

Growing Like India: The Unequal Effects of Service-Led Growth.

Tianyu Fan, Michael Peters, and Fabrizio Zilibotti

Yale University*

March 2022

Abstract

We construct a spatial equilibrium model where agents have nonhomothetic preferences over final goods that differ in the intensity of use of local services as production inputs. Over time, the expansion of employment in local services is both a consequence (income effects) and a cause (productivity growth) of the development process. We estimate the model using household data from the Indian NSS exploiting granular (district-level) information on sectors of employment and individual data for consumption expenditure. By way of counterfactuals, we assess the welfare effects of different sources of technical progress for different group of agents in the economy (rich vs. poor; rural vs. urban residents). We find that productivity growth in consumer services is an important driver of rising living standards between 1987 and 2011 accounting for 1/3 of aggregate welfare gains. However, these gains are heavily skewed toward high-income households living in cities. Productivity growth in the service sector is also a powerful driver of the process of structural change shifting employment out of agriculture into the service sector with only limited industrialization.

1 Introduction

Urbanization and structural change are transforming the lives of hundreds of million of people throughout the globe. Consider India, the second most populous country in the world: Thirty years ago, only a quarter of the population resided in urban areas and almost two thirds of the labor force was employed in agriculture. Today, the share of people living in urban areas has increased by 10 percentage points while the employment share of agriculture is down to 42%. While economic development has improved living conditions across the board, the sources of welfare gains are diverse. In rural areas, poverty has fallen mainly owing to productivity growth in agriculture. Meanwhile, the urban bourgeoisie has benefited not only from the availability of better and cheaper goods but also from the growing supply of local services that have changed the face of urban life.

This paper provides a framework to quantify the heterogeneous welfare effects of structural change across localities and the income distribution ladder, building a bridge between economic growth and economic development.

*This study is part of a broader research project also involving Philippe Aghion and Robin Burgess. A previous version of this article was circulated under the title: *Service-Led or Service-Biased Growth? Equilibrium Development Accounting across Indian Districts*. We are grateful to our discussants Sebastian Sotelo and Klaus Desmet for their thoughtful suggestions at the AEA meeting and the NBER Summer Institute. We thank Daron Acemoglu, Treb Allen, Fabian Eckert, Reto Foellmi, Doug Gollin, Cormac O’Dea, Peter Klenow, Samuel Kortum, Rachel Ngai, Richard Rogerson, María Sáez Martí, Danyang Xie, and seminar participants at the ASSA Meeting 2021, Cowles Macro Conference, the STEG Workshop, Dartmouth College, Penn State University, Peking University, RIDGE, the University of Sankt Gallen, and Yale University. We also thank Haonan Li, Sarah Moon, Shengqi Ni, Pariroo Rattan, and Huihuang Zhu for excellent research assistance. Fabrizio Zilibotti acknowledges financial support from the NSF grant 1949228. Michael Peters and Fabrizio Zilibotti thank the “Minnesota Opportunity and Inclusive Growth Institute” for the hospitality.

We abandon the straightjacket of aggregate representative agent models and construct a multisectoral model where people with heterogeneous purchasing power reside in different locations and consume different baskets of goods and services. This allows us to associate changes in the economic environment with their effects on the welfare of people with diverse socioeconomic characteristics.

Geography and Production Structure. We construct a spatial equilibrium theory whose main building blocks are nonhomothetic preferences and the assumption that—while manufacturing and agricultural goods are traded across regions—some services are of a local nature. We assume labor is perfectly mobile across industries while labor mobility is subject to frictions across geographical locations. We make two key assumptions: first, some of the services produced in the economy must be provided locally; second, final goods that use intensively such local services are luxuries, while goods with a low local service content are necessities. For instance, a large share of the value added of fashionable restaurants consists of the labor services of cooks, waiters, etc. In contrast, a basic home-made meal consists mainly of tradable goods.

Because the service sector is broad and heterogeneous, we split it into two parts. On the one hand, we label *consumer services* (henceforth, CS) services that contribute to households’ access to consumption goods (e.g., restaurants or retail shops) or directly enter their consumption basket (e.g., health or leisure services). On the other, we label *producer services* (PS) services that are used as inputs to the goods-producing industry, such as business services, corporate law services, and part of transport services. Consumers’ preferences can be represented both on final goods and over the value added of three grand sectors: *food*, *industrial goods*, and *CS*. In this classification, we view PS as an integral part of the industrial goods sector. As such, their value added can be shipped across locations. In contrast, CS are neither directly nor indirectly traded. While the assumption that CS are consumed locally is 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.

Then, we use both micro and macro data to estimate the spatial and time variation of productivity in each sector. Our approach is in the wave of the development accounting literature. We do not attempt to explain the determinants of productivities but retrieve them from the equilibrium condition of a structural model. Finally—by means of counterfactual analysis—we evaluate the importance of different source of productivity changes on structural change and on the welfare of people with different earnings living in different locations.

PIGL Preferences: We assume preferences parametrized by an indirect utility function in the PIGL class. This preference class was first introduced by Muellbauer (1976) and has been recently popularized in the literature on growth and structural change by Boppart (2014). PIGL has two important features. First, it has nice 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. Second, it enables us to estimate the income elasticity using household expenditure data—a point to which we return below.

The aggregation properties of PIGL preferences enable us to perform a variety of counterfactual welfare calculations based on the estimated model. In particular, because agents have nonhomothetic preferences and CS must be provided locally, productivity growth in different sectors benefits people differentially—rich versus poor as well as urban versus rural residents. Our estimated model allows us to quantify the heterogeneous welfare gains and development effects of service-led growth both across the income distribution and across space.

Identification: The estimation of the productivity vector is subject to an identification problem. An increase

in the employment share of CS could stem from demand forces (i.e., local income growth under nonhomothetic preferences) and supply forces (i.e., local productivity of CS). We refer to these two channels as *service-biased* and *service-led* growth, respectively. The identification of their relative importance hinges on the income elasticity of demand. To this aim, we estimate Engel curves using household expenditure data. This step is potentially treacherous. Our estimation of productivities uses sectoral employment data and a demand system defined over sectoral value added aggregates. As Herrendorf et al. (2013) show, the parameters of this demand system are in general different from those derived from preferences over final goods and services. The mapping between the two set of parameters depends in the general on the input-output matrix. We formally establish that—under PIGL preferences—the income elasticity we estimate is a common parameter to the value added and the final expenditure demand system, irrespective of input-output relations. While this equivalence does not extend to other parameters, we only use household expenditure to retrieve the elasticity that is common in the two systems.

Application: Service-led Growth in India: We apply our methodology to India, a fast-growing economy, with an average annual 4.2% growth rate during 1987–2011, for which individual of good quality are available. In this period, the lion’s share of the process of structural change was a shift from agriculture to services with only a minor role of the manufacturing sector. India is not an exception in this respect. In recent years, structural change has similar features in in many developing economies.

Our estimation exploits individual geolocalized consumption and employment data, and we estimate sectoral productivity growth for ca. 400 Indian districts. The 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 final goods are provided more efficiently (e.g., because of a better division of labor or for a market size effect that allows the entry of large more efficient retailers.)

Second, we document an important role for service-led growth for economic development. At the aggregate level, rising efficiency in the provision of consumer services accounts for almost one third of the increase in living standard (i.e., welfare) 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. 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. Finally, our approach is close in spirit to Burstein et al. (2005) who emphasize in a different context the nontradable nature of consumer services and the large value added share of these services in final expenditure goods.

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 productivity. 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 in the benchmark model, though we study the case of labor mobility as an extension. 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-biased 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 the left panel of Figure 1, we show the time-series evolution of sectoral employment shares.¹ 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 whose employment shares increased, respectively, by nine and seven percentage points. By contrast, employment in the manufacturing sector is stagnant. Today, the service sector accounts for about one third of aggregate employment and almost one half of total value added in India.

The service sector encompasses a set of heterogeneous activities. In the right panel of Figure 1, we decompose it into five subsectors: (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.² 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.

The figure 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. But service employment grew substantially in both localities.

Was the rapid expansion of the service sector shown in Figure 1 a source or a corollary of Indian growth? In the remainder of the paper, we argue that rising productivity in consumer-oriented services like retail and hospitality was indeed an important source.

¹ The figure is constructed using micro data on employment from the NSS whose description is deferred to Section 4.

² 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.

3 Theory

The model economy comprises R regions. Within each region there are three broad sectors of activity: agriculture (F for *food*), the industrial sector (G for *goods*) and consumer services (CS). Consumers' preferences are defined over a continuum of products, and each product is a combination of the output of these three sectors. The main distinction between food and goods on the one hand and CS on the other hand is their tradability: While food and goods are tradable across regions (subject to iceberg costs), CS must be locally provided.³

In our benchmark model, we assume that the aggregate supply of labor is inelastically provided in each region and that workers' human capital is perfectly substitutable across the sectors of activity. Below we extend this treatment of labor supply in two ways. First, we allow for spatial mobility and rationalize the geographic allocation of labor through the lens of a geography model with local amenities. This extension allows us to run experiments in which labor can move across regions in response to counterfactual changes in the environment. Second, we also consider an extension where we allow for imperfect sectoral substitutability through distinguishing workers by their skills and sectors by their skill requirement. In that case, regional differences in the supply of skills emerge as a source sectoral of comparative advantage and hence specialization. Throughout our analysis we assume that markets are frictionless and competitive.

