

Service-Led or Service-Biased Growth?

Equilibrium Development Accounting across Indian Districts

Tianyu Fan, Michael Peters, and Fabrizio Zilibotti

Yale University*

March 2021

Abstract

In many developing countries today, the structural transformation takes the form of a shift of employment out of agriculture into the service sector. By contrast, industrial employment is mostly stagnant. Is the service sector an engine of growth and hence growth service led? Or is its expansion a mere corollary of growth, where rising incomes stemming from productivity growth in goods-producing industries increase the demand for services? To determine whether growth is service led or service biased, we estimate a spatial equilibrium model with nonhomothetic preferences. Our methodology is in the spirit of development accounting and lends itself to a quantitative assessment of both the aggregate and the heterogeneous welfare effects of sectoral productivity growth. In an application to India, we find that productivity growth in consumer services such as retail and hospitality was an important driver of rising living standards between 1987 and 2011 and accounted for 1/3 of aggregate welfare gains. However, such benefits were highly skewed and accrued mostly to high-income households living in urbanized locations.

1 Introduction

Most industrialized countries have undergone a similar pattern of structural transformation. At an early stage of development, a growing industrial sector drew labor from a declining agricultural sector. At a later stage, the employment shares of both agriculture and manufacturing fell, while the service sector became the main source of employment growth. This pattern fits the experience of most Western and East-Asian economies. However, in other parts of the globe, economic development appears to have taken a different turn. Over the last four decades, the share of manufacturing jobs has barely grown in most developing countries, including fast-growing economies such as India and sub-Saharan Africa. There, the structural transformation has taken the form of a shift from agriculture to services.

To many scholars (e.g., McMillan and Rodrik (2011) or Rodrik (2015)), the absence of employment growth in the manufacturing sector, also dubbed "premature de-industrialization" is a cause of concern, because technical progress in manufacturing is seen as the engine of growth. By contrast, the expansion of the service sector is often interpreted as a by-product of growth driven by rising incomes. Hence, growth is service biased, but not service led.

*This study is part of a broader research project also involving Philippe Aghion and Robin Burgess. We are extremely grateful to our discussant Sebastian Sotelo for his thoughtful suggestions. We thank Fabian Eckert, Reto Foellmi, Doug Gollin, Cormac O'Dea, Peter Klenow, Samuel Kortum, Rachel Ngai, Richard Rogerson, María Sáez Martí, and seminar participants at the ASSA Meeting 2021, Cowles Macro Conference, the STEG Workshop, Penn State University, Peking University, RIDGE, the University of Sankt Gallen, and Yale University. We also thank Sarah Moon, Shengqi Ni, Pariroo Rattan, and Huihuang Zhu for excellent research assistance. Financial support from SNSF project "Inequality, Cultural Transmission, and Human Capital Accumulation" (grant 100018 165616) is gratefully acknowledged. Michael Peters thanks the "Minnesota Opportunity and Inclusive Growth Institute" for the hospitality.

A key hurdle to resolving this controversy is measurement: quantifying productivity growth in service industries is difficult. In this paper, we propose a novel structural methodology that allows us to estimate productivity in the service sector and quantitatively assess its importance for the development process. We then apply this methodology to the growth process of India since 1987. Our results go against this pessimistic view. We find substantial productivity growth in consumer services such as retail or hospitality that can account for about one third of aggregate welfare gains. In short: Indian growth was to a large extent service led.

Methodology: Equilibrium Development Accounting. Our approach follows the tradition of the development accounting literature. We do not attempt to explain the determinants of productivity growth but provide a method to measure sectoral productivity through the lens of a structural model. The estimation is disciplined by a theory that stands on two building blocks: (i) nonhomothetic preferences, and (ii) a spatial multisector equilibrium model with inter-regional trade where firms have heterogeneous productivities in different locations. These two ingredients allow us to capture two important features of services: the possibility of service-biased growth if services are luxuries and the fact that services are non-tradable and hence depend on local demand.

Because the service sector is broad and heterogeneous, we split it into two parts. Many services improve households' access to consumption goods (e.g., restaurants or retail) or directly enter their consumption basket (e.g., leisure services). We call these services *consumer services (CS)*. Other services are predominantly inputs to the production of goods and include, among others, business services, corporate law services, and part of transport services. We refer to such services as *producer services (PS)*. To capture this distinction, we model CS as directly entering the households' consumption basket so that consumers' preferences are defined over three final items: *food*, *industrial goods*, and *CS*. By contrast, PS are inputs into the production of industrial goods so that people can work in four sectors of activity: agriculture, manufacturing, PS and CS.

To highlight the fact that restaurant visits and personal services are local goods, we assume food and industrial goods are freely traded, whereas CS have to be purchased locally. Local CS productivity therefore directly affects the price and availability of CS in each market. Similarly, local income directly affects the demand for the local provision of CS. By contrast, PS are embedded in industrial goods so that their value added is tradable. To close the model, we assume labor is the only productive factor and is perfectly mobile across industries although immobile across locations. Thus, there is a single wage per effective unit of labor in each location, although wages differ across locations. These assumptions are extreme but can be relaxed by introducing non-prohibitive labor-mobility frictions across both sectors and space.

Our model allows us to disentangle the relative importance of service-biased versus service-led growth by estimating the variation in sectoral productivity across space and time. Conditional on a set of preference and technology parameters, regional endowments of human capital, and sectoral labor productivity in each region, the model yields a unique equilibrium vector of wages and sectoral labor allocations. Conversely, if we have data for the allocation of labor across sectors in each district and for local earnings, we can retrieve a unique set of labor productivity for each region-sector pair that rationalizes the data.

The main identification challenge emerges because of non-homothetic preferences. Imagine, for example, two small districts R (rich) and P (poor), which are part of a large multi-district economy, and abstract for simplicity from PS. Let A_{rs} denote labor productivity in district r and sector s . Consider two polar opposite scenarios. In the first scenario, $A_{RF} > A_{PF}$ and $A_{RM} > A_{PM}$; that is, the productivity of the tradable sectors food and manufacturing are larger in R than in P . Instead, the productivity of CS is the same in both districts, $A_{RCS} = A_{PCS}$. In equilibrium, workers in district R earn a higher wage. If food is a necessity and CS is a luxury, consumers

in R spend a higher share of their income on CS, whereas consumers in P spend a higher share of their income on food. Because CS is nontradable, district R has a larger employment share in CS, whereas P specializes in the production of goods.¹ In this scenario, spatial differences in expenditure and employment shares are entirely driven by income effects in demand. We refer to this case as *service-biased growth*.

The second scenario is one in which productivity in the tradable sectors is identical in R and P , whereas $A_{RCS} > A_{PCS}$. As before, district R is richer, has a larger service sector, and consumers spend a smaller share of their budget on food. However, in this case, any difference stems from a technological gap in the CS sector. We refer to this scenario as *service-led growth*.

Although the example highlights the difference across space, the same argument applies to the analysis of a given district at two points in time. Under *service-biased growth*, the growth of the service sector would be entirely a consequence of the productivity growth in the goods-producing sectors. Under *service-led growth*, productivity growth in the service sector would be the sole cause of productivity growth and structural change.

To solve this identification problem, we estimate the aggregate demand system for consumer services and directly measure the strength of the income effect. The key aspect of the demand system is the income elasticity. If service demand is very income elastic, growth tends to be service biased, because rising incomes sharply increase service demand. If, by contrast, this income elasticity is small, changes in employment indicate changes in productivity.

Given the pivotal role of this elasticity, we estimate it from Engel curves using microdata on income and expenditure. To link this empirical estimate to the aggregate demand system of our theory, we have to address two challenges. First, in the presence of nonhomothetic preferences, the income elasticity at the micro level is not necessarily informative about the elasticity of aggregate demand. Second, as stressed in Herrendorf et al. (2013), in general, no direct mapping exists between the preference parameters estimated from data on final expenditure to the ones of the value-added demand system used in our theory.

The preference specification we opt for overcomes both of these difficulties. We model preference as belonging to the PIGL class, which was first introduced by Muellbauer (1976) and has been recently popularized in the literature on growth and structural change by Boppart (2014). The PIGL preference class has a crucial aggregation property: the choice of a set of agents endowed with PIGL preferences facing a common price vector can be rationalized as the choice of a representative agent whose preferences also fall into the PIGL class. In particular, the parameter governing the income elasticity for aggregate value added coincides with the micro elasticity. Moreover, we also show this parameter can be estimated from data on final expenditure.

Our methodology allows us quantify the extent to which the growth of the service sector is either a source or a consequence of the development process. First, we estimate sectoral productivity growth at the regional level and thus we can directly determine whether productivity in the service sector actually grew. Second, it lends itself to a quantitative assessment of both the aggregate and the heterogeneous welfare impact of service-led growth. In particular, when agents have nonhomothetic preferences and live in different locations, the growth of different sectors benefits people differentially—rich versus poor as well as urban versus rural residents. Using our model, we can quantify the heterogeneous welfare gains of service-led growth both across the income distribution and across space. We view this as important, because growth theory is often criticized for glossing over the distributional implications of economic growth.

Application: Service-led growth in India (1987–2011). We apply our methodology to India, a fast-growing economy, with an average annual growth rate of 4.2% during 1987–2011. In this period, the lion’s share of the

¹ Because food and industrial goods are tradable, whether P specializes in agriculture or industry depends on its comparative advantages.

process of structural change was a shift from agriculture to services. Our estimation exploits individual geolocalized consumption and employment data, and we estimate sectoral productivity growth for almost 400 Indian districts.

Our results are interesting in several respects. First, at the spatial level, there are large sectoral productivity differences. In particular, the CS sector features a large productivity gap between urban and rural districts. Thus, urban districts have a higher service employment share not only because their inhabitants are richer, but also because the market provides them more efficiently.

Second, we document a very important role for service-led growth. At the aggregate level, rising efficiency in the provision of consumer services accounts for almost one third of the increase in living standard since 1987. For comparison, the impact of agricultural productivity growth is roughly similar, but growth in the industrial sector was substantially less important. In fact, using a non-parametric bootstrap procedure to estimate the sampling uncertainty in our estimates, we show that the difference between service-led and industrial growth is statistically different. To the best of our knowledge, this paper is the first to quantify the importance of the consumer-service sector for a developing economy such as India.

Third, service-led growth is very unequal. Productivity growth in CS is the main source of welfare gains for richer households, especially those in urbanized districts. The residents in the top quintile of urbanization would have been better off taking a 41% income cut in 2011 than moving back to the productivity that the CS sector had in 1987. Similarly, service-led growth accounts for the vast majority of welfare gains the richest 20% of the Indian population experienced since 1987. By contrast, for poorer households living in rural districts, improvements in living standards hinge mostly on productivity growth in agriculture.

Finally, productivity growth in CS turns out to also be the key driver of structural change. Had productivity in the service sector stagnated, the employment share of agriculture would not have declined. By contrast, the effect of agricultural productivity growth is negligible.

Related Literature. Our paper contributes to the macroeconomic literature on the structural transformation including, among others, Ngai and Pissarides (2007), Herrendorf et al. (2013, 2014, and 2020), Gollin et al. (2014), Hobijn et al. (2019), and Garcia-Santana et al. (2020).

A recent literature focuses on the service sector, however, mostly with a focus on developed economies such as the US. Buera and Kaboski (2012) emphasize the importance of the (demand-driven) growth of a skill-intensive service industry in the post-1950s US economy. Hsieh and Rossi-Hansberg (2019) argue that in more recent years, ICT has triggered an industrial revolution and has been a major source of productivity growth. Their view is echoed by Eckert et al. (2020). An exception to this rich-country focus is Duarte and Restuccia (2010), who document large cross-country productivity differences in service industries, a finding broadly in line with our results across locations within India and Gollin et al. (2015), who emphasize how urbanization often goes hand in hand with a booming consumption of non-tradable services, although their focus on such “consumption cities” in resource-rich African economies is very different from ours. Desmet et al. (2015) and Dehejia and Panagariya (2016) also study the role of the service sector in India and document an important role for cities, in particular in the provision of PS. Finally, our finding that service growth was decidedly pro-rich and pro-urban is consistent with Chatterjee and Giannone (2021), who use data on regional income growth for a large number of countries and document that rising productivity in services is associated with regional divergence.

On the methodological side we build on the large literature on development accounting; see, for example, Caselli (2005) and Hall and Jones (1999). This literature postulates aggregate production functions and uses information on the accumulation of productive factors to fit the data. Our methodology is closer to the structural development

accounting of Gancia et al. (2013), who exploit the restrictions imposed by an equilibrium model to identify sectoral productivities. Similarly, Cheremukhin et al. (2015) and Cheremukhin et al. (2017) use an accounting approach in conjunction with a neoclassical growth to study the determinants of growth in China and Russia.

We perform our accounting exercise in the context of a model with inter-regional trade linkages, commonly used in the economic geography literature; see, for example, Redding and Rossi-Hansberg (2017) or Allen and Arkolakis (2014). In contrast to these papers, we abstract from labor mobility, even though one could extend our methodology to allow for it. Cravino and Sotelo (2019) is a recent example of an analysis of the structural transformation in the context of a model with international trade.

Non-homothetic preferences play a key role in our analysis. The classic reference for the service bias (and resulting cost disease) of economic growth is Baumol (1967). Earlier papers emphasizing their importance for the growth process include Foellmi and Zweimueller (2006), Kongsamut et al. (2001), and Matsuyama (2000). The more recent literature on structural change with nonhomothetic preferences includes, among others, Boppart (2014) and Alder et al. (2019) who, like us, propose generalizations of the PIGL preferences class proposed in Muellbauer (1976). Eckert and Peters (2020) is the first paper to incorporate these preferences in a spatial model of structural change. In contrast to us, they focus on the interaction between spatial mobility and the structural transformation. Instead, Matsuyama (2019) and Comin et al. (2020) use a class of generalized CES preferences related to Sato (2014). The authors show these preferences can account accurately for the patterns of structural transformation across several countries. In our paper, we use PIGL preferences because their tractable and transparent aggregation properties are especially suitable. Our results on the unequal gains from service growth are reminiscent to Fajgelbaum and Khandelwal (2016), who measure the unequal gains from trade in a setting with nonhomothetic preferences.

We also contribute to the vast literature on economic development of India including, among others, Aghion et al. (2005, 2008), Akcigit et al. (2021), Basu (2008), Basu and Maertens (2007), Foster and Rosenzweig (1996, 2004), Goldberg et al. (2010), Kochhar et al. (2006), and Martin et al. (2017).

Road Map. The structure of the paper is as follows. Section 2 presents some stylized facts of the role of services in India. Section 3 lays out the theoretical framework. Section 4 describes the data and the main empirical patterns in India. Section 5 discusses the estimation method. Section 6 discusses the main results. Section 7 performs robustness analysis. Section 8 concludes. The Appendix contains technical details, a description of the data sources, and additional tables and figures.


2 Structural Change and Service Growth in India: 1987-2011

Between 1987 and 2011, the Indian economy experienced a remarkable transformation. Not only did income per capita grow by a factor of 3, but the employment structure also changed markedly. In Figure 1, we show the time-series evolution of sectoral employment shares.² Panel (a) is based on the standard ISIC classification. Two facts are apparent: First, agriculture is the largest employment source, accounting for almost half of total employment in 2011. Second, the structural transformation in India is mostly an outflow out of agriculture and an inflow into services and construction. Today, the service sector accounts for one third of aggregate employment. By contrast, employment in the manufacturing sector is stagnant. Panel (b) relies on a model-based classification of activities to which we return in Section 4.

The service sector encompasses a set of heterogeneous activities. In Figure 2, we decompose it into five subsectors:

² The figure is constructed using micro data on employment from the NSS whose description is deferred to Section 4.

PANEL a: STRUCTURAL CHANGE IN INDIA (ISIC CLASS.)



PANEL b: STRUCTURAL CHANGE IN INDIA (OUR SAMPLE)




Figure 1: STRUCTURAL CHANGE IN INDIA: 1987–2011. The figure shows the evolution of sectoral employment shares over time. Panel (a) is based on a standard ISIC classification. Panel (b) is based on the classification we use in our analysis—see Section 4.

(1) Wholesale, Retail, Hotel and Restaurants; (2) Health and Community Services; (3) Financial, Business, and Transport; (4) ICT, and (5) Education and Public Administration (PA). The first and second subsectors, which serve mostly consumers, employed well over half of all Indian service workers in 2011. The third and fourth subsectors sell part of their services to industrial firms. Finance, business, and transport services accounted for about a quarter of service sector employment in 2011. Although the growth rate of employment in the ICT sector was especially fast, this sector accounts for a mere 2.5% of service employment in 2011. Education and PA are mostly government-run activities. The share of the Indian labor force employed in this subsector is constant over time—in contrast to all other subsectors that grew rapidly during the period studied.³ Thus, the expansion of services in India is not confined to business-oriented service industries, such as finance and ICT services. The vast majority of employment gains are found in consumer services such as retail, hospitality, and health.

Figure 2 also highlights important differences across local labor markets in India. We split India into rural and urban districts, broken down so that approximately half of the workers belong to rural districts and urban districts, respectively. Service activities are much more prevalent in urban than in rural areas, especially so in business-oriented activities such as financial services and ICT.

Was the rapid expansion of the service sector shown in Figures 1 and 2 a source or a corollary of Indian growth? In the remainder of the paper, we argue that rising service productivity was indeed an important source.

3 Theory

We consider a general equilibrium environment with R regions. Consumers have preferences over three goods: agricultural goods (F for *food*), industrial goods (G for *goods*), and consumer services (CS). A single factor of production—labor—is inelastically supplied. Whereas goods and food are tradable, CS are not traded and must be provided locally. All markets are frictionless and competitive.

Our theory reflects the three salient aspects highlighted in Figures 1 and 2. First, we allow consumer preferences to be non-homothetic. Hence, the increase in retail employment can be driven both by income effects and by rising productivity in the retail sector. Second, we break down services into CS that—like retail—directly enter consumers’ utility functions and PS that—like corporate legal services—are inputs to the production of goods. Finally, we

³ In absolute terms, employment in the first and second subsectors increased by approximately 32 million in 1987–2011. Employment in the third and fourth subsectors increased by approximately 20 million. Finally, employment in education and PA increased by approximately 7 million—proportionally to the growth of the labor force.




Figure 2: SERVICE INDUSTRIES. The figure displays the employment shares in 1987 and 2011 of five groups of service industries: (1) Wholesale, Retail, Hotel, and Restaurants; (2) Health and Community Services; (3) Financial, Business, and Transport Services; (4) ICT Services; and (5) Education and Public Administration.

assume CS to be non-tradable across space so that some locations (e.g., cities) specialize in services because they are efficient in providing them or because their inhabitants are rich (or both).

3.1 Technology and Preferences

Technology. All goods and services are produced under constant return technologies such that

$$Y_{rst} = A_{rst}H_{rst} \quad \text{for } s = F, CS, G, \quad (1)$$

where H_{rst} denotes the amount of human capital employed in sector s in region r at time t .

This "three sector structure" has important implications for how we take our model to the data on sectoral employment. In particular, total employment in the industrial sector H_{rGt} encapsulates both workers in manufacturing and workers providing PS, as such services are an input into the production of goods. Hence, $H_{rGt} = H_{rMt} + H_{rPSt}$.

That we can write this production structure in the form of (1) is a result, not an assumption. In Section 3.5, we provide a micro-founded model of the industrial sector where firms produce using blue-collar workers and PS and show that the equilibrium gives rise to $Y_{rGt} = A_{rGt}H_{rGt}$, where A_{rGt} is a function of structural parameters that does not depend on equilibrium prices. Therefore, we can characterize the general equilibrium taking A_{rGt} as if it were a primitive, and then solve for the equilibrium in the industrial sector. By contrast, total productivity in agriculture, A_{rFt} , and CS, A_{rCSt} , are fundamental parameters that we take as exogenous.

Preferences. We assume consumers' preferences to be in the PIGL class. PIGL preferences have two important advantages that we rely on extensively. First, they allow us to parameterize the extent of nonhomotheticity in a flexible way, which we can estimate from individual data. Second, they have simple and transparent aggregation properties that allow us to take a district-level demand system to the data.

PIGL preferences do not have an explicit utility representation but are represented by an indirect utility function of the form

$$V(e, \mathbf{p}) = \frac{1}{\varepsilon} \left(\frac{e}{B(\mathbf{p})} \right)^\varepsilon - D(\mathbf{p}), \quad (2)$$

where e denotes total spending and $\mathbf{p}$ the vector of prices. The functions $D(\mathbf{p})$ and $B(\mathbf{p})$ are homogeneous of degree zero and one, respectively.

For the traded commodities, consumers buy a CES aggregate of differentiated regional varieties with an elasticity of substitution σ . Hence, the appropriate price is the respective CES price index $p_s = (\sum_r p_{rs}^{1-\sigma})^{1/(1-\sigma)}$. We parameterize the functions $B(\mathbf{p})$ and $D(\mathbf{p})$ in (2) as

$$B(\mathbf{p}) = \prod_{s \in \{F, G, CS\}} p_s^{\omega_s} \quad \text{and} \quad D(\mathbf{p}) = \left(\sum_{s \in \{F, G, CS\}} \tilde{\nu}_s \ln p_s \right),$$

with $\sum_s \omega_s = 1$ and $\sum_s \tilde{\nu}_s = 0$. This specification yields the indirect utility function

$$V(e, \mathbf{p}) = \frac{1}{\varepsilon} \left(\frac{e}{\prod_s p_s^{\omega_s}} \right)^{\varepsilon} - \sum_s \tilde{\nu}_s \ln p_s. \quad (3)$$

Roy's Identity implies the sectoral expenditure shares associated with $V(e, p)$ (see Section A-1 in the Appendix) are given by

$$\vartheta_s^h(e, \mathbf{p}) = \omega_s + \tilde{\nu}_s \left(\frac{e}{\prod_s p_s^{\omega_s}} \right)^{-\varepsilon}. \quad (4)$$

Equation (4) highlights the presence of both income and price effects. The expenditure share on s is increasing in income (i.e., s is a luxury) if and only if $\tilde{\nu}_s < 0$. For instance, if food is a necessity, $\tilde{\nu}_F > 0$ and its expenditure share converges to ω_F from above. Likewise, if CS is a luxury, $\tilde{\nu}_{CS} < 0$ and ϑ_{CS}^h increases toward ω_{CS} from below. The strength of nonhomotheticities is governed by the parameter ε , which, with a slight abuse of terminology, we label the “income elasticity”. This income elasticity ε turns out to be the crucial parameter in our analysis.

Importantly, because CS are a luxury, the welfare benefits of an increase in the productivity of CS are skewed toward the rich. In fact, the expenditure share $\vartheta_{CS}(e, \mathbf{p}_{rt})$ exactly measures the welfare exposure of a change in prices at the individual level.⁴ Below, we show the heterogeneity in the welfare impact across the income distribution is quantitatively large and makes service-led growth significantly pro rich.

3.2 Heterogeneity, Aggregate Demand, and Utilitarian Welfare

As already mentioned, PIGL preferences have attractive aggregation properties. Suppose individuals have heterogeneous abilities, and let w_{rt} denote the wage per efficiency unit of labor. Income (and expenditure) for individual h is then given by $e_{rt}^h = q^h w_{rt}$, where q^h is the number of efficiency units of labor. Let $F_{rt}(q)$ denote the distribution function of q in region r at the t . Empirically, we will relate the spatial variation in the distribution of q to observable differences in human capital. Using (4), the *aggregate* spending share on goods in sector s in region r is given by

$$\vartheta_{rs}(w_{rt}, \mathbf{p}_{rt}) \equiv \frac{L_{rt} \int \vartheta_s^h(q w_{rt}, p_{rt}) q w_{rt} dF_{rt}(q)}{L_{rt} \int q w_{rt} dF_{rt}(q)} = \omega_s + \nu_{rs} \left(\frac{E_{rt}[q] w_{rt}}{p_F^{\omega_F} p_G^{\omega_G} p_{CSr}^{\omega_{CS}}} \right)^{-\varepsilon}, \quad (5)$$

⁴ Formally, letting $e(\mathbf{p}_{rt}, V)$ denote the expenditure function of achieving a utility level of V given prices $\mathbf{p}_{rt}$, $\partial \ln e(\mathbf{p}_{rt}, V) / \partial \ln p_{rCS} = \vartheta_{CS}(e, \mathbf{p}_{rt})$.

where

$$\nu_{rs} \equiv \frac{E_{rt} [q^{1-\varepsilon}]}{E_{rt} [q]^{1-\varepsilon}} \tilde{\nu}_s. \quad (6)$$

Comparing (5) and (4) highlights in what sense PIGL allows for a representative household: the aggregate demand system in (5) is isomorphic to that of a representative consumer in region r who earns the average income $E_{rt} [q] w_{rt}$, and has the inequality-adjusted preference parameter ν_{rs} in (6) instead of the primitive parameter $\tilde{\nu}_s$. The correct inequality-adjustment to go from the micro preferences $\tilde{\nu}_s$ to the market demand parameters ν_{rs} is given by $E_{rt} [q^{1-\varepsilon}] / E_{rt} [q]^{1-\varepsilon}$ and hence depends, in general, on the local income distribution.

The analysis simplifies if we assume q follows a Pareto distribution with c.d.f. $F_{rt}(q) = 1 - (\underline{q}_{rt}/q)^\zeta$, with a region-invariant tail parameter ζ . Then, $E_r [q] = \frac{\zeta}{\zeta-1} \underline{q}_r$ and $E_r [q^{1-\varepsilon}] = \frac{\frac{\zeta}{1-\varepsilon}}{\frac{\zeta}{1-\varepsilon}-1} \underline{q}_r^{1-\varepsilon}$, so that equation (6) simplifies to $\nu_{rs} = \nu_s \equiv \frac{\zeta^\varepsilon (\zeta-1)^{1-\varepsilon}}{\zeta+\varepsilon-1} \tilde{\nu}_s$. Thus, if income is Pareto distributed with a common tail parameter, all regions have the same “aggregate” parameter ν_s , which is proportional to the primitive individual preference parameter $\tilde{\nu}_s$. In this case, regional demand differences are solely driven by local prices, wages, and human capital $E_{rt} [q] \propto \underline{q}_{rt}$, and we can write $\vartheta_{rs}(w_{rt}, \mathbf{p}_{rt}) = \vartheta_s(\underline{q}_{rt} w_{rt}, \mathbf{p}_{rt})$. Hence, despite heterogeneity in income and non-homothetic preferences, we can express aggregate demand as a function of wages, prices, and the average level of efficiency.

A similar aggregation property also allows us to calculate utilitarian welfare in location r . Given the local skill distribution F_{rt} and a vector of local wages and prices, utilitarian welfare is defined as $\mathcal{U}_{rt}(w_{rt}, \mathbf{p}_{rt}) \equiv \int V(qw_{rt}, p_{rt}) dF_{rt}(q)$. Using the indirect utility function in (3), $\mathcal{U}_{rt}(w_{rt}, \mathbf{p}_{rt})$ has the simple closed-form expression

$$\mathcal{U}_{rt}(w_{rt}, \mathbf{p}_{rt}) = \mathcal{U}(E_{rt}[q] w_{rt}, \mathbf{p}_{rt}) = \frac{\zeta^{1-\varepsilon} (\zeta-1)^\varepsilon}{\zeta-\varepsilon} \times \left(\frac{1}{\varepsilon} \left(\frac{E_{rt}[q] w_{rt}}{p_{Ft}^{\omega_F} p_{Gt}^{\omega_G} p_{rCS}^{\omega_{CS}}} \right)^\varepsilon - \sum_s \nu_s^\mathcal{U} \ln p_{rst} \right), \quad (7)$$

where $\nu_s^\mathcal{U} \equiv \nu_s \times ((\zeta-\varepsilon)(\zeta-(1-\varepsilon)))/(\zeta(\zeta-1))$. Hence, utilitarian welfare is akin to the indirect utility of a representative agent with average income $E_{rt}[q] w_{rt}$ and a scaled taste parameter $\nu_s^\mathcal{U}$ that accounts for the income distribution (ζ) and the income elasticity (ε). Note that given this scaled taste parameter, the distribution F_{rt} only enters through the average income term $E_{rt}[q] w_{rt}$. Hence, we also write $\mathcal{U}_{rt}(w_{rt}, \mathbf{p}_{rt}) = \mathcal{U}(E_{rt}[q] w_{rt}, \mathbf{p}_{rt})$ to stress that welfare differences across time and space depend entirely on wages w_{rt} , prices p_{rt} , and average human capital $E_{rt}[q]$.

3.3 Equilibrium

To characterize the equilibrium, note that the CES demand structure and the fact that the competitive prices are given by $p_{rst} = A_{rst}^{-1} w_{rt}$ implies a set of nationwide market-clearing conditions for the two tradable goods

$$w_{rt} H_{rst} = \left(\frac{w_{rt}^{1-\sigma} A_{rst}^{\sigma-1}}{\sum_{j=1}^R w_{jt}^{1-\sigma} A_{jst}^{\sigma-1}} \right) \times \sum_{j=1}^R \vartheta_s(\underline{q}_{jt} w_{jt}, \mathbf{p}_{rt}) w_{jt} H_{jt} \quad \text{for } s = F, G, \quad (8)$$

and a set of district-specific market-clearing conditions for nontradable CS,

$$w_{rt} H_{rCS} = \vartheta_{CS}(\underline{q}_{rt} w_{rt}, \mathbf{p}_{rt}) w_{rt} H_{rt}. \quad (9)$$

Together with the market-clearing conditions for local labor markets, $H_{rF} + H_{rG} + H_{rCS} = H_r$, these equations fully describe the equilibrium.

Equations (8) and (9) highlight how our framework allows for differences in tradability between PS and CS. CS value added has to be consumed locally. By contrast, PS value added is embodied in goods and hence tradable. Cities like Delhi or Mumbai, for example, which might have a comparative advantage in PS like finance or ICT, are thus able to export such service inputs to the rest of India. While our treatment of the difference in tradability between PS (which is freely tradable) and CS (which is entirely local) is obviously stark, it is qualitatively consistent with Gervais and Jensen (2019), who estimate sector-specific trade costs and conclude that PS are as tradable as tangible goods whereas trade costs in CS activities are substantially higher.

