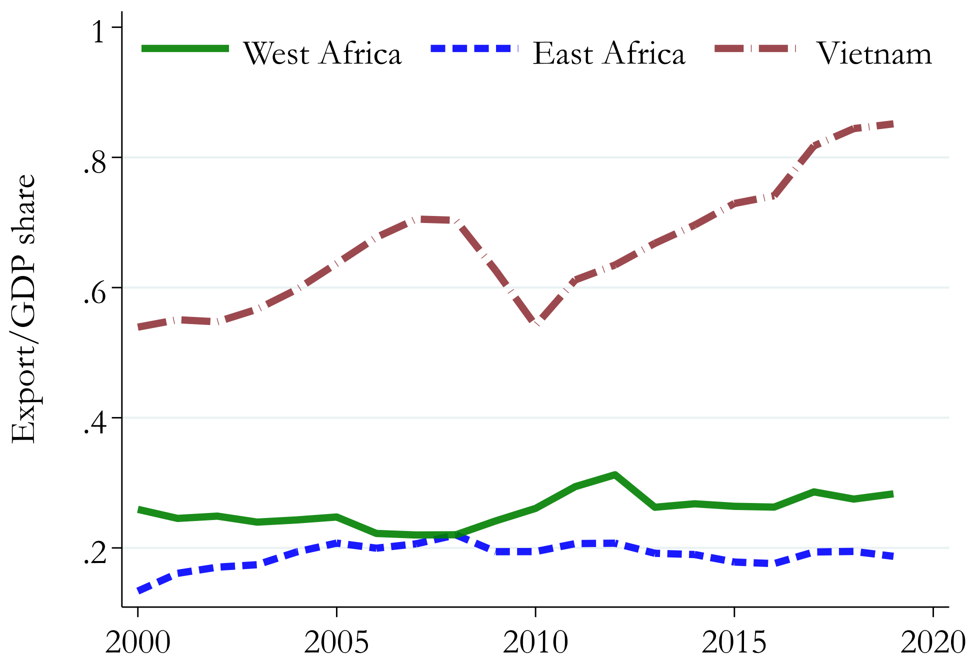

Notes: This figure plots the time series of export shares in GDP for 2000–2019. Malawi and Ethiopia are excluded from the East Africa series due to missing data. Regional averages are population-weighted using 2002 population shares. Data source: WDI.