3.1 Technology and Preferences

Technology: Each region r produces a measure one continuum of differentiated final products that enter consumers' utility. Each good is produced using two physical inputs—food and goods—and local CS workers. For instance, a restaurant meal is a combination of food, kitchen tools, and the service provided by cooks and waiters. Formally, the production function for final good $n \in [0, 1]$ in region r at time t is given by

$$Y_{rnt} = \tilde{\lambda}_n x_{rFt}^{\lambda_{nF}} x_{rGt}^{\lambda_{nG}} (\mathcal{A}_{rnt} H_{rCS} t)^{\lambda_{nCS}}, \quad (1)$$

where x_F , x_G , denote the inputs of food and goods in the production of commodity n , $H_{rCS} t$ is the mass of consumer service workers allocated to the production of good n and \mathcal{A}_{rnt} reflects the efficiency of providing the consumer service content for product n in region r . The scalar $\tilde{\lambda}_n$ is an inconsequential constant to simplify some expressions.⁴

As Equation (1) shows, the productions functions for non-tradable consumer goods are parametrized by the elasticities λ_{ns} and the productivity of CS \mathcal{A}_{rnt} . The parameters λ_{ns} determine the intensity of food, goods and CS in the production of product n and are thus akin to input-output coefficients. Intuitively, a home-cooked meal is a product with a large food content ($\lambda_{nF} \approx 1$) and a low content of CS (the retail store). A restaurant meal also requires food but has a larger CS content. Finally, personal services like haircuts or nanny services consist almost entirely of CS ($\lambda_{nCS} \approx 1$). The productivity term \mathcal{A}_{rnt} governs the efficiency of the provision of CS for good n in region r and captures regional differences in the quality of retail stores or the variety of available restaurants. Quantifying the welfare impact of increases in \mathcal{A}_{rnt} on consumer welfare is the main aim of our paper.

As mentioned above food and industrial goods in equation (1) are tradable. Formally, x_F and x_G are CES

³ The industrial sector employs both production workers and workers producing production services (PS). Because the value added of corporate lawyers and consultants is embodied in industrial goods, PS are ultimately tradable.

⁴ In particular, $\tilde{\lambda}_n = \lambda_{nF}^{-\lambda_{nF}} \lambda_{nG}^{-\lambda_{nG}} \lambda_{nCS}^{-\lambda_{nCS}}$.

aggregates of differentiated varieties, each produced in a different region:

$$x_s = \left(\sum_{r=1}^R y_{rs}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad \text{for } s \in \{F, G\}. \quad (2)$$

The technology of the two tradable varieties produced in region r features constant return to labor:

$$y_{rFt} = A_{rFt} H_{rFt} \quad \text{and} \quad y_{rGt} = A_{rGt} H_{rGt}. \quad (3)$$

Here, A_{rFt} denotes sectoral TFP, while H_{rst} denotes the amount of human capital employed in the production of food and industrial goods. In particular, $H_{rGt} = H_{rBCt} + H_{rCSt}$, i.e., H_{rGt} includes labor services provided by both blue-collar workers and PS workers. However, (3) does *not* restrict blue-collar workers and corporate lawyers to be perfect substitute. Indeed, one can provide microfoundation to this reduced-form representation of the industrial goods technology in terms of a production function where the services provided by blue-collar workers and PS workers are imperfect substitutes. More formally, one can postulate $y_{rGt} = g_{rt}(H_{rBCt}, H_{rPS t})$, where g_{rt} are generic linearly homogeneous functions that can vary at the region-time level in an unrestricted way. Such technological variation captures both factor-neutral productivity differences and differences in the relative efficiency of producer services across space and time. Thus, a location can be highly productive in the production of industrial goods either because they have access to a good productive technology of blue collar workers or because their corporate lawyers are very efficient. As long as firms maximize profits, one can always express the constant-return technology in the form $y_{rGt} = A_{rGt} H_{rGt}$, where the TFP parameter A_{rGt} summarizes the regional heterogeneity in the technology.⁵

Equations (1), (2), and (3) highlight the special role of the CS sector in our theory: its value added is combined with that of tradable commodities to turn the latter into final goods, local consumers can enjoy. In particular, note that we refer to \mathcal{A}_{rnt} as CS productivity even though (given the Cobb-Douglas nature of the production function in (1)), \mathcal{A}_{rnt} applies to all inputs and can as well be thought of as the productivity of the local final good sectors. We show below that the assumption that consumer goods must be supplied locally, while food and goods can be purchased in nationwide markets, allows us to separately identify \mathcal{A}_{rnt} from A_{rGt} and A_{rFt} .

Tradability is also the critical difference between PS and CS in our theory. While CS value added can only be supplied locally, PS value added embodied in goods is ultimately tradable. Moreover, the way how PS enter the production of industrial goods can vary freely across locations. Hence, cities like Delhi or Mumbai, for example might have a comparative advantage in PS like finance or ICT and can then export the value added of such services to the rest of India (and in on of our extension to the rest of the world).

Preferences: Following Boppart (2014) and—more specifically—Alder et al. (2019), we assume consumers’ preferences over the continuum of final products n are in the PIGL class. PIGL preferences have some appealing properties for our purposes. First, they have simple and transparent aggregation properties that allow us to take a spatial demand system to the data. Second, they allow us to derive analytic expressions for individual and aggregate

⁵ Total production can be written as $y_{rGt} = g_{rt}(1 - s_{rPS t}, s_{rPS t}) H_{rGt}$, where $H_{rGt} = H_{rBCt} + H_{rPS t}$ and $s_{rPS t} = H_{rPS t} / H_{rGt}$ denotes the human capital share allocated towards PS. Total factor productivity A_{rGt} is then given by $A_{rGt} \equiv \max_{s_{PS}} g_{rt}(1 - s_{PS}, s_{PS})$. Note that A_{rGt} is a primitive and fully determined from the production function g_{rt} . As a particular example, suppose g_{rt} takes the CES form with elasticity of substitution ρ and labor augmenting productivity A_{rBCt} and $A_{rPS t}$ for blue-collar and PS workers respectively. Then $A_{rGt} = \left(A_{rBC}^{\rho-1} + A_{rPS}^{\rho-1} \right)^{1/(\rho-1)}$. For the purpose of our analysis we do not have to take a stand but we will directly estimate the distribution of A_{rGt} from the data.

welfare effects. Third, they provide a simple mapping of preferences over final goods into preferences over value added.⁶

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

$$\mathcal{V}(e, \mathbf{p}_r) = \frac{1}{\varepsilon} \left(\frac{e}{B(\mathbf{p}_r)} \right)^\varepsilon - D(\mathbf{p}_r), \quad (4)$$

where e denotes total spending and \mathbf{p}_r the vector of prices in region r . The functions $D(\mathbf{p})$ and $B(\mathbf{p})$ are homogeneous of degree zero and one, respectively. We parametrize the functions $B(\mathbf{p})$ and $D(\mathbf{p})$ in (4) as

$$B(\mathbf{p}_r) = \exp \left(\int_{n=0}^1 \beta_n \ln p_{rn} dn \right) \quad \text{and} \quad D(\mathbf{p}_r) = \left(\int_{n=0}^1 \kappa_n \ln p_{rn} dn \right),$$

where $\int_{n=0}^1 \beta_n dn = 1$ and $\int_{n=0}^1 \kappa_n dn = 0$. This specification yields the indirect utility function

$$\mathcal{V}(e, \mathbf{p}_r) = \frac{1}{\varepsilon} \left(\frac{e}{\exp \left(\int_n \beta_n \ln p_{rn} dn \right)} \right)^\varepsilon - \int_n \kappa_n \ln p_{rn} dn. \quad (5)$$

Roy's Identity implies that the expenditure share individual h allocates to final good n given prices \mathbf{p}_r and spending e , $\vartheta_n^{FE,h}(e, \mathbf{p}_r)$, is given by (see Section A-1 in the Appendix):

$$\vartheta_{FE,n}^h(e, \mathbf{p}_r) = \beta_n + \kappa_n \left(\frac{e}{\exp \left(\int_n \beta_n \ln p_{rn} dn \right)} \right)^{-\varepsilon}. \quad (6)$$

where the subscript FE is a mnemonic for final expenditure. In this specification, prices \mathbf{p} and spending e conveniently enter only through a single summary statistic, that resembles a notion of real income. Moreover, household h 's Engel curve for product n , i.e. the change in the spending share as a function of income, is given by

$$\frac{\partial \vartheta_n^{FE,h}(e, \mathbf{p}_r)}{\partial e} = -\varepsilon \kappa_n \left(\frac{e}{\exp \left(\int_n \beta_n \ln p_{rn} dn \right)} \right)^{-\varepsilon} \frac{1}{e}. \quad (7)$$

Our class of preferences encompasses Cobb-Douglas preferences as a special case when $\kappa_n = 0$. A good n is luxury if $\kappa_n < 0$ and a necessity if $\kappa_n > 0$. Without any loss of generality, we order products inversely to the parameter κ_n so that good with a high index n are luxuries and goods with a low index are necessities. Note that, because $\int_{n=0}^1 \kappa_n dn = 0$, if some goods are necessities, some other goods must be luxuries.

Equation (7) also highlights that the slope of the Engel curves and the strength of income effects is governed by the parameter ε . This elasticity—that we label the *Engel elasticity*—plays a key role in our analysis.