3.4 Service-Led or Service-Biased Growth?

The equilibrium condition in (9) highlights the focal point of our analysis: Is growth service led or service biased? Using the definition of the expenditure share ϑ_{CS} in (5), the equilibrium CS employment share is given by

$$\frac{H_{rCS}}{H_{rt}} = \omega_{CS} + \nu_{CS} p_F^{\varepsilon \omega_F} p_G^{\varepsilon \omega_G} \times \left(\underbrace{E_{rt}[q]}_{\text{Skills}} \times \underbrace{w_{rt}^{1-\omega_{CS}}}_{\text{Wages}} \times \underbrace{A_{rCS}^{\omega_{CS}}}_{\text{Productivity}} \right)^{-\varepsilon}, \quad (10)$$

where we have eliminated p_{rCS} using the competitive equilibrium condition $p_{rCS} = w_{rt}/A_{rCS}$.

Equation (10) highlights how our theory encompasses both the possibility that growth was service-led and that growth was service-biased. Because equilibrium employment depends on the local supply of skills ($E_{rt}[q]$), local wages (w_{rt}), and local productivity (A_{rCS}), locations such as Delhi or Bangalore might have a large retail sector either because consumers living there are, on average, rich or because the local retail sector is highly productive. Similarly, rising educational attainment and factor prices or rising productivity can increase CS employment over time. The former channel describes a *service-biased* growth scenario: the rise of service employment is due to income effects. The latter channel captures a *service-led* growth scenario where rising service employment is driven by rising productivity and the CS sector is a source rather than a corollary of growth.

To solve this identification problem, we leverage both the structure imposed by our theory and additional data. First, the data on earnings, schooling, and an estimate of the returns to schooling allow us to measure local skills and their price. Given an income elasticity ε and the (endogenous) prices of tradable goods p_F and p_G , we can use (10) to identify A_{rCS} . Similar to the traditional approach in development accounting, we use a set of structural parameters to identify productivity in a model-consistent way. However, our inference hinges on solving for a set of equilibrium prices, p_F , p_G , and w_r . For this reason, we label our methodology *equilibrium development accounting*.

3.5 Equilibrium in the Industrial Sector

Thus far, we have taken the technology of the goods-producing sector $Y_{rGt} = A_{rG}(H_{rMt} + H_{rPS})$ as exogenous. In this section we provide a micro-founded model of the industrial sector that generates such an aggregate production function as an equilibrium outcome. We assume each district's industrial good is produced by a continuum of firms using production workers and PS as inputs. PS are provided by a separate service sector comprising corporate law services, accounting, financial advising, and so on.

Environment. We construct a model with heterogeneous firms where the demand of PS is related to firm size. Intuitively, in small firms, activities such as accounting are carried out by the manager, whereas large firms outsource them to professional providers. We formalize this idea by positing a nonhomothetic production function of the following form:

$$y_i = z_i^{1-\alpha-\beta} H_{PMi}^\alpha (A_{rPS} H_{PSi} + \kappa)^\beta. \quad (11)$$

Here, z_i is firm i 's productivity, and H_{PMi} and H_{PSi} denote the inputs of manufacturing production workers and PS, respectively. A_{rPS} denotes the productivity of the PS sector in region r . The parameter $\kappa \geq 0$ governs the nonhomotheticity of firms' technology—if $\kappa > 0$, the input share of PS increases with firms' size. We assume $\alpha + \beta < 1$, which captures a limited span of managerial control. Decreasing returns to scale guarantee firms earn profits in equilibrium.

The mass of active firms is endogenously determined via free entry. To enjoy the opportunity of drawing a realization from the productivity distribution, an entrant firm must pay a sunk labor cost of f_E manufacturing workers. The productivity z_i is drawn from a Pareto distribution $F_{rt}(z) = 1 - (A_{rMt}/z)^\lambda$, where A_{rM} is a lower-bound productivity that parametrizes the state of technology in region r .

Active firms must hire f_O manufacturing workers in order to start production. Some low- z_i firms therefore might opt to remain inactive. In particular, we assume

$$\frac{f_O}{f_E} > \frac{\beta + (1 - \alpha)(\lambda - 1)}{1 - \alpha - \beta}, \quad (12)$$

which ensures some low-productivity firms will choose not to be active. Although inessential, it simplifies the analysis and avoids a taxonomic presentation.

Equilibrium. In this section, we summarize—with the aid of Figure 3—the properties of the equilibrium, whose analytical characterization is deferred to Appendix Section A-2. Under condition (12), there are two productivity thresholds, z^* and $z_L(A_{PS})$, with $A_M < z^* \leq z_L(A_{PS})$. These thresholds partition the set of firms in three groups: low-productivity firms with $z \in [A_M, z^*]$ are inactive; medium-productivity firms with $z \in [z^*, z_L(A_{PS})]$ produce using only production workers;⁵ high-productivity firms with $z > z_L(A_{PS})$ demand both production workers and PS. Furthermore, high-productivity firms also specialize on PS on the intensive margin. In particular, the relative PS employment share of a firm with productivity z is given by

$$\frac{H_{PS}(z)}{H_{PM}(z)} = \frac{\beta}{\alpha} \times \left(1 - \frac{z}{z_L}\right),$$

and hence increasing in z . For very productive firms, the PS share approaches the "Cobb-Douglas" limit β/α because the nonhomotheticity κ ceases to be important. The following proposition summarizes the main properties of the equilibrium.

Proposition 1. *The equilibrium production function of the industrial-goods sector in region r is given by $Y_{rG} =$*

⁵ The medium-productivity range may be empty. In particular, if $A_{PS} > \frac{1-\alpha}{\beta} \frac{\kappa}{f_O}$, all active firms hire PS and $z^* = z_L(A_{PS})$.




Figure 3: EQUILIBRIUM IN THE INDUSTRIAL SECTOR. The figure displays the main qualitative features of the industrial equilibrium. All firms with $z \geq z^*(A_{PS})$ are active. All firms with $z \geq z_L(A_{PS})$ demand PS. The shaded blue area represents the productivity distribution.

$A_{rG}H_{rG}$, where $H_{rG} = H_{rM} + H_{rPS}$ and $A_{rG} = A_G(A_{rM}, A_{rPS})$ is given by

$$A_G(A_M, A_{PS}) := \begin{cases} Q_1 A_M^{(1-\alpha-\beta)} \left(\left[1 + \frac{1}{\lambda-1} \frac{\beta}{1-\alpha} \left(\frac{1-\alpha}{\beta} \varsigma(A_{PS}) \right)^{-\frac{(1-\alpha)(\lambda-1)+\beta}{1-\alpha-\beta}} \right] \right)^{\frac{1-\alpha-\beta}{\lambda}} & \text{if } \varsigma(A_{PS}) > \frac{\beta}{1-\alpha}, \\ Q_2 A_M^{(1-\alpha-\beta)} A_{PS}^\beta (1 - \varsigma(A_{PS}))^{\frac{(1-\lambda)(1-\alpha-\beta)}{\lambda}} & \text{if } \varsigma(A_{PS}) \leq \frac{\beta}{1-\alpha} \end{cases},$$

where $\varsigma(A_{PS}) \equiv \frac{\kappa}{f_0 A_{PS}}$, Q_1 and Q_2 are constants, and $A_G(A_M, A_{PS})$ is continuous and increasing in both arguments.

The employment share of PS is given by

$$\frac{H_{PS}}{H_G} = \begin{cases} \frac{\beta}{\lambda} \frac{\beta + (\lambda-1)(1-\alpha)}{\beta + (1-\alpha)(\lambda-1) \left(\frac{1-\alpha}{\beta} \varsigma(A_{PS}) \right)^{\frac{(\lambda-1)(1-\alpha)+\beta}{1-\alpha-\beta}}} & \text{if } \varsigma(A_{PS}) > \frac{\beta}{1-\alpha} \\ \left(\beta - (1-\alpha-\beta) \frac{\lambda-1}{\lambda} \frac{\varsigma(A_{PS})}{1-\varsigma(A_{PS})} \right) & \text{if } \varsigma(A_{PS}) \leq \frac{\beta}{1-\alpha} \end{cases}. \quad (13)$$

$H_{PS}/H_G \in [0, \beta]$ is continuous and strictly increasing in A_{PS} .

Proof. See Section A-2 in the Appendix. $\square$

Proposition 1 contains three main results. First, it shows the aggregate production function of the goods-producing sectors has indeed the simple form used in our analysis above: it is linear in total employment, that is, PS and manufacturing employment, with a constant productivity A_G . Second, it provides a closed-form expression for A_G that is increasing in both A_M and A_{PS} . Hence, a highly productive PS sector acts as a source of comparative advantage for tradable goods. Note that this comparative static of A_{PS} reflects both the direct effect of PS inputs being more productive and the indirect effect of selection: an increase in A_{PS} increases z^* and decreases z_L , and hence induces a reallocation of resources toward more productive firms. Third, Proposition 1 yields a closed-form expression for the aggregate PS employment share H_{PS}/H_G . Interestingly, the sole determinant of structural change within the industrial sector (i.e., of PS deepening) is the local productivity of the PS sector A_{rPS} . By contrast,

manufacturing productivity A_{rM} does not affect the composition of employment. Free entry is key for this stark result. If manufacturing productivity A_{rM} (or the demand for the regional industrial variety) increases, only the mass of manufacturing firms adjusts. The technology choice of active firms remains unaffected.

In summary, our model has a tractable recursive structure. The trade equilibrium pins down the employment share of the industrial sector as a function of A_{rM} and A_{rPS} , given preferences and technology in the other sectors. The employment breakdown into manufacturing and PS is then determined by A_{rPS} only.

4 Empirical Analysis

We now turn to the empirical analysis. We first describe our main data sources. We then discuss our measurement choices to relate this data to our theory.

4.1 Data

Our analysis relies on four datasets.⁶

1. The NSS Employment-Unemployment Schedule for the years 1987 and 2011, henceforth, the “NSS data.”
2. The Economic Census for the years 1990, and 2013, henceforth, the “EC.”
3. A Special Survey of the Indian Service Sector for the year 2006, henceforth, the “Service Survey.”
4. The NSS Consumer-Expenditure Schedule, henceforth, the “Expenditure data.”

The NSS data, which form the backbone of our analysis, are a household survey with detailed information on employment characteristics and households’ location of residence. We use data for 1987 and 2011. The NSS data allows us to measure sectoral employment shares and average income at the district-year level. Consistent with our theory, we measure employment shares in four sectors: agriculture, manufacturing, PS, and CS. For agriculture and manufacturing, we follow the ISIC sectoral classification in the NSS data. As highlighted in Figure 2, the situation is more complicated in the service industry. Whereas, for example, retail workers are arguably part of the CS sector, the distinction is less clear for lawyers because this category includes both corporate lawyers and divorce lawyers, providing, respectively, PS and CS. To solve this problem, we combine information from the Economic Census and the Service Survey to estimate the extent to which particular subsectors within the service sector provide their value added to firms rather than consumers. We describe this procedure in detail in Section 4.2 below. To measure income, we proxy earnings by average expenditure. We prefer this measure to direct information on wages to also capture informal employment.

The EC is a complete count of all establishments engaged in the production or distribution of goods and services in India. It covers all sectors except crop production and plantation and collects information on each firm’s location, industry, employment, and the nature of ownership. It contains approximately 24 million and 60 million establishments in 1990 and 2013, respectively.⁷ The relatively unexplored Service Survey was conducted in 2006 and is designed to be representative of India’s service sector. It covers almost 200,000 private enterprises subdivided into six service industries. In Appendix Section B-1, we compare it with the EC and document that it is representative

⁶ A more detailed description of these datasets is deferred to Appendix Section B-1. Here, we highlight the main features.

⁷ As shown in earlier studies, see for example Hsieh and Klenow (2014) or Akcigit et al. (2021), most Indian firms are very small, with an average size ranging between two and three employees, over half having a single employee, and only one in 1,000 firms employing more than 100 workers.

of the distribution of firm size in India. We use the EC and the Service Survey to classify service employment into CS and PS.

Finally, we use the NSS Consumer-Expenditure Schedule. This dataset contains detailed information on households’ expenditure allocation across narrowly defined goods, and thus allows us to measure expenditure shares on food and CS. We refer to Section B-1.5 in the Appendix for details on the product classification. We leverage this information to estimate the income elasticity ε , which is the key preference parameter for our analysis.

To compare spatial units over time, we create a time-invariant definition of geography. We define regions as Indian districts. Because the boundaries of several districts changed over time, we harmonized them using GIS software, relying on maps for the years 1987, 1991, 2001, and 2011. We define regions so that they have the same boundaries over time. To keep the number of regions as large as possible, a region is always the smallest area that covers a single district or a set of districts with consistent boundaries over time. There are two small districts that did not exist in 1987 (but do in 2011) and we exclude these from our analysis. We also exclude districts with less than 50 observations as these do not allow us to estimate sectoral employment shares in a credible way. In the end, we obtain 360 regions that cover all of India. For simplicity, we refer to these regions as “districts.” Section B-3 in the Appendix describes in detail how we construct this crosswalk.

4.2 Measurement

Consumer vs. Producer Services. We aim to distinguish between PS and CS in a way that is consistent with our theory. Ideally, we would want to measure employment in PS and CS with the help of detailed input-output matrices so as to associate the value added of each firm to the identity of the buyers (either private individuals or firms). To the best of our knowledge, this information is not available.

We therefore leverage micro data on the firms’ downstream trading partners contained in the Service Survey. Specifically, this data report whether a firm is selling mostly to consumers or to other firms. We could thus, in principle, calculate the share of employment in every service industry-district cell distinguishing between firms selling to other firms and those serving consumers. In practice, this procedure is not feasible, because the Service Survey contains too few firms to precisely estimate these employment shares for each service industry-district cell. Instead, we rely on the fact that the probability of a firm selling to other firms rather than to consumers is highly correlated with firm size—larger firms are more likely to sell to firms. We show this pattern in Table 1, which displays the share of firms that mainly sell to other firms by employment size. A clear pattern emerges: small firms with one or two employees sell almost exclusively to final consumers, whereas a significant share of large firms sell to other firms.

	Firm size: Number of employees								
	1	2	3	4	5	6-10	11-20	21-50	51+
PS share	5.0%	3.8%	6.2%	8.5%	11.5%	12.6%	11.8%	27.6%	42.5%
Number of firms	97337	46571	13227	5156	2777	4841	2830	601	403

Table 1: SHARE OF PRODUCER SERVICES BY FIRM SIZE. The table reports the share of firms selling to firms (rather than private individuals) in different size categories.

We use the pattern reported in Table 1 in the following way. First, we estimate the PS employment share by firm size for different industries within the service sector. We then use the *district*-specific size distribution from the EC

to infer the aggregate PS employment share in district r . More formally, the PS employment share (relative to the total service sector) in subsector k in region r is given by $s_{rk}^{PS} = \sum_b \omega_{kb}^{PS} \ell_{kbr}$, where ω_{kb}^{PS} is the share of employment in firms selling to firms in sector k in size class b , and ℓ_{kbr} is the employment share of firms of size b in sector k in region r . Note this procedure assumes the structure of production for firms of equal size to not vary across Indian districts. The regional variation in PS and CS employment stems from differences in (i) total service employment, (ii) the relative share of different service industries, and (iii) the distribution of firm size. We exclude from the analysis a subset of service industries for which the categorization into PS and CS is ambiguous. These industries include public administration and defense, compulsory social security, education, and extraterritorial organizations and bodies. In Section B-2.2 in the Appendix, we describe this procedure in more detail.

In Figure 4, we display the result of this exercise for different subsectors within the service sector. Within the retail and restaurant sector, only a few establishments cater to other firms. Hence, we estimate that more than 97% of employment in that industry is engaged in the production of CS. The situation is very different in the financial or the ICT sector, where roughly 26% of 53% of employment caters mainly to other firms.




Figure 4: PRODUCER VS CONSUMER SERVICES IN DIFFERENT INDUSTRIES. The figure shows the share of producer and CS in 2011 in different industries within the service sector.

Finally, we merge construction and utilities with the service sector. Although the construction sector is sometimes included in the industrial sector, the key distinction in our theory is that goods are tradable whereas services are nontradable. Because construction and utilities are local goods, we find it natural to merge them with services.⁸ The construction sector serves both consumers (e.g., residential housing) and firms (e.g., business construction). To break these activities into PS and CS, we follow a procedure similar to that used for services. We exploit information from the “Informal Non-Agricultural Enterprises Survey 1999–2000” (INAES) dataset, which also reports whether a firm sells to consumers or other firms and which covers the construction sector. Given the sample size, splitting the destination of construction activities is possible only at the national, not the district level. We obtain the following breakdown. First, we remove 9.1% of the construction activity from the sample, which corresponds to the share of government activity (infrastructure and public goods). Then, based on the INAES data, we attribute 87.1% of what is left to CS and 12.9% to PS in every district-year. For more details, see Section B-2.2 in the Appendix.

⁸ In Section 7, we redo our analysis when we include construction in the manufacturing sector and show our results do not depend on this particular choice.

Given these measurement choices, we are now in the position to quantify the structural transformation in India, across both time and space. Panel (b) of Figure 1 uses the sectoral classification we adopt in our analysis. Relative to Panel (a), we exclude the public sector, merge services with construction and utilities, and break down services into CS and PS, as discussed above. The time-series evolution of agricultural and manufacturing employment is essentially unchanged. Within the service sector, CS grow particularly quickly.⁹

In Figure 5, we turn to the spatial heterogeneity across Indian districts. We focus on urbanization as our measure of spatial heterogeneity. This as a mere descriptive device. In Section OA-1 in the Online Appendix, we show a strong positive correlation between urbanization and the expenditure per capita in the NSS data for 2011. Thus, we take the urbanization rate as a proxy for economic development across Indian districts. Figure 5 displays sectoral employment shares by urbanization quintiles. The average urbanization rates of the five quintiles are, respectively, 6.4%, 12.1%, 19.5%, 29.2%, and 56.4%. Richer, urban locations have lower employment shares in agriculture and specialize in the production of services and industrial goods. Over time, the share of agriculture declines. Between 1987 and 2011 the structural transformation was especially fast in more urbanized districts. In 1987, agriculture was the main sector of activity even in the top quintile of urbanization. By contrast, in 2011, more than half of the working population was employed in CS and PS. This difference is even starker when one looks at earnings instead of employment (see Section OA-1 in the Online Appendix.).




Figure 5: SECTORAL EMPLOYMENT OVER TIME AND SPACE. The figure plots the sectoral employment shares by urbanization quintile in 1987 and 2011.

Human Capital. To be consistent with our theory, we measure each district's endowment of human-capital units $F_{rt}(q)$ and its distribution across sectors in terms of efficiency units of labor. To measure the distribution of human capital across sectors within a district, we rely on the sectoral distribution of earnings, because differences in earnings reflect heterogeneity in the endowment of effective units of labor.¹⁰ To measure the distribution of human capital across districts, we follow the approach in the development accounting literature and leverage data on the regional distribution of schooling. We assume individual human capital q_i as a function of schooling s_i is given by $q_i = \exp(\rho s_i) \times v_i$, where s_i denotes the number of years of education, ρ is the annual return to schooling, and

⁹ Panel (b) of Figure 1 shows the employment share of education and PA remains constant over time at a 5% level. This finding suggests our choice to exclude them is largely inconsequential.

¹⁰ In Section 7.3.2 below, we extend our model to allow for imperfect substitution of skills across sectors.

v_i is an idiosyncratic shock, which we assume to be iid across districts and years and which satisfies $E[v_i] = 1$. To measure schooling attainment s_i , we classify people into four educational groups: (i) less than primary school, (ii) primary and upper primary/middle school, (iii) secondary school, and (iv) more than secondary school. We associate each step in the education ladder with three extra years of education, consistent with the organization of schools in India.

We estimate ρ using Mincerian regressions—see Section 5.1. Given an estimate of ρ , we then calculate the average amount of human capital per region as $E_{rt}[q] = \sum_e \exp(\rho \times e) \ell_r(e)$, where $\ell_r(e)$ denotes the share of people in region r with e years of education. Hence, the distribution of educational attainment across space determines the spatial distribution of human capital. Finally, consistent with our assumption that q follows a Pareto distribution with lower bound $\underline{q}_{rt}$, $\underline{q}_{rt}$ satisfies $E_{rt}[q_i] = \frac{\zeta}{\zeta-1} \underline{q}_{rt}$.

Table 2 shows why it is important to allow for human capital differences across years, sectors, and space. First, the level of schooling increased markedly between 1987 and 2011 and is itself a source of growth. Second, educational attainment differs across sectors. That agriculture is the least skill-intensive industry and educational attainment is highest in PS is not surprising. However, note the CS sector also employs lots of skilled individuals and is more skill intensive than the manufacturing sector.¹¹ Through the lens of our model, these patterns imply that the average number of efficiency units differs across sectors, and by using earnings shares rather than employment shares, our methodology takes such differences into account. Finally, there are large spatial differences whereby city dwellers are much more educated than the rural population. By explicitly measuring the local supply of human capital, we refrain from attributing these differences to differences in local TFP.

	Less than primary	Primary, upper primary, and middle	Secondary	More than secondary
<i>Aggregate Economy (1987 - 2011)</i>				
1987	66.79%	22.03%	7.99%	3.19%
2011	40.32%	30.10%	18.79%	10.79%
<i>By sector (2011)</i>				
Agriculture	53.72%	29.23%	14.45%	2.60%
Manufacturing	32.63%	35.31%	20.68%	11.39%
CS	25.16%	31.99%	27.94%	14.90%
PS	17.38%	26.58%	26.29%	29.74%
<i>By Urbanization (2011)</i>				
Rural	46.97%	30.00%	16.30%	6.84%
Urban	33.69%	30.30%	21.27%	14.73%

Table 2: EDUCATIONAL ATTAINMENT. The table shows the distribution of the educational attainment. Wholesale, Retail, Hotel, Restaurants, Health, and Community Service are classified as CS. Financial, Business, Transport, and ICT Services are classified as PS. The breakdown of rural and urban districts is chosen in a way that approximately half of the population lives in rural and urban districts.

¹¹ For ease of comparison with Figure 2, we classify CS and PS according to the NIC classification, that is, assign wholesale, retail, hotels, restaurants, health, and community services to CS and financial, business, transport and ICT services to PS.

5 Estimation: Identification and Results

With the aforementioned data at hand, we now turn to the estimation of our model. Our approach is in the tradition of development accounting, which has a long history in macro and development economics (see, e.g., Caselli (2005), Hall and Jones (1999), and Gancia et al. (2013)). Whereas these studies infer productivity at the country-level from an aggregate production function, we estimate the entire distribution of productivity $\{A_{rst}\}$ across sectors and space. We do so by relying on the entire equilibrium structure of our model, and hence refer to our method as *equilibrium development accounting*.

The centerpiece of our methodology is the distinction between structural parameters and local productivity. Our model is characterized by 14 structural parameters describing preferences, technologies, and the distribution of skills

$$\Omega = \left\{ \underbrace{\varepsilon, \nu_{CS}, \nu_F, \omega_{CS}, \omega_F, \sigma}_{\text{Preference parameters}}, \underbrace{\lambda, \beta, \alpha, f_O, f_E, \kappa}_{\text{Manufacturing technology}}, \underbrace{\rho, \zeta}_{\text{Human capital}} \right\}.$$

In terms of local productivity, each region is characterized by a 3-tuple of regional productivity levels in agriculture, CS, and the goods-producing industry:

$$\mathbf{A}_t = \{A_{rFt}, A_{rCSt}, A_{rGt}\}.$$

The industry productivity A_{rGt} in turn depends on manufacturing and PS productivity A_{rMt} and A_{rPSt} .

In Section 5.1, we describe how we estimate the structural parameters in Ω . Given the parameters in Ω , a unique mapping exists from the equilibrium skill prices $\{w_{rt}\}$ and sectoral employment allocations $\{H_{rst}\}$ to the underlying productivity fundamentals in $\mathbf{A}_t$. In Section 5.2, we describe this procedure and our estimates of $\mathbf{A}_t$.

5.1 Estimation of Structural Parameters Ω

The Income Elasticity ε .

The crucial parameter in our analysis is the income elasticity ε , which determines how quickly demand shifts away from agricultural goods as incomes rise. To estimate ε , we use the cross-sectional relationship between income and expenditure shares at the household level and estimate ε via indirect inference. In particular, letting $\vartheta_{F,h}$ denote the observed expenditure share of food of household h and e_h total household spending, we estimate the following Engel curve using the data on household expenditure:

$$\ln \vartheta_{F,h} = \delta_r + \beta \times \ln e_h + x'_h \gamma + u_h, \quad (14)$$

where δ_r is a region fixed effect and x_h contains household characteristics that could induce a correlation between total spending $\ln e_h$ and food shares. We then estimate ε via indirect inference; that is, we estimate (14) in our model using the estimated coefficient $\hat{\beta}$ as a moment.¹²

¹² To estimate (14) in our model, we randomly draw a sample of 1 million individuals from the region-specific income distribution $F_{rt}(q)$, calculate the model-implied food shares, and then run a regression of log food share against log income and region fixed effects. We draw our sample in a way to replicate the relative size of each district; that is, the share of observations from district r is the same as observed in the data.

	ln food expenditure share			
ln expenditure	-0.332*** (0.008)	-0.318*** (0.007)	-0.331*** (0.008)	-0.366*** (0.010)
Winsorized (2%)		✓		
Addtl. Controls			✓	
IV				✓
F-stat				2093.77
N	101,654	101,654	101,601	94,436
R^2	0.476	0.462	0.486	0.478

Table 3: INCOME ELASTICITY FOR FOOD. The table shows the estimated coefficient β of the regression (14). In all specifications, we control for region fixed effects, an urban/rural dummy, a full set of fixed effects for household size, and the number of workers within the household. In column 2, we topcode household expenditure at the 98% quantile. In column 3, we control for the type of the household, the religion, the social class and whether the household receives rationing cards. In column 4, we instrument household expenditure with a set of three-digit occupation fixed effects. Standard errors clustered at the district level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Although β is not an explicit structural parameter in our theory, it is tightly connection to structural income elasticity ε . According to our theory, the log spending share on food is given by

$$\ln \vartheta_F^h(e, \mathbf{p}_r) = \ln \left(\omega_F + \nu_F^h \left(\frac{e}{p_F^{\omega_F} \times p_G^{\omega_G} \times p_{CSr}^{\omega_{CS}}} \right)^{-\varepsilon} \right).$$

Hence, if $\omega_F \approx 0$ (which is the case in our structural estimation), our theory implies

$$\ln \vartheta_F^h(e, \mathbf{p}_r) = \ln(p_F^{\omega_F} p_G^{\omega_G} p_{CSr}^{\omega_{CS}})^{\varepsilon} + \ln \nu_F^h - \varepsilon \times \ln e;$$

that is, the estimated income elasticity β directly coincides with the structural parameter ε . Note that the region-specific price of CS, p_{rCSr} , is absorbed in the district fixed effect δ_r in (14) and that the additional household level controls x_h aim to capture variation in preferences (ν_F^h) which we abstract from in our theory.

Table 3 reports the results of estimating (14). The first column contains our baseline specification, where we control for a region fixed effect, a dummy for whether the household lives in an urban or rural area (within districts), a full set of fixed effects for household size and the number of workers within the household. We cluster standard errors at the region level to account for the correlation of spending shares through regional prices. We estimate an empirical elasticity of 0.332 that is precisely estimated.

In the remaining columns of Table 3, we report additional specifications to show the robustness of this estimate. In column 2, we winsorize the expenditure variable $\ln e_h$ at the top 2% level to limit the importance of measurement error. In column 3, we control for additional household-level controls. In particular, we control non-parametrically for differences in the household type, that is, whether the household is self-employed (in agriculture or non-agriculture), a regular wage earner or a casual laborer (in agriculture or non-agriculture), the household's religion, and social group and whether the household is eligible to purchase subsidised food grain from the Indian government. In both specifications, the estimate is indistinguishable from our baseline estimate.

In column 4, we present the results from an IV specification. We do so to address concerns about measurement error in $\ln e_h$ that would make our estimate for β downward biased. We instrument total expenditure with a full

set of three-digit occupation fixed effects.¹³ Expectedly, these fixed effects strongly predict total expenditure as evidenced by the large F-statistic. The resulting estimate is slightly larger but quantitatively very similar to the OLS specifications reported in columns 1 - 3. In Section B-5 in the Appendix, we report additional specifications and also show the constant elasticity between expenditure and the expenditure share on food is a good approximation for a large part of the expenditure distribution.

For our baseline results we take an empirical elasticity of -0.33 as our target moment. In Section 7, we show that our results are robust to other choices for ε in line with the results reported in Table 3.

Value Added vs. Final Expenditure To link the structural parameter in our theory (ε) to these empirical estimates, note that our theory specifies consumers' preferences over sectoral value added, whereas the data on household expenditure pertain to final expenditure. As stressed in Herrendorf et al. (2013), the structural parameters of the demand system based on final goods differ from the one based on value added. Although true in general, this is fortunately less of a concern for the elasticity parameter ε , the only parameter we estimate from the data on expenditure shares.

To see this formally, suppose the final-good demand system takes the form

$$\vartheta_s^{FE}(e, q) = \omega_s^{FE} + \tilde{\nu}_s^{FE} \left(\frac{e}{\prod_j q_j^{\omega_j^{FE}}} \right)^{-\varepsilon^{FE}},$$

where q_j denotes the price of the final good of sector j , and we use the superscript “FE” to highlight that the respective structural parameters correspond to the final-expenditure demand system. Furthermore, suppose a unit of the final good in sector s is produced with a Cobb-Douglas production function of sectoral value added:

$$y_s^{FE} = \prod_{j \in (A, CS, G)} (y_{js}^{VA})^{\lambda_{js}}. \quad (15)$$

Hence, the matrix $[\lambda_{js}]_{js}$ describes the input-output matrix of the economy.

As we show in Section B-5 in the Appendix, these assumptions imply that the value-added share of sector s is

$$\vartheta_s^{VA} = \omega_s^{VA} + \tilde{\nu}_s^{VA} \times \left(\frac{e}{\prod_j (p_j^{VA})^{\omega_j^{VA}}} \right)^{-\varepsilon^{FE}},$$

where $\omega_s^{VA} = \sum_j \lambda_{js} \omega_j^{FE}$ and $\tilde{\nu}_s^{VA} = \sum_j \lambda_{js} \tilde{\nu}_j^{FE}$. Hence, the implied value-added demand system is *exactly* consistent with our theory. It is of the PIGL form and the final-good expenditure elasticity ε^{FE} coincides with the value-added expenditure elasticity ε . This is different for the other parameters ω_j^{VA} and $\tilde{\nu}_j^{VA}$, both of which are “input-output” weighted averages of the final-good structural parameters ω_j^{FE} and $\tilde{\nu}_j^{FE}$. Hence, as stressed by Herrendorf et al. (2013), in general, knowledge of the entire input-output structure $[\lambda_{js}]_{js}$ is required to identify the parameters of the value-added demand system. This, however, is not the case for the expenditure elasticity ε , which is the only parameter we estimate from the data on final-good expenditure.

In Section B-5 in the Appendix, we show that this intuition is more general than for the particular case of Cobb-Douglas aggregation embedded in (15) and extend this argument for a general CES production function.

¹³ The survey assigns the occupation of the household member with the highest earnings to the entire household.

Parameter	Target	Value
<i>Preference parameters</i>		
ε	Engel curve	0.34
ω_F	Agricultural spending share US	0.01
ω_{CS}	Agricultural Employment share 2011	0.69
ν_F	Agricultural Employment share 1987	1.28
ν_{CS}	Normalization	-1
σ	Set exogenously	3
<i>Production function parameters</i>		
λ	Tail of the employment distribution	1.42
β	Employment share of lawyers in the US	0.7
α	Profit share	0.158
f_O	Normalization	1
f_E	Normalization	0.11
κ	Normalization	1
<i>Skill parameters</i>		
ρ	Mincerian schooling returns	0.056
ζ	Earnings distribution within regions	3

Table 4: STRUCTURAL PARAMETERS. The table summarizes the estimated structural parameters. The details of the estimation are discussed in the text.

Other Preference Parameters $\nu_{CS}, \nu_F, \omega_{CS}, \omega_F$ and σ .

The market-level demand system depends on the aggregate preference parameters ν_{CS} and ν_F , which are in turn related to the primitive micro-level preference parameters $\tilde{\nu}_{CS}$ and $\tilde{\nu}_F$ —cf. equation (6). We estimate ν_s directly from the data and infer the structural micro parameters ν_s^h given an estimate of the inequality parameter ζ . Identifying ν_s^h separately from ν_s is only required to quantify the welfare consequences of service-led growth, not to estimate the model.

Equation (5) shows the taste shifters ν_{CS} and ν_F determine sectoral spending and employment when holding income and prices constant. In Section A-3 in the Appendix, we prove that ν_{CS} is not separately identified from the productivity in CS, A_{rCS} . Hence, without loss of generality, we can normalize it to -1. The taste shifter for agricultural products, ν_F , can then be directly identified from the aggregate agricultural employment share in a given year. We opt to match it in the year 1987. Doing so yields $\nu_F = 1.276$. Given the normalization of $\nu_{CS} = -1$, $\nu_M = -(\nu_F + \nu_{CS}) = -0.276$. Hence, manufacturing products are also luxury goods, because their expenditure share is increasing in income. However, their income elasticity is below the one for CS.

To identify the share parameters ω_{CS} and ω_F , recall that $\vartheta_F^h(e, p) > \lim_{e \rightarrow \infty} \vartheta_F^h(e, p) = \omega_F$ and that $\vartheta_{CS}^h(e, p) < \lim_{e \rightarrow \infty} \vartheta_{CS}^h(e, p) = \omega_{CS}$. Hence, the expenditure share on food (CS) approaches ω_F (ω_{CS}) from above (below) as income increases. In the US, which we take an example of a rich economy, where nonhomothetic demand is less important, the agricultural employment share is about 1%. Hence, we take $\omega_F = 0.01$. For ω_{CS} , we follow a similar strategy as for ν_F and match the aggregate agricultural employment share in a given year. Given our interest in the long-run growth experience of India, we opt to match agricultural employment in 2011, which implies $\omega_{CS} = 0.69$.¹⁴ Finally, we set the inter-regional trade elasticity σ to a consensus estimate in the literature and assume $\sigma = 3$.

¹⁴ Our model implies regional *employment* shares in CS are bounded by ω_{CS} from above. As we discuss in more detail in Section B-2.4 in the Appendix, our Indian data contains seven districts that feature employment shares in CS that exceed ω_{CS} . Because these districts are very small and account for less than 1% of employment, we drop them from our analysis.

Skill Parameters ζ and ρ .

Our specification of skills $q_i = \exp(\rho s_i)v_i$ implies log earnings of individual i in region r at time t , y_{irt} are given by the usual Mincerian regression $\ln y_{irt} = \ln w_{rt} + \rho s_i + \ln v_i$. Hence, we can estimate ρ from the within-region variation between earnings and education, which we can measure from the NSS data. We estimate an average annual rate of return of 5.6%. Although this estimate is on the lower end of standard Mincerian regressions, recall that we are using data on consumption rather than income. In Section 7, we discuss the robustness of our results with respect to the Mincerian estimate.

We also estimate the tail parameter of the skill distribution ζ . This parameter does not affect the equilibrium conditions given that we estimate the aggregate preference parameter ν_s directly. Hence, our estimate of regional productivity does not depend on the value of zeta. An estimate of ζ is only required once we want to calculate welfare. To estimate ζ , recall that the distribution of income in region r is given by $G_r(y) = 1 - \left(\frac{q_r w_r}{y}\right)^\zeta$, implying $\ln(1 - G_r(y)) = \zeta \ln\left(\frac{q_r w_r}{y}\right) - \zeta \ln y$. We therefore estimate ζ from a cross-sectional regression $\ln(1 - G_r(y_i)) = \delta_r + \beta \ln y_i + u_{ir}$, where δ_r is a district fixed effect. In practice, we consider a support of regional incomes above the median, because the Pareto distribution is a better fit to the right tail of the income distribution. This procedure yields an estimate of $\zeta \approx 2$ (see Appendix Section B-7).

Technology parameters: $\lambda, \beta, \alpha, f_O, f_E$, and κ .

To decompose A_{rG} into the A_{rM} and A_{rPS} , we need to know the underlying parameters of the industrial sector. Proposition 1 establishes that all allocations only depend on $A_{rPS} \times f_O/\kappa$. Hence, the parameters κ and f_O are not separately identified from A_{rPS} , and we normalize $f_O = \kappa = 1$. This normalization entails no loss of generality, because the scale factor $f_O\kappa$ has no bearing on any of our results. Similarly, the entry cost f_E is not separately identified from the level of productivity A_{rMt} as long as some firms are “discarded” after their efficiency draw z is observed, that is, condition (12) is satisfied. We therefore chose f_E to satisfy (12). This leaves us with three parameters: λ, β , and α .

We identify the tail of the productivity distribution λ from the employment distribution of the EC. As we show in detail in Section B-8 in the Appendix, our model implies that, as for the distribution of productivity, the distribution of employment for large firms is also Pareto with shape λ . We find an estimate of $\lambda = 1.42$, which is very precisely estimated.

We then pick α and β to jointly match a profit share of 10% and the long-run share of PS workers within the industrial sector. Our model implies $\lim_{A_{PS} \rightarrow \infty} H_{PS}/H_G = \beta$. In the US, PS account for about 28% of employment and production workers for 12%. This observation suggests $\beta = \frac{0.28}{0.28+0.12} = 0.7$. Given λ and β , the parameter α is tied to the profit share because $\alpha + \beta$ determines the returns to scale and hence the share accruing to the fixed factor. In particular, our model implies the profit share equals $(1 - \alpha - \beta)/\lambda$. For $\beta = 0.7$ and $\lambda = 1.42$, a profit rate of 10% requires that $\alpha = 0.158$.

The Demand for Consumer Services: Direct Evidence

In Table 3, we used data on food shares to estimate the income elasticity ε . In principle, we could use data on the expenditure share of CS. We choose food expenditures for two reasons. First, food expenditures might be better measured if these expenditures are more salient. Second, as argued above, the log-linear specification in (14) is particularly informative about ε if ω_s is small, because our theory then exactly implies a log-linear relationship,

	ln consumer service expenditure share			
ln expenditure	0.268*** (0.022)	0.270*** (0.021)	0.214*** (0.021)	0.571*** (0.031)
Winsorized (2%)		✓		
Addtl. Controls			✓	
IV				✓
F-stat				722.50
N	100,383	100,383	100,334	93,571
R^2	0.240	0.240	0.250	0.213

Table 5: INCOME ELASTICITY FOR CONSUMER SERVICES. The table shows the estimated income elasticity for the CS expenditure share (see (14)). In all specifications, we control for region fixed effects, an urban/rural dummy, a full set of fixed effects for household size and the number of workers within the household. In column 2, we topcode household expenditure at the 98% quantile. In column 3, we control for the type of household, the religion, the social class, and whether the household receives rationing cards. In column 4, we instrument household expenditure with a set of three-digit occupation fixed effects. Standard errors clustered at the district level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

and the distinction between final expenditure and value added becomes less important. Because CS are a luxury, $\omega_F \approx 0$ but $\omega_{CS} > 0$.

Reassuringly, our estimated model yields a nontargeted expenditure elasticity for CS that is broadly in line with the empirical evidence. Specifically, we run—in the model and in the data—the same specification as in (14) except that we use households’ expenditure share on CS, $\ln \vartheta_{CS,h}$, as the dependent variable. We follow the official classification of the NSS expenditure module to assign expenditures to CS. These expenditures include, for example, domestic servants, barber shops, or tailor services. We also add entertainment expenses such as movie theaters, club fees, and cable TV. Again, we refer to Section B-1.5 in the Appendix for the exact list of expenses.

In Table 5, we report the results. The structure of the table is exactly the same as in Table 3; that is, in column 2, we winsorize the data; in column 3, we control for additional household characteristics; and in column 4, we report the IV specification, where we instrument total household expenditure with three-digit occupation dummies. Again, we cluster standard errors at the district level. Empirically, CS are indeed luxuries, because their spending shares are increasing in income. Quantitatively, this elasticity is estimated to be between 0.25 and 0.3 for the OLS specification and around 0.55 in the IV case. When we estimate this specification in our model, we estimate a coefficient of about 0.4. Hence, even though we do not use the data on CS spending to estimate the model, the implied elasticities are consistent with what we see in the household data.

Our model also implies the price of CS varies across space and reflects differences in wages w_r and CS productivity A_{rCS} . In particular, conditional on total expenditure $\ln e_h$, CS shares are predicted to be large in regions where prices are low, that is, where A_{rCS} is large relative to the local wage. In the context of the regression in Table 5, this variation is captured by the regional fixed effects.

In Figure 6, we plot the correlation between our estimates of regional fixed effects and the regional urbanization rate. To visualize the relative size of districts, the size of the markers reflects the size of the population. Figure 6 shows a robust positive relationship: cities are particularly productive in CS (relative to the prevailing wage). As we show below, we find the same qualitative patterns in our structural analysis even though we do not use the information from the expenditure data.

Our estimated demand system also delivers estimates of the elasticity of substitution that are quantitatively broadly consistent with findings in the literature. For the class of PIGL preferences, the elasticity of substitution is




Figure 6: URBANIZATION AND REGIONAL FIXED EFFECTS OF CONSUMER SERVICE SPENDING. The figure displays the correlation with the region fixed effects stemming from Table 5 (column 1) against the urbanization rate.

not a structural parameter but depends on relative prices and total expenditure.¹⁵ As we show in detail in Section B-6 in the Appendix, we find services and goods are complements, with an average elasticity of substitution of 0.7. By contrast, food and CS are, on average, substitutes with a substitution elasticity of around 1.3.

5.2 Estimation of Productivity Fundamentals $\mathbf{A}_t$

Given the structural parameter vector Ω , data on local wages and sectoral employment allocations as well as time-series data on relative prices and aggregate income, the equilibrium conditions uniquely identify a set of local productivity fundamentals $\mathbf{A}_t$. We refer to Section A-3 in the Appendix for details, but we describe the main intuition here. Consider first the identification of $A_{rCS,t}$, which we discussed in Section 3.4. Equation (10) implies we can uniquely solve for $A_{rCS,t}$ as

$$A_{rCS,t} = \left(\frac{(-\nu_{CS})}{\omega_{CS} - \frac{H_{rCS,t}}{H_{r,t}}} \right)^{\frac{1}{\omega_{CS}} \frac{1}{\varepsilon}} p_F^{\frac{\omega_F}{\omega_{CS}}} p_G^{\frac{\omega_G}{\omega_{CS}}} (E_{rt}[q] \times w_{rt}^{1-\omega_{CS}})^{-\frac{1}{\omega_{CS}}}. \quad (16)$$

Controlling for the level of human capital $E_{rt}[q]$ and the equilibrium factor price w_{rt} , CS productivity is increasing in the observed employment share $\frac{H_{rCS,t}}{H_{r,t}}$.¹⁶ Conversely, holding the employment share $\frac{H_{rCS,t}}{H_{r,t}}$ constant, CS productivity $A_{rCS,t}$ is decreasing in both human capital and factor prices. Structurally decomposing the observed variation in employment shares into the part that is service led (i.e., $A_{rCS,t}$) versus the part that is service biased because of income effects (i.e. $E_{rt}[q] w_{rt}^{1-\omega_{CS}}$) is a key aspect of our equilibrium accounting methodology.

The procedure to estimate productivity in tradable sectors is different. Equation (8) implies relative productivity across two locations is given by

$$\frac{A_{rs}}{A_{js}} = \left(\frac{H_{rs}}{H_{js}} \right)^{\frac{1}{\sigma-1}} \left(\frac{w_r}{w_j} \right)^{\frac{\sigma}{\sigma-1}} \quad \text{for } s = F, G. \quad (17)$$

¹⁵ The Allen Uzawa elasticity of substitution between goods s and k is given by $EOS_{sk} = 1 - \varepsilon \frac{(\vartheta_s - \omega_s)(\vartheta_k - \omega_k)}{\vartheta_s \vartheta_k}$. See Section A-4 in the Appendix.

¹⁶ Recall that if CS are a luxury, $\nu_{CS} < 0$ and $\frac{H_{rCS,t}}{H_{r,t}} < \omega_{CS}$.

Hence, sectoral productivity differences can be inferred from relative skill prices and relative factor inputs (in units of human capital) given the elasticity of substitution σ . No other preference parameters are involved in this estimation, because food and industrial goods are tradable so that local demand is dissociated from local income.

Although we can use (17) to estimate relative sectoral productivity, we still need additional restrictions to estimate the level of productivity in agriculture and industry. As we show in the Appendix, we can exploit the time-series data on the relative price of food (relative to goods) and on aggregate GDP.¹⁷

To finally decompose our estimates of productivity in the goods-producing sector A_{rGt} into the manufacturing (A_{rM}) and PS (A_{rPS}) component, we rely on the regional employment share in PS. Specifically, equation (13) in Proposition 1 shows that—given α, β , and λ —we can infer $\varsigma(A_{rPS})$ (and hence A_{rPS}) from the observed relative employment share of PS relative to manufacturing.¹⁸

This discussion underscores the sense in which our methodology is an accounting procedure: for given parameters, we estimate sectoral productivity that exactly rationalizes the observed data on wages and sectoral factor inputs as equilibrium outcomes. Our identification strategy leverages the recursive structure of Proposition 1: we can identify productivity in agriculture, CS, and the industrial sector independently of the particular microstructure of the industrial sector. Such structure is only required to decompose the overall productivity of the industrial sector into the part stemming from PS and the part stemming from manufacturing workers.




Figure 7: ESTIMATED SECTORAL PRODUCTIVITIES. The figure shows a bin scatter plot of the estimated sectoral labor productivities in agriculture, CS, and industry across urbanization-rate bins. Each plot is constructed by pooling the estimates for 1987 and 2011 after absorbing year effects.

In Figure 7, we summarize the main cross-sectional pattern of our productivity estimates by displaying a bin scatter plot of the (logarithm of the) estimated labor productivities in the agricultural, industrial, and CS sector as functions of the urbanization rate. The relationship between productivity and urbanization is increasing for CS (Panel (b)) and in the industrial sector (Panel (c)). For agriculture, the relationship is relatively flat and slightly hump shaped. The declining portion corresponding to districts with an urbanization rate above 50% likely reflects the scarcity of land (a factor of production from which we abstract) in urban areas.

¹⁷ We measure GDP in terms of the numeraire industrial good. Because of nonhomothetic preferences, we cannot define a standard consumption price index. For comparison, we calculated wage growth for a fictitious agent endowed with the median wage and living in a district in which the supply of CS is at the median level. Based on the consumption basket of such an individual in 1987 and 2011, we calculated real wage growth using a Laspeyres and a Paasche index. The resulting real wage growth in the two cases is 1.82 and 4.86, respectively. Our calibration yields a wage growth factor of 2.60, which is in between.

¹⁸ Our identification relies on the nonhomothetic factor demand functions. If $\kappa = 0$, equation (13) implies $\varsigma(A_{PS}) = 0$ irrespective of A_{PS} and that the PS employment share would be constant and equal to β . Intuitively, with a Cobb-Douglas technology, the relative employment and expenditure shares would be independent of productivity. Moreover, A_M and A_{PS} would both be factor neutral, and the aggregate TFP A_G would only depend on $A_M^{1-\alpha-\beta} A_{PS}^\beta$. Hence, A_M and A_{PS} could not be independently identified.

Remarkably, the productivity dispersion and its correlation with urbanization is strongest in the CS sector. Hence, the large employment share of CS in urbanized districts is not only a consequence of high wages (the Baumol effect) or of an abundance of human capital, but also of high CS productivity relative to rural areas. Qualitatively, this pattern is consistent with the estimated fixed effects in Figure 6 stemming from the CS-expenditure data: conditional on income, individuals in cities spend a larger share of their resources on CS.

Among the tradable goods, productivity is significantly more dispersed in the industrial than in the agricultural sector. To understand why, note a district’s relative productivity is identified by its sectoral earning share relative to its skill price (see (17)). The “compressed” productivity distribution in agriculture reflects the observation that wages are negatively correlated with the employment share of agriculture across districts. By contrast, wages are positively correlated with the employment share of industry, implying a wider productivity dispersion.

6 The Importance of Service Led Growth

We now turn to our main question of interest: Was Indian growth service led? If so, did productivity growth in the provision of CS play a quantitatively important role for rising living standards and the structural transformation?

To answer these questions, we first use our sectoral productivity estimates A_{rst} and calculate sectoral productivity growth between 1987 and 2011 for each district. We summarize these distributions of annualized productivity growth in Table 6. In the first row, we focus on CS-productivity growth; in the remaining rows, we report the distributions of growth rates in the tradable sectors for comparison. Two salient facts emerge. First and foremost, productivity in the CS sector grew in the majority of districts. Hence, the rise in CS employment was not merely driven by changes in demand driven by rising incomes or human-capital accumulation. Second, productivity growth was unequal across space and particularly so in the CS sector.¹⁹

	Sectoral productivity growth					
	10th	25th	50th	75th	90th	Aggregate
<i>Service-led growth</i>						
Consumer Services (g_{rCS})	-1.3	0.8	3.5	8.1	14.0	5.0
<i>Growth in other sectors</i>						
Agriculture (g_{rF})	0.02	1.1	2.1	3.1	4.0	2.2
Industry (g_{rG})	0.7	2.1	3.3	4.6	5.9	3.5

Table 6: REGIONAL DISTRIBUTION OF SECTORAL PRODUCTIVITY GROWTH. The table reports different moments of the distribution of growth rates in the different industries between 1987 and 2011. These growth rates are annualized and calculated as $g_{rs} = \frac{1}{2011-1987} (\ln A_{rs2011} - \ln A_{rs1987})$. Columns 1 - 5 report different quantiles. The “Aggregate” column reports the population-weighted (2011) average. All distributions are truncated at the top and bottom 3%.

To quantify the macroeconomic impact of these growth estimates reported in Table 6, we compute counterfactual equilibria where we set the respective sector’s productivity growth to zero in all districts. The resulting changes in wages and employment allocations thus reflect the effect of sectoral productivity growth holding constant productivity growth in all other sectors. Our model then allows us to compute the welfare effects for consumers and how these effects vary across space and the income-distribution ladder. As we shall see in Section 6.1, our analysis uncovers a great deal of heterogeneity in both dimensions. In addition, we can also compute the implications for

¹⁹ To account for measurement error, we winsorize the top and bottom 3% of the estimated productivity distributions. The details are discussed in the Appendix, where we also report robustness results to these choices (see Section B-9).

the structural transformation, and we do so in Section 6.2.

These implications of service-led growth stem from our estimated model and are therefore associated with sampling uncertainty. Intuitively, because the underlying micro data is a sample of individuals, the measured sectoral employment shares in each district are random variables. And given the accounting nature of our analysis, our estimates of productivity fundamentals $\mathbf{A}_t$ (and, in turn, the counterfactual exercises of shutting down sectoral productivity growth) inherit this sampling uncertainty.

To quantify the extent of this uncertainty, we estimate the *distribution* of both the welfare effects and the sectoral reallocation of employment using a non-parametric bootstrap procedure (Horowitz, 2019). While we refer to Section B-11 in the Appendix for the details of the implementation, the idea is conceptually simple. The non-parametric bootstrap treats the empirical distribution of the data as if it were the indeed the underlying population distribution. We can then construct a bootstrap sample with the same sample size from the Indian micro data by drawing households with replacement and redo our analysis. When we repeat this step B times, we can calculate each statistic of interest B times and hence estimate the entire distribution. Using this procedure we can calculate confidence intervals for all the outcomes we report. In practice, we take $B = 200$.

6.1 The Welfare Implications of Service Led Growth

To measure changes in welfare, we calculate equivalent variations relative to the *status quo* in 2011. We focus on three layers of heterogeneity: (i) across individuals differentiated by income, (ii) across districts differentiated by their rate of urbanization, and (iii) for the aggregate Indian economy.

As already discussed in Section 3.2 above, the PIGL demand system allows us to capture such heterogeneous welfare impacts in a tractable way. More specifically, suppose we want to compare the two vectors of wages and prices $\{w_r, \mathbf{p}_r\}_r$ and $\{w_r^{CF}, \mathbf{p}_r^{CF}\}_r$, where *CF* stands for *counterfactual*. Let $\bar{w}^h(w_r^{CF}, \mathbf{p}_r^{CF} | \mathbf{p}_r, q)$ be the income individual h with skill level q facing prices $\mathbf{p}_r$ requires to achieve the same level of utility as under $\{w_r^{CF}, \mathbf{p}_r^{CF}\}_r$. Using the indirect utility function V given in (3), $\bar{w}^h$ is implicitly defined by

$$V(\bar{w}^h(w_r^{CF}, \mathbf{p}_r^{CF} | \mathbf{p}_r, q), \mathbf{p}_r) = V(qw_r^{CF}, \mathbf{p}_r^{CF}). \quad (18)$$

Thus, we can directly calculate the welfare-equivalent income $\bar{w}^h$ from $\{w_r^{CF}, \mathbf{p}_r^{CF}, \mathbf{p}_r\}_r$ for each level of human capital q . In a similar vein, we can calculate the utilitarian welfare consequences at the district level. Exploiting the aggregation result from Section 3.2, the appropriate representative agent in district r facing prices p_r would require a level of *regional* spending power $\bar{w}(w_r^{CF}, \mathbf{p}_r^{CF} | \mathbf{p}_r, E_r[q])$ defined by

$$\mathcal{U}(\bar{w}(w_r^{CF}, \mathbf{p}_r^{CF} | \mathbf{p}_r, E_r[q]), \mathbf{p}_r) = \mathcal{U}(E_r[q]w_r^{CF}, \mathbf{p}_r^{CF}), \quad (19)$$

where $\mathcal{U}$ is defined in (7).²⁰

Given $\bar{w}^h$ and $\bar{w}$ defined in (18) and (19), we calculate the respective welfare changes relative to 2011 as

²⁰ Using equations (18) and (3) and (19) and (7) we get that

$$\bar{w}^h(w_r^{CF}, \mathbf{p}_r^{CF} | \mathbf{p}_r, q) = \left(\left(\frac{qw_r^{CF}}{\prod (p_{rs}^{CF})^{\omega_s}} \right)^{\varepsilon} - \left(\prod p_{rs}^{\omega_s} \right)^{\varepsilon} \left(\sum \tilde{\nu}_s \ln \frac{p_{rs}^{CF}}{p_{rs}} \right) \right)^{1/\varepsilon}.$$

The expression for the aggregate variation $\bar{w}(w_r^{CF}, \mathbf{p}_r^{CF} | \mathbf{p}_r, E_r[q])$ differs only in two ways: it uses $E_r[q]$ instead of q and is evaluated using the scaled preference parameter $\nu_s^{\mathcal{U}}$ in lieu of the primitive parameter $\tilde{\nu}_s$.

$$\Delta \mathcal{W}_r^h(q) \equiv \frac{\bar{\omega}^h(w_r^{CF}, \mathbf{p}_r^{CF} | \mathbf{p}_r, q)}{qw_{r,2011}} - 1 \quad \text{and} \quad \Delta \mathcal{W}_r \equiv \frac{\bar{\omega}(w_r^{CF}, \mathbf{p}_r^{CF} | \mathbf{p}_r, E_r[q])}{E_r[q]w_{r,2011}} - 1. \quad (20)$$

Hence, $\Delta \mathcal{W}_r^h(q)$ is the change in income that an individual with human capital q living in district r in 2011 would require to achieve the same level of utility in the counterfactual allocation. If, for example, $\Delta \mathcal{W}_r^h(q) = -20\%$, the consumer would be indifferent between giving up 20% of her income in 2011 and an allocation in which productivity in a particular sector did not grow between 1987 and 2011. Similarly, $\Delta \mathcal{W}_r$ is the change in regional income required to achieve the same utilitarian welfare in district r .

Because these expressions vary across individuals and locations, they allow us to quantify the unequal effects of sectoral productivity growth. To gauge the welfare consequences at the aggregate level, we also report the change in aggregate welfare $\Delta \mathcal{W}$, which we compute as

$$\Delta \mathcal{W} = \frac{\sum_r \bar{\omega}(w_r^{CF}, \mathbf{p}_r^{CF} | \mathbf{p}_r, E_r[q]) L_{r,2011}}{\sum_r E_r[q] w_{r,2011} L_{r,2011}} - 1 = \sum_r \Delta \mathcal{W}_r \frac{E_r[q] w_{r,2011} L_{r,2011}}{\sum_r E_r[q] w_{r,2011} L_{r,2011}}.$$

Results: Sources of Welfare Growth in India (1987 - 2011)

We now set to zero—for each sector—the distribution of productivity growth since 1987 shown in Table 6, recalculate the equilibrium, and then use the expressions above to calculate the implied welfare changes at the aggregate ($\Delta \mathcal{W}$), regional ($\Delta \mathcal{W}_r$), and individual ($\Delta \mathcal{W}_r^h(q)$) level.

Aggregate Effects. We first discuss the aggregate effects, which we display in Figure 8. More precisely, we report the change in aggregate welfare attributable to the different sectors and, as discussed above, calculate the whole distribution of these welfare measures to take the sampling variation into account. We report these distributions as a boxplot. The “box” covers the 25%-75% quantiles of the distribution of aggregate welfare gains, the line within the box indicates the median and the two vertical lines on the top and the bottom indicate the 5% and 95% quantile.

Figure 8 highlights the first important results of our analysis: a substantial part of economic development in India since 1987 was indeed service led. On average, the Indian population would have been willing to reduce their income in 2011 by 26% in lieu of giving up the observed productivity growth originating in the CS sector. Furthermore, with 90% probability the welfare gains of service-led growth are between 22% and 29%. To put this number into perspective, the equivalent variation of the entirety of Indian income growth since 1987 is 64%. Hence, productivity growth in the CS sector accounts for roughly one third of the entire increase in economic well-being.

Figure 8 also shows that agricultural productivity was an important source of welfare improvement between 1987 and 2011. The salience of agriculture is hardly surprising given its large employment share in India. The relatively small welfare effects of productivity growth in the industrial sector is more surprising. The corresponding equivalent variation of productivity growth amounts to about 17% and is very precisely estimated.²¹ Hence, service-led growth has a larger welfare effect than productivity growth in the industrial sector. In fact, our estimates of the distribution of such gains do not even overlap. Finally, for comparison, we also report the welfare consequences of human-capital accumulation. At least as measured in the quantity of schooling, these welfare gains are relatively modest, namely a mere 9% of 2011 income.

In sum, Figure 8 shows service-led growth played an important role for economic development in India since

²¹ We return below to the decomposition of the effects within the industrial sector—manufacturing versus PS.




Figure 8: AGGREGATE WELFARE EFFECTS. The figure displays the percentage welfare losses $\Delta\mathcal{W}$ associated with counterfactually setting productivity in agriculture, CS, and industry, as well as the level of human capital, to their respective levels in 1987 in all Indian districts. We compute the distribution of such welfare losses using a non-parametric bootstrap. The respective boxes cover the 25%-75% quantile of the bootstrap distribution. The horizontal lines on the top and bottom refer to the 5% and 95% quantile of the bootstrap distribution. For comparison, the figure also shows the welfare loss of resetting all productivities and human capital to their 1987 level as the horizontal line at the bottom.

1987. In Section 7, we scrutinize this finding through a battery of sensitivity checks and show that the importance of service-led growth is a robust result.