3.2 Final Expenditure and Value Added

Equation (6) defines the expenditure share over final products $n \in [0, 1]$. Because the estimation methodology we adopt relies on sectoral employment data, it is useful to define a demand system in terms of the value added

⁶ An alternative class of nonhomothetic preferences recently proposed by Comin et al. (2020) has attractive properties for explaining long-run trends in the data. However, these preferences are less tractable when it comes to aggregation and moving between a value added to an expenditure approach. We leave to future research to explore how this class of preferences can be incorporated into the analysis.

produced by the three grand sectors F, G, and CS. In this section, we show how the PIGL preference specification in (5) allows us to seamlessly go back and forth between preferences and demand defined over final expenditure or value added. Most important, the Engel elasticity ε can be estimated using any of two approaches—a result that will prove very handy when we take the model to the data. In the rest of the paper, we often add time subscripts to emphasize exogenous and endogenous variables that are allowed to change over time in our application.

To link the final expenditure and the value added demand system, we establish relationships between the price vector for final goods and price indexes for the tradable inputs. The price of final good n in region r is given by

$$p_{rnt} = P_{rFt}^{\lambda_n F} P_{rGt}^{\lambda_n G} (\mathcal{A}_{rnt}^{-1} w_{rt})^{\lambda_n CS}. \quad (8)$$

Here, P_{rFt} and P_{rGt} are the prices in region r of the tradable food and industrial goods, and w_{rt} is the wage per efficiency unit of human capital in region r at t . Note that $\mathcal{A}_{rnt}^{-1} w_{rt}$ is the unit cost of the local CS input for producing the final good n .

Given perfect competition, the CES aggregation of regional varieties in (2), and the presence of iceberg trade costs, the price of the tradable goods s can be written as

$$P_{rst}^{1-\sigma} = \sum_{j=1}^R \tau_{rj}^{1-\sigma} A_{jst}^{\sigma-1} w_{jt}^{1-\sigma}, \quad \text{for } s \in \{F, G\}, \quad (9)$$

where $\tau_{rj} \geq 1$ captures the iceberg cost of shipping the variety j to region r . Note that, absent trade costs, the price of tradable goods would be equalized across regions and the regional variation in final good prices p_{rnt} would only come from differences in local wages and in CS productivity. However, the presence of iceberg costs increases the price faced by consumers in remote locations *ceteris paribus*.

Combining (5) with (8) allows us to represent consumers' preferences directly over sectoral value added. The following proposition can be established.

Proposition 1. *The indirect utility function of consumers in region r at time t as a function of household expenditure e (see (5)) can be written as*

$$\mathcal{V}_{rt}^{VA} (e, \mathbf{P}_{rt}) = \frac{1}{\varepsilon} \left(\frac{e}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} (\mathcal{A}_{rCS}^{-1} w_{rt})^{\omega_{CS}}} \right)^{\varepsilon} - \tilde{\nu}_F \ln P_{rFt} - \tilde{\nu}_G \ln P_{rGt} - \tilde{\nu}_{CS} \ln (\mathcal{A}_{rCS}^{-1} w_{rt}), \quad (10)$$

where $\mathbf{P}_{rt} = (P_{rFt}, P_{rGt}, w_{rt})$, P_{rFt} and P_{rGt} are given by (9),

$$\omega_s \equiv \int_{n=0}^1 \lambda_{ns} \beta_n \, dn, \quad \tilde{\nu}_s \equiv \int_{n=0}^1 \lambda_{ns} \kappa_n \, dn, \quad \text{for } s \in \{F, G, CS\}, \quad (11)$$

and

$$\mathcal{A}_{rCS} \equiv \exp \left(\int_n \frac{\beta_n \lambda_{CSn}}{\omega_{CS}} \ln \mathcal{A}_{rnt} \, dn \right). \quad (12)$$

The associated expenditure shares over sectoral value added aggregates—cf. (6)—are given by

$$\vartheta_{rst}^{VA,h} (e, \mathbf{P}_{rt}) = \omega_s + \tilde{\nu}_s \left(\frac{e}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} (\mathcal{A}_{rCS}^{-1} w_{rt})^{\omega_{CS}}} \right)^{-\varepsilon}. \quad (13)$$

Proposition 1 characterizes consumers' preferences and demand over sectoral value added. The indirect utility function (10) has the same functional form as in (5), which was defined over final goods. Equally important, the demand function over value added (13) features the same Engel elasticity parameter ε as the demand function (6) over final goods. This property of PIGL preferences will enable us to estimate ε from micro-data on household expenditure shares (on final goods) and then use it in the demand system defined over sectoral value added that is the base of our estimation.

The latter result may come as surprise because—as Herrendorf et al. (2013) point out—the mapping from the parameters of a final-expenditure demand to those of a value-added demand system generally involves the input-output matrix. In general, this is true in our model: while the two demand systems share the same elasticity ε , the parameters ω_s and $\tilde{\nu}_s$ are input-output weighted averages of the underlying final good demand parameters β_n and κ_n —see equation (11). More specifically, whether the demand for sectoral value added is income elastic depends on the correlation of the good-specific demand parameters with their factor intensities λ_{ns} . As seen in equation (13), the demand for sectoral value added is rising in income if $\tilde{\nu}_s < 0$. Equation (11) shows that this in turn is the case if income elastic *products* have a large *sectoral* input requirement. Hence, following our convention that products are inversely ordered by their $\tilde{\kappa}_n$, rich individuals buy a large CS value added content, if λ_{nCS} tends to be large for high- n products. By contrast, if all goods were produced with equal factor proportions, i.e. $\lambda_{ns} = \lambda_s$, or more generally if λ_{ns} were orthogonal to $\tilde{\kappa}_n$ for all s , the demand for sectoral value added would be homothetic and independent of prices (i.e., Cobb Douglas) even though the underlying demand for final goods was nonhomothetic. Formally, $\tilde{\nu}_s = \int_{n=0}^1 \lambda_{ns} \kappa_n dn = \lambda_s \int_{n=0}^1 \kappa_n dn = 0$, so that $v_{rst}^{VA,h}(e) = \omega_s$ —see equation (13).

Finally, Proposition 1 shows that, holding constant the prices of tradable goods and the spending e , the local CS sector enters the indirect utility function only through the summary statistic A_{rCS_t} . This is an index constructed as the (geometric) average productivity of the technologies of all final goods weighted by their local CS content and by the demand share of each (nontradable) final good. In this sense, A_{rCS_t} is a productivity index for the local CS sector, and a sufficient statistics to compute the welfare consequences of changes in the productivity of the CS sector. In other words, knowing the evolution of A_{rCS_t} over time allows us to quantify the welfare consequences of productivity growth in the service sector using information on sectoral aggregates.

3.3 Heterogeneity and Aggregate Demand

As the next step toward characterizing equilibrium, we derive the aggregate demand system. Under general preferences, aggregating non-homothetic preferences of individuals with heterogeneous income requires knowledge of the entire income distribution. However, PIGL preferences grant an analytical derivation of the aggregate demand system only in terms of prices, wages and structural parameters.

Suppose individuals have heterogeneous human capital that can be supplied to all sectors of production. Individual h 's income 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 time t . Differences in human capital can reflect differences in both ability and education. Empirically, we will relate the spatial variation in the distribution of q to observable differences in educational attainments.

Because our analysis abstracts from savings and capital accumulation, income equals expenditure. Then, equation (6) implies that the *aggregate* spending share on value added produced in sector s by consumers located in

region r is given by

$$\vartheta_{rst}^{VA} \equiv \frac{L_{rt} \int \vartheta_{rst}^{VA,h} (qw_{rt}) qw_{rt} dF_{rt}(q)}{L_{rt} \int qw_{rt} dF_{rt}(q)} = \omega_s + \nu_{rst} \left(\frac{A_{rCS}^{\omega_{CS}} \mathbb{E}_{rt}[q] w_{rt}^{1-\omega_{CS}}}{P_{rF}^{\omega_F} P_{rG}^{\omega_G}} \right)^{-\varepsilon}, \quad (14)$$

where

$$\nu_{rst} \equiv \frac{\mathbb{E}_{rt}[q^{1-\varepsilon}]}{\mathbb{E}_{rt}[q]^{1-\varepsilon}} \tilde{\nu}_s, \quad (15)$$

having defined—with slight abuse of notation—the expectation operator $\mathbb{E}_{rt}[x] \equiv \mathbb{E}[x; F_{rt}(x)]$. Comparing (14) with (6) highlights in what sense PIGL allows for a representative household: the *aggregate* demand system in (14) is isomorphic to that of a representative consumer in region r who earns the average income $\mathbb{E}_{rt}[q] w_{rt}$, and has the inequality-adjusted preference parameter ν_{rst} in (15) instead of the primitive parameter $\tilde{\nu}_s$. The inequality adjustment is the term $\mathbb{E}_{rt}[q^{1-\varepsilon}] / \mathbb{E}_{rt}[q]^{1-\varepsilon}$ which depends, in general, on the local distribution of efficiency units F_{rt} and thus can in principle vary across time and space.

The analysis further 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 . In this case, $\mathbb{E}_r[q] = \frac{\zeta}{\zeta-1} \underline{q}_r$ and $\mathbb{E}_r[q^{1-\varepsilon}] = \frac{\zeta}{\frac{\zeta}{1-\varepsilon}-1} \underline{q}_r^{1-\varepsilon}$, so that equation (15) boils down to

$$\nu_{rst} = \nu_s = \frac{\zeta^\varepsilon (\zeta - 1)^{1-\varepsilon}}{\zeta + \varepsilon - 1} \tilde{\nu}_s. \quad (16)$$

Thus, if income is Pareto distributed with a common tail parameter and an intercept \underline{q}_{rt} that can vary across space and time, all regions have the same “aggregate” parameter ν_s , which is proportional to the primitive individual preference parameter $\tilde{\nu}_n$. In this case, regional demand differences are solely driven by local prices, wages, human capital $\mathbb{E}_{rt}[q] \propto \underline{q}_{rt}$ and CS productivity A_{rCS} and despite heterogeneity in income and non-homothetic preferences, we can express the aggregate demand function of region r as a stable function of these objects. In addition, (14) stresses that they only affect aggregate spending as a combined region-specific aggregate statistic.