Heterogeneous Effects. A centerpiece of our contribution is the quantification of the unequal effects of economic growth. Our analysis captures this inequality in two ways. First, as shown in Table 6, we estimate that regions differed in the productivity growth they experienced since 1987. Second, the non-homothetic nature of preferences implies consumers on different levels of the income ladder care differently about sectoral productivity growth: growth in CS and industrial goods is particularly beneficial for the rich, whereas growth in the agricultural sector mostly benefits the poor. Using the notation in (20), the former is mostly reflected in the regional welfare change $\Delta\mathcal{W}_r$, the latter mostly in the welfare change at the individual level $\Delta\mathcal{W}_r^h(q)$.

We first look at the spatial dimension and group districts by quintiles of the urbanization rate in 2011. We then calculate the (income-weighted) average welfare changes $\Delta\mathcal{W}_r$ within each urbanization quintile. These results are shown in the left panel of Figure 9. As in Figure 8, we again display our results using boxplots to illustrate the uncertainty in our estimates.

The welfare consequences of productivity growth vary widely across space. Unsurprisingly, agricultural productivity growth is pro-rural. On average, households in the lowest quintile of urbanization are prepared to sacrifice 24% of their 2011 income to avoid going back to the 1987 productivity level in agriculture. This positive welfare change declines sharply in the top quintile, where productivity growth in agriculture is only worth 16% of the 2011 income. By contrast, productivity growth in CS and the industrial sector were decidedly pro-urban. This pattern is most pronounced for the CS sector whose productivity growth is worth 42% of the 2011 income for the most urbanized quintile. Our estimates of the distributions of these welfare gains make the “urban-rural” split of India also statistically precise. While we cannot reject that the welfare consequences of sectoral productivity growth are




Figure 9: THE HETEROGENEOUS WELFARE IMPACT OF SERVICE-LED GROWTH. The figure displays the average percentage welfare losses associated with counterfactually setting productivity in agriculture, CS, and industry, as well as human capital, at the respective 1987 level, broken down by urbanization quintile in 2011 (Panel (a)) and by the 10th, 20th, 50th, 75th, 90th, 95th, and 99th percentile of the income distribution in 2011 (Panel (b)). We compute the distribution of such welfare losses using a non-parametric bootstrap. The respective boxes cover the 25%-75% quantile of the bootstrap distribution. The horizontal lines on the top and bottom refer to the 5% and 95% quantile of the bootstrap distribution

the same across the lower four quintiles of the distribution, the top urban quintile seems to be qualitatively different: a large part of the realized welfare gains were service-led and agricultural productivity is of limited importance.

Although these differences in the spatial incidence of sectoral productivity growth are partly driven by differences in productivity growth, they also reflect differences in the income distribution. Because the population of cities is, on average, richer, their welfare is particularly reliant on the price of CS. The right panel of Figure 9 decomposes the welfare effects across the Indian income distribution. We focus on the 10th, 20th, 50th, 75th, 90th, 95th, and 99th percentiles. As expected, the benefits of productivity growth in CS and (to a lesser extent) industry are sharply increasing in income, whereas the opposite is true for agriculture. Interestingly, the welfare change for the top 99% attributable to CS productivity growth is smaller than for the average of the top quintile of the urbanization distribution, because not all the rich people live in cities. Furthermore, our methodology uncovers statistically meaningful differences in the sources of welfare growth between the top 25% and the bottom 75% of the population. For the bottom 75% of the population, the welfare effects of productivity growth in agriculture, CS and the industrial sector are roughly of the same size. For the top 25%, service-led growth was quantitatively much more important.

In summary, the welfare effects of growth are heavily skewed. In urban areas and for rich households, the standards of living grew mostly because of productivity growth in CS and—to a lesser extent—in the industrial sector. By contrast, technical progress in agriculture is the main source of welfare gains for the poor, living in rural districts.

Decomposing the Effects of Productivity Growth within the Industrial Sector.

We have so far presented the welfare consequences of resetting productivity in the industrial sector to its 1987 level. As highlighted in Proposition 1, we can decompose the effect of industrial productivity growth (A_{rGt}) into that of manufacturing (A_{rMt}) and PS ($A_{rPS t}$). As we show in more detail in Section B-10 in the Appendix, when

we implement this decomposition, we find productivity growth in the manufacturing sector accounts for the vast majority of the welfare impact of industrial growth. In our baseline calibration, productivity growth originating in the PS sector plays a minor role and accounts for less than 5% of the aggregate welfare gain of industrial productivity growth.

Our methodology infers a small role for PS productivity for two reasons. First, in the Indian micro data, the PS employment share is relatively small. Hence, productivity growth originating in that sector tends to have a quantitatively small effect. We expect this result to be different in more advanced economies, where PS play a quantitatively more important role (Eckert et al., 2020). Second, the decomposition depends on the parameters of the production function, in particular, β , which (see equation (11)) governs the importance of PS as an input of production and determines the share of PS workers in the long run.

Our results are robust with respect to both of these concerns. First, in Section 7.2, we explicitly address the concern that our methodology underestimates the employment share in PS, given the observation of a fast development in service industries such as ICT and show all our results are qualitatively robust to reasonable measurement choices that give a more prominent role for PS employment. Second, in Section B-10 in the Appendix, we study the extent to which our results depend on different calibrations of the production function, in particular, β . The effect of PS productivity growth remains small and accounts for at most one fifth of growth in the industrial sector.

6.2 Service Led Growth and Structural Change

Sectoral productivity growth is not only an important driver of welfare growth, but is also at the heart of the sectoral reallocation of employment, that is, the structural transformation. We report these employment effects in Figure 10. Each of the three panels focuses on one sector and depicts the counterfactual sectoral employment share if productivity growth in agriculture (green bars), CS (orange bars), and the industrial sector (blue bars) had been zero since 1987. The dashed horizontal lines show the actual employment share in 1987 and 2011, for reference.²² Like in Figures 8 and 9 above, we again display the distribution of these effects through a box plot. In contrast to these welfare effects, however, the aggregate employment effects are precisely estimated. Hence, sampling variation plays a minor role in as far as the implications for the structural transformation are concerned.

Figure 10 has a clear message: productivity growth in CS was responsible for the largest part of the observed structural transformation. As seen in the left panel, in the absence of productivity growth in CS, the agricultural employment share would have been 60% instead of 50%. Hence, CS productivity growth accounts for more than half of the decline in agricultural employment between 1987 and 2011. The other panels show that employment in both CS and industry would have been lower if productivity had not grown in the CS sector.

At the same time, Figure 10 also highlights an important role for service-*biased* growth. As shown in panel (b), even in the absence of productivity growth in the CS sector, the employment share of CS would have grown by five percentage points between 1987 and 2011. However, its expansion would have been less spectacular than observed in the data.

The reason productivity growth in CS markedly affects agricultural employment is the following. In the absence of productivity growth, Indian consumers would be poorer and CS would be relatively more expensive. Given our estimated demand system, both forces push toward an increase in the demand for agricultural goods. The income

²² The figure shows results for employment in effective units of labor, which we label employment with a slight abuse of terminology.




Figure 10: SECTORAL PRODUCTIVITY GROWTH AND STRUCTURAL CHANGE. Each panel shows the counterfactual employment share in the respective sector when we set productivity in agriculture, CS, and industry to their 1987 level. The dashed horizontal lines show employment in 1987 and 2011, for reference. We compute the distribution of such welfare losses using a non-parametric bootstrap. The respective boxes cover the 25%-75% quantile of the bootstrap distribution. The horizontal lines on the top and bottom refer to the 5% and 95% quantile of the bootstrap distribution.

effect increases agricultural demand because food is a necessity. The substitution effect complements this force because we estimate food and CS to be slight substitutes.

By contrast, productivity growth in agriculture (green bars) appears to have marginally *increased* employment in agriculture and slowed down employment growth in industry and CS. This finding runs against the view that productivity growth in agriculture is a precondition for industrialization. It is instead in line with the findings of Foster and Rosenzweig (2004) on the effects of the Green Revolution and those of Kelly et al. (2020), who document a negative effect of agricultural productivity on the Industrial revolution across British regions.

In conclusion, service-led growth explains the lion's share of India's structural transformation between 1987 and 2011. Not only would India's consumers be substantially worse off in welfare terms, but India would also still resemble a much more agricultural economy, with industrial and service employment playing a less important role.

7 Robustness

In this section, we perform a robustness analysis of our results. We focus on the results concerning welfare reported in Figures 8 and 9.²³ We focus on three aspects. First, in Section 7.1, we study the sensitivity of the results to changes in the structural parameters. Next, in Section 7.2, we address some measurement issues. Finally, in Section 7.3, we discuss alternative modeling choices.


7.1 Sensitivity to Structural Parameters

Consider, first, the parameters governing preferences and skills. On the preference side, we focus on the asymptotic expenditure share on food ω_F and the income elasticity ε .²⁴ For the distribution of skills, we focus on the Mincerian returns ρ and the tail parameter of the skill distribution ζ . All results are based on re-estimating the entire model.


²³ The results for the structural transformation are available upon request.

²⁴ The other parameters ν_F and ω_{CS} are point identified in our theory.


PANEL a: LONG-RUN SHARE OF AGRICULTURE ω_F



PANEL b: INCOME ELASTICITY ε



PANEL c: RETURN TO EDUCATION ρ



PANEL d: SKILL DISTRIBUTION ζ




Figure 11: ROBUSTNESS ANALYSIS. Panels (a), (b), (c), and (d) show the aggregate welfare effects as a function of the preference parameters ω_F , ε , the Mincerian rates of return to education ρ , and the tail parameter of the skill distribution ζ . The vertical dashed line corresponds to the parameter value in our benchmark analysis.

Preferences. The parameter ω_F was calibrated so that the asymptotic GDP share of agriculture is 1%, corresponding to the output share of the US farming sector in 2017. However, the GDP share of agriculture is larger than 1% in many European countries, at about 2% in Italy and France, and 3% in Spain. Therefore, considering a range of larger values of ω_F is useful. Panel (a) of Figure 11 shows the implied welfare impact of sectoral productivity growth is essentially independent of ω_F .

Panel (b) of Figure 11 focuses on the income elasticity ε . We expect our results to be sensitive to this parameter. In particular, a high income elasticity attributes a large share of the growth of the CS sector to income effects. Conversely, a low income elasticity would require large productivity growth to explain the observed expansion of the CS sector. Consequently, we expect the welfare effects of service-led growth to decrease in ε . The results shown in Panel (b) of Figure 11 confirm our expectation and show that changing ε yields significant quantitative differences. For instance, if we set $\varepsilon = 0.7$, the aggregate welfare effect falls to a mere 10%. However, recall that the highest estimate of the food income elasticity was 0.366 (see Table 3) and that the parameter ε approximately coincided with this elasticity. Figure 11 shows that for any $\varepsilon < 0.5$, a large share of Indian growth is service led. Hence, for growth to be preeminently service biased, the income elasticity should be much higher than what the household-level data suggest.

Skills. In the lower panels of Figure 11, we focus on the determinants of human capital. Our estimate of the return to education ρ based on micro data is an annual 5.6% return. This estimate is on the lower end of typical

Mincerian regressions. A potential concern is that we use data on consumption that might reflect consumption sharing within households with different skills and education levels. This might lead to attenuation bias. For this reason, we consider alternative calibrations in which the return to education is higher, up to an annual 10% that is an upper bound to the range of the typical estimates. As seen in Panel (c) of Figure 11, our main results are not sensitive to this parameter.

Panel (d) of Figure 11 shows the effect of the tail of the skill distribution ζ . This parameter mostly affects our decomposition of productivity growth into agriculture and CS: the higher the ζ , the higher the importance of CS growth relative to agricultural productivity. This result reflects the importance of nonhomothetic demand. The smaller ζ , the higher the income inequality. And because higher inequality increases aggregate demand for CS for a given average wage, less productivity growth is “required” to explain the increase in CS employment if ζ were small. Figure 11 shows this intuition is borne out but that the effects are quantitatively moderate.

	Aggregate Effects				Effects by Urbanization Quantile					
	Agriculture	CS	Industry	HC	Agriculture		CS		Industry	
	1st	5th	1st	5th	1st	5th	1st	5th	1st	5th
Baseline	-20.9	-26.4	-17.4	-8.7	-23.6	-16.2	-17.2	-41.4	-12.7	-23.2
<i>Alternative measurement choices (Section 7.2)</i>										
Double PS	-20.9	-22.4	-20.1	-8.7	-22.6	-17.7	-19.8	-28.3	-14.4	-27.1
ICT & Business to PS	-21.2	-20.2	-18.8	-8.6	-23.3	-17.6	-20.0	-23.8	-13.3	-25.7
Construction to manufacturing	-20.2	-23.5	-22.5	-8.7	-26.5	-12.2	-3.7	-49.5	-12.8	-33.5
<i>Alternative modelling assumptions (Section 7.3)</i>										
Open economy	-21.0	-23.2	-14.3	-8.3	-23.5	-16.6	-16.8	-34.9	-10.6	-18.4
Open economy (large ICT)	-20.5	-19.3	-14.2	-8.3	-22.8	-16.4	-18.5	-22.7	-10.5	-17.8
Imperfect skill substitution	-24.9	-26.5	-16.5	-16.2	-29.5	-18.3	-12.1	-45.3	-11.6	-22.6

Table 7: THE IMPORTANCE OF SERVICE-LED GROWTH: ROBUSTNESS. In this table, we report a summary of our results from the robustness tests described in more detail in the main text. In the first four columns, we report the aggregate welfare loss in the absence of productivity growth (cols 1 - 3) or human-capital accumulation (col 4). In the remaining columns, we report the welfare loss for the 1st and 5th quintile of the urbanization distribution.

7.2 Measurement: Revisiting the PS-CS Split

Our classification of service employment into PS and CS hinges on whether firms in the service sector sell mostly to firms or consumers. For our baseline analysis, we use firm-level information contained in the service survey in this regard. According to this classification, the vast majority of service employment indeed caters to consumers. Even though sectors that sell in significant proportions to firms—such as ICT and business services—grow very quickly, the majority of the service sector continues to be in consumer-oriented industries such as wholesale, retail, and restaurants.²⁵

Although we consider our data-driven approach an accurate way to separate CS from PS, our procedure could underestimate the PS sector if firms report sales to small firms as sales to individuals. To gauge the quantitative importance of such measurement concerns, we consider two alternative classifications. First, we assume the (human-

²⁵ To corroborate our results, we also measured aggregate employment from the Economic Census 2013; that is, we focused on the industry of firms rather than of the employees. In the Economic Census, industries such as wholesale, retail, restaurants, health and community services account for 37.9% of total employment, which compares with approximately 6.5% for financial, business, and ICT services. Note that even these sectors serve in part consumers as many lawyers (who are part of the business service industries) and banks sell their services to households.

capital-adjusted) employment share of PS is twice as large as in our benchmark estimate in each service industry shown in Figure 4. Second, we assume the entire ICT and business service industries serve manufacturing firms (while retaining our baseline approach for the remaining service industries). We regard both alternatives as generous upper bounds for the importance of PS employment.

The results are shown in rows 2 and 3 of Table 7. The first four columns report the aggregate welfare effect ($\Delta\mathcal{W}$), shown in Figure 8. The last six columns focus on the spatial heterogeneity ($\Delta\mathcal{W}_r$), shown in Figure 9. For parsimony, we only report the top and bottom urbanization quantiles. As expected, the importance of productivity growth in CS decreases when we attribute a larger share of the expanding service sector to PS. This is especially important for the most urban locations, because the spatial concentration of PS exceeds the one of CS. However, in all cases, Indian growth continues to be preeminently service led.

Finally, we explore the importance of the construction sector. Recall that we attributed construction to the service sector, given its non-tradable nature. Because, traditionally, construction is absorbed in the manufacturing sector, we also redid our analysis under this alternative measurement choice. We report the result in row 4 of Table 7. Although this reclassification reduces the importance of CS and increases the importance of the industrial sector, we still find CS to be the most important contributor to Indian growth. Interestingly, construction plays a particularly important role for the spatial heterogeneity, because it is relatively pro rural. If we do not count the construction sector as part of the service sector, the spatial incidence of service-led growth is even more pro urban than in our baseline estimate. Specifically, the welfare effect of productivity growth in CS remains the same in the most urbanized districts, whereas it turns minuscule in the most rural districts.

7.3 Alternative Modeling Assumptions

Given the accounting nature of our methodology, a natural question concerns the extent to which our results are sensitive to our specific modeling assumptions. In this section, we consider two alternatives. First we extend our model to allow for international trade. Second, we consider an environment where skills are imperfectly substitutable.

7.3.1 Open Economy

Thus far, we have treated India as a closed economy. However, international trade, in particular exports of ICT services, has become increasingly important for India. In this section, we therefore extend our model to an open-economy environment. For brevity, we only summarize the main features of the extended model. The technical analysis can be found in Appendix Section A-5.

We assume consumers, both in India and in the rest of the world, consume industrial goods sourced from many countries. Different national varieties, which are in turn CES aggregates of regional varieties, enter into a CES utility function as imperfect substitutes. To capture that India might have a specific comparative advantage in ICT services, we assume India exports both domestic goods and ICT services. For simplicity, we assume ICT services are not sold in the Indian domestic market. In our estimation, we assume balanced trade, but we allow India to run a trade deficit in goods and a surplus in ICT services, which is in line with the empirical observation.

To calibrate this model, we need information on the revenue of ICT services, the exports and imports of goods, and an estimate of the trade elasticity. We measure ICT revenue from the income share of ICT workers. We classify as ICT service workers all those employed in the following service industries: (i) telecommunications, (ii) computer programming, (iii) consultancy and related activities software publishing, and (iv) information-service activities. In our NSS data, these activities constitute 0.72% of total employment in 2011 (in 1987, it was a less than 0.1%). ICT

workers earn, on average, higher wages than other workers. When one considers the earning share, they account for 1.56% of total earnings in 2011 (in 1987, it was 0.11%). Given the small size of the ICT sector in 1987, we assume it was zero and target the earnings share in 2011. In terms of exports, according to the World Bank, the export of goods and merchandise increased from 11.3 billions (4.1% of GDP) in 1987 to 302.9 billions (16.6% of GDP) in current USD. The manufacturing sector accounted for 66% of such merchandise exports in 1987 and for 62% in 2011. According to the OECD, the domestic value added in gross exports amounts to 83.9% of exports for India and we assume this percentage to be constant over time. In accordance with these data, we assume the value-added export of trade increased from 13.9% in 1987 to 53.6% in 2011 as a share of the GDP in the manufacturing sector. Finally, we set the trade elasticity to 5 (Simonovska and Waugh, 2014).

The results of quantifying the sources of growth in this context are contained in rows 6 and 7 of Table 7. In row 6, we report the results from the measurement choices outlined above. In row 7, we report the results when the ICT sector is twice as large as actually observed. Expectedly, such choices reduce the importance of the CS, because they reduce measured employment growth in these industries. Again, this is particularly relevant for cities, which saw the fastest increase in ICT employment. Nevertheless, CS continue to play an important role for aggregate growth and for urban areas in particular. Adding foreign trade does not alter the result that Indian growth is largely service led.

7.3.2 Imperfect Substitution and Skill Bias in Technology

In our model, we allow for individual heterogeneity in human capital but maintain that workers endowed with different efficiency units are perfect substitutes for one another. In this section, we generalize our model by assuming workers with different educational attainments are imperfect substitutes in production (see Section B-12 in the Appendix for details). As we showed in Table 2, agricultural workers have, on average, lower educational attainment than those employed in service industries. Thus, an increase in the skill endowment could be responsible for the reallocation of workers from agriculture to CS (see, e.g., Porzio et al. (2020) or Schoellman and Hendricks (2020)). By ignoring such skill-based specialization, our Ricardian model could exaggerate the importance of technology for the development of the service sector.

For simplicity, we work with two skill groups and define workers to be skilled if they have completed secondary school. We assume the production functions to be of the usual CES form:

$$Y_{rs} = A_{rst} \left((H_{rst}^-)^{\frac{\rho-1}{\rho}} + (Z_{rst} H_{rst}^+)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}} \quad \text{for } s = F, CS, G,$$

where H^+ and H^- denote high- and low-skilled workers, respectively. Note that the technology admits differences in both TFP A_{rst} and skill bias Z_{rst} across sector-districts and time.²⁶ We assume the elasticity of substitution ρ to be constant across sector-districts and externally calibrate $\rho = 1.8$, which is in the consensus region (see, e.g., Ciccone and Peri (2005) and Gancia et al. (2013)). Our conclusions do not hinge on the particular calibration of ρ .

We continue to allow for heterogeneous productivities across workers of the same educational group. A worker's wage is a draw from a skill-specific Pareto distribution with the same tail parameter as in our baseline analysis.²⁷

²⁶ Allowing the skill bias of technology to vary across space is important. If Z were constant across districts, the model would predict skill premia to be lower in skill-rich regions. However, this assumption contradicts the observation that both the relative supply of skills and the skill premium are positively correlated with urbanization.

²⁷ Separately identifying the lower bound of the Pareto distribution of human-capital draws from the level of the technology parameters is impossible. Therefore, we normalize the lower bound to unity for both skill groups. Because we are only interested in changes over time in TFP, this normalization is immaterial.

As in our baseline analysis, this model is exactly identified, and for given structural parameters, we can rationalize the data of sectoral earnings shares by skill group and average earnings by skill group for each region in India by choice of A_{rst} and Z_{rst} (see Section B-12 in the Appendix).

The results of this extension are reported in the last row of Table 7. Because productivity growth is now parametrized by changes in factor-neutral productivity A_{rst} and the skill bias Z_{rst} , we set the respective growth rates in both sectors to zero. The quantitative role for the CS sector is very similar to the one of our baseline calibration. Interestingly, human capital now plays a more important role, owing to the increasing supply of high-skilled labor over time.

This extension also allows us to uncover additional facts about the skill bias in technology. In Figure 12, we plot our estimates of the skill bias Z_{rst} as binned scatter plots. First, across districts, Z_{rs} increases in the level of urbanization for all sectors. This increase reflects the empirical observation that the skill premium is higher in urban than in rural districts. Second, we find evidence for skill-biased technical change: over time, Z_{rs} increases in all sectors. Although our accounting approach cannot uncover causal links, these patterns are consistent with models of directed technical change and directed technology adoption such as Acemoglu and Zilibotti (2001) and Gancia et al. (2013), where firms adopt more skill-intensive technologies in response to the wider availability of skilled workers.




Figure 12: SKILL BIAS OF TECHNOLOGY. The figure shows a binned scatter plot of Z_{rF} , Z_{rCS} , and Z_{rG} as a function of the urbanization rate in 1987 (dashed line and “+” markers) and 2011 (solid line dots).

8 Conclusion

Although an expanding service sector is often seen as a rich-country phenomenon, tertiarization is well underway in most developing countries. In particular, rising employment in consumer services such as retail and restaurants accounts for the bulk of the decline in agricultural employment while industrial employment is often stagnant. Are these patterns a sign that services can be a source of productivity growth even at low levels of economic development? Or is rising service employment simply a corollary of rising incomes if services are luxury goods? In short: Is growth service led or service biased?

In this paper, we developed a methodology to answer this question. Our approach is in the spirit of development accounting but uses the restrictions imposed by a spatial equilibrium model. The estimated model allows us to determine the importance of different sectors as an engine of growth and structural transformation. Moreover, it lends itself to a quantitative analysis of both the aggregate welfare effects of growth and its distributional consequences.

At the core of our identification strategy are consumers’ preferences, in particular, the income elasticity of aggregate service demand. The higher this elasticity, the more service-biased economic growth is. Conversely, if the income elasticity of consumer demand is limited, rising employment in the consumer service sector is a sign that growth was service led. Given the importance of this parameter, we infer it directly using Indian household data. Importantly, we show that the income elasticity of consumers’ observable demand system over final expenditure coincides with the one defined over value added that is relevant in our theory.

Our analysis delivers two main results. First and foremost, Indian growth was to a large extent service led. Quantitatively, productivity growth in sectors such as retail, hospitality, or transportation account for one third of welfare growth between 1987 and 2011. Second, the welfare impact of service-led growth was strikingly unequal and benefitted mostly wealth individuals in urbanized locations. The reasons are that service productivity grew particularly fast in urban areas and that richer consumers care more about the consumption of services owing to nonhomothetic preferences.

We also document that productivity growth in consumer services was the main driver of the structural transformation and accounts for almost half of the decline in agricultural employment. By contrast, technical progress in agriculture, did not promote structural change. This result is in line with a growing body of literature – including Kelly et al. (2020), Moscona (2019), or Foster and Rosenzweig (2004) – who document similar findings for India, within and across countries, and for the British Industrial Revolution.

Our approach has several limitations that we hope to overcome in future research. Two are particularly important. First, owing to our accounting approach, we took consumer service productivity as exogenous. Understanding the exact nature of productivity growth and how it materializes seems to us a question of first-order importance, in particular as far as potential policy-implications are concerned. Second, knowing whether service-led growth is unique to the Indian experience or also important in other developing countries would be interesting. If service-led growth is indeed an integral part of the growth trajectory of developing countries today, the absence of employment growth in the manufacturing sector might be less concerning than previously thought.

References

- Acemoglu, D. and F. Zilibotti: 2001, ‘Productivity Differences’. *Quarterly Journal of Economics* **116**(2), 563–606.
- Aghion, P., R. Burgess, S. Redding, and F. Zilibotti: 2005, ‘Entry Liberalization and Inequality in Industrial Performance’. *Journal of the European Economic Association* **3**(2-3), 291–302.
- Aghion, P., R. Burgess, S. Redding, and F. Zilibotti: 2008, ‘The Unequal Effects of Liberalization: Evidence from Dismantling the License Raj in India’. *American Economic Review* **98**(4), 1397–1412.
- Akcigit, U., H. Alp, and M. Peters: 2021, ‘Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries’. *American Economic Review* **111**(1), 231–275.
- Alder, S., T. Boppart, and A. Muller: 2019, ‘A Theory of Structural Change that Can Fit the Data’. CEPR DP13469.
- Allen, T. and C. Arkolakis: 2014, ‘Trade and the Topography of the Spatial Economy’. *Quarterly Journal of Economics* **129**(3), 1085–1140.
- Basu, K.: 2008, ‘The Enigma of India’. *Journal of Economic Literature* **46**(2), 396–406.
- Basu, K. and A. Maertens: 2007, ‘The Pattern and Causes of Economic Growth in India’. *Oxford Review of Economic Policy* **23**(2), 143–167.
- Baumol, W.: 1967, ‘Macroeconomics of Unbalanced Growth: The Anatomy of the Urban Crisis’. *American Economic Review* **57**(3), 415–426.
- Boppart, T.: 2014, ‘Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences’. *Econometrica* **82**(6), 2167–2196.

- Buera, F. J. and J. P. Kaboski: 2012, 'The Rise of the Service Economy'. *American Economic Review* **102**(6), 2540–69.
- Caselli, F.: 2005, 'Accounting for Cross-Country Income Differences'. In: P. Aghion and S. Durlauf (eds.): *Handbook of Economic Growth*, Vol. 1A. Elsevier B.V., Chapt. 9, pp. 679–741.
- Chatterjee, S. and E. Giannone: 2021, 'Unequal Global Convergence'. Working Paper.
- Cheremukhin, A., M. Golosov, S. Guriev, and A. Tsyvinski: 2015, 'The economy of People's Republic of China from 1953'. NBER Working Paper 21397.
- Cheremukhin, A., M. Golosov, S. Guriev, and A. Tsyvinski: 2017, 'The industrialization and economic development of Russia through the lens of a neoclassical growth model'. *Review of Economic Studies* **84**(2), 613–649.
- Ciccone, A. and G. Peri: 2005, 'Long-Run Substitutability Between More and Less Educated Workers: Evidence from U.S. States, 1950-1990'. *Review of Economics and Statistics* **87**(4), 652–663.
- Comin, D., D. Lashkari, and M. Mestieri: 2020, 'Structural Change with Long-Run Income and Price Effects'. Forthcoming in *Econometrica*.
- Cravino, J. and S. Sotelo: 2019, 'Trade-induced structural change and the skill premium'. *American Economic Journal: Macroeconomics* **11**(3), 289–326.
- Dehejia, R. and A. Panagariya: 2016, 'The link between manufacturing growth and accelerated services growth in India'. *Economic Development and Cultural Change* **64**(2), 221–264.
- Desmet, K., E. Ghani, S. O'Connell, and E. Rossi-Hansberg: 2015, 'The spatial development of India'. *Journal of Regional Science* **55**(1), 10–30.
- Duarte, M. and D. Restuccia: 2010, 'The Role of the Structural Transformation in Aggregate Productivity*'. *Quarterly Journal of Economics* **125**(1), 129–173.
- Duarte, M. and D. Restuccia: 2019, 'Relative Prices and Sectoral Productivity'. *Journal of the European Economic Association* **18**(3), 1400–1443.
- Eckert, F., S. Ganapati, and C. Walsh: 2020, 'Skilled Scalable Services: The New Urban Bias in Economic Growth'. Minneapolis Reserve Bank of Minneapolis, Institute Working Paper 25.
- Eckert, F. and M. Peters: 2020, 'Spatial Structural Change'. Working Paper.
- Fajgelbaum, P. D. and A. K. Khandelwal: 2016, 'Measuring the unequal gains from trade'. *The Quarterly Journal of Economics* **131**(3), 1113–1180.
- Foellmi, R. and J. Zweimüller: 2006, 'Income Distribution and Demand-Induced Innovations'. *Review of Economic Studies* **73**(4), 941–960.
- Foster, A. and M. Rosenzweig: 1996, 'Technical Change and Human-Capital Returns and Investments: Evidence from the Green Revolution'. *American Economic Review* **86**(4), 931–953.
- Foster, A. and M. Rosenzweig: 2004, 'Agricultural Productivity Growth, Rural Economic Diversity, and Economic Reforms: India, 1970–2000'. *Economic Development and Cultural Change* **52**(3), 509–542.
- Gancia, G., A. Müller, and F. Zilibotti: 2013, 'Structural Development Accounting'. In: D. Acemoglu, M. Arellano, and E. Dekel (eds.): *Advances in Economics and Econometrics: Theory and Applications (Tenth World Congress of the Econometric Society)*. Cambridge University Press, pp. 373–418.
- Garcia-Santana, M., J. Pijoan-Mas, and L. Villacorta: 2020, 'Investment Demand and Structural Change'. Mimeo, CEMFI.
- Gervais, A. and J. B. Jensen: 2019, 'The tradability of services: Geographic concentration and trade costs'. *Journal of International Economics* **118**, 331–350.
- Goldberg, P., A. K. Khandelwal, N. Pavcnik, and P. Topalova: 2010, 'Imported Intermediate Inputs and Domestic Product Growth: Evidence from India'. *Quarterly Journal of Economics* **125**(4), 1727–1767.
- Gollin, D., R. Jedwab, and D. Vollrath: 2015, 'Urbanization with and without industrialization'. *Journal of Economic Growth* **21**, 1573–7020.
- Gollin, D., D. Lagakos, and M. Waugh: 2014, 'The Agricultural Productivity Gap'. *Quarterly Journal of Economics* **129**(2), 939–993.