3.4 Equilibrium

Given the aggregate demand functions characterized in (14), we are now in the position to derive the equilibrium. The equilibrium can then be characterized in the following way.

Proposition 2. *Let P_{rF} and P_{rG} be given by (19). The sectoral labor allocation $\{H_{rF}, H_{rG}, H_{rCS}\}_r$ and local wages $\{w_{rt}\}$ are determined by the following equilibrium conditions:*

1. *Market clearing for local consumer services*

$$w_{rt} H_{rCS} = \left(\omega_{CS} + \nu_{CS} \left(\frac{A_{rCS}^{\omega_{CS}} \mathbb{E}_{rt}[q] w_{rt}^{1-\omega_{CS}}}{P_{rF}^{\omega_F} P_{rG}^{\omega_G}} \right)^{-\varepsilon} \right) w_{rt} H_{rt}, \quad (17)$$

2. *Market clearing for tradable goods F and G*

$$w_{rt} H_{rst} = \sum_{j=1}^R \pi_{rsjt} \left(\omega_s + \nu_s \left(\frac{A_{jCS}^{\omega_{CS}} \mathbb{E}_{jt}[q] w_{jt}^{1-\omega_{CS}}}{P_{jF}^{\omega_F} P_{jG}^{\omega_G}} \right)^{-\varepsilon} \right) w_{jt} H_{jt} \quad \text{for } s = F, G, \quad (18)$$

where

$$\pi_{rsot} = \tau_{ro}^{1-\sigma} A_{ost}^{\sigma-1} w_{ot}^{1-\sigma} / P_{rst}^{1-\sigma}. \quad (19)$$

3. Labor market clearing

$$H_{rFt} + H_{rGt} + H_{rCSt} = H_{rt}.$$

Proof. See Appendix. □

Proposition 2 fully characterizes the sectoral employment allocations across space and the local distribution of wages. The contrast between equations (17), on the one hand, and (18) on the other hand reflects the tradable nature of food and goods and the nontradable nature of CS. The local demand for value added in the CS sector hinges on local demand forces such as wages, human capital and local productivity. Instead, the demand for tradable goods originates from all localities. The spending share on food and industrial goods in Equation (19) follows from the CES demand structure across regional varieties subject to iceberg trade costs.

Proposition 2 also highlights that *sectoral* value added and employment are fully determined by the parameters ν_s and ω_s in Equation (11) and by the aggregate CS TFP index A_{rCSt} given in (12). They do not independently depend on product-specific preference parameters $[\beta_n, \kappa_n]_{n=0}^1$ nor on the efficiency terms $[A_{rnt}]_{n=0}^1$. Similarly, the local size of the industrial sector H_{rGt} only depends on A_{rGt} and we do not need to impose more structure on how producer services and blue collars interact in production. Crucially, $\nu_s, \omega_s, A_{rCSt}$ and A_{rGt} are not only sufficient to compute the equilibrium, but they also determine all our outcomes of interest, in particular individual and aggregate welfare. Below we show how we can recover these objects directly from our data.

3.5 Welfare and Inequality

The focal point of our contribution is a quantification of the welfare consequences of productivity growth sourcing from different sectors. Because preferences are nonhomothetic and CS are provided locally, productivity growth affects heterogeneously the utility of consumers with different income and residing in different regions. If goods with a high CS content have a high income elasticity, the welfare effects of productivity growth in CS are skewed toward the rich. More precisely, the expenditure share $\vartheta_{CS}(e, \mathbf{P}_{rt})$ exactly measures the welfare exposure of a change in prices at the individual level.⁷ Similarly, if (say) cities experience faster growth in CS productivity A_{rCSt} , city dwellers are going to be the main beneficiaries. In contrast, large part of the benefits from productivity growth in tradable sectors is exported.

For our purposes, we need to evaluate welfare at the individual level and at the regional level. At the household-level, we apply directly the indirect utility in equation (10). Substituting away the equilibrium price for CS, we can express the indirect utility of an individual living in region r as a function of the local wage w_{rt} , the local productivity of consumer services A_{rCSt} and the prices of the two tradable goods. At the regional level, we exploit the aggregation properties of PIGL preferences to calculate $\mathcal{U}_{rt}(w_{rt}, \mathbf{P}_{rt}) \equiv \int \mathcal{V}_r^{VA}(qw_{rt}, \mathbf{P}_{rt}) dF_{rt}(q)$. Note that this expression depends on the local skill distribution F_{rt} and a vector of local wages and prices. Plugging in the indirect utility function in (10) yields

$$\mathcal{U}_{rt}(w_{rt}, \mathbf{P}_{rt}) = \frac{\zeta^{1-\varepsilon} (\zeta - 1)^\varepsilon}{\zeta - \varepsilon} \times \left(\frac{1}{\varepsilon} \left(\frac{\mathbb{E}_{rt}[q] w_{rt}}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} (A_{rCSt}^{-1} w_{rt})^{\omega_{CS}}} \right)^\varepsilon - \sum_s \nu_s^\mu \ln P_{rst} \right), \quad (20)$$

⁷ Formally, letting $e(\mathbf{P}_{rt}, V)$ denote the expenditure function associated with the utility level V given the price vector \mathbf{P}_{rt} , $\partial \ln e(\mathbf{P}_{rt}, V) / \partial \ln P_{rst} = \vartheta_{rst}^h(e, \mathbf{P}_{rt})$.

where $\nu_s^\mu \equiv \nu_s \times ((\zeta - \varepsilon)(\zeta - (1 - \varepsilon)))/(\zeta(\zeta - 1))$. Hence, utilitarian welfare is akin to the indirect utility of a representative agent with average income $\mathbb{E}_{rt}[q]w_{rt}$ and a scaled taste parameter ν_s^μ that accounts for the income distribution (ζ) and the income elasticity (ε). Given this scaled taste parameter, the distribution F_{rt} only enters through the average income term $\mathbb{E}_{rt}[q]w_{rt}$.

3.6 Service-Led or Service-Biased Growth? Measuring Productivity in CS

In order to quantify the welfare impact of productivity growth in the consumer-oriented services, we need to estimate by how much CS productivity grew in a certain time interval. The equilibrium condition (17) contains the key for our identification strategy. In particular, it implies that the local CS employment share is given by

$$\frac{H_{rCS_t}}{H_{rt}} = \omega_{CS} + \nu_{CS} P_{rF_t}^{\varepsilon\omega_F} P_{rG_t}^{\varepsilon\omega_G} \times \left(\underbrace{\mathbb{E}_{rt}[q]}_{\text{Skills}} \times \underbrace{w_{rt}^{1-\omega_{CS}}}_{\text{Wages}} \times \underbrace{A_{rCS_t}^{\omega_{CS}}}_{\text{Productivity}} \right)^{-\varepsilon}. \quad (21)$$

Equation (21) highlights how our theory accommodate both service-led and service-biased growth. The equilibrium employment depends on the local supply of skills ($\mathbb{E}_{rt}[q]$), local wages (w_{rt}), and local productivity (A_{rCS_t}). The retail sector could then be large in urban districts such as Delhi or Bangalore either because local consumers are, on average, more educated and richer, or because (in addition) the productivity of CS is larger than in less developed areas in the country. For instance, more efficient department store chains open branches in large cities but not in rural districts where consumers must resort to smaller private retailers.

How does one separate the two growth sources? To attain identification, we leverage both the structure imposed by our theory and a variety of micro data. 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 (estimated below) and the prices of tradable goods that we retrieve from equilibrium conditions, we can use (21) to identify A_{rCS_t} . 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 simultaneously for a set of equilibrium prices, P_{rF} , P_{rG} , and w_r .

4 Empirical Analysis

In this section, we describe our data sources and discuss measurement issues.

4.1 Data and Geography

We use four datasets.⁸

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.”

⁸ A more detailed description of these datasets is deferred to Appendix Section B-1. Here, we highlight the main features.

The NSS, which provides the backbone for our analysis, is a household survey with detailed information on employment characteristics and households' location of residence. We use data for 1987 and 2011. The NSS yields measures of sectoral employment shares and average consumption (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. In the service sector, we must separate CS from PS. For some service industry, the assignment is clearcut. For instance, it seems natural to classify hotels and restaurants as part of the CS sector. In other industries, the distinction is less sharp. For instance, within legal services corporate lawyers provide PS whereas divorce lawyers provide CS. To solve this problem in a consistent way, we combine information from the Economic Census and the Service Survey to estimate the extent to which each service industry provides services to firms or to 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.⁹ 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 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 use this information to estimate the income elasticity ε , which is the key preference parameter for our analysis.

Geography: 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. We exclude two small districts that exist in 2011 but did not exist in 1987. We also exclude districts with less than 50 observations because they do not allow us to precisely estimate sectoral employment shares. In the end, we obtain 360 regions that cover the vast majority of the Indian territory. We label these regions districts. Section B-3 in the Appendix describes in detail how we construct this crosswalk.

4.2 Measurement

Consumer vs. Producer Services: We distinguish between PS and CS in a way that is consistent with our theory. Our preferred criterion is to classify as PS firms selling to other firms and as CS firms selling to consumers.¹⁰

⁹ As shown by 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.

¹⁰ We recognize that this is an imperfect approximation. In particular, in our theory PS are only purchased by firms, while CS can also be purchased as inputs for the production of other CS. For instance, wholesale firms serving local retailers should be part of the CS sector even though they sell to firms. Measurement error is also a concern. Some firms might report selling to individuals,

We perform robustness analysis to alternative classification criteria.

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 2, 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, respectively, 26% and 53% of employment caters mainly to other firms.

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 provided locally, we find it natural to merge them with services.¹¹ The construction sector serves both consumers (e.g., residential housing) and firms (e.g., business con-

although some of these individuals are small entrepreneurs using the service as inputs to production activities. In spite of these issues, we regard our criterion as a reasonable proxy measure to distinguish CS from PS in the data.

¹¹ In Section 7, we redo our analysis when we include construction in the manufacturing sector and show our results do not depend on

struction). 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.

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.¹²

In Figure 3, 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 3 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 districts 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.).

A concern is that our methodology might underestimate employment in PS relative to CS. In Section 7.2, we address this concern by showing that all our results are qualitatively robust to reasonable alternative measurement choices that give a more prominent role for PS employment.

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, which reflect differences in the endowment of effective units of labor.¹³ 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 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.

this particular choice.

¹² 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.

¹³ In Section 7.3.2 below, we extend our model to allow for imperfect substitution of skills across sectors.

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 $\mathbb{E}_{rt}[q] = \sum_e \exp(\rho \times e) l_r(e)$, where $l_r(e)$ 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 $\mathbb{E}_{rt}[q_i] = \frac{\zeta}{\zeta-1} \underline{q}_{rt}$.

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.¹⁴ 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.

5 Estimation: Identification and Results

We now turn to the estimation of the model. Our approach is in the tradition of development accounting—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. Because we rely on the equilibrium structure of our model, we refer to our method as *equilibrium development accounting*.

¹⁴ For ease of comparison with Figure 1, 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.

The centerpiece of the methodology is the distinction between structural parameters and local productivities. The model has eight structural parameters describing preferences and the distribution of skills

$$\mathbf{\Omega} = \left\{ \underbrace{\varepsilon, \nu_{CS}, \nu_F, \omega_{CS}, \omega_F, \sigma}_{\text{Preference parameters}}, \underbrace{\rho, \zeta}_{\text{Human capital}} \right\}.$$

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

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

Given the parameter vector $\mathbf{\Omega}$, there exists a one-to-one mapping from the equilibrium skill prices $\{w_{rt}\}$ and sectoral employment allocations $\{H_{rst}\}$ to the underlying productivity fundamentals in \mathbf{A}_t . In Section 5.1, we describe how we estimate the vector of structural parameters $\mathbf{\Omega}$. In Section 5.2, we discuss the estimation procedure for \mathbf{A}_t and its results.

5.1 Estimation of Preference and Human Capital Parameters

The Income Elasticity ε : The crucial parameter in our analysis is the income elasticity ε , which determines how fast the demand for agricultural goods shrinks as income rises. To estimate ε , we use the cross-sectional relationship between income and expenditure shares at the household level and estimate ε via indirect inference. We recall that this approach, which involves expenditure data, differs from the value added approach based on employment data adopted thus far. While in general the two demand systems have different parameters, Proposition 1 establishes that under PIGL preferences the two demand systems share the same Engel elasticity ε .

In terms of our theory, let the set \mathcal{F} denote the subset of the product space $[0, 1]$, that contains all products classified as food. According to our theory, the spending share on food is then given by

$$\vartheta_{FE, \mathcal{F}}^h(e, \mathbf{p}_r) = \int_{n \in \mathcal{F}} \vartheta_{FE, n}^h(e, \mathbf{p}_r) = \beta_{\mathcal{F}} + \kappa_{\mathcal{F}} \left(\frac{e}{\exp(\int_n \beta_n \ln p_{rn} dn)} \right)^{-\varepsilon}, \quad (22)$$

where $\beta_{\mathcal{F}} \equiv \int_{n \in \mathcal{F}} \beta_n dn$ and $\kappa_{\mathcal{F}} \equiv \int_{n \in \mathcal{F}} \kappa_n dn$. Note that $\beta_{\mathcal{F}}$ is the asymptotic expenditure share on food as income grows, i.e. $\lim_{e \rightarrow \infty} \vartheta_{FE, \mathcal{F}}^h(e, \mathbf{p}_r) = \beta_{\mathcal{F}}$. If $\beta_{\mathcal{F}}$ is small, Equation (22) yields a log-linear relationship between household income and food shares

$$\ln \vartheta_{FE, \mathcal{F}}^h(e, \mathbf{p}_r) = \varepsilon \left(\int_n \beta_n \ln p_{rn} dn \right) - \varepsilon \times \ln e + \ln \kappa_{\mathcal{F}}. \quad (23)$$

We can then estimate ε from the regression

$$\ln \vartheta_{FE, \mathcal{F}rt}^h = \delta_r + \gamma \times \ln e_h + x_h' \psi + u_h, \quad (24)$$

where $\vartheta_{FE, \mathcal{F}rt}^h$ denotes the observed food expenditure share of household h in region r , e_h denotes total household spending, 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.

	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 (24). 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Comparing (24) with (23) shows that the estimated parameter γ coincides with the structural parameter ε while the term $(\int_n \beta_n \ln p_{rn} dn)$ is absorbed in the district fixed effect δ_r . Moreover, the additional household level controls x_h capture cross-sectional variation in preferences (κ_F) which we abstract from in our theory but which might be correlated with household income e . The linear regression was derived by setting β_F to zero. To avoid to rely on this approximation, we do not simply equate ε to $\hat{\gamma}$ in Equation (23); rather, we estimate ε via indirect inference, namely, we estimate Equation (24) using the empirically observed coefficient $\hat{\gamma}$ as a moment. To implement this approach, we compute expenditure shares on food at the individual level, i.e. $\vartheta_{FE,\mathcal{F}}^h(e, \mathbf{p}_r)$. Note that both $\kappa_{\mathcal{F}}$ and $\beta_{\mathcal{F}}$ are unknown, because we do not estimate preferences at the level of the continuum of final goods (i.e. β_n and κ_n). To address this problem, we exploit the relationship between value added share and final expenditure shares highlighted in Proposition (1). As we show below, we can estimate ω_F and $\tilde{\nu}_F$ from value added data—recall that $\omega_F = \int_n \lambda_{nF} \beta_n$ and $\tilde{\nu}_F = \int_n \lambda_{nF} \kappa_n$ from (11). If, by virtue of being classified as food, the final goods in the set \mathcal{F} have a large value added content of agricultural goods, whereas goods not classified as food have a low value added content of agricultural goods—i.e. $\lambda_{nF} \approx 1$ for $n \in \mathcal{F}$ and $\lambda_{nF} \approx 0$ for $n \notin \mathcal{F}$ —then, $\beta_{\mathcal{F}} \approx \omega_F$ and $\kappa_{\mathcal{F}} \approx \tilde{\nu}_F$. We can then compute $\vartheta_{FE,\mathcal{F}}^h(e, \mathbf{p}_r)$ in the model and estimate the empirical specification (24).¹⁵

Table 3 reports the estimation results. We cluster standard errors at the region level to account for the correlation of spending shares through regional prices. The first column refers to a baseline specification which controls 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 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 add 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

¹⁵ To estimate (24) 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.

the Indian government. In both specifications, the estimate is very close to the baseline estimate.

In column 4, we present the results from an IV specification addressing concerns about measurement error in $\ln e_h$ that could bias downward the estimated Engel elasticity. We instrument total expenditure with a full set of three-digit occupation fixed effects.¹⁶ Expectedly, these fixed effects strongly predict total expenditure as shown by the large F-statistic. The exclusion restriction is that occupational choices only affect spending shares through their effect on income. The IV estimate is only slightly larger than the OLS estimate. In Appendix Section B-5 we report the results from additional specification 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 the results are robust to all estimates reported in Table 3.

Other Preference Parameters: For the remaining parameters of the demand system, we follow the value added approach. 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 (15). We estimate ν_s directly from the data and infer the structural micro parameters $\tilde{\nu}_s$ given an estimate of the inequality parameter ζ . Identifying $\tilde{\nu}_s$ separately from ν_s is only required to quantify the welfare consequences of service-led growth, not to estimate the model.

We turn then to the identification of the parameters ω_s and ν_s for $s \in \{F, G, CS\}$. In Appendix Section A-2, we show that the set of market clearing conditions and the Walras' law yield the following equation:

$$\sum_{r=1}^R w_{rt} H_{rFt} = \omega_F \sum_{r=1}^R w_{rt} H_{rt} + \nu_F \sum_{r=1}^R \left(\omega_{CS} - \frac{H_{rCS t}}{H_{rt}} \right) w_{rt} H_{rt}, \quad (25)$$

which must hold for $t = 1987$ and $t = 2011$, assuming that preference are stable over time. Since earnings and labor allocations are observable, this yields two equations in three unknown parameters ω_F , ν_F , and ω_{CS} . We externally calibrate ω_F and then use (25) to identify ν_F and ω_{CS} .

The parameter ω_F pins down the asymptotic expenditure share on agriculture goods to which ϑ_F declines as total expenditure grows. In the US, the agricultural employment share (as well as its value added share) is about 1%. Hence, we set $\omega_F = 0.01$. Incidentally, this implies that the structural Engel elasticity ε is close to the empirical estimate $\hat{\gamma}$ in (23). Having set $\omega_F = 0.01$, Equation (25) yields then $\nu_F = 1.276$ and $\omega_{CS} = 0.69$.¹⁷ This implies an asymptotic expenditure share on CS of 69% that we view as reasonable. For instance, the value added share of the service sector in the US (that is not a targeted moment) has averaged at 77% throughout the last decade. This empirical shares includes both PS and CS.