- Hall, R. and C. Jones: 1999, ‘Why Do Some Countries Produce So Much More Output Per Worker Than Others?’. *Quarterly Journal of Economics* **114**(1), 83–116.
- Herrendorf, B., R. Rogerson, and A. Valentinyi: 2013, ‘Two Perspectives on Preferences and Structural Transformation’. *American Economic Review* **103**(7), 2752–89.
- Herrendorf, B., R. Rogerson, and A. Valentinyi: 2014, ‘Growth and Structural Transformation’. In: *Handbook of economic growth*, Vol. 2. Elsevier, pp. 855–941.
- Herrendorf, B., R. Rogerson, and A. Valentinyi: 2020, ‘Structural Change in Investment and Consumption—A Unified Analysis’. *Review of Economic Studies*.
- Hobijn, B., T. Schoellman, and A. Vindas: 2019, ‘Structural Transformation by Cohort’. Mimeo, Arizona State University.
- Horowitz, J. L.: 2019, ‘Bootstrap methods in econometrics’. *Annual Review of Economics* **11**, 193–224.
- Hsieh, C.-T. and P. J. Klenow: 2014, ‘The life cycle of plants in India and Mexico’. *The Quarterly Journal of Economics* **129**(3), 1035–1084.
- Hsieh, C.-T. and E. Rossi-Hansberg: 2019, ‘The Industrial Revolution in Services’. NBER Working Paper No. 25968.
- Kelly, M., J. Mokyr, and C. O’Grada: 2020, ‘The Mechanics of the Industrial Revolution’. CEPR DP14884.
- Kochhar, K., U. Kumar, R. Rajan, A. Subramanian, and I. Tokatlidis: 2006, ‘India’s Pattern of Development: What Happened, What Follows?’. *Journal of Monetary Economics* **53**(5), 981–1019.
- Kongsamut, P., S. Rebelo, and D. Xie: 2001, ‘Beyond balanced growth’. *The Review of Economic Studies* **68**(4), 869–882.
- Martin, L., S. Nataraj, and A. Harrison: 2017, ‘In with the Big, Out with the Small: Removing Small-Scale Reservations in India’. *American Economic Review* **107**(2), 354–386.
- 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**(6), 1093–1120.
- Matsuyama, K.: 2019, ‘Engel’s Law in the Global Economy: Demand-Induced Patterns of Structural Change, Innovation, and Trade’. *Econometrica* **87**(2), 497–528.
- McMillan, M. and D. Rodrik: 2011, ‘Globalization, Structural Change, and Productivity Growth’. In: M. Bacchetta and M. Jansen (eds.): *Making Globalization Socially Sustainable*. International Labor Organization, pp. 49–84.
- Moscona, J.: 2019, ‘Agricultural Development and Structural Change, Within and Across Countries’. Working Paper, MIT.
- Muellbauer, J.: 1976, ‘Community Preferences and the Representative Consumer’. *Econometrica* **44**(5), 979–999.
- Ngai, L. R. and C. A. Pissarides: 2007, ‘Structural Change in a Multisector Model of Growth’. *American Economic Review* **97**(1), 429–443.
- Porzio, T., F. Rossi, and G. Santangelo: 2020, ‘The Human Side of Structural Transformation’. CEPR DP15110.
- Redding, S. J. and E. Rossi-Hansberg: 2017, ‘Quantitative Spatial Economics’. *Annual Review of Economics* **9**(1), 21–58.
- Rodrik, D.: 2015, ‘Premature Industrialization’. Mimeo, Harvard Kennedy School.
- Sato, R.: 2014, ‘The Most General Class of CES Functions’. *Econometrica* **43**, 999–1003.
- Schoellman, T. and L. Hendricks: 2020, ‘Skilled Labor Productivity and Cross-Country Income Differences’. Mimeo, Federal Reserve of Minneapolis.
- Simonovska, I. and M. E. Waugh: 2014, ‘The elasticity of trade: Estimates and evidence’. *Journal of international Economics* **92**(1), 34–50.

APPENDIX A: TECHNICAL DETAILS

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

A-1 Derivation of the expenditure shares $\vartheta_s(e, \mathbf{p})$

Roy's Identity implies that the expenditure share on sector s is given by

$$\vartheta_s(e, \mathbf{p}) = - \frac{\frac{\partial V(e(\mathbf{p}, u), \mathbf{p})}{\partial p_s} p_s}{\frac{\partial V(e(\mathbf{p}, u), \mathbf{p})}{\partial e} e}.$$

Using the indirect utility function in (3), $\vartheta_s(e, \mathbf{p})$ is given by

$$\vartheta_s(e, \mathbf{p}) = \omega_s + \nu_s \left(\frac{e}{\prod_s p_{st}^{\omega_s}} \right)^{-\varepsilon}$$

A-2 Equilibrium in the Industrial Sector

In this section we characterize the equilibrium in the industrial sector. The technical details are contained in the Online Appendix. As highlighted in Proposition 1, we have to distinguish two cases. In particular, recall the definition

$$\varsigma(A_{PS}) \equiv \frac{\kappa}{f_O A_{PS}}. \quad (\text{A-1})$$

Henceforth, we simply write ς . Below we will show that some active firms do not hire lawyers if and only if

$$\varsigma \geq \frac{\beta}{1 - \alpha}. \quad (\text{A-2})$$

Note that ς is decreasing in A_{PS} (see (A-1)). Hence, condition (A-2) requires the productivity of lawyers A_{PS} to be low enough.

Firm-level allocations

We first solve for the firm-level allocations, i.e. firm profits, firm employment and the productivity cutoff z^* . Let p_G denote the price of the industrial good. If active, firm z_i solves the maximization problem

$$\pi(z_i) = \max_{H_{PMi}, H_{PSi} \geq 0} \left\{ p_G z_i^{1-\alpha-\beta} H_{PMi}^\alpha (A_{PS} H_{PSi} + \kappa)^\beta - w(H_{PMi} + H_{PSi}) - f_{OW} \right\}. \quad (\text{A-3})$$

where f_{OW} denotes the overhead costs. Note that we explicitly impose the constraint that $H_{PSi} \geq 0$. Firms operate if and only if $\pi(z_i) \geq 0$. We denote the productivity threshold by z^* , i.e., $\pi(z^*) = 0$. Under condition (12), $z^* > A_M$, that is there is a range of low-productivity firms that choose to be inactive.

Proposition 2. *Suppose that $\varsigma \geq \frac{\beta}{1-\alpha}$, where ς is given in (A-1). Let z^* denote the endogenous productivity threshold, such that firms with $z_i \geq z^*$ produce in equilibrium. Let z_L denote the cutoff where firms start hiring producer services. Then:*

1. *The productivity threshold z^* is given by*

$$z^* = \left(\frac{w}{p_G} \frac{1}{\kappa^\beta \alpha} \left(\frac{\alpha}{1-\alpha} f_O \right)^{1-\alpha} \right)^{\frac{1}{1-\alpha-\beta}}. \quad (\text{A-4})$$

2. The producer service cutoff z_L is given by

$$z_L = z^* \left(\frac{1-\alpha}{\beta} \varsigma \right)^{\frac{1-\alpha}{1-\alpha-\beta}} > z^* \quad (\text{A-5})$$

3. Optimal factor demands are given by

$$H_{PS}(z_i) = \begin{cases} 0 & \text{if } z_i < z_L \\ \varsigma^{\frac{z_i - z_L}{z_L}} f_O & \text{if } z_i \geq z_L \end{cases}, \quad (\text{A-6})$$

and

$$H_{PM}(z_i) = \begin{cases} \frac{\alpha}{\beta} f_O \varsigma \left(\frac{z_i}{z_L} \right)^{\frac{1-\alpha-\beta}{1-\alpha}} & \text{if } z_i < z_L \\ \frac{\alpha}{\beta} \varsigma^{\frac{z_i}{z_L}} f_O & \text{if } z_i \geq z_L \end{cases}. \quad (\text{A-7})$$

4. Firm-level profits are given by

$$\pi(z_i) = \begin{cases} \left(\left(\frac{1-\alpha}{\beta} \varsigma \left(\frac{z_i}{z_L} \right)^{\frac{1-\alpha-\beta}{1-\alpha}} - 1 \right) f_O w & \text{if } z_i < z_L \\ \left(\varsigma \left(1 + \left(\frac{1-\alpha-\beta}{\beta} \right) \frac{z_i}{z_L} \right) - 1 \right) f_O w & \text{if } z_i \geq z_L \end{cases}. \quad (\text{A-8})$$

Proof. See Section OA-2.1 in the Online Appendix. $\square$

Note that (A-5) determines z_L directly as a function of z^* . Moreover, under our assumption that $\varsigma > \frac{\beta}{1-\alpha}$ indeed $z^* < z_L$ and all firms with $z_i \in [z^*, z_L]$ do not hire lawyers. As $\varsigma \rightarrow \frac{\beta}{1-\alpha}$, we have $z^* \rightarrow z_L$. Note also that the profit function in (A-8) is concave in z as long as firms do not hire lawyers but linear in z once they hire lawyers.

Proposition 3. Suppose that $\varsigma < \frac{\beta}{1-\alpha}$, where ς is given in (A-1). Let $\tilde{z}$ denote the endogenous productivity threshold, such that firm with $z_i \geq \tilde{z}$ will produce in equilibrium. Then:

1. The productivity threshold is given by

$$\tilde{z} = \frac{1}{1-\alpha-\beta} \left(\frac{w}{p_G} \right)^{\frac{1}{1-\alpha-\beta}} \left(\frac{1}{\alpha} \right)^{\frac{1-\alpha}{1-\alpha-\beta}} \left(\frac{1}{\beta A_{PS}} \right)^{\frac{\beta}{1-\alpha-\beta}} f_O (1-\varsigma). \quad (\text{A-9})$$

2. Optimal factor demands are given by

$$\begin{aligned} H_{PMi} &= \frac{\alpha}{1-\alpha-\beta} f_O (1-\varsigma) \frac{z_i}{\tilde{z}} \\ H_{PSi} &= \frac{\beta}{1-\alpha-\beta} f_O (1-\varsigma) \frac{z_i}{\tilde{z}} - \frac{\kappa}{A_{PS}}. \end{aligned}$$

3. Firm-level profits are given by

$$\pi(z_i) = \pi(z_i) = \left(\frac{z - \tilde{z}}{\tilde{z}} \right) f_O (1-\varsigma) w. \quad (\text{A-10})$$

Proof. See Section OA-2.1 in the Online Appendix. $\square$

Free Entry and the Equilibrium Wage

Free entry requires that the cost of entry are equal to the expected profits, i.e.

$$f_E w = E[\pi] = \int \pi(x) f(x) dx.$$

This condition allows us to solve for the equilibrium real wage $\frac{w}{p_G}$.

Proposition 4. Suppose that $\varsigma \geq \frac{\beta}{1-\alpha}$. Then

$$\left(\frac{z_L}{A_M}\right)^\lambda = \frac{(1-\alpha-\beta)}{\beta+(1-\alpha)(\lambda-1)} \left[\left(\varsigma \frac{1-\alpha}{\beta}\right)^{\lambda \frac{1-\alpha}{1-\alpha-\beta}} + \frac{\varsigma}{\lambda-1} \right] \frac{f_O}{f_E} \quad (\text{A-11})$$

$$\left(\frac{z^*}{A_M}\right)^\lambda = \frac{(1-\alpha-\beta)}{\beta+(1-\alpha)(\lambda-1)} \left[1 + \frac{1}{\lambda-1} \left(\frac{\beta}{1-\alpha}\right)^{\frac{(1-\alpha)\lambda}{1-\alpha-\beta}} \varsigma^{-\frac{(1-\alpha)(\lambda-1)+\beta}{1-\alpha-\beta}} \right] \frac{f_O}{f_E}, \quad (\text{A-12})$$

and

$$\frac{w}{p_G} = \left(\frac{z_L}{\kappa}\right)^{1-\alpha-\beta} \alpha^\alpha (\beta A_{PS})^{1-\alpha} \quad (\text{A-13})$$

$$= \left(\frac{(1-\alpha-\beta)}{\beta+(1-\alpha)(\lambda-1)} \left[1 + \frac{1}{\lambda-1} \left(\frac{\beta}{1-\alpha}\right)^{\frac{(1-\alpha)\lambda}{1-\alpha-\beta}} \varsigma^{-\frac{(1-\alpha)(\lambda-1)+\beta}{1-\alpha-\beta}} \right] \frac{f_O}{f_E} \right)^{\frac{1-\alpha-\beta}{\lambda}} A_M^{1-\alpha-\beta} \kappa^\beta \alpha^\alpha \left(\frac{1-\alpha}{f_O}\right)^{1-\alpha} \quad (\text{A-14})$$

Suppose that $\varsigma < \frac{\beta}{1-\alpha}$. Then

$$\left(\frac{\tilde{z}}{A_M}\right)^\lambda = \frac{1}{\lambda-1} \frac{f_O}{f_E} (1-\varsigma), \quad (\text{A-15})$$

and

$$\frac{w}{p_G} = \left(\frac{\tilde{z}(1-\alpha-\beta)}{f_O(1-\varsigma)}\right)^{1-\alpha-\beta} (\beta A_{PS})^\beta (\alpha)^\alpha \quad (\text{A-16})$$

$$= (1-\alpha-\beta)^{1-\alpha-\beta} \left(\frac{1}{\lambda-1} \frac{1}{f_E}\right)^{\frac{1-\alpha-\beta}{\lambda}} \left(\frac{1}{f_O(1-\varsigma)}\right)^{\frac{\lambda-1}{\lambda}(1-\alpha-\beta)} A_M^{1-\alpha-\beta} \left(\beta \frac{\kappa}{f_O \varsigma}\right)^\beta (\alpha)^\alpha. \quad (\text{A-17})$$

Proof. See Section OA-2.3 in the Online Appendix. $\square$

Proposition 4 characterizes the cutoffs and the real wage in terms of parameters. In particular, all cutoffs z_L are independent of the wage w . Note also that ς is decreasing in A_{PS} so that $\frac{\partial z_L}{\partial A_{PS}} < 0$, i.e. if lawyers become more productive, the cutoff to hire lawyers declines. Moreover, $\frac{\partial z^*}{\partial A_{PS}} > 0$.¹ Because ς is decreasing in A_{PS} , the real wage is increasing in A_{PS} (see Section OA-2.5 in the Online Appendix).

Aggregate Labor Allocations

Now consider aggregate employment. In our economy, workers in the manufacturing sector are used for production work (H_{PM}), to provide PS (H_{PS}), pay for overhead (H_{OM}) and generate new business ideas (H_{EM}). Hence, labor market clearing requires that

$$H_G = H_{PS} + H_{PM} + H_{EM} + H_{OM}.$$

¹ Note that we assumed that $z^* > A_M$. Hence, we need to impose that

$$\frac{(1-\alpha-\beta)}{\beta+(1-\alpha)(\lambda-1)} \left[1 + \frac{1}{\lambda-1} \left(\frac{\beta}{1-\alpha}\right)^{\frac{(1-\alpha)\lambda}{1-\alpha-\beta}} \varsigma^{-\frac{(1-\alpha)(\lambda-1)+\beta}{1-\alpha-\beta}} \right] \frac{f_O}{f_E} > 1$$

Proposition 5. *The number of entry and production workers is given by*

$$\begin{aligned} H_{EM} &= \frac{1 - \alpha - \beta}{\lambda} H_G \\ H_{PM} &= \alpha H_G \end{aligned}$$

independent of ς . The number of firms, M , is given by

$$M = \frac{1 - \alpha - \beta}{\lambda} \frac{H_G}{f_E}.$$

independent of ς . The number of lawyers and overhead workers is given by

$$H_{PS} = \begin{cases} \frac{\sigma(\varsigma)}{1 + \sigma(\varsigma)} \frac{\beta + (\lambda - 1)(1 - \alpha)}{\lambda} H_G & \text{if } \varsigma \geq \frac{\beta}{1 - \alpha} \\ \left(\beta - (1 - \alpha - \beta) \frac{\lambda - 1}{\lambda} \frac{\varsigma}{1 - \varsigma} \right) H_G & \text{if } \varsigma < \frac{\beta}{1 - \alpha} \end{cases}$$

and

$$H_{OM} = \begin{cases} \frac{1}{1 + \sigma(\varsigma)} \times \frac{\beta + (\lambda - 1)(1 - \alpha)}{\lambda} H_G & \text{if } \varsigma \geq \frac{\beta}{1 - \alpha} \\ (1 - \alpha - \beta) \frac{\lambda - 1}{\lambda} \frac{1}{1 - \varsigma} H_G & \text{if } \varsigma < \frac{\beta}{1 - \alpha} \end{cases},$$

where

$$\sigma(\varsigma) \equiv \frac{1}{\lambda - 1} \left(\frac{\beta}{1 - \alpha} \right)^{\lambda \frac{1 - \alpha}{1 - \alpha - \beta}} \varsigma^{-\frac{(\lambda - 1)(1 - \alpha) + \beta}{1 - \alpha - \beta}}.$$

Proof. See Section OA-2.4 in the Online Appendix. □

Proposition 5 implies that H_{PS} is increasing in A_{PS} , whereas H_{OM} is decreasing in A_{PS} . Interestingly, their sum is independent of A_{PS} , i.e.

$$\frac{H_{OM} + H_{PS}}{M} = \left(\frac{\beta + (\lambda - 1)(1 - \alpha)}{1 - \alpha - \beta} \right) f_E.$$

Proposition 5 also implies that the endogenous number of firms that are created, M , is independent of A_{PS} conditional on total industrial employment H_G . But the number of ideas, that actually produce, $M(1 - F(z^*))$, is decreasing in A_{PS} as the production cutoff z^* is increasing in A_{PS} . Hence, improvements in the productivity of lawyers induce selection by truncating the productivity distribution. Note that all these allocations are independent of aggregate manufacturing productivity A_M . This is in contrast to the micro-level, where employment shares vary systematically with firm productivity z . In particular, (A-6) and (A-7) imply that for firms that hire lawyers (i.e. $z_i \geq z_L$) we have

$$\frac{H_{PS}(z)}{H_{PM}(z)} = \frac{\beta}{\alpha} \left(1 - \frac{z_i}{z_L} \right),$$

i.e. more productive firms have a higher employment share of lawyers relative to production workers. However, the aggregate employment share of production workers hired by (large) firms who hire a positive share of lawyers (i.e. $z_i > z_L$) is given by

$$\frac{\int_{z \geq z_L} H_{PM}(z) dG(z)}{\int_{z \geq z_L} H_{PS}(z) dG(z)} = \lambda \frac{\alpha}{\beta}.$$

Hence, even though there is micro-heterogeneity in the intensity of hiring lawyers, the aggregate employment share of lawyers (among firms who hire lawyers) is constant and depends on the tail of the productivity distribution λ . A thicker tail, i.e. λ smaller, *increases* the aggregate employment share of lawyers by shifting resources towards large firms, which are lawyer intensive.

Figure A-1 depicts the allocation of employment as a function of $\varsigma = \frac{\kappa}{A_{PS} f_O}$ for both cases discussed above. Note that all employments are continuous at the threshold $\varsigma = \frac{\beta}{1 - \alpha}$.




Figure A-1: Aggregate labor allocations

Aggregate Manufacturing Productivity

The free entry condition ensures that the industrial sector as a whole does not generate any profits. Hence, aggregate revenue is equal to aggregate labor payments

$$p_G Y_G = w H_G.$$

Proposition 6. *Let aggregate productivity in the goods producing sector A_G be defined by*

$$A_G \equiv \frac{Y_G}{H_G}.$$

Then,

$$A_G = \begin{cases} Q_2 \left(\frac{1}{1-\varsigma} \right)^{\frac{\lambda-1}{\lambda}(1-\alpha-\beta)} \left(\frac{1}{\varsigma} \right)^\beta A_M^{1-\alpha-\beta} & \text{if } \varsigma < \frac{\beta}{1-\alpha} \\ Q_1 \left(1 + \frac{1}{\lambda-1} \frac{\beta}{1-\alpha} \left(\frac{1-\alpha}{\beta} \varsigma \right)^{-\frac{(1-\alpha)(\lambda-1)+\beta}{1-\alpha-\beta}} \right)^{\frac{1-\alpha-\beta}{\lambda}} A_M^{1-\alpha-\beta} & \text{if } \varsigma \geq \frac{\beta}{1-\alpha} \end{cases},$$

where

$$Q_1 = \left(\frac{(1-\alpha-\beta)}{\beta + (1-\alpha)(\lambda-1)} \frac{f_O}{f_E} \right)^{\frac{1-\alpha-\beta}{\lambda}} \kappa^\beta \alpha^\alpha \left(\frac{1-\alpha}{f_O} \right)^{1-\alpha}$$

and

$$Q_2 = \alpha^\alpha (1-\alpha-\beta)^{1-\alpha-\beta} \left(\frac{1}{\lambda-1} \frac{1}{f_E} \right)^{\frac{1-\alpha-\beta}{\lambda}} (\beta \kappa)^\beta \left(\frac{1}{f_O} \right)^{\frac{(\lambda-1)(1-\alpha)+\beta}{\lambda}}.$$

Proposition 6 follows directly from the fact that $A_G = w/p_G$ and the solution for w/p_G from Proposition 4. The importance of Proposition 6 is that it shows that the manufacturing sector is characterized by an aggregate production function for the industrial good sector, where total productivity in industrial production A_G is fully determined from parameters: the productivity of lawyers A_{PS} (encapsulated in ς), the level of productivity A_M , the overhead cost f_O and the entry cost f_E . Note that A_G is continuous in ς and satisfies

$$\lim_{\varsigma \rightarrow \infty} A_G = \left(\frac{(\lambda-1)(1-\alpha)}{\beta + (1-\alpha)(\lambda-1)} \right)^{\frac{1-\alpha-\beta}{\lambda}} \kappa^\beta \alpha^\alpha \left(\frac{1-\alpha-\beta}{\lambda-1} \frac{1}{f_E} \right)^{\frac{1-\alpha-\beta}{\lambda}} \left(\frac{1-\alpha}{f_O} \right)^{\frac{(\lambda-1)(1-\alpha)+\beta}{\lambda}} A_M^{1-\alpha-\beta}$$




Figure A-2: Aggregate productivity A_G

and

$$A_G \left(\varsigma = \frac{\beta}{1-\alpha} \right) = \alpha^\alpha \kappa^\beta \left(\frac{1-\alpha-\beta}{\lambda-1} \frac{1}{f_E} \right)^{\frac{1-\alpha-\beta}{\lambda}} \left(\frac{1-\alpha}{f_O} \right)^{\frac{(\lambda-1)(1-\alpha)+\beta}{\lambda}} A_M^{1-\alpha-\beta}.$$

Figure A-2 depicts aggregate productivity A_G as a function of $\varsigma = \frac{\kappa}{A_{PS} f_O}$ for both cases discussed above.

A-3 Estimation of Productivity Fundamentals and Structural Parameters (Sections 5.2 and 5.1)

In this section we describe the details of our strategy to estimate the productivity fundamentals $\{A_{rst}\}$ and two structural parameters, ω_{CS} and ν_F .

Consider a single time period. Given the regional distribution of human capital and the sectoral distribution of earnings we can calculate $\{[w_r]_r, H_{rF}, H_{rG}, H_{rCS}\}_r$ in a model consistent way. In particular, the regional supply of human capital in location r is given by

$$H_{rt} = L_{rt} \sum_e \exp(\rho \times s) \ell_r(e).$$

where ρ is the return to education and $\ell_r(e)$ denotes the share of people in region r with e years of education.

We use the observed sectoral earnings shares to measure sectoral labor supplies. Specifically, for each sector s , we calculate

$$H_{rst} = \frac{\sum_i 1[i \in s] w_i}{\sum_i w_i} \times H_{rt}$$

where w_i is the wage of individual i . The average regional skill price w_r can be calculated as $w_r = (\sum_{i \in r} w_i) / H_{rt}$.

Given, this data we can then infer sectoral productivity in the following way:

Step 1: Estimating *relative* food productivity and *relative* manufacturing productivity It is useful to write productivities as

$$A_{rF} = A_F a_{rF} \text{ with } \sum_{r=1}^R a_{rF}^{\sigma-1} = 1$$

$$A_{rG} = A_G a_{rG} \text{ with } \sum_{r=1}^R a_{rG}^{\sigma-1} = 1.$$

Here A_s denotes the level of sectoral productivities and a_{rs} denotes the relative productivity of region r . Using the market clearing conditions in (8), it is easy to show that

$$a_{rF} = \left(\frac{H_{rF} w_r^\sigma}{\sum_{r=1}^R (H_{rF}) w_r^\sigma} \right)^{\frac{1}{\sigma-1}} \quad (\text{A-18})$$

$$a_{rG} = \left(\frac{H_{rG} w_r^\sigma}{\sum_{r=1}^R (H_{rG}) w_r^\sigma} \right)^{\frac{1}{\sigma-1}} \quad (\text{A-19})$$

Note a_{rF} and a_{rG} are insensitive to the level of w_r .

Step 2: Estimating A_F and A_G The two prices of the tradable goods are

$$p_{Ft} = \left(\sum_r \left(\frac{w_r}{A_{rF}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} = \frac{1}{A_F} \left(\sum_r \left(\frac{w_r}{a_{rF}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$$

$$p_{Gt} = \left(\sum_r \left(\frac{w_r}{A_{rG}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} = \frac{1}{A_G} \left(\sum_r \left(\frac{w_r}{a_{rG}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}.$$

Note that $\left(\sum_r \left(\frac{w_r}{a_{rF}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$ and $\left(\sum_r \left(\frac{w_r}{a_{rG}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$ are observable from (A-18) and (A-19). Hence, let us write

$$\Lambda_s^w \equiv \left(\sum_r \left(\frac{w_r}{a_{rs}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \quad (\text{A-20})$$

where Λ_s^w is known. The superscript “ w ” indicates that this object is homogenous of degree one in the level of wages. To determine A_F and A_G we use two restrictions:

1. First we choose the manufacturing good as the numeraire. This determines A_G because $p_{Gt} = 1$ implies that $A_G = \Lambda_G^w$. Note that an increase in w by a common factor increases A_G by the same amount.
2. Let the relative price of agricultural goods relative to manufacturing goods be p_t^{AG} . Then

$$p_t^{AG} = \frac{p_t^A}{p_t^G} = \frac{A_F^{-1} \Lambda_F^w}{A_G^{-1} \Lambda_G^w} = A_F^{-1} \Lambda_F^w.$$

Given p_t^{AG} , we can thus identify A_F .

Step 3: Estimating $\{A_{rCS}\}_{r=1}^R$ and the structural parameters ω_{CS} and ν_F To derive (A-18) and (A-19) we used $R - 1$ equations for each sector. This means that we still have R equations for the non-tradable CS sector

and two aggregate resource constraints for the tradable goods. These are

$$\frac{H_{rCS}}{H_r} = \omega_{CS} + \nu_{CS} (p_{At}^{\omega_A} p_{rCS}^{\omega_{CS}} p_{Gt}^{\omega_G})^\varepsilon (E_{rt}[q]w_{rt})^{-\varepsilon} \quad (\text{A-21})$$

$$\sum_r w_r H_{rF} = \sum_{j=1}^R \left(\omega_A + \nu_A \left(p_{At}^{\omega_A} p_{CSjt}^{\omega_{CS}} p_{Gt}^{\omega_G} \right)^\varepsilon (E_{rt}[q]w_{rt})^{-\varepsilon} \right) w_j H_j \quad (\text{A-22})$$

$$\sum_r w_r H_{rG} = \sum_{j=1}^R \left((1 - \omega_A - \omega_{CS}) - (\nu_A + \nu_{CS}) \left(p_{At}^{\omega_A} p_{CSjt}^{\omega_{CS}} p_{Gt}^{\omega_G} \right)^\varepsilon (E_{rt}[q]w_{rt})^{-\varepsilon} \right) w_j H_j \quad (\text{A-23})$$

Equation (A-23) is redundant, it is implied by (A-21) and (A-22) due to Walras' Law. Substituting the numeraire assumption $p_{Gt} = 1$, $p_{At} = p_t^{AG}$, where p_t^{AG} is the observed relative price and $p_{rCS} = \frac{w_{rt}}{A_{rCS}}$ yields

$$\frac{H_{rCS}}{H_r} = \omega_{CS} + \nu_{CS} (p_t^{AG})^{\varepsilon\omega_A} \left(\frac{w_{rt}}{A_{rCS}} \right)^{\varepsilon\omega_{CS}} (E_{rt}[q]w_{rt})^{-\varepsilon} \quad (\text{A-24})$$

$$\sum_r w_r H_{rF} = \omega_A \sum_{j=1}^R w_j H_j + \nu_A (p_t^{AG})^{\varepsilon\omega_A} \sum_{j=1}^R \left(\frac{w_{jt}}{A_{CSj}} \right)^{\varepsilon\omega_{CS}} (E_{rt}[q])^{-\varepsilon} w_{jt}^{-1-\varepsilon} H_j. \quad (\text{A-25})$$

For a given year these are $R+1$ equations in R productivities $\{A_{rCS}\}$ and 4 structural parameters $(\omega_{CS}, \omega_F, \nu_F, \nu_{CS})$ (recall that we take ε as given because we estimate it from the expenditure shares). If we have T years, we have $TR + 4$ unknowns and $TR + T$ equations. Thus, by insisting that preferences are constant over time, we add over-identifying restrictions to our analysis as we add additional years.