Consider, next, ν_{CS} . In Appendix Section A-2, we show that ν_{CS} is not separately identified from the productivity $A_{rCS t}$.¹⁸ Hence, without loss of generality, we normalize it to -1. Given these estimates, the restrictions imposed by PIGL preferences identify ω_G and ν_G . The asymptotic share of the good producing sector (that, recall, includes both manufacturing production and PS) is 30%. Moreover, $\nu_M = -(\nu_F + \nu_{CS}) = -0.276$. This implies that industrial goods are also luxury goods, although their income elasticity is smaller than for the CS.

¹⁶ The survey assigns the occupation of the household member with the highest earnings to the entire household.

¹⁷ 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.

¹⁸ This means that we cannot identify the level of $A_{rCS t}$. However, this is not important for our goals. Under the assumption of stable preferences, we can calculate the growth over time of $A_{rCS t}$ that is a focal point of our analysis.

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>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.

Finally, we set the inter-regional trade elasticity σ to a consensus estimate in the literature and assume $\sigma = 3$.

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 a standard Mincerian regression $\ln y_{irt} = \ln w_{rt} + \rho s_i + \ln v_i$. Thus, we can estimate ρ from the within-region variation between earnings and education, which we can measure from the NSS data. This yields an average annual rate of return of 5.6%, which is on the lower end of standard Mincerian regressions, although broadly in line with the findings of recent studies for India using the NSS—see Singhari and Madheswaran (2016). In Section 7, we discuss the robustness of our results to using a higher return to education.

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 $\zeta \approx 3$ (see Appendix Section B-7).

5.2 Estimation of Productivity Fundamentals \mathbf{A}_t

Given the structural parameter vector $\mathbf{\Omega}$, 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-2 in the Appendix for details, but we describe the main intuition here. Consider first the identification of A_{rCS_t} . Equation (21) implies we can uniquely solve for A_{rCS_t} as

$$A_{rCS_t} = \left(\frac{(-\nu_{CS})}{\omega_{CS} - \frac{H_{rCS_t}}{H_{rt}}} \right)^{\frac{1}{\omega_{CS}} \frac{1}{\varepsilon}} P_{rFt}^{\frac{\omega_F}{\omega_{CS}}} P_{rGt}^{\frac{\omega_G}{\omega_{CS}}} (\mathbb{E}_{rt}[q] \times w_{rt}^{1-\omega_{CS}})^{-\frac{1}{\omega_{CS}}}. \quad (26)$$

Controlling for the level of human capital $\mathbb{E}_{rt}[q]$ and the equilibrium factor price w_{rt} , CS productivity is in-

creasing in the observed employment share $\frac{H_{rCS_t}}{H_{rt}}$.¹⁹ Conversely, holding the employment share $\frac{H_{rCS_t}}{H_{rt}}$ 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 (i.e., driven by income effects—cf. $\mathbb{E}_{rt}[q] w_{rt}^{1-\omega_{CS}}$) is a key contribution of our equilibrium accounting methodology.

The procedure to estimate productivity in tradable sectors is different. Equation (18) implies relative productivity across two locations is given by

$$\frac{A_{rs}}{A_{js}} = \left(\frac{H_{rs}}{H_{js}} \right)^{\frac{1}{\sigma-1}} \left(\frac{w_r}{w_j} \right)^{\frac{\sigma}{\sigma-1}} \quad \text{for } s = F, G. \quad (27)$$

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

Equation (27) allows us to estimate the cross-region distribution of productivities in the two tradable sectors. In addition, we need to determine the average productivity levels in agriculture and industry in 1987 and 2001. To this aim, we first choose the average producer price in the industrial sector to unity as the numeraire: $\sum_r (p_{rGt})^{1-\sigma} = 1$, for $t \in \{1987, 2011\}$, where $p_{rGt} = \frac{w_r}{A_{rGt}}$ is the producer price of region r 's local variety. This choice is unimportant. Then, we set $\sum_r \frac{e_{rt}}{P_{rGt}} = GDP_t$, for $t \in \{1987, 2011\}$, where GDP_t is the measured real GDP for India and P_{rGt} are theoretical industrial good prices given by (9).²⁰ In words, we equate the total expenditure deflated using model-generated regional industrial prices to the real GDP of India. This implies we measure GDP growth in terms of the industrial good.²¹ Finally, we must determine the average productivity in agriculture. To this aim, we impose that the (weighted average across Indian districts) relative price of agriculture to industrial goods in the model matches the data counterpart in both 1987 and 2011. More details are provided in Appendix B-4.

Figure 4 summarizes the 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 observed urbanization rate in 1987. 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.

Both 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 (income effect) or of an abundance of human capital, but also of high CS productivity relative to rural areas. 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 (27)). 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

¹⁹ Recall that, if CS are a luxury, $\nu_{CS} < 0$ and $\frac{H_{rCS_t}}{H_{rt}} < \omega_{CS}$.

²⁰ The unit of measure for GDP is an additional degree of freedom. In practice, we set $GDP_{1987} = 1$ and $GDP_{2011} = 1 + g_{87,11}$, where $g_{87,11}$ is the observed growth rate of real GDP in India between 1987 and 2011.

²¹ Because of nonhomothetic preferences, we cannot define a standard consumption price index. For comparison, we calculated income growth for a fictitious agent endowed with the median earnings 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 income growth using a Laspeyres and a Paasche index. The resulting real income growth in the two cases is 1.82 and 4.86, respectively. Our calibration yields an income growth factor of 2.60, which is in between.

with the employment share of industry, implying a wider productivity dispersion.

Figure 4 describes the spatial variation in the *level* of sectoral productivity. We are equally interested in the distribution of sectoral productivity growth between 1987 and 2011. Using our estimates A_{rst} we can calculate sectoral productivity growth between 1987 and 2011 for each district. We summarize the distributions of annualized productivity growth in Table 5.

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. Two facts are salient. First and foremost, productivity in the CS sector grew in the vast majority of districts. Second, productivity growth was unequal across space, particularly so in the CS sector.²²

	Sectoral productivity growth					
	10th	25th	50th	75th	90th	Aggregate
	<i>Service-led growth</i>					
Consumer Services (g_{rCS})	-1.4	0.8	3.6	8.6	15.0	8.2
	<i>Growth in other sectors</i>					
Agriculture (g_{rF})	0.2	1.2	2.1	3.1	4.0	2.5
Industry (g_{rG})	0.8	2.1	3.3	4.7	6.0	4.5

Table 5: 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%.

5.3 Non-targeted Moments

In this section, we highlight some implications of the estimated model for non-targeted data moments. In particular, we document that (i) the quantitative predictions are consistent with the empirical regional variation of food prices, (ii) the Engel elasticity ε estimated from the food expenditure shares is also consistent with the expenditure shares of CS, (iii) the predictions of our model for average productivity growth at the sectoral level line up well with existing estimates from the Groningen Productivity Database and (iv) the implied elasticities of substitution for the value added of different sectors are in line with existing estimates from other studies.

Regional Food Prices: Our estimated model predicts local prices that can be compared with the data. The expenditure survey reports both the total expenditure and the total quantity bought for a variety of food items.²³ We can compute the average price of variety v in region r , p_{vr} as the ratio between total expenditure and quantity.

Given this information, we compute the average food price in region r as the regional fixed effect δ_r in the regression

$$\ln p_{vr} = \delta_r + \delta_v + u_{rj}. \quad (28)$$

Note that the variety-specific fixed effect δ_v controls for differences in the unit of measurements across districts. The left panel of Figure 5 shows the correlation between the estimated $\hat{\delta}_r$ and the regional price of agricultural goods

²² 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).

²³ In Section B-1.5 in the Appendix, we report the distribution of expenditure shares across the different varieties. The largest 10 categories account for 50% of household spending. These contain among others rice, cereal, oil, and milk.

in the model, that is $\ln p_{rFt}$. According to our model, this is the local price of final goods that consists mostly out of agricultural inputs (i.e. $\lambda_{nF} \approx 1$).²⁴ The two measures are strongly positively correlated, even though we do not use the data on local food prices as targets of our estimation. In the model, the variation in local food prices reflects local agricultural productivity, local wages and food prices of close-by locations (which have low transport costs). In the Appendix, we also show [TO BE ADDED] that local food prices are correlated with the local agricultural employment share both in the data and the model.

Spending on Consumer Services: We have used data on food shares to estimate the Engel elasticity ε . Alternatively, we could have used data on the expenditure share of CS. We prefer food expenditures for two reasons. First, expenditure on food items is likely to be better measured. Second, the log-linear specification in (24) is particularly informative about ε if ω_s is small, because our theory then exactly implies a log-linear relationship, and the distinction between final expenditure and value added becomes less important. The evidence from advanced economies suggest that, while $\omega_F \approx 0$, ω_{CS} is significantly larger than zero as rich economies spend large share of their income on services.

Reassuringly, the results described in Appendix Section B-1.5 show that one would obtain similar results if one had used existing data for CS spending instead of food expenditure. Specifically, we run—in the model and in the data—the same specification as in (24) 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.

We find that CS are luxuries: high-income households spend a higher share on CS. Quantitatively, we find the elasticity between the spending share on consumer service goods and individual income to be between 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 obtain a coefficient of about 0.4. Hence, even though we do not use the data on CS spending to estimate the model, the implied empirical elasticities are broadly 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 expenditure regression, this variation is captured by the regional fixed effects (see (24)).

In the right panel of Figure 5, 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. There is a robust positive relationship: cities are particularly productive in CS. Hence, even though we do not use the information from the CS expenditure data, this is qualitatively consistent with our structural estimation, which also implied a positive correlation between A_{rCS} and urbanization - see Figure 4.

Alternative Estimates of Sectoral Productivity Growth: While we are not aware of productivity growth estimates at the regional level, the Groeningen Productivity Database provides estimates of nationwide productivity growth at the sector level that can be compared with our estimates. In the first two columns of Table 6 we report annual growth in value added per worker between 1987 and 2011 for a variety of subsectors as reported in the Groeningen database. In columns 3 and 4 we report the results from our model. More specifically, we report the average regional growth rate of value added per worker weighted by the relative share of aggregate expenditure in

²⁴ The local price of good n is given by $\ln p_{rnt} = \lambda_{Fn} \ln p_{rFt} + \lambda_{Gn} \ln p_{Grt} + \lambda_{CSn} \ln \frac{w_{rt}}{A_{rnt}}$. For $\lambda_{nF} \approx 1$ this reduces $\ln p_{rnt} = \lambda_{Fn} \ln p_{rFt}$.

each district.

The comparison shows that our model-based accounting approach yields aggregate productivity growth rates that are similar to the corresponding Groeningen figures. First and foremost, the 8.7% estimated aggregate productivity in the CS is in the ball park of the productivity growth of the service sector reported in the Groeningen data (7.9%). The exact comparison is complicated because in our classification the service sector includes CS, PS, and Government Services—from which we abstract, and whose productivity growth is lower than for other service categories. But even at the industry category level, the results are broadly consistent. For instance, the category Trade, Restaurant, and Hotel—which includes mostly CS—has a productivity growth above 7% in the Groeningen data. The category Finance, Insurance, etc.—which includes both PS and CS in our classification—attains the highest productivity growth in the Groeningen data.

Our estimated productivity growth for the two tradable industries also accord well with the Groeningen data. Productivity growth in agriculture is the lowest both in the Groeningen data and by our estimate, although the latter is significantly larger: 2.5% vs. 1%. Finally, our 5% estimated productivity growth in the industrial sector is close to the estimated productivity growth in manufacturing in the Groeningen data (5.6%). Because our industrial sector also includes part of the construction, utilities, and service activities, our estimates imply a somewhat lower productivity growth than does the Groeningen data.

While this aggregate evidence is, by construction, silent on the extent to which productivity growth is unbalanced across space, we find it reassuring that our model delivers productivity growth estimates that are broadly consistent with existing aggregate data.

Groeningen Data	Value Added pw Growth	Our Model	Value Added pw Growth
Agriculture	0.98	Agriculture	2.83
Manufacturing	5.63	Industrial goods	4.91
Mining	3.42		
Construction	6.44		
Utilities	5.68		
All services	7.90	Consumer services	8.67
Finance/insurance/ real estate/business services	10.22		
Trade/restaurants/hotels	7.22		
Transport/storage/communication	7.89		
Community/social/personal services	4.97		
Government services	5.08		

Table 6: SECTORAL VALUE ADDED GROWTH: The left panel of the table reports sectoral productivity (value added per worker) growth from the Groeningen Productivity Database for years 1987-2011. The right panel reports the average growth rate of sectoral productivity (value added per worker) as implied by our model.

Elasticities of Substitution: Finally, we can calculate the elasticity of substitution between the value added of different sectors with existing estimates in the literature based on alternative methodologies. For the class of PIGL preferences, the elasticity of substitution is not a constant structural parameter but varies with relative prices and total expenditure.²⁵ We defer the detail of the analysis to Section B-6 in the Appendix. The main finding is that services and industrial goods are complements, with an elasticity of substitution ranging between 0.6 and

²⁵ 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-3 in the Appendix.

0.9 across the urbanization quantiles. A relationship of complementarity is consistent with the common wisdom in the literature. On the other hand, we find that food and CS are, on average, substitutes with an elasticity of substitution ranging between 1.6 (least urbanized districts) and 1.2 (most urbanized districts). Finally, food and industrial goods are also substitutes but with lower elasticities ranging between 1.1 and 1.2.

6 The Unequal Welfare Effects of Service-Led Growth

This section contains the main results of the paper. We address the following related questions: (i) How important a driver of rising living standards was productivity growth in the service sector? (ii) How skewed were the benefits of service-led growth across different socioeconomic groups? (iii) How important was productivity growth in CS in promoting structural change in India?

To quantify the macroeconomic impact of these growth estimates reported in Table 5, 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 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, we uncover a great deal of heterogeneity in both dimensions. In addition, we can also compute the implications for the structural transformation.

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 Methodology

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

As discussed in Section 3.3, the PIGL demand system allows us to capture such heterogeneous welfare effects 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 $\{w_r^{CF}, \mathbf{P}_r^{CF}\}_r$. In our experiments below, changes in wages and prices are caused by counterfactual changes in sector-region-specific productivities.

Using the indirect utility function V given in (5), \bar{w}^h is implicitly defined by

$$\mathcal{V}^{VA}(\bar{w}^h(w_r^{CF}, \mathbf{P}_r^{CF} | \mathbf{P}_r, q), \mathbf{P}_r) = \mathcal{V}^{VA}(qw_r^{CF}, \mathbf{P}_r^{CF}). \quad (29)$$

Note that we can calculate the welfare-equivalent income \bar{w}^h from $\{w_r^{CF}, \mathbf{P}_r^{CF}, \mathbf{P}_r\}_r$ conditional on the individual human capital q . In a similar vein, we can calculate the utilitarian welfare effects at the district level. Exploiting the aggregation properties of PIGL, we can determine the level of *regional* spending power $\bar{w}(w_r^{CF}, \mathbf{P}_r^{CF} | \mathbf{P}_r, \mathbb{E}_r[q])$ the representative agent in district r facing prices P_r would require to attain indifference. This is implicitly defined by

$$\mathcal{U}(\bar{w}(w_r^{CF}, \mathbf{P}_r^{CF} | \mathbf{P}_r, \mathbb{E}_r[q]), \mathbf{P}_r) = \mathcal{U}(\mathbb{E}_r[q]w_r^{CF}, \mathbf{P}_r^{CF}), \quad (30)$$

where \mathcal{U} is defined in (20).²⁶

Given \bar{w}^h and \bar{w} defined in (29) and (30), we can calculate the equivalent variations relative to 2011 as

$$\Delta\mathcal{W}_r^h(q) \equiv \frac{\bar{w}^h(w_r^{CF}, \mathbf{P}_r^{CF} | \mathbf{P}_r, q)}{qw_{r,2011}} - 1 \quad \text{and} \quad \Delta\mathcal{W}_r \equiv \frac{\bar{w}(w_r^{CF}, \mathbf{P}_r^{CF} | \mathbf{P}_r, \mathbb{E}_r[q])}{\mathbb{E}_r[q]w_{r,2011}} - 1. \quad (31)$$

Here, $\Delta\mathcal{W}_r^h(q)$ is the change in income that an individual with human capital q living in district r in 2011 would require to attain the same level of utility as 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 2011 income and a counterfactual allocation in which productivity in a particular sector is reset to the 1987 level. Similarly, $\Delta\mathcal{W}_r$ is the analogue equivalent variation for the representative agent in region r . The heterogeneity in the the equivalent variations across individuals and districts allow us to quantify the unequal effects of sectoral productivity growth.

We can also calculate the equivalent variation at the national level by averaging across the regional income variations using regional employments as weights:

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

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

This section provides the main results in the paper. We construct counterfactual scenarios by sequentially setting to zero the productivity growth of each sector in the period 1987–2011. In each experiment, we calculate the counterfactual equilibrium prices and income levels and use them to calculate the equivalent variations using to the methodology discussed above. We calculate equivalent variations 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 with the aid of Figure 6. We report the welfare change

²⁶ Using equations (29) and (5) and (30) and (20) we get that

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

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

(equivalent variation) associated with counterfactually shutting down productivity growth in each sector. We also construct a counterfactual in which we set to the 1987 level the human capital in each region, as measured by educational attainment and by the estimated return to education. We calculate the distribution of these welfare changes taking into account the sampling variation. We report these distributions as a boxplot. Each box shows 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. Note that the more negative the welfare loss in the plot, the larger is the welfare gain associated with productivity growth in a sector. The solid horizontal line marked 1987 yields the welfare effect of simultaneously shutting down productivity growth in all sector and human capital accumulation.

The first important result is that a substantial part of the total welfare gains appear to be service led. On average, the Indian population would have been willing to reduce its 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 6 also shows that agricultural productivity was an important source of welfare improvement. The salience of agriculture is hardly surprising given its large employment share in India. The smaller welfare effects of productivity growth in the industrial sector is perhaps more surprising. The equivalent variation amounts to 17% and is very precisely estimated. Hence, productivity growth in CS is more important in welfare terms than productivity growth in the industrial sector. The welfare consequences of human-capital accumulation are modest. These welfare gains are relatively modest, namely a mere 9% of 2011 income. Note these differences are based on private returns to education. To the extent to which there is a wedge between the private and social return to education (e.g., a better-educated labor force favors technical progress in some sector) this is not part of our calculation. Overall, Figure 6 shows that service-led growth plays an important role in India since 1987. In Section 7, we scrutinize this finding through a battery of sensitivity checks and show that the result is very robust.