Note that (A-24) implies that

$$(p_t^{AG})^{\varepsilon\omega_A} \left(\frac{w_{rt}}{A_{rCS}} \right)^{\varepsilon\omega_{CS}} = -\frac{1}{\nu_{CS}} \left(\omega_{CS} - \frac{H_{rCS}}{H_r} \right) (E_{rt}[q]w_{rt})^\varepsilon.$$

Substituting this into (A-25) yields

$$\sum_r w_r H_{rF} = \omega_F \sum_{r=1}^R w_r H_r - \frac{\nu_F}{\nu_{CS}} \sum_{r=1}^R \left(\omega_{CS} - \frac{H_{rCS}}{H_r} \right) w_{rt} H_r. \quad (\text{A-26})$$

Given the data on $\{w_r, H_{rs}\}$, this is - for a given year - a single equation in three unknowns: ω_F , $\frac{\nu_F}{\nu_{CS}}$ and ω_{CS} . This leaves us with R equations for consumer service employment. From (A-24) it is apparent that ν_{CS} is not separately identified from the *level* of productivity in the consumer service sector: holding ω_{CS} and ε fixed, the data on wages w_{rt} and employment shares $\frac{H_{rCS}}{H_r}$ identifies $\nu_{CS} A_{rCS}^{-\varepsilon\omega_{CS}}$. Hence, under the assumption that consumer services are a luxury, we can wlog normalize $\nu_{CS} = -1$. For a given choice of ω_F we can therefore use (A-26) in 1987 and 2011 to uniquely solve for ω_{CS} and ν_F .

A-4 The Elasticity of Substitution

The Allen Uzawa elasticity of substitution between goods s and k is given by

$$EOS_{sk} = \frac{\frac{\partial^2 e(p,V)}{\partial p_s \partial p_k} e(p,V)}{\frac{\partial e(p,V)}{\partial p_s} \frac{\partial e(p,V)}{\partial p_k}}.$$

The expenditure function is given by

$$e(p, V) = \left(V + \sum_s \tilde{\nu}_s \ln p_s \right)^{1/\varepsilon} \varepsilon^{1/\varepsilon} \prod_{s \in \{F, G, CS\}} p_s^{\omega_s}.$$

This implies that

$$\begin{aligned} \frac{\partial e(p, V)}{\partial p_s} &= \left(V + \sum_s \tilde{\nu}_s \ln p_s \right)^{1/\varepsilon} \varepsilon^{1/\varepsilon} \prod_{s \in \{F, G, CS\}} p_s^{\omega_s} \left(\frac{\frac{1}{\varepsilon} \tilde{\nu}_s}{V + \sum_s \tilde{\nu}_s \ln p_s} + \omega_s \right) \frac{1}{p_s} \\ &= e(p, V) \left(\frac{\frac{1}{\varepsilon} \tilde{\nu}_s}{V + \sum_s \tilde{\nu}_s \ln p_s} + \omega_s \right) \frac{1}{p_s}. \end{aligned}$$

This also implies that

$$\begin{aligned} \frac{\partial^2 e(p, V)}{\partial p_s \partial p_k} &= \frac{\partial e(p, V)}{\partial p_k} \left(\frac{\frac{1}{\varepsilon} \tilde{\nu}_s}{V + \sum_s \tilde{\nu}_s \ln p_s} + \omega_s \right) \frac{1}{p_s} - e(p, V) \frac{\frac{1}{p_s} \frac{1}{\varepsilon} \tilde{\nu}_s \tilde{\nu}_k \frac{1}{p_k}}{(V + \sum_s \tilde{\nu}_s \ln p_s)^2} \\ &= e(p, V) \frac{1}{p_k} \frac{1}{p_s} \left\{ \left(\frac{\frac{1}{\varepsilon} \tilde{\nu}_k}{V + \sum_s \tilde{\nu}_s \ln p_s} + \omega_k \right) \left(\frac{\frac{1}{\varepsilon} \tilde{\nu}_s}{V + \sum_s \tilde{\nu}_s \ln p_s} + \omega_s \right) - \varepsilon \frac{\frac{1}{\varepsilon} \tilde{\nu}_s \frac{1}{\varepsilon} \tilde{\nu}_k}{(V + \sum_s \tilde{\nu}_s \ln p_s)^2} \right\}. \end{aligned}$$

Now note that

$$\begin{aligned} \frac{\frac{1}{\varepsilon} \tilde{\nu}_k}{V + \sum_s \tilde{\nu}_s \ln p_s} + \omega_k &= \tilde{\nu}_k \frac{1}{\varepsilon} \left(V + \sum_s \tilde{\nu}_s \ln p_s \right)^{-1} + \omega_k \\ &= \tilde{\nu}_k \left(\frac{e(p, V)}{\prod_{s \in \{F, G, CS\}} p_s^{\omega_s}} \right)^{-\varepsilon} + \omega_k = \vartheta_k. \end{aligned}$$

Hence,

$$\begin{aligned} \frac{\partial e(p, V)}{\partial p_s} &= e(p, V) \vartheta_s \frac{1}{p_s} \\ \frac{\partial^2 e(p, V)}{\partial p_s \partial p_k} &= e(p, V) \frac{1}{p_k} \frac{1}{p_s} \left\{ \vartheta_k \vartheta_s - \varepsilon \frac{\frac{1}{\varepsilon} \tilde{\nu}_s}{V + \sum_s \tilde{\nu}_s \ln p_s} \frac{\frac{1}{\varepsilon} \tilde{\nu}_k}{V + \sum_s \tilde{\nu}_s \ln p_s} \right\} \\ &= e(p, V) \frac{1}{p_k} \frac{1}{p_s} \{ \vartheta_k \vartheta_s - \varepsilon (\vartheta_s - \omega_s) (\vartheta_k - \omega_k) \}. \end{aligned}$$

This implies that

$$\begin{aligned} EOS_{sk} &= \frac{e(p, V) \frac{1}{p_k} \frac{1}{p_s} \{ \vartheta_k \vartheta_s - \varepsilon (\vartheta_s - \omega_s) (\vartheta_k - \omega_k) \} e(p, V)}{e(p, V) \vartheta_s \frac{1}{p_s} e(p, V) \vartheta_k \frac{1}{p_k}} \\ &= 1 - \varepsilon \frac{(\vartheta_s - \omega_s) (\vartheta_k - \omega_k)}{\vartheta_s \vartheta_k}. \end{aligned}$$

A-5 Open economy

In this model we present the formal analysis for the open economy extension discussed in Section 7.3.1.

A-5.1 Environment and Equilibrium

We assume that the consumption of the physical good of consumers in India is a combination of domestic and imported goods with a constant elasticity of substitution η :

$$C_G = \left(C_{G,D}^{\frac{\eta-1}{\eta}} + \varphi C_{G,ROW}^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}.$$

Here, $C_{G,D}$ and $C_{G,ROW}$ are the physical quantities of the domestic and imported physical good, φ is a taste parameter capturing the preference for the imported good, and η is the elasticity of substitution that we interpret as a trade elasticity.

Letting $p_{G,D}$ and $p_{G,ROW}$ denote the respective prices, the price index of the bundle C_G is given by

$$P_G = \left(p_{G,D}^{1-\eta} + \varphi^\eta p_{G,ROW}^{1-\eta} \right)^{\frac{1}{1-\eta}}. \quad (\text{A-27})$$

The expenditure share on Indian goods is $\frac{p_{G,D} C_{G,D}}{P_G C_G} = \left(\frac{P_{G,D}}{P_G} \right)^{1-\eta}$. Combining this expression with Equation (A-27) yields the expenditure shares

$$\begin{aligned} \frac{p_{G,D} C_{G,D}}{P_G C_G} &= \frac{\varphi^{-\eta} \left(\frac{P_{G,D}}{p_{G,ROW}} \right)^{1-\eta}}{1 + \varphi^{-\eta} \left(\frac{P_{G,D}}{p_{G,ROW}} \right)^{1-\eta}}, \\ \frac{p_{G,ROW} C_{G,ROW}}{P_G C_G} &= \frac{1}{1 + \varphi^{-\eta} \left(\frac{P_{G,D}}{p_{G,ROW}} \right)^{1-\eta}}. \end{aligned}$$

For simplicity we subsume trade costs in the relative price of foreign goods.

The Indian economy is assumed to export both domestic goods and a special category of services that is traded internationally: ICT exports. Consider first the export of goods. We model total spending on Indian goods (in term of domestic goods) from the rest of the world (ROW) as

$$X_{G,D} = \frac{\varphi^{-\eta} \left(\frac{P_{G,D}}{p_{G,ROW}} \right)^{1-\eta}}{1 + \varphi^{-\eta} \left(\frac{P_{G,D}}{p_{G,ROW}} \right)^{1-\eta}} \Upsilon_G,$$

i.e. $X_{G,D}$ are total exports from India, Υ_G is a demand shifter (for goods) and $p_{G,ROW}$ denotes the price of goods in the ROW. For simplicity we assume the price elasticity of exports and imports to be the same and equal to η .

Consider, next, the exported ICT services.² We assume that the ROW buys a bundle of regional varieties ICT services

$$Y_{ICTt} = \left(\sum_{r=1}^R (y_{rICTt})^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

where y_{rICTt} denotes the quantity of services produced in region r and exported to the rest of the world. ICT services are produced in region r according to the production function $y_{rICTt} = A_{rICTt} H_{rt}$. Hence, the price of ICT services is given by

$$p_{ICTt} = \left(\sum_r p_{rICTt}^{1-\sigma} \right)^{\frac{1}{1-\sigma}} = \left(\sum_r \left(\frac{w_{rt}}{A_{rICTt}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}.$$

² For simplicity, we assume that ICT services are not sold in the domestic market but only internationally.

As we do for goods, we model the import demand for ICT services as

$$X_{ICT} = p_{ICT}^{1-\eta} \Upsilon_{ICT}.$$

Again, any trade costs are subsumed in the demand shifter Υ_{ICT} .

Equilibrium

The equilibrium with trade is pinned down by the following equilibrium conditions:

1. Market clearing for agricultural goods:

$$w_r H_{rF} = \left(\frac{w_r^{1-\sigma} A_{rF}^{\sigma-1}}{\sum_{j=1}^R w_j^{1-\sigma} A_{jF}^{\sigma-1}} \right) \times \left(\sum_{j=1}^R \vartheta_{jF} w_j H_j \right)$$

2. Market clearing for manufacturing goods:

$$w_r H_{rG} = \left(\frac{w_r^{1-\sigma} A_{rG}^{\sigma-1}}{\sum_{j=1}^R w_j^{1-\sigma} A_{jG}^{\sigma-1}} \right) \times \underbrace{\left(\underbrace{\frac{\varphi^{-\eta} \left(\frac{(\sum_j w_j^{1-\sigma} A_{jG}^{\sigma-1})^{\frac{1}{1-\sigma}}}{p_{G,ROW}} \right)^{1-\eta}}{\varphi^{-\eta} \left(\frac{(\sum_j w_j^{1-\sigma} A_{jG}^{\sigma-1})^{\frac{1}{1-\sigma}}}{p_{G,ROW}} \right)^{1-\eta} + 1}}_{\text{Domestic spending}} \sum_{j=1}^R \vartheta_{jG} w_j H_j + \underbrace{\left(\sum_j w_j^{1-\sigma} A_{jG}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_G}_{\text{Total exports}} \right)}_{\text{Aggregate demand for physical goods}}$$

3. Market clearing for local CS:

$$w_r H_{rCS} = \vartheta_{rCS} w_r H_r.$$

4. Market clearing for local ICT services:

$$w_r H_{rICT} = \left(\frac{w_r^{1-\sigma} A_{rICT}^{\sigma-1}}{\sum_{j=1}^R w_j^{1-\sigma} A_{jICT}^{\sigma-1}} \right) \times \underbrace{\left(\sum_j w_j^{1-\sigma} A_{jICT}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICT}}_{\text{ICT exports}}$$

5. Labor market clearing:

$$H_{rF} + H_{rG} + H_{rCS} + H_{rICT} = H_r.$$

6. Balanced Trade:

$$\underbrace{\left(\left(\sum_j w_j^{1-\sigma} A_{jG}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_G + \left(\sum_j w_j^{1-\sigma} A_{jICT}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICT} \right)}_{\text{Exports}} = \underbrace{\frac{\sum_{j=1}^R \vartheta_{jG} w_j H_j}{\varphi^{-\eta} \left(\frac{(\sum_j w_j^{1-\sigma} A_{jG}^{\sigma-1})^{\frac{1-\eta}{1-\sigma}}}{p_{G,ROW}} \right)^{1-\eta} + 1}}_{\text{Imports}}$$

Letting $x \equiv \varphi^\eta p_{G,ROW}^{1-\eta}$ denote the (scaled) terms of trade, these are $5R + 1$ equations in $5R + 1$ unknowns $\{x, \{w_r, H_{rF}, H_{rG}, H_{rCS}, H_{rICT}\}_r\}$. Again, we can pick a numeraire

$$p_{G,IND} = \left(\sum_r \left(\frac{w_r}{A_{rG}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} = 1.$$

Given the productivities $\{A_{rF}, A_{rG}, A_{rCS}, A_{rICT}\}_r$, the population distribution $\{H_r\}_r$, the demand shifters of the foreign sector $(\Upsilon_{ICT}, \Upsilon_G)$ and the other preference parameters of the model, we can calculate $\{x, \{w_r, H_{rF}, H_{rG}, H_{rCS}, H_{rICT}\}_r\}$.

A-5.2 Identification of Productivity Fundamentals in the Open Economy Model

For the economy with trade we need to identify the following additional objects:

$$\left\{ [A_{rICT}]_{r=1}^R, \Upsilon_G, \Upsilon_{ICT} \right\}.$$

There are $R + 2$ unknowns. For these $R + 2$ unknowns we have the following conditions

1. Relative ICT payments across localities for ICT exports:

$$\frac{w_r H_{rICT}}{w_j H_{jICT}} = \frac{w_r^{1-\sigma} A_{rICT}^{\sigma-1}}{w_j^{1-\sigma} A_{jICT}^{\sigma-1}}$$

These are $R - 1$ equations to determine A_{rICT} up to scale, i.e.

$$A_{rICT} = A_{ICT} a_{rICT} \text{ with } \sum_r a_{rICT}^{\sigma-1} = 1$$

yields

$$a_{rICT} = \left(\frac{H_{rICT} w_r^\sigma}{\sum_j H_{jICT} w_j^\sigma} \right)^{\frac{1}{\sigma-1}}$$

Because the level of ICT productivity A_{ICT} is not separately identified from the aggregate demand shifter Υ_{ICT} , without loss of generality we can set $A_{ICT} = 1$.³

³ To see this, note that the equilibrium condition for ICT exports implies that

$$w_r H_{rICT} = \left(\frac{w_r^{1-\sigma} A_{rICT}^{\sigma-1}}{\sum_j w_j^{1-\sigma} A_{jICT}^{\sigma-1}} \right) \left(\sum_j w_j^{1-\sigma} A_{jICT}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICT} = \left(\frac{w_r^{1-\sigma} a_{rICT}^{\sigma-1}}{\sum_j w_j^{1-\sigma} a_{jICT}^{\sigma-1}} \right) \left(\sum_j w_j^{1-\sigma} a_{jICT}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} A_{ICT}^{\eta-1} \Upsilon_{ICT}$$

Hence, Υ_{ICT} and A_{ICT} are not separately identified.

2. To identify Υ_{ICT} we use that

$$\begin{aligned}\sum_r w_r L_{rICT} &= \sum_r \left(\frac{w_r^{1-\sigma} A_{rICT}^{\sigma-1}}{\sum_{j=1}^R w_j^{1-\sigma} A_{jICT}^{\sigma-1}} \right) \left(\sum_j w_j^{1-\sigma} A_{jICT}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICT} \\ &= \left(\sum_j w_j^{1-\sigma} A_{jICT}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICT}.\end{aligned}\tag{A-28}$$

The RHS is total value added of the ICT sector, which we can calculate directly in the data. Given that w_j and A_{jICT} is observed, we can calculate Υ_{ICT} .

3. To identify Υ_G we use a moment about the share of manufacturing value added that is exported. Our model implies that:

$$\text{Total value added in manufacturing} = \sum_r w_r H_{rG}$$

and

$$\text{Total value added of exports} = \left(\sum_j w_j^{1-\sigma} A_{jG}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_G$$

Hence, the share of value added in the manufacturing sector is

$$M_1 = \frac{\left(\sum_j w_j^{1-\sigma} A_{jG}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_G}{\sum_r w_r H_{rG}} = \frac{P_{G,IND}^{1-\eta} \Upsilon_G}{\sum_r w_r H_{rG}} = \frac{\Upsilon_G}{\sum_r w_r H_{rG}}\tag{A-29}$$

Hence, for a given moment of the export share of manufacturing M_1 and data on $\{w_j, L_{jG}\}_j$ we can solve for Υ_G .

A-6 The Model with Imperfect Skill Substitution

In Section 7.3.2 we extended our analysis to a more general production function, where high and low skilled workers are imperfect substitutes. In this section we describe the details of this exercise.

A-6.1 Environment and Equilibrium

Suppose that the technology in sector s in region r is given by

$$Y_{rs} = A_{rs} \left((H_{rs}^-)^{\frac{\rho-1}{\rho}} + (Z_{rs} H_{rs}^+)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}},$$

where A_{rs} denotes factor neutral productivity, Z_{rs} denotes the skill bias and H_{rs}^- (H_{rs}^+) are the quantities of human capital of low (high) skilled individuals. Again we assume that individuals are heterogenous. Specifically, people of skill type $j \in \{-, +\}$ draw their efficiency level from a pareto with the same shape, i.e.

$$P\left(q_i^j \leq k\right) = 1 - \left(\frac{q_{rt}^j}{k}\right)^\zeta \equiv F_{rt}^j(k).$$

Total income of an individual i of skill type j in region r at time t is therefore given by $y_{rt}^{i,j} = w_{rt}^j q_i^j$, where the skill price w_{rt}^j is now skill-specific. The aggregate expenditure share on goods from sector s goods in region r is then

given by

$$\vartheta_{rst} \equiv \frac{L_{rt}^- \int \vartheta_s^h(qw_{rt}^-, p_{rt}) qw_{rt}^- dF_{rt}^-(q) + L_{rt}^+ \int \vartheta_s^h(qw_{rt}^+, p_{rt}) qw_{rt}^+ dF_{rt}^+(q)}{L_{rt}^- \int qw_{rt}^- dF_{rt}^-(q) + L_{rt}^+ \int qw_{rt}^+ dF_{rt}^+(q)},$$

where $\vartheta_s^h(qw_{rt}^-, p_{rt})$ denotes the sectoral expenditure share at the individual level. Substituting the expression for $\vartheta_s^h(qw_{rt}^-, p_{rt})$ and using the fact that $y_{rt}^{i,j}$ is also pareto distributed yields

$$\vartheta_{rst} = \omega_s + \tilde{\nu}_s \frac{\zeta - 1}{\zeta - (1 - \varepsilon)} \left(\frac{1}{\prod_s p_s^{\omega_s}} \right)^{-\varepsilon} \left(s_{rt}^{Y,-} \left(w_{rt}^- q_{rt}^- \right)^{-\varepsilon} + \left(1 - s_{rt}^{Y,+} \right) \left(w_{rt}^+ q_{rt}^+ \right)^{-\varepsilon} \right).$$

where $s_{rt}^{Y,-} = \frac{L_{rt}^- w_{rt}^- q_{rt}^-}{L_{rt}^- w_{rt}^- q_{rt}^- + L_{rt}^+ w_{rt}^+ q_{rt}^+}$ is the income share of low skilled individuals in region r at time t . Hence, the sectoral expenditure share is given by

$$\vartheta_{rst} = \vartheta_s \left(q_{rt}^- w_{rt}^-, q_{rt}^+ w_{rt}^+, s_{rt}^{Y,-}, \mathbf{p}_{rt} \right),$$

i.e. sectoral spending varies at the regional level because of (i) differences in regional factor prices w_{rt}^- and w_{rt}^+ , (ii) differences in the prices of non-tradable goods p_{rCst} and (iii) differences in the skill composition $s_{rt}^{Y,-}$.

Equilibrium The equilibrium is characterized by the following conditions. The CES structure and perfect competition imply that prices are given by

$$p_{rs} = \frac{1}{A_{rs}} \left((w_{rt}^-)^{1-\rho} + Z_{rs}^{\rho-1} (w_{rt}^+)^{1-\rho} \right)^{\frac{1}{1-\rho}}.$$

The relative skill demand for sector s in region r is given by

$$\frac{w_{rt}^+ H_{rst}^+}{w_{rt}^- H_{rst}^-} = Z_{rt}^{\rho-1} \left(\frac{w_{rt}^+}{w_{rt}^-} \right)^{1-\rho}.$$

The CES demand system across regional varieties implies the market clearing conditions

$$w_{rt}^- H_{rst}^- + w_{rt}^+ H_{rst}^+ = \left(\frac{p_{rs}}{p_s} \right)^{1-\sigma} \times \sum_{j=1}^R \vartheta_s \left(q_{jt}^- w_{jt}^-, q_{jt}^+ w_{jt}^+, s_{jt}^{Y,-}, \mathbf{p}_{jt} \right) \bar{w}_{rt} L_{rt},$$

where $\bar{w}_{rt}$ denotes average income and $p_s = \left(\sum_{r=1}^R p_{rs}^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$. The market clearing condition for non-tradable CS implies

$$w_{rt}^- H_{rCst}^- + w_{rt}^+ H_{rCst}^+ = \vartheta_{CS} \left(q_{jt}^- w_{jt}^-, q_{jt}^+ w_{jt}^+, s_{jt}^{Y,-}, \mathbf{p}_{jt} \right) \bar{w}_{rt} L_{rt}. \quad (\text{A-30})$$

Finally, labor market clearing implies

$$H_{rF}^j + H_{rG}^j + H_{rCS}^j = H_r^j \text{ for } j \in \{-, +\}.$$

These equations uniquely determine the regional wages $\{w_{rt}^-, w_{rt}^+\}$ and the sectoral labor allocations $\{H_{rst}^-, H_{rst}^+\}$.

A-6.2 Measurement and Equilibrium Accounting

As before we use these equations and the observable data to infer the productivity vector $\{A_{rst}, Z_{rst}\}$ for each region-sector pair. To connect our data to the objects in the model, we make the following measurement choices:

1. We classify individuals into high and low skill workers by their years of schooling. We assume workers with at least secondary schooling are high skill workers. In Figure A-3 we show the share of high skilled employment




Figure A-3: SHARE OF HIGH SKILLED EMPLOYMENT BY URBANIZATION. The figure shows the share of employment with at least secondary schooling for different quantiles of urbanization.

as a function of the urbanization rate. In rural regions, only 20% of workers are high-skilled. In cities, this share is twice as large.

2. As in our baseline model, we assume a Mincerian return $\rho = 5.6\%$ per year of schooling within skill groups. This allows us to measure the aggregate skill supplies H_{rt}^- and H_{rt}^+ for each region.
3. As in our baseline model, we use the observed sectoral earnings shares by skill group to measure sectoral labor supplies. Specifically, for each skill group $j = \{-, +\}$ and sector s , we calculate

$$H_{rst}^j = \frac{\sum_i 1[i \in j \text{ and } i \in s] w_i}{\sum_i 1[i \in j] w_i} \times H_{rt}^j$$

where w_i is the wage of individual i .

4. We then calculate the regional skill prices as $w_r^j = \frac{1}{L_{rt}^j} \sum_{i=1}^{L_{rt}^j} y_{rti}^j$ where y_{rti}^j denotes total income of individual i in region r at time t in skill group j .

These data are sufficient to uniquely solve for $\{A_{rst}, Z_{rst}\}$ and to perform the counterfactual analysis reported in Section 7.3.2.

APPENDIX B: DATA AND MEASUREMENT

In this section, we discuss the data and empirical issues discussed in the text.

B-1 Data Sources

Our analysis relies on four data sets:

1. The national sample survey (NSS)
2. The Economic Census (EC)
3. The Service Sector in India: 2006-2007
4. The Informal Non-Agricultural Enterprises Survey 1999-2000 (INAES)
5. The Household Expenditure survey

In this section we describe these datasets in detail.

B-1.1 National Sample Survey (NSS)

The National Sample Survey (NSS) is a representative survey conducted by the Government of India to collect socio-economic data at the household level since 1950. Each round of the survey consists of several schedules that cover different topics like “consumer expenditure”, “employment and unemployment”, “participation in education”, etc. We focus on the “consumer expenditure” module and the “employment and unemployment” module and use data from rounds 43, 55, 60, 64, 66 and 68 of NSS, which span the years 1987 to 2011. The survey covers the entirety of India except a few regions due to unfavorable field conditions.⁴

We use the “employment and unemployment” module to measure sectoral employment shares and total earnings. An individual is defined as being employed if his/her usual principal activity is one of the following: (i) worked in household enterprises (self-employed); (ii) worked as helper in household enterprises; (iii) worked as regular salaried/wage employee; (iv) worked as casual wage labour in public works; (v) worked as casual wage labour in other types of work. We describe the details of our sectoral employment classification in Section B-2 below.

As our measure of income we focus on total expenditure. More specifically, we measure total household expenditure and divide it by household size. We do so to properly account for the relative income of self-employed of informally employed employees. In the main analysis, we winsorize the expenditure data at 98th percentiles to account for measurement error.

As we describe in more detail in Section B-1.5, the NSS provides two measures of expenditure. The so-called uniform reference period (URP) measure simply measures total expenditure as expenditure within the last 30 days. The mixed reference period (MRP) measure asks respondents for the total expenditure within the last year for a subset of durable goods to account for the lumpiness of purchases. For all years except 1987, expenditure is reported using the MRP classification. To make the results comparable across years, we merge the expenditure module (described in Section B-1.5) with the employment module in 1987 at the household-level and use the MRP measure contained in the expenditure module. In practice, this choice is inconsequential because the different measures are highly correlated. In Table B-1 we report the correlation between the monthly per capita expenditure (MPCE) measure reported in the employment module, the MPCE URP measure reported in the expenditure module and the URP and MRP measures after winsorizing. This correlation exceeds 0.9 for all measures and our results do not hinge on which measure we use for 1987.

To measure human capital, we utilize information on educational attainment. We classify individual’s education into four levels: (i) less than primary; (ii) primary, upper primary, and middle; (iii) secondary; (iv) more than

⁴ For example, Ladakh and Kargil districts of Jammu & Kashmir, some interior villages of Nagaland, and villages in Andaman and Nicobar Islands are not covered in some rounds of the survey.

	MPCE Employment module	MPCE Expenditure module	MPCE.URP Authors' calculation	MPCE.MRP Authors' calculation
MPCE (Employment module)	1			
MPCE.URP (Expenditure module)	0.968	1		
MPCE.URP (authors' calculation)	0.967	0.998	1	
MPCE.MRP (authors' calculation)	0.916	0.939	0.941	1

Table B-1: CORRELATION MATRIX OF DIFFERENT EXPENDITURE MEASURES. The table shows the correlation between household expenditure reported in the NSS employment schedule, NSS expenditure schedule, and calculated by authors. We trim the top 1% and bottom 1% of observations.

secondary. We then associate different years of schooling to each category to estimate annual returns. Building on the official classification in India, we attribute 0, 3, 6, and 9 years respectively.

The “consumer expenditure” module collects information on households’ consumption of various kinds of food, entertainment, sundry articles, consumer services and housing expenses during last 30 days and consumption of clothing, bedding, footwear, education and medical goods and services, and various durable goods during last 365 days. We measure total monthly household consumer expenditure as the sum of all monthly-based expenditures and 30/365 of yearly-based expenditure.

In Table B-2 we report the summary statistics about the sample size of the NSS in the different years. Depending on the year, our data comprises between 60,000 and 120,000 household and between 300,000 and 600,000 individuals.

Round	Year	Households	Individuals
43	1987-1988	126,353	654,903
55	1999-2000	107,215	596,688
60	2004	59,042	303,233
64	2007-2008	125,578	572,254
66	2009-2010	100,957	459,784
68	2011-2012	101,717	456,970

Table B-2: NATIONAL SAMPLE SURVEY: SUMMARY STATISTICS.

B-1.2 Economic Census

The India Economic Census (EC) is a complete count of all establishments, i.e. production units engaged in production or distribution of goods and services not for the purpose of sole consumption, located within the country. The Censuses were conducted in the years 1977, 1980, 1990, 1998, 2005, 2013, 2019. The micro-level data in 1990, 1998, 2005, 2013 are publicly available.

The EC collects information such as firms’ location, industry, ownership, employment, source of financing and the owner’s social group. It covers all economic sectors excluding crop production and plantation. The EC in 2005 and 2013 exclude some public sectors like public administration, defense and social security. In terms of geography, the EC covers all States and Union Territories of the country except for the year 1990, which covers all states except Jammu and Kashmir.

In Table B-3 we report some summary statistics of the EC in various years. In the most recent year, 2013, the EC has information on almost 60m firms. As expected, the majority of such firms is very small. Average employment is around 2 and 55% firms have a single employee. The share of firms with more than 100 employees is 0.06%.

Year	Number of firms	Total employment	Employment distribution			
			Avg.	1	empl. < 5	> 100
1990	24216790	74570280	3.08	53.77%	91.24%	0.13%
1998	30348881	83308504	2.75	51.18%	91.71%	0.11%
2005	41826989	100904120	2.41	55.76%	93.17%	0.12%
2013	58495359	131293872	2.24	55.47%	93.44%	0.06%

Table B-3: THE ECONOMIC CENSUS: SUMMARY STATISTICS. The table report the number of firms, total employment, average employment and the share of firms with 1, less than 5 and more than 100 employees.

B-1.3 Service Sector in India: 2006-2007

The Service Sector in India (2006-2007) dataset is part of an integrated survey by the NSSO (National Sample Survey Organisation) in its 63rd round. In the 57th round (2001-2002), the dataset was called “Unorganized Service Sector”. With the inclusion of the financial sector and large firms, the dataset was renamed as “Service Sector in India” and is designed to be representative for India’s service sector. In Table B-4 we compare this Service Survey with the Economic Census for a variety of subsectors within the service sector. Table B-4 shows that the service survey is consistent with the EC, i.e. average firm size and the share of firms with less than 5 employees are quite comparable in most subsectors.

The Service Survey covers a broad range of service sectors, including hotels and restaurants (Section H of NIC 04); transport, storage and communication (I); financial intermediation (J); real estate, renting and business activities (K); education (M); health and social work (N) and other community, social and personal service activities (O). Excluded are the following subsectors: railways transportation, air transport, pipeline transport; monetary intermediation (central banks, commercial banks, etc); trade unions; government and public sector enterprises and firms that appeared in the Annual Survey of Industries frame (ASI 2004-2005). In terms of geography, the survey covers the whole of Indian Union except for 4 districts and some remote villages.⁵ The survey was conducted in a total number of 5573 villages and 7698 urban blocks. A total of 190,282 enterprises were ultimately surveyed.

For our analysis we use two pieces of information: the number of employees and whether the main customer is another firm or a household.

B-1.4 Informal Non-Agricultural Enterprises Survey 1999-2000 (INAES)

We use this dataset to allocate employment in the construction sector to either consumer or producer services. The Informal Non-Agricultural Enterprises Survey is part of the 55th survey round of the NSSO. It covers all informal enterprises in the non-agricultural sector of the economy, excluding those engaged in mining, quarrying and electricity, gas and water supply.⁶ The Informal Non-Agricultural Enterprises Survey collects information on operational characteristics, expenses, value-added, fixed asset, loan and factor income. For our analysis we use two pieces of information: the number of employees and whether the main customer is another firm or a household.

B-1.5 Household Expenditure Survey

To estimate the expenditure elasticity ε we rely on data on consumer expenditure. This data is contained in the National Sample Survey, Round 68, Schedule 1.0. The dataset reports detailed information on a large set of

⁵ The survey covered the whole of India except (i) Leh (Ladakh), Kargil, Punch and the Rajauri districts of Jammu & Kashmir, (ii) interior villages situated beyond 5 km of a bus route in Nagaland, (iii) villages of the Andaman and Nicobar Islands, that remain inaccessible throughout the year.

⁶ The organized sector comprises all factories registered under Sections 2(m)(i) and 2(m)(ii) of the Factories Act, 1948, where 2(m)(i) includes manufacturing factories which employ 10 or more workers with electric power, and 2(m)(ii) includes manufacturing factories which 20 or more without electric power. The unorganized sector comprises all factories not covered in the organized sector. The informal sector is a subset of the unorganized sector. The unorganized sector includes four types of enterprises: (i) unincorporated proprietary enterprises; (ii) partnership enterprises; (iii) enterprises run by cooperative societies, trusts, private; (iv) public limited companies. The informal sector only includes firms in category (i) and (ii).

NIC'2004	Sector	Number of firms		Average employment		Less than 5 employees	
		EC	Service Survey	EC	Service Survey	EC	Service Survey
55	Hotels and restaurants	1499101	30744	2.52	2.49	90%	91%
60	Land transport; transport via pipelines	1317904	41065	1.67	1.24	97%	99%
61	Water transport	7914	174	4.35	1.92	0.90	0.98
63	Transport activities; travel agencies	188474	2101	3.40	3.33	86%	85%
64	Post and telecommunications	723119	22885	2.06	1.41	96%	99%
65-67	Financial intermediation	293489	16331	5.61	3.81	69%	82%
70	Real estate activities	70128	3648	2.18	1.64	93%	96%
71	Renting of machinery and household goods	365246	5387	2.00	1.77	94%	97%
72	Computer and related activities	66414	1060	6.01	13.45	83%	86%
73	Research and development	2097	5	16.66	4.58	66%	89%
74	Other business activities	519696	10610	2.81	1.92	90%	95%
85	Health and social work	783644	11930	3.39	1.99	88%	95%
91	Activities of membership organizations	1002996	2837	1.82	1.32	94%	98%
92	Recreational, cultural and sporting activities	222061	2698	2.95	2.91	85%	82%
93	Other service activities	1419685	26132	1.74	1.54	97%	99%

Table B-4: ECONOMIC CENSUS AND SERVICE SURVEY. The table reports statistics about firms' number and employment from the Economic Census 2005 and Service Survey 2006.

spending categories. In Table B-5 we report the broad classifications. The data also contains a finer allocation of spending within each category. For the purpose of this paper, we rely only on the classification in Table B-5.

No.	Description	No.	Description	No.	Description
1	Cereals	13	Served processed food	25	Conveyance
2	Cereal substitute	14	Packaged processed food	26	Rent
3	Pulses and products	15	Pan	27	Consumer taxes and cesses
4	Milk and milk products	16	Tobacco	28	Sub-total (1 – 27)
5	Salt and sugar	17	Intoxicants	29	Clothing
6	Edible oil	18	Fuel and light	30	Bedding
7	Egg, fish and meat	19	Medical (non-institutional)	31	Footwear
8	Vegetables	20	Entertainment	32	education
9	Fruits (fresh)	21	Minor durable-type goods	33	Medical (institutional)
10	Fruits (dry)	22	Toilet articles	34	Durable goods
11	Spices	23	Other household consumables	35	Sub-total (29 – 34)
12	Beverages	24	Consumer services excl. conveyance		

Table B-5: BROAD CLASSIFICATION OF NSS EXPENDITURE SURVEY. The table report the classification of broad expenditure items in the Expenditure Survey.

We classify consumers' spending on food as categories 1 - 17. We classify spending on consumer services as all spending in the consumer service category (category 24) and on entertainment (category 20). In Tables B-6 and B-7 we report the more detailed classification of the consumer service and entertainment spending categories.

Spending on category c is measured as spending within a particular reference period. For all categories, subject report total spending during the last 30 days. For durable goods, and medical and educational spending (i.e. categories 29-34), the subjects additionally report total spending in the last year. This second concept of expenditure aims to account for the lumpiness of purchases. For this group we therefore take 1/12 of annual spending as our measure of monthly expenditure. We measure total spending as the sum of all spending across all categories to calculate the spending share on food and consumer services.

In Table B-8 we report a selected set of summary statistics for the main variables of interest. In total we have expenditure data for slightly more than 100,000 households. In the first two rows we show the distribution of household expenditure for the case of measuring durable spending at the monthly frequency (the uniform reference

No.	Description	No.	Description
480	Domestic servant/cook	490	Postage and telegram
481	Attendant	491	Miscellaneous expenses
482	Sweeper	492	Priest
483	Barber, beautician, etc.	493	Legal expenses
484	Washerman, laundry, ironing	494	Repair charges for non-durables
485	Tailor	495	Pet animals (incl. birds, fish)
486	Grinding charges	496	Internet expenses
487	Telephone charges: landline	497	Other consumer services excluding conveyance
488	Telephone charges: mobile		

Table B-6: EXPENDITURE ITEMS WITHIN CONSUMER SERVICES. This table reports the detailed expenditure items within the category consumer services (category 24 in Table B-5)

No.	Description	No.	Description
430	Cinema, theatre	435	Photography
431	Mela, fair, picnic	436	VCD/ DVD hire (incl. instrument)
432	Sports goods, toys, etc.	437	Cable TV
433	Club fees	438	Other entertainment
434	Goods for recreation and hobbies		

Table B-7: EXPENDITURE ITEMS WITHIN ENTERTAINMENT. This table reports the detailed expenditure items within the category entertainment (category 20 in Table B-5)

period *URP*) and or at the annual frequency (the mixed reference period *MRP*). Table B-8 shows that the dispersion in spending is much higher for *URP* case, especially in the right tail. We therefore use the *MRP* measure as our measure of total expenditure.

Table B-8 also reports a set of statistics for the distribution of food shares and consumer service spending shares. The full distribution is shown in Figure B-1. There is ample cross-sectional dispersion. Through the lens of our theory, this dispersion is generated through heterogeneity in income and relative prices.

	N	mean	sd	min	median	p90	p95	max
Household expenditure (<i>URP</i>)	101,662	8,226	12,784	40	6,264	14,475	19,081	1,239,930
Household expenditure (<i>MRP</i>)	101,662	8,316	7,438	44	6,572	14,960	19,433	339,832
Household size	101,662	4.57	2.25	1	4	7	9	39
Food expenditure share	101,662	0.49	0.13	0	0.50	0.64	0.68	1
CS expenditure share	101,662	0.06	0.04	0	0.06	0.11	0.14	0.67

Table B-8: NSS EXPENDITURE SURVEY: SUMMARY STATISTICS. The table reports selected summary statistics from the NSS expenditure survey.

For our regression analysis reported in Tables 3 and 5 we control for additional household level of covariates. We control for the size of the household and the number of (potential) workers in the household which we define as all individuals between age 15 and 65. We also control for additional household demographics, namely

- the type of the household, which for rural areas is one of (i) self-employed in agriculture, (ii) self-employed in non-agriculture, (iii) regular wage/salary earner, (iv) casual worker in agriculture and (v) casual worker in non-agriculture, (vi) other and in urban areas one of (i) self-employed (ii) regular wage/salary earner, (iii) casual worker, (vi) other;




Figure B-1: DISTRIBUTION OF FOOD AND CONSUMER SERVICE EXPENDITURE SHARES. The figure shows the unconditional distribution of the expenditure shares for food (left panel) and consumer services (right panel).

- the household's religion, which is one of Hinduism, Islam, Christianity, Sikhism, Jainism, Buddhism, Zoroastrianism and other;
- the household's social group, which is one of the following: scheduled tribe, scheduled case, backward class and other.

Finally, the survey also reports whether the household is eligible to receive a rationing card.

B-2 Classification of Industries

At the heart of our analysis is the sectoral composition of regional employment. In our theory we distinguish between four sectors: Agriculture, Manufacturing, Consumer Services and Producer Services. To map these concepts to sectors in the data, we first construct six broad industries (see Section B-2.1). In a second step we then attribute employment in services and construction to consumer and producer services respectively - see Section B-2.2.

B-2.1 Broad Industry Classification

We initially divide economic activities into six industries:

1. Agriculture
2. Manufacturing
3. Construction and Utilities
4. Services
5. Information and Communications Technology (ICT)
6. Public Administration and Education.

To do so we rely on India's official classification system National Industrial Classification (NIC). We report our classification of industries in Table B-9.

Because the NIC classification system changes over time, we construct a concordance table between 2-digit industries of different versions of the NIC based on official NIC documents and detailed sector descriptions. This

Industry	NIC 2008	Description
Agriculture	01-03	Agriculture, forestry and fishing
Manufacturing	05-09	Mining of coal and lignite
	10-33	Manufacturing
Construction & Utilities	35	Electricity, gas, steam and air conditioning supply
	36-39	Water supply; sewerage, waste management and remediation activities
	41-43	Construction
Services	45-47	Wholesale and retail trade; repair of motor vehicles and motorcycles
	49-53	Transportation and storage
	55-56	Accommodation and Food service activities
	581	Publishing of books, periodicals and other publishing activities
	64-66	Financial and insurance activities
	68	Real estate activities
	69-75	Professional, scientific and technical activities
	77-82	Administrative and support service activities
	86-88	Human health and social work activities
	90-93	Arts, entertainment and recreation
	94-96	Other service activities
	97	Activities of households as employers of domestic personnel
ICT	582-63	Information and communication
Public Administration	84	Public administration and defence; compulsory social security
&	85	Education
Education	99	Activities of extraterritorial organizations and bodies

Table B-9: INDUSTRIAL CLASSIFICATION. The table reports the industrial classification into six broad sectors.

concordance system allows to compare sectoral employment patterns over time. Our crosswalk is reported in Table B-10.

B-2.2 Attributing Employment to Producer and Consumer Services

Our theory highlights the difference between PS, that are inputs in the production of goods, and CS, that are bought directly by consumers and are luxuries. In terms of the data, see Table B-9, financial and insurance activities (NIC codes 64-66) are examples of the former while retail trade (NIC codes 45-47) is an example of the latter.

In order to arrive at a systematic classification, we rely on the Service Survey (see Section B-1.3) that reports the identify of the main *buyer* of a given firm. We therefore refer to firms that sell to other firms as PS-firms and firms that sell to consumers as CS-firms.

Ideally, we would calculate the employment share of PS-firms in each subsector of the service sectors and in each region. The regional variation is important because our theory stresses that CS and PS productivity varies at the regional level. Given the large number of regions and subsection, the sample size of the Service Survey is not sufficiently large to estimate these averages precisely.

We therefore generate the regional variation in employment shares by using regional variation in the firm-size distribution and differences in the employment share of PS firms by firm size. Empirically, large firms are - within their subsector - much more likely to sell to firms. To see this, consider Figure B-2 where we depict the employment share of PS firms as a function of firm size in the raw data. In Table B-11 we show that the same pattern is presents within two- and three-digit industries and whether or not we use sampling weights. In particular, we regress a dummy variable for whether the firm sells mainly to other firms on different firm size dummies. The coefficients are generally positive and increasing.

To exploit this size-dependence, we adopt the following procedure:

1. For each 2-digit subsector k within the service sector listed in Table B-9 and size bin b we calculate the

sector	NIC-1987	NIC-1998 & NIC-2004	NIC-2008
Agriculture			
Agriculture and hunting	00-04	01	01
Forestry and logging	05	02	02
Fishing and aquaculture	06	05	03
Manufacturing			
Coal, lignite, and peat	10	10	05, 0892
Crude petroleum and natural gas	11,19	11	06, 091
Metal ores	12, 13, 14	12,13	07
Other mining and quarrying	15	14	08(except0892), 099
Food products	20,21, 220-224	15	10, 11
Tobacco products	225-229	16	12
Textiles and wearing apparel	23-24	17, 18	13, 14
Leather products	29(except 292)	19	15
Wood products	27(except 276-277)	20	16
Paper products, printing and publishing	28	21, 22	17, 18, 581
Refined petroleum	314-319	23	19
Chemicals	30	24	20, 21
Rubber and plastics products	310-313(except3134)	25	22
Other non-metallic mineral products	32	26	23
Basic metals	33(except338)	27	24
Fabricated metal	34(except342), 352, 391	28, 2927	25, 3311
Machinery and equipment	35-36(except352), 390, 392, 393, 395, 396, 399	29-32 (except2927)	261-264, 268, 27, 28, 3312, 3314, 3319, 332, 9512
Medical, precision and optical instruments	380-382	33	265-267, 325, 3313
Transport equipment	37, 397	34, 35	29, 30, 3315
Furniture	276, 277, 3134, 342	361	31
Other manufacturing	383-389	369	32(except325)
Construction & Utilities			
Electricity, gas, steam supply	40, 41, 43	40	35
Water supply	42	41	36
Sewerage and waste treatment	338, 6892, 91	37,90	37, 38, 39
Construction	50, 51	45	41, 42, 43
Services			
Wholesale	398, 60-64, 682, 686, 890, 974	50, 51(except51901)	45, 46
Retail	65-68(except682,686,6892)	52(except526,52591)	47
Repair services	97(except974)	526	952
Land transport	70	60	49
Water transport	71	61	50
Air transport	72	62	51
Supporting and auxiliary transport activities	730, 731, 732, 737, 738, 739, 74	63	52, 79
Post and telecommunications	75	64	53, 61
Hotels	691	551	55
Restaurants	690	552	56
Computer and related activities	394, 892, 897	72, 922	582, 62, 63, 9511
Financial service	80	65, 67	64, 66
Insurance and pension	81	66	65
Real estate activities	82	70	68
Legal activities	83	7411	691
Accounting	891	7412	692
Business and management consultancy	893	7413, 7414	70, 732
Architecture and engineering	894, 895	742	71
Research and development	922	73	72
Advertising	896	743	731
Other business activities	898, 899	749	74, 78, 80, 81, 82
Renting	733, 734, 735, 736, 85	71	77
Health and social work	93, 941	85	75, 86, 87, 88
Recreational cultural and sporting activities	95	92(except922)	59, 60, 90, 91, 93
Gambling	84	51901, 52591	92
Membership organizations	94(except941)	91	94
Personal service	96, 99	93, 95	96, 97
goods-producing activities for own use	#N/A	96	981
services-producing activities for own use	#N/A	97	982
Public Administration & Education			
Public administration and defence	90	75	84
Education	920-921	80	85
Extraterritorial organizations	98	99	99

Table B-10: CONCORDANCE BETWEEN 2-DIGIT INDUSTRY CLASSES. The table report the classification of NIC codes in different years to the broad sectoral categories of Table B-9.




Figure B-2: PRODUCER SERVICE SHARE BY FIRM SIZE. The figure shows the share of service firms whose main customer are other firms (as opposed to private individuals) with breakdown by firm size.

Probability of selling to firms				
2 employees	0.013*** (0.001)	0.014*** (0.002)	0.014*** (0.001)	0.016*** (0.002)
3 employees	0.030*** (0.002)	0.028*** (0.006)	0.028*** (0.002)	0.029*** (0.005)
4 employees	0.055*** (0.004)	0.063*** (0.011)	0.049*** (0.004)	0.059*** (0.011)
5 employees	0.080*** (0.006)	0.074*** (0.011)	0.070*** (0.006)	0.072*** (0.010)
6-10 employees	0.090*** (0.005)	0.062*** (0.007)	0.080*** (0.005)	0.057*** (0.007)
11-20 employees	0.085*** (0.006)	0.042*** (0.008)	0.074*** (0.006)	0.039*** (0.008)
21-50 employees	0.192*** (0.016)	0.106*** (0.026)	0.164*** (0.016)	0.099*** (0.025)
more than 50 employees	0.345*** (0.023)	0.159*** (0.044)	0.304*** (0.022)	0.137*** (0.034)
Industry FE (2 digit)	Yes	Yes		
Industry FE (3 digit)			Yes	Yes
Sampling weights	No	Yes	No	Yes
N	173743	173743	173743	173743
R ²	0.100	0.077	0.133	0.104

Table B-11: CORPORATE CUSTOMERS AND FIRM SIZE. *Notes:* Columns 1 and 2 (3 and 4) control for two (three) digit industry fixed effects. Columns 2 and 4 weigh each observation by the sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

employment share of PS firms as

$$\omega_{kb}^{PS} = \frac{\sum_{f \in (k,b)} 1\{f \in PS\} l_f}{\sum_{f \in (k,b)} l_f}.$$

Here, f denotes a firm, $1\{f \in PS\}$ is an indicator that takes the value 1 if firm f is a PS firm and l_f denotes firm employment. In practice we take three size bins, namely “1 or 2 employees”, “3-20 employees” and “more than 20” employees. We always weigh observations with the sampling weights provided in the Service Survey.⁷

2. We then use the Economic Census (see Section B-1.2) and calculate the share of employment of firms in size bin b in subsector k in region r as

$$\ell_{kbr} = \frac{\sum_{f \in (k,b,r)} l_f}{\sum_{f \in (k,r)} l_f}$$

3. We then combine these two objects to calculate the share of employment of PS firms in region r in subsector k as

$$s_{rk}^{PS} = \sum_b \ell_{kbr} \omega_{kb}^{PS}$$

4. Finally, we use s_{rk}^{PS} to calculate the share of employment in PS and CS in region r as

$$\varpi_r^{PS} = \frac{\sum_k s_{rk}^{PS} l_{rk}^{NSS}}{l_{rk}^{NSS}} \quad \text{and} \quad \varpi_r^{CS} = \frac{\sum_k (1 - s_{rk}^{PS}) l_{rk}^{NSS}}{l_{rk}^{NSS}},$$

where l_{rk}^{NSS} denotes total employment in subsector k in region r as measured from NSS.

Five subsectors within the service sector are not covered by the Service Survey. Table B-12 reports our approach to attribute the employment in these subsectors to the PS or CS sector respectively.

NIC2004	Industry	Approach
22	Publishing, printing and reproduction of recorded media	Attribute all employment to PS
50	Sale, maintenance and repair of motor vehicles and motorcycles; retail sale of automotive fuel	Use average PS share (at firm-size bin level) from other sectors for which we have information
51	Wholesale trade and commission trade, except of motor vehicles and motorcycles	Use average PS share (at firm-size bin level) from other sectors for which we have information
52	Retail trade, except of motor vehicles and motorcycles; repair of personal and household goods	Attribute all employment to CS
62	Air transport	Attribute all employment to PS

Table B-12: IMPUTATION OF PS EMPLOYMENT. This table reports our imputation of PS and CS employment for subsectors that are not covered by the service survey.

B-2.3 Attributing employment in Construction & Utilities to PS and CS and the Public Sector

As explained in the text we also attribute employment in construction and utilities to either CS or PS. We follow a similar strategy as for the service sector. To do so, we use the Informal Non-Agricultural Enterprises Survey

⁷ In some industries, there are not enough firms with more than 20 employees to estimate ω_{kb}^{PS} precisely. If there are less than 5 firms and ω_{kb}^{PS} is smaller than ω_{kb}^{PS} in the pre-ceding size bin (i.e. $\omega_{k3}^{PS} < \omega_{k2}^{PS}$), we set $\omega_{k3}^{PS} = \omega_{k2}^{PS}$. Hence, for cells with few firms we impose the the share of PS firms is monotonic in firm size.

1999-2000 (see Section B-1.4).

From the description of National Industry Classification, some subsectors are clearly for public purposes. We therefore classify 5-digit level industries within the construction sector into public and private. The results are reported in Table B-13.

NIC-2004	Description	Public/Private
45101	Site preparation in connection with mining	Public
45102	Site preparation other than in connection with mining	Public
45201	General construction (including alteration, addition, repair and maintenance) of residential buildings.	Private
45202	General construction (including alteration, addition, repair and maintenance) of non-residential buildings.	Private
45203	Construction and maintenance of roads, rail-beds, bridges, tunnels, pipelines, rope-ways, ports, harbours and runways etc.	Public
45204	Construction/erection and maintenance of power, telecommunication and transmission lines	Public
45205	Construction and maintenance of waterways and water reservoirs	Public
45206	Construction and maintenance of hydro-electric projects	Public
45207	Construction and maintenance of power plants, other than hydro-electric power plants	Public
45208	Construction and maintenance of industrial plants other than power plants	Private
45209	Construction n.e.c. including special trade construction	Private
45301	Plumbing and drainage	Private
45302	Installation of heating and air-conditioning systems, antennas, elevators and escalators	Private
45303	Electrical installation work for constructions	Private
45309	"Other building installation n.e.c.	Private
45401	Setting of wall and floor tiles or covering with other materials like parquet, carpets, wall paper etc.	Private
45402	Glazing, plastering, painting and decorating, floor sanding and other similar finishing work	Private
45403	Finish carpentry such as fixing of doors, windows, panels etc. and other building finishing work n.e.c.	Private
45500	Renting of construction or demolition equipment with operator	Private

Table B-13: CLASSIFICATION OF THE CONSTRUCTION SECTOR. The table reports how we classify different subsectors in the construction sector as either public or private sectors.

We drop for our analysis all subsectors, that we classify as public. These account for roughly 9.2% of total construction employment (see below). For all subsectors attributed to the private sector, we estimate the CS and PS share based on the information in the Informal Non-Agricultural Enterprises Survey. The survey has information on firms in the construction sector and reports the identify of the main buyer of the firm. In particular, we observe in the data whether the firm sells to (i) the government, (ii) a cooperative or marketing society, (iii) a private enterprise, (iv) a contractor or intermediary, (v) a private individual and (vi) others. We associate all firms that answer (ii), (iii) or (iv) with PS firms and all firms that answer (v) with CS firms. We then calculate the PS share of a given private subsector as total PS employment relative to total CS and PS employment in the respective subsector, i.e. for subsector k we calculate the PS share as

$$\omega_k^{PS} = \frac{\sum_{f \in k} 1\{f \in PS\} l_f}{\sum_{f \in k} l_f},$$

where l_f denotes firm employment and $1\{f \in PS\}$ is an indicator for whether firm f is a PS firm.

In Table B-14 we report the relative employment shares of public employment (as classified in Table B-13), CS and PS in the construction sector as a whole. The share of public employment is around 10% with a slight bump in 2009, presumably a consequence of the financial crisis in 2008. Among the private subsectors, 12.9% of employment is associated with the provision of producer services.

To calculate total employment in PS and CS industries within the private subsectors of the construction sector at the regional level, we apply the 5-digit PS shares ω_k^{PS} to the NSS employment data and calculate total as

$$\varpi_r^{PS} = \frac{\sum_k \omega_k^{PS} l_{rk}^{NSS}}{l_{rk}^{NSS}} \quad \text{and} \quad \varpi_r^{CS} = \frac{\sum_k (1 - \omega_k^{PS}) l_{rk}^{NSS}}{l_{rk}^{NSS}}.$$

Note that the regional variation in PS and CS shares within the construction sector only arises because regions differ in the relative size of the different private subsectors listed in Table B-13.

B-2.4 Top-coding of CS and PS employment shares

Our model implies that regional CS expenditure shares (and hence employment shares) are bounded from above by ω_{CS} . Similarly, our model implies that the share of producer service employment ("lawyers") relative to production

	1999	2004	2007	2009
Public employment	0.073	0.102	0.073	0.136
CS employment share	0.806	0.781	0.809	0.755
PS employment share	0.121	0.116	0.118	0.109
PS/(PS+CS)	0.131	0.130	0.127	0.126

Table B-14: COMPOSITION OF THE CONSTRUCTION SECTOR. The table shows the relative employment shares of PS, CS and public employment in the construction sector in different years. We associate public employment to sectors classified as "public" in Table B-13. The classification of employment in the private subsectors to CS and PS is explained in the main text. The last row reports the relative employment share of PS within the private subsectors.

workers is bounded from above by β . Given our estimated of $\omega_{CS} = 0.69$ and $\beta = 0.7$ reported in Table 4, 20 districts violate this requirement and we topcode them in our analysis. Topcoding such districts is entirely inconsequential because these districts are very small. In Table B-15 we report the aggregate population share of such districts by year. As can be seen, in almost all years, such districts account for less than 1% of the total population in India in almost all years.

Total share of population			
Year	CS empl. share $>\omega_{CS}$	PS empl. share $>\beta$	
1987	0.0012	0.0074	
2011	0.0018	0.0092	

Table B-15: EMPLOYMENT SHARES OF TOP-CODED DISTRICTS. The table reports the aggregate population share of all districts whose employment share in CS exceeds ω_{CS} (column 1) or β (column 2).

B-3 Geography: Harmonizing Regional Borders

In this section we describe our procedure to harmonize the geographical boundaries to construct a consistent panel of time-invariant localities. This need arises because the borders of numerous Indian districts have changed between 1987 and 2011. This is for example seen in the left panel of Figure B-3 that plots the districts' boundaries in 2001 and 2011. The purple line represents the boundaries in 2001, the red line represents the boundaries in 2011.

The most common type of regional re-districting is a *partition* where one district has been separated into several districts in the subsequent years. The second type is a *border move*, where the shared border between two districts changes. The third is a *merge*, where two districts are merged into a single district.

In order to carry out the analysis on a panel of districts with a consistent geography, we construct regions that have consistent borders in 1987 and 2011. To keep the number of regions as large as possible, a region is always the smallest area that covers a single district or a set of districts with consistent borders over time. For instance, in the case of a partition, the region is constructed as the district in the pre-partition year. In the case of a border move, a region is constructed as the union of two districts. We construct a regional map with consistent borders from 1987 to 2011. The right panel of Figure B-3 shows the official Indian districts in the year 2011 (dashed red lines) and the time-invariant geographical units, that we for simplicity also refer to as districts (solid blue lines). There are two small districts that did not exist in 1987 (but do in 2011) and we exclude these from our analysis. Furthermore, because our methodology requires us to calculate sectoral employment shares at the district level, we exclude districts with less than 50 observations as these do not allow us to estimate such shares in a credible way.




Figure B-3: DISTRICT BORDERS IN INDIA: 1987 - 2011. The left figure plots the districts' boundaries in 1987 and 2011. The purple line represents the boundaries in 1987 and the dashed red line represents the boundaries in 2011. The right figure shows the official Indian districts in the year 2011 (dashed red lines) and the time-invariant geographical units we construct (solid blue lines) upon which our analysis is based.

B-4 The relative price of agricultural goods

Our estimation uses the relative price of agricultural goods (relative to manufacturing goods) to identify the relative productivity in the agricultural sector (relative to manufacturing). The Ministry of Planning and Program Implementation (MOSPI) of the Government of India reports value added by 2-digit sectors at current prices and constant prices from 1950-2013⁸ We then construct the sectoral price index by

$$p_i = \frac{\text{GDP at current price}_i}{\text{GDP at constant price}_i} \quad (\text{B-1})$$

and normalize both price indexes in the year 2005 to unity. We then calculate the relative price of agricultural products as

$$p_{\text{relative}} = \frac{p_{\text{agri}}}{p_{\text{manu}}} \quad (\text{B-2})$$

To check the validity of our results, we also use two additional data sources to calculate the relative price. The first is the GGDC 10-Sector Database⁹, which provides long-run data on sectoral productivity performance in Africa, Asia, and Latin America. This data set reports the annual series of value-added at current national prices and value-added at constant 2005 national prices. We follow the same procedures to calculate the relative price.

The second is the Wholesale Price Index (WPI) reported by Office of the Economic Advisor¹⁰. The WPI tracks ex-factory price for manufactured products and market price for agricultural commodities. One issue with this is that the base year (and the basket of goods) changes during different time periods. Two series are relevant to our research. The first one is the series with the base year 1993, which is available from 1994 to 2009. The second

⁸ The data is available at <http://www.mospi.gov.in/data>. See "Summary of macro economic aggregates at current prices, 1950-51 to 2013-14" and "Summary of macro economic aggregates at constant(2004-05) prices, 1950-51 to 2013-14".

⁹ The data is available at <https://www.rug.nl/ggdc/productivity/10-sector>

¹⁰ The data is available at <https://eaindustry.nic.in/>




Figure B-4: RELATIVE PRICE OF AGRICULTURAL TO MANUFACTURING GOODS. The figure shows the relative price of agricultural products as calculated in (B-1) and (B-2) from the different sources mentioned in the main text. “MOSPI” refers to the data from the Indian Government is used in our analysis. “GGDC” stems from GGDC 10-Sector Database. “WPI (1993)” and “WPI (2004)” are based on the Wholesale Price Index with a 1993 base year and a 2004 base year respectively.

one is the series with the base year 2004, which is available from 2005 to 2016. Again, we use the same method to calculate the relative price, and normalized the relative prize in the year 2005 to 1.

In Figure B-4 we plot the relative price of agricultural goods to manufacturing goods. Since the pattern from the different data sources are very similar, we use the results based on MOPSI in our analysis.

B-5 Estimating the expenditure elasticity (ε)

In Table 3 in the main text we reported the estimated elasticity of agricultural expenditure shares with respect to total expenditure. In Figure B-5 we show the estimated Engel curve for the year 2011 graphically. As implied by the PIGL preference specification, the relationship is approximately linear for a large part of the expenditure distribution.




Figure B-5: ENGEL CURVES IN INDIA. The figure shows a binscatter plot of the log food shares and log expenditure for the year 2011 at the individual household level after absorbing district fixed effects.

Value added versus Final Expenditure

Our estimation of ε relies on individual-level expenditure data. This data refers to consumer spending on final goods. Our theory, in contrast, is described in terms of value added. This raises the question whether we can identify the preference parameters of our value-added preferences from the data on final expenditure.

As shown in Herrendorf et al. (2013), the relevant parameters of consumers' preferences over final goods are in general distinct from the preference parameters consistent with the value-added interpretation. However, as we show now, this is less of a concern for the case of the income elasticity ε , which is the only parameter we estimate using data on expenditure.

To see this, suppose the consumer preferences over final goods are also of the PIGL form and given

$$V^{FE}(e, q) = \frac{1}{\varepsilon^{FE}} \left(\frac{e}{\prod_s q_j^{\omega_s^{FE}}} \right)^{\varepsilon^{FE}} - \sum_s \tilde{\nu}_s^{FE} \ln q_j. \quad (\text{B-3})$$

Here q_j denotes the price of the final good of sector $s \in \{F, CS, G\}$ and we denote the parameters with the superscript “FE” to highlight that these are parameters of the final expenditure specification. The associated expenditure share to (B-3) is given by

$$\vartheta_s^{FE}(e, q) = \omega_s^{FE} + \tilde{\nu}_s^{FE} \left(\frac{e}{\prod_j q_j^{\omega_j^{FE}}} \right)^{-\varepsilon^{FE}}. \quad (\text{B-4})$$

The micro-data on expenditure in principle identifies the parameters $\{[\omega_j^{FE}]_j, \varepsilon^{FE}\}$.

To map these preference parameters to the demand system in terms of value added, we have the technology to transform sectoral value added into final goods. Suppose first that the mapping from value added to final expenditure is given by a Cobb-Douglas technology

$$y_s^{FE} = \prod_{j \in (A, CS, G)} (y_{js}^{VA})^{\lambda_{js}}, \quad (\text{B-5})$$

where $\sum_j \lambda_{js} = 1$ for all s .¹¹ Here y_s^{FE} is the total quantity of final goods in sector s and y_{js}^{VA} is the value added content of sector j in the production of sector s goods. Hence, the matrix $[\lambda_{js}]_{js}$ describes the input output matrix of the economy.

The specification in (B-5) implies that λ_{js} is the cost share of sector j value added in the production of final goods in sector s

$$p_j^{VA} y_{js}^{VA} = \lambda_{js} q_s^{FE} y_s^{FE}.$$

Furthermore, final good prices q_s are given by $q_s = \prod_j (p_j^{VA})^{\lambda_{js}}$. Total value added of sector j is therefore given by

$$p_j^{VA} c_j^{VA} = \sum_s p_j^{VA} c_{js}^{VA} = \sum_s \lambda_{js} q_s^{FE} c_s^{FE}.$$

¹¹ Below we generalize our argument to a more general CES technology.

This implies that the value added share of sector j is given by

$$\begin{aligned}
\vartheta_j^{VA} = \frac{p_j^{VA} c_j^{VA}}{e} &= \sum_s \lambda_{js} \frac{q_s^{FE} c_s^{FE}}{e} = \sum_s \lambda_{js} \vartheta_j^{FE} \\
&= \sum_s \lambda_{js} \left(\omega_s^{FE} + \tilde{\nu}_s^{FE} \left(\frac{e}{\prod_m q_m^{\omega_m^{FE}}} \right)^{-\varepsilon^{FE}} \right) \\
&= \sum_s \lambda_{js} \omega_s^{FE} + \sum_s \lambda_{js} \tilde{\nu}_s^{FE} \times \left(\frac{e}{\prod_m \prod_j (p_j^{VA})^{\lambda_{jm} \omega_m^{FE}}} \right)^{-\varepsilon^{FE}} \\
&= \sum_s \lambda_{js} \omega_s^{FE} + \sum_s \lambda_{js} \tilde{\nu}_s^{FE} \times \left(\frac{e}{\prod_s (p_s^{VA})^{\sum_m \lambda_{sm} \omega_m^{FE}}} \right)^{-\varepsilon^{FE}}.
\end{aligned}$$

Now define

$$\begin{aligned}
\omega_j &\equiv \sum_s \lambda_{js} \omega_s^{FE} \\
\tilde{\nu}_j &\equiv \sum_s \lambda_{js} \tilde{\nu}_s^{FE} \\
\varepsilon &= \varepsilon^{FE}
\end{aligned}$$

as the cost-share weighted averages of the respective final good parameters. Then,

$$\vartheta_j^{VA} = \omega_j + \tilde{\nu}_j \left(\frac{e}{\prod_s (p_s^{VA})^{\omega_j}} \right)^{-\varepsilon}. \quad (\text{B-6})$$

Note that (B-6) is exactly the value added demand system of our theory. In particular, the elasticity parameter ε coincides exactly with the elasticity parameter of the final expenditure demand system in (B-4) and hence can be estimated from the data on final expenditure. This is not the case for the other preference parameters $\tilde{\nu}_j$ and ω_j . In particular, knowledge of the input-output matrix $[\lambda_{js}]$ is required to translate the final good share $[\omega_s^{FE}]_s$ into the value added shares $[\omega_s]$. We therefore do not attempt to also estimate ω_F or $\tilde{\nu}_F$ from the expenditure data but only estimate ε .

The Cobb-Douglas specification in (B-5) is particularly convenient because it gives rise to the exact value added demand system of our theory. However, the results that the expenditure income elasticity ε^{FE} and the value added elasticity ε are closely linked is more general. Suppose that the final good production function in (B-5) is given by the more general CES form

$$y_s^{FE} = \left(\sum_j (b_{js})^{1/\eta_s} (c_{js})^{\frac{\eta_s-1}{\eta_s}} \right)^{\frac{\eta_s}{\eta_s-1}}.$$

Following the same steps as above, we then find that

$$\begin{aligned}
\vartheta_j^{VA} &= \sum_{s \in} \frac{b_{js} (p_j^{VA})^{1-\eta_s}}{\sum_m b_{ms} (p_m^{VA})^{1-\eta_s}} \vartheta_s^{FE} \\
&= \sum_s \frac{b_{js} (p_j^{VA})^{1-\eta_s}}{\sum_m b_{ms} (p_m^{VA})^{1-\eta_s}} \left(\omega_s^{FE} + \tilde{\nu}_s^{FE} \left(\frac{e}{\prod_j q_j^{\omega_j^{FE}}} \right)^{-\varepsilon^{FE}} \right) \\
&= \left(\frac{\sum_s b_{js} (p_j^{VA})^{1-\eta_s} \omega_s^{FE}}{\sum_m b_{ms} (p_m^{VA})^{1-\eta_s}} \right) + \left(\frac{\sum_s b_{js} (p_j^{VA})^{1-\eta_s} \tilde{\nu}_s^{FE}}{\sum_m b_{ms} (p_m^{VA})^{1-\eta_s}} \times \left(\frac{1}{\prod_j q_j^{\omega_j^{FE}}} \right)^{-\varepsilon^{FE}} \right) \times e^{-\varepsilon^{FE}} \\
&= \left(\frac{\sum_s b_{js} (p_j^{VA})^{1-\eta_s} \omega_s^{FE}}{\sum_m b_{ms} (p_m^{VA})^{1-\eta_s}} \right) + \left(\frac{\sum_s b_{js} (p_j^{VA})^{1-\eta_s} \tilde{\nu}_s^{FE}}{\sum_m b_{ms} (p_m^{VA})^{1-\eta_s}} \times \left(\frac{1}{\prod_j q_j^{\omega_j^{FE}}} \right)^{-\varepsilon^{FE}} \right) \times \left(\frac{e}{\prod_j q_j^{\omega_j^{FE}}} \right)^{-\varepsilon^{FE}} \\
&= \left(\frac{\sum_s b_{js} (p_j^{VA})^{1-\eta_s} \omega_s^{FE}}{\sum_m b_{ms} (p_m^{VA})^{1-\eta_s}} \right) + \left(\frac{\sum_s b_{js} (p_j^{VA})^{1-\eta_s} \tilde{\nu}_s^{FE}}{\sum_m b_{ms} (p_m^{VA})^{1-\eta_s}} \times \prod_j \left(\sum_s b_{sj} \left((p_s^{VA})^{\omega_j^{FE}} \right)^{1-\eta_j} \right)^{\frac{\varepsilon^{FE}}{1-\eta_j}} \right) \times e^{-\varepsilon^{FE}}.
\end{aligned}$$

Hence, define

$$\begin{aligned}
\bar{\omega}_j([p_s^{VA}]) &\equiv \bar{\omega}_{jrt} = \frac{\sum_s b_{js} (p_j^{VA})^{1-\eta_s} \omega_s^{FE}}{\sum_m b_{ms} (p_m^{VA})^{1-\eta_s}} \\
\bar{\nu}_j([p_s^{VA}]) &\equiv \bar{\nu}_{jrt} = \frac{\sum_s b_{js} (p_j^{VA})^{1-\eta_s} \tilde{\nu}_s^{FE}}{\sum_m b_{ms} (p_m^{VA})^{1-\eta_s}} \times \prod_j \left(\sum_s b_{sj} \left((p_s^{VA})^{\omega_j^{FE}} \right)^{1-\eta_j} \right)^{\frac{\varepsilon^{FE}}{1-\eta_j}},
\end{aligned}$$

where the dependence on sectoral prices $[p_{srt}^{VA}]$ implies that these functions vary across sectors, region and time. The value added expenditure shares are therefore given by

$$\vartheta_j^{VA} = \bar{\omega}_{jrt} + \bar{\nu}_{jrt} \times e^{-\varepsilon^{FE}}. \quad (\text{B-7})$$

Hence, the derived demand system has the same functional form as the value added demand system in our theory. In particular, the final good expenditure elasticity ε^{FE} again emerges as the value added expenditure elasticity. However, in contrast to (B-6), the “parameters” $\bar{\omega}_{jrt}$ and $\bar{\nu}_{jrt}$ explicitly depend on value added prices and hence vary across locations and time. Recall that we target the parameter β from the regression

$$\ln \vartheta_{Frt}^i = \delta_{Frt} + \beta \times \ln e_{irt} + u_{irt},$$

i.e. we also identify the parameter β from the cross-sectional variation within locations as dictated by (B-7).

B-6 The Elasticity of Substitution and the Income Elasticity of Expenditure Shares

For the class of PIGL preferences, the elasticity of substitution is not a structural parameter but depends on relative prices and total expenditure. The Allen Uzawa elasticity of substitution between goods s and k is given by

$$EOS_{sk} = \frac{\frac{\partial^2 e(p,V)}{\partial p_s \partial p_k} e(p,V)}{\frac{\partial e(p,V)}{\partial p_s} \frac{\partial e(p,V)}{\partial p_k}},$$

where $e(p,V)$ denotes the expenditure function. As we show in Section A-4 in the Appendix, our preference specification implies that

$$EOS_{sk} = 1 - \varepsilon \frac{(\vartheta_s - \omega_s)(\vartheta_k - \omega_k)}{\vartheta_s \vartheta_k}.$$

In the left panel of Figure B-6 we report the implied elasticities of substitution as a function of the regional urbanization.¹² The substitution elasticities are relative to close to unity. Goods and consumer services are complements, in particular in poor, rural districts. Food and consumer services are slightly more substitutable than implied by a Cobb Douglas utility function. In the right panel we depict the elasticity of sectoral expenditure shares with respect to income. In our model, this elasticity is given by

$$\frac{\partial \ln \vartheta_s}{\partial \ln e} = -\varepsilon \frac{\vartheta_s - \omega_s}{\vartheta_s}.$$

Quantitatively, our estimated model predicts that the expenditure elasticity for agricultural products is close to -0.34. This is expected because $\varepsilon \approx 0.34$. The expenditure elasticities on goods and consumer services are both positive and between 0.2 and 0.5. The consumer service elasticity is particularly large in rural regions, that are on average poor and unproductive in consumer services.




Figure B-6: ELASTICITIES OF SUBSTITUTION AND EXPENDITURE ELASTICITIES. The figure plots the elasticity of the substitution between the three goods (left panel) and the elasticity of expenditure shares with respect to expenditure (right panel) by urbanization quintile in 2011.

¹² More specifically, for each urban quintile we calculate the population weighted average of the respective regional elasticity of substitution for the region's representative household.

B-7 Estimating the shape of the human capital distribution (ζ)

We estimate the tail parameter of the distribution of efficiency units ζ from the distribution of income. Our model implies that total income and expenditure of individual h is given by $e_{rt}^h = q^h w_{rt}$, where q follows a Pareto distribution

$$f_{rt}(q) = \frac{\zeta q_{rt}^\zeta}{q^{\zeta+1}}$$

This implies that

$$\ln(f_{rt}(q)) = \ln(\zeta q_{rt}^\zeta) - (\zeta + 1) \ln(q). \quad (\text{B-8})$$

Hence, we estimate ζ from a regression of the (log of the) upper tail density on log efficiency units which we calculate as $q_{rt}^h = \frac{e_{rt}^h}{w_{rt}}$. In Table B-16 we report the estimated ζ based on (B-8). We report both the estimate based on the full sample (column 1) and the estimates by urbanization quintile (columns 2 - 6). We also report our estimates based on two measures of earnings: total expenditure by capita (which we use as our earnings measure for our main analysis) and total incomes, which is also reported in the NSS data.

Table B-16 contains two results. First, the estimated tail parameter for the aggregate economy is slightly below 3, stable across years and does not depend on the exact measure of earnings. Second, the estimated tail parameter is declining in the urbanization rate indicating that urban locations have higher inequality. Our estimates also indicate that inequality was lower in 2011 than in 1987. For our quantitative model we set ζ to an average value of 3. In Section 7 we show that our results are robust to a variety of choices for ζ . Hence, for simplicity, we abstract from the heterogeneity in ζ across urbanization quantiles.

		Variable	Full Sample	Quintiles of Urbanization				
				1st	2nd	3rd	4th	5th
1987	Income		2.82	3.11	3.06	3.25	2.93	2.92
	Expenditure		2.84	3.64	3.57	3.21	3.03	2.79
2011	Income		2.85	4.04	3.47	3.13	2.90	2.71
	Expenditure		2.90	3.80	3.57	3.16	2.96	2.63

Table B-16: IDENTIFICATION OF ζ . The table reports the estimate of ζ based on (B-8). In the first columns we report the estimates for the years 1987 and 2011. In the remaining columns we perform our estimation separately for different quantiles of the urbanization distribution.

B-8 Estimating the shape of the productivity distribution (λ)

We identify the shape parameter λ of the productivity distribution from the tail of the employment distribution. Our model implies that total employment of a firm with productivity z is given by

$$l(z) = L_{PS}(z) + L_{PM}(z) = \frac{\beta + \alpha}{\beta} \varsigma(A_{PS}) \frac{z}{z_L} f_0 - \varsigma(A_{PS}) f_0.$$

This implies that the employment distribution is given by

$$F_l(l) = P(l(z) \leq l) = 1 - \left(\frac{A_{rMt} \frac{\alpha + \beta}{\beta} \frac{\varsigma(A_{PS}) f_0}{z_L}}{l + \varsigma(A_{PS}) f_0} \right)^\lambda.$$

Hence, for large firms (i.e. $l \rightarrow \infty$), the tail of the employment distribution is exactly equal to λ .

To estimate λ , note that

$$\ln(1 - F_l(l)) = C_0 - \lambda \ln l, \quad (\text{B-9})$$




Figure B-7: IDENTIFICATION OF λ . The figures depicts the relationship between $\ln(1 - F_l(x))$ and $\ln x$ for different cutoffs of the employment distribution. We consider all firms with employment exceeding the 50%, 70%, 80% and 90% quantile. For each sample we consider a grid of 200 points the log employment distribution. The slope of the relationship coincides with λ (see (B-9)).

where $C_0 = \ln \left(A_{rMt} \frac{\alpha+\beta}{\beta} \frac{\zeta f_0}{z_L} \right)^\lambda$. Hence, we estimate λ from a regression of the log of the upper tail probability on log employment.

In Figure B-7 we depict the empirical relationship of (B-9) for the different sub samples of tail of the employment distribution. Specifically, we consider a grid of 200 points of the log employment distribution $\ln l$, calculate $F_l(l)$ for these grid points and then plot $\ln(1 - F_l(l))$ against $\ln l$. We consider four samples, namely all firms with employment exceeding the 50%, 70%, 80% and 90% quantiles.

Equation (B-9) implies that the relationship should be linear and that the slope should be equal $-\lambda$. Figure B-7 shows that the employment distribution in India indeed has a pareto tail and that the estimated slope does not depend markedly on the employment cutoff for the employment distribution.

In Table B-17 we report these results in a regression format. The four columns refer to the different subsamples of the tail of the employment distribution. We estimate a pareto tail of 1.42. Reassuringly, the slope is very precisely estimated and the estimates are almost identical across the different samples. We also find the economic magnitude plausible. The firm size distribution in the US is often found to have a tail of around 1.1. Given that the importance of large firms is larger in rich countries, an estimate of 1.42 strikes us as plausible.

B-9 Outliers In Quantitative Analysis

For our quantitative analysis in Section 6 we winsorize a small number of outliers. For a small number of regions we estimate very large changes in CS productivity. Intuitively, because CS employment in our model is bounded by ω_{CS} from above, our theory can only rationalize employment shares close to ω_{CS} with an exceedingly high level

	(1)	(2)	(3)	(4)
	> 50%	> 70%	> 80%	> 90%
log employment	-1.426*** (0.010)	-1.420*** (0.011)	-1.418*** (0.011)	-1.414*** (0.012)
Industry FE	Yes	Yes	Yes	Yes
R^2	0.991	0.989	0.989	0.987

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B-17: ESTIMATION OF λ . The table reports the estimates of the tail of the productivity distribution λ . In the different columns we restrict the sample to firms with a size exceeding the 50%, 70%, 80% and 90% quantile of the size distribution.

of CS productivity.

This is seen Table B-18 where we report the upper and lower quantiles of the regional distribution of welfare changes for the different counterfactuals. Consider for example the agricultural sector. If agricultural productivity had not grown since 1987, the most adversely affected region would have seen its welfare decline by 65.6% in terms of an equivalent variation. Conversely, some regions would have gained. The region benefitting the most would have seen an increase in welfare by 54.8%. The last row of Table B-18 shows that some regions would have seen very large gains if CS productivity had not grown. These are regions where CS productivity *declined* between 1978 and 2011. As explained above, this pattern is entirely driven by few districts being close to the theoretical threshold of ω_{CS} . For comparison, in the last row we also report the estimated distribution of the welfare effects in our baseline analysis, where we truncate the productivity growth distribution at the top and bottom 3%. This has large effects on the distribution of welfare effects in the right tail of the distribution.

	Regional welfare changes (%)									
	Min	1%	2%	3%	5%	95%	97%	98%	99%	Max
Agriculture	-65.6	-53.7	-49.4	-46.7	-45.6	3.0	8.2	10.3	19.2	54.8
Industry	-44.9	-33.1	-29.0	-27.4	-25.2	-5.1	-4.0	-2.4	3.2	26.8
Cons. Serv.	-98.9	-97.9	-96.0	-92.9	-86.1	26.3	68.9	386.2	2575.4	6.8×10^4
Cons. Serv. (Baseline)	-97.0	-96.7	-95.1	-92.8	-86.1	26.3	60.3	72.2	121.6	156.4

Table B-18: DISTRIBUTION OF WELFARE LOSSES. The table reports the lower and upper percentiles of the regional distributions of sectoral welfare losses.

These extreme values at the bottom of the regional productivity growth distribution have aggregate effects. For our baseline analysis we trim the top and bottom 3% of the productivity growth distribution and set regional productivity growth in such regions to the 3% and 97% quantile respectively. In Table B-19 we report the change in aggregate in the absence of CS productivity growth as a function of this trimming cutoff. Without any trimming, the aggregate effect is 32%, i.e. is positive due to the extreme outliers reported in Table B-19. Once such outliers are truncated, we recover our baseline results of a welfare loss of about -26%. In the last row of Table B-19 we report the aggregate employment share of the affected districts. The changes in the aggregate effects of CS growth are not driven by few large districts but by a small number of small districts with very large changes in CS productivity.

B-10 Decomposing the Effects of Productivity Growth Within the Industrial Sector: Details

In this section we provide more details concerning the decomposition of productivity growth in the industrial sector into its two components: manufacturing productivity A_{rMt} and PS productivity A_{rPSt} . In the left panel of Figure

	No trimming	Trimming cutoff				
		1%	2%	3%	4%	5%
Welfare Loss	31.9%	-6.3%	-21.9%	-26.4%	-27.1%	-27.6%
Employment Share	0	1.2%	1.4%	3.2%	4.2%	6.7%

Table B-19: WELFARE LOSSES WITH DIFFERENT TRIMMING CUT-OFFS. The table reports the aggregate welfare effects of productivity growth in the CS sector for different trimming rules. A trimming cutoff $x\%$ means that we set the $x\%$ highest and lowest productivity growth rates to $x\%$ and $1 - x\%$ respectively.




Figure B-8: DECOMPOSITION OF INDUSTRIAL GROWTH: MANUFACTURING VS PRODUCER SERVICES. The figure decomposes the effect of industrial productivity growth ($d \ln A_{rGt}$) into manufacturing ($d \ln A_{rMt}$) and producer service (A_{rPS}) growth. In the left panel we display the decomposition in our baseline calibration. In the right panel we display the decomposition as a function of the PS share in the production function β .

B-8 we report the composition of the aggregate effect into the respective components in our baseline calibration. Clearly, manufacturing productivity accounts for the lion share of industrial growth.¹³

In the right panel, we study the sensitivity of our results with respect to β . In our baseline calibration, we choose $\beta = 0.7$ based on the observation of a large employment share of PS in the United States. However, we recognize that the calibration of $\beta = 0.7$ is subject to some uncertainty. For instance the asymptotic share of the PS sector could be smaller in India than in the United States—as it is in other large economies such as China and Germany. In the right panel of Figure B-8 we therefore replicate our decomposition as a function of β . The lower β , the higher the role for producer services. Intuitively: the closer the observed equilibrium to its asymptotic share of PS employment, the more sensitive are the employment allocations to A_{PS} . Hence, even modest increases in the PS intensity of production cause us to infer relative large changes in PS productivity. Still, manufacturing productivity accounts for more than 80% of productivity growth in the industrial sector.

For completeness, In Figure B-9, we decompose the welfare effect of the industrial sector into the manufacturing and PS component by urbanization quintile. As with the aggregate results reported in Figure B-8, the quantitative importance of producer service growth is small.

B-11 The Bootstrap Procedure

In this section we describe the implementation of our bootstrap procedure. We rely on a non-parametric bootstrap, which treats the observed empirical distribution of the data as the population (see, for example, Horowitz (2019)).

¹³ In this decomposition there is also a covariance component. Because the importance of this component is minuscule, we do not show it in Figure B-8.




Figure B-9: DECOMPOSITION OF INDUSTRIAL GROWTH BY URBANIZATION. This figure report the decomposition of the aggregate welfare effect of productivity growth in the industrial sector between manufacturing (blue bars) and PS (red bars) for different urbanization quintiles.

We implement this procedure in the following way:

1. From the underlying micro data of the NSS we draw households randomly with replacement and we sample, within each district, the same number of households as the current dataset.¹⁴
2. Given this bootstrap sample, we recalculate all statistics used in our accounting procedure, i.e. sectoral employment shares, sectoral income shares and the supply of human capital at the district level.
3. We then redo our entire analysis on this bootstrap sample, i.e.
 - (a) we re-estimate the structural parameters that rely on this data, i.e. the income elasticity ε (by targeting the estimated income elasticity of the expenditure of food reported in Table 3) and the preference parameters ν_F and ω_{CS} (as explained in Section 5),
 - (b) we re-estimate the productivity fundamentals $\mathbf{A}_t$,
 - (c) we calculate our counterfactuals by setting sectoral productivity growth between 1987 and 2011 to zero,
4. This procedure provides us with alternative estimates of the welfare effects and the impact on the structural transformation. Let $\Delta\mathcal{W}^{(b)}$, $\Delta\mathcal{W}_r^h(q)^{(b)}$ and $\Delta\mathcal{W}_r^{(b)}$ denote the aggregate, individual and regional welfare impact from bootstrap iteration b . Similarly, let $L_{s2011}^{CF_F, (b)}$, $L_{s2011}^{CF_{CS}, (b)}$ and $L_{s2011}^{CF_I, (b)}$ denote counterfactual employment share in sector s in bootstrap iteration (b) in 2011 if productivity in agriculture (F), CS (CS) and Industry (I) had not grown since 1987. We always use the same choices to treat outliers as in our baseline analysis (see Section B-9)
5. We replicate this procedure B times and hence arrive at the vector

$$\left\{ \Delta\mathcal{W}^{(b)}, \Delta\mathcal{W}_r^h(q)^{(b)}, \Delta\mathcal{W}_r^{(b)}, L_{s2011}^{CF_F, (b)}, L_{s2011}^{CF_{CS}, (b)}, L_{s2011}^{CF_I, (b)} \right\}_{b=1}^B \quad (\text{B-10})$$

In practice we take $B = 200$.

¹⁴ We decided to sample individuals *within* districts for two reasons. First, we wanted to ensure the regional population shares (which we take as exogenous in our theory) are relatively constant across bootstrap iterations (they are not exactly constant because different households have different sampling weights). Second, some districts are small. By fixing the number of sampled households within each districts we ensure a comparable sample size with our baseline analysis.




Figure B-10: BOOTSTRAP DISTRIBUTION OF AGGREGATE EMPLOYMENT SHARES. The figure shows the bootstrap distribution of the aggregate sectoral employment share in 1987 (left panel) and 2011 (right panel). The vertical dashed line corresponds to the empirically observed value.

6. From B-10 we can estimate the distribution of the statistics of interest. For example, the τ th quantile of the distribution of aggregate welfare gains, $q_{\Delta\mathcal{W}}^\tau$, can be estimated from the empirical distribution

$$\frac{1}{B} \sum_{b=1}^B 1 \left[\Delta\mathcal{W}^{(b)} \leq q_{\Delta\mathcal{W}}^\tau \right] \leq \tau$$

The quantiles for the other objects of interest are calculated in a similar way.

7. In the box plots in Figures 8, 9 and 10 we plot the 5%, 25%, 50%, 75% and 95% quantiles of the respective distribution.

Note that, for simplicity, this procedure only captures the sampling variation stemming from the NSS micro data. Hence, we do not for example resample firms in the Economic Census or the firm survey to re-estimate the relative weights of PS versus CS employment within the different sub sectors of the service sector (see Section B-2).

In Figure B-10 we show the bootstrap distribution of the aggregate sectoral employment shares in 1987 (left panel) and 2011 (right panel). Expectedly, the sampling variation in these aggregate statistics is very small and the distribution is close to the value of our baseline analysis which is shown as a dashed vertical line.

In Figure B-11 we show the estimated distribution of the welfare losses depicted in Figures 8 and 9. We show the losses attributable to productivity growth in agriculture (panel a), in CS (panel b) and in the industrial sector (panel c). For each case we depict the aggregate welfare losses and the losses for the 1st and 5th urbanization quintile on the left and for different quantiles of the income distribution on the right. The distributions are well-behaved and do not seem to be driven by extreme outliers.

B-12 Details of Robustness Analysis (Section 7)

This section contains additional results for our robustness analysis in Section 7.

B-12.1 Sensitivity to Structural Parameters

In Figure B-12 we show the robustness of our results with respect to the elasticity of substitution across traded varieties σ (left panel) and the trade elasticity η used in the open economy extension of our theory (right panel). The structure of the graphs is the same as in Figure 11 in the main text.




Figure B-11: BOOTSTRAP DISTRIBUTION OF WELFARE LOSSES. The figure shows the bootstrap distribution of the welfare losses when we counterfactually set sectoral productivity in 2011 to its level in 1987. In panel (a) we shut down productivity growth in agriculture, in panel (b) in CS and in panel (c) in the industrial sector. Within each panel, on the left we show the aggregate welfare losses and the losses for the 1st and 5th urbanization quintile. On the right we show the losses for the different quantiles of the income distribution.




Figure B-12: ROBUSTNESS ANALYSIS. Panels (a) and (b) show the welfare effects as a function of the variety elasticity σ and the trade elasticity η . The vertical dashed line corresponds to the parameter value in our benchmark analysis.

The variety elasticity σ has - quantitatively - a negligible effect on our results: the aggregate welfare effects of sectoral productivity do not depend much on the assumed value of σ . As far as the effects of consumer services is concerned, they are - if anything - increasing in σ , i.e. our benchmark value of $\sigma = 3$ (which is on the low end of the usual estimates) is conservative. The same is true for the trade elasticity η , shown in the right panel. Again, a higher level of η increases the welfare gains of sectoral productivity growth but the quantitative effects are small. We therefore conclude that our main results are robust to our choices of σ and η .