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 5, we estimate that productivity growth differs across regions. Second, the non-homothetic preferences imply that consumers on different levels of the income ladder care differently about different sectoral productivity growth: growth in CS and (to a lesser extent) industrial goods is particularly beneficial for the rich, whereas growth in the agricultural sector mostly benefits the poor. We first focus on the spatial dimension. We 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.²⁷ These results are shown in the left panel of Figure 7. As in Figure 6, we display our results using boxplots which also illustrate the uncertainty in our estimates.

The welfare consequences of productivity growth vary widely across space. Unsurprisingly, the benefits of agricultural productivity growth is skewed toward rural areas. 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. The equivalent variation declines sharply in the top quintile, where productivity growth in agriculture is only worth 16% of the 2011 income. By contrast, the benefits from productivity growth in CS and the industrial

²⁷ In all experiments we perform, we shut down sectoral productivity growth simultaneously in all locations. Thus, part of the benefits spread around India through trade across districts.

sector are skewed toward urban locations. This pattern is most pronounced for the CS sector whose productivity growth is worth 41% 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 the same across the lower four quintiles of the distribution, the top urban quintile seems to be qualitatively different: there, a welfare gains were mostly service led while the benefits from agricultural productivity growth in India are modest.

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 7 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:

The industrial good sector comprises both manufacturing production and PS. It is possible to disentangle the welfare effects of productivity growth in these two sectors. Although productivity growth is faster in PS, the employment share of these services is small. As a result, the welfare gains accrue mostly from productivity growth in the manufacturing sector. For instance, in our baseline specification, the average welfare effect of producer services productivity growth accounts for a mere 2.4% of the average welfare effect of industry productivity growth. We study below the robustness of this result to alternative ways of splitting the service sector into CS and PS.

6.3 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 8. 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.²⁸ Like in Figures 6 and 7 above, we again display the distribution of these effects though 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.

²⁸ The figure shows results for employment in effective units of labor, which we label “employment” with a slight abuse of terminology.

Figure 8 shows that productivity growth in CS was responsible for the lion’s share of the observed structural transformation. The left panel shows that absent productivity growth in CS the agricultural employment share would have been 60% instead of 50%. Thus, 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 had productivity not grown in the CS sector.

Figure 8 also highlights an important role of income effects (*service-biased* growth.) Panel (b) shows that even without any productivity growth in the CS sector the employment share of CS would have grown by five percentage points between 1987 and 2011. Yet, 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 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 modest effects on structural change. If anything, it marginally *increased* employment in agriculture and slowed down employment growth in industry and CS. This result is in line with the findings of Foster and Rosenzweig (2004) on the effects of the Green Revolution and those of Kelly et al., 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 be a more rural economy.

7 Robustness

In this section, we perform robustness analysis for the welfare effects reported in Figures 6 and 7. We consider three sets of issues. 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 study generalizations of the model. In particular, we extend our model to an open economy setting, we consider a production structure where skills are imperfectly substitutable, and we allow for workers to be mobile across space so that the spatial distribution of the population endogenously responds to changes in local wages and prices.

7.1 Sensitivity to Structural Parameters

We consider the parameters governing preferences and skills. All results are based on re-estimating the entire model.

Preferences. We focus on the parameters ω_F and ε that we, respectively, calibrate and estimate outside of the theory. The other preference parameters are either point identified in our theory or pinned down by matter of normalization.

We calibrated ω_F to 1% so as to match the value added (and employment) share of the US farming sector in 2017. However, the value added share of agriculture is larger than 1% in many industrial countries (e.g. 2% in Italy and France, 3% in Spain.) Therefore, we consider a range of larger values of ω_F . Panel (a) of Figure 9 shows the implied welfare impact of sectoral productivity growth is essentially independent of ω_F .

Panel (b) of Figure 9 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, the welfare effects of service-led growth are decreasing in ε . The results shown in Panel (b) of Figure 9 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 9 shows that for any $\varepsilon < 0.5$, a large share of productivity growth sources in the service sector. Hence, for the growth of the service sector to be preeminently driven by income effect, we should believe in a much higher income elasticity than is indicated by the household-level data.

Skills. In the lower panels of Figure 9, 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 9, our main results are not sensitive to this parameter.

Panel (d) of Figure 9 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 9 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
Baseline	-21.2	-26.5	-17.3	-8.7	-24.0	-16.5	-17.3	-41.3	-12.0	-23.3
	<i>Alternative measurement choices (Section 7.2)</i>									
Double PS	-21.2	-22.3	-20.0	-8.7	-23.7	-18.0	-19.6	-28.1	-13.5	-27.2
ICT & Business to PS	-21.5	-20.3	-18.7	-8.6	-24.1	-17.8	-19.9	-23.7	-12.6	-25.8
Construction to manufacturing	-20.4	-22.7	-22.4	-8.7	-25.5	-12.5	-3.1	-48.5	-12.9	-33.3
	<i>Alternative modelling assumptions (Section 7.3)</i>									
Open economy	-20.3	-23.0	-19.5	-8.8	-24.0	-16.9	-15.9	-34.6	-15.2	-24.3
Open economy (large ICT)	-20.8	-18.9	-19.5	-8.0	-23.8	-16.7	-17.4	-22.2	-19.5	-14.9
Imperfect skill substitution	-25.2	-26.7	-16.5	-16.2	-29.6	-18.6	-12.0	-45.4	-11.0	-22.6
Spatial labor mobility $\eta = \frac{2}{3}$	-21.1	-28.4	-17.3	-8.7	-23.5	-16.9	-19.6	-41.8	-12.3	-22.7
Spatial labor mobility $\eta = \frac{4}{3}$	-21.0	-29.1	-17.2	-8.6	-23.0	-17.1	-20.5	-41.5	-12.6	-22.2
Spatial labor mobility $\eta = 2$	-20.8	-29.2	-17.2	-8.6	-22.5	-17.3	-20.8	-41.0	-12.8	-21.8

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.²⁹

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-capital-adjusted) employment share of PS is twice as large as in our benchmark estimate in each service industry shown in Figure 2. Second, we assume the entire ICT and business service industries serve manufacturing firms (while retaining our baseline approach for the remaining service industries).

The results are shown in rows 2 and 3 of Table 7. The first four columns report the aggregate welfare effect (ΔW), shown in Figure 6. The last six columns focus on the spatial heterogeneity (ΔW_r), shown in Figure 7. 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, productivity growth in CS continues to be a large driver of welfare changes.

Finally, we turn to 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 Generalizations of the Theory

In this section, we consider three generalizations of the theory. First, we incorporate international trade. Second, we consider an environment where skills are imperfectly substitutable and skill-intensity varies across sectors. Third, we allow for workers to be spatially mobile.

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 extend our model to an open-economy

²⁹ 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.

environment. For brevity, we only summarize the main features of the extended model. The technical analysis can be found in Appendix Section A-4.

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.³⁰ 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.³¹ 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 is now pinned down by two parameters, we set both A_{rs} and Z_{rs} to the respective 1987 level when running counterfactuals. 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 10, 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.

7.3.3 Spatially Mobile Workers

In our baseline model, workers are exogenously assigned to regions. In the counterfactual analysis, we assumed people to be spatially immobile. However, people could decide to leave urban areas in response to sector-region productivity changes. For instance, a counterfactual decline in CS productivity could lead people to move out of cities. To gauge the quantitative importance of labor mobility, we reestimate our model in the presence of a migration choice. We model migration as a discrete choice problem, where individuals receive idiosyncratic preference shocks and locations differ in a scalar amenity. Formally, we assume that individual h 's value of living in location r at time t is given by

$$Q_{rt}^h = \mathcal{B}_{rt} \bar{\omega}^h(w_{rt}, \mathbf{P}_{rt} | \mathbf{P}_{rt}, \bar{q}_{rt}) u_{rt}^h. \quad (32)$$

Here, \mathcal{B}_r denotes the value of regional amenities, $\bar{\omega}^h(w_{rt}, \mathbf{P}_{rt} | \mathbf{P}_{rt}, \bar{q}_{rt})$ describes the average utility of being in region r in monetary terms and u_{rt}^h is a preference shock idiosyncratic to individual h and location r , which we assume

³⁰ 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.

³¹ 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.

to be Fréchet distributed with parameter θ .³² Given these assumptions, the share of people living in region r at time t is given by

$$\frac{L_{rt}}{L_t} = \frac{(\mathcal{B}_{rt}\bar{w}^h(w_{rt}, \mathbf{P}_{rt}|\mathbf{P}_{rt}, \bar{q}_{rt}))^\theta}{\sum_j (\mathcal{B}_{jt}\bar{w}^h(w_{jt}, \mathbf{P}_{jt}|\mathbf{P}_{jt}, \bar{q}_{jt}))^\theta} \quad (33)$$

Except for the presence of non-homothetic preferences, this setup is standard in most models of economic geography (see Redding and Rossi-Hansberg (2017)).

In Section A-6 in the Appendix we discuss the solution of this model in more detail. We first show that all our estimates of both structural parameters and sectoral productivities are the same as in the model with immobile labor. Intuitively, given the observed population, we can estimate the model exactly as in our baseline analysis. We can then residually estimate the spatial distribution of amenities \mathcal{B}_{rt} to rationalize the observed population distribution as an equilibrium outcome.

To perform counterfactuals, we need an estimate of the spatial labor supply elasticity θ , which in our context captures a long-run migration elasticity. In the absence of exogenous variation in local wages, this elasticity is hard to directly estimate. Therefore, we consider two scenarios: in our baseline scenario we pick θ , such that in our CS counterfactual, the amount of spatial reallocation is as high as what occurred in India between 1987 and 2011 holding local amenities fixed. For robustness, we also consider a higher-elasticity scenario. Given that the empirical literature finds relatively low levels of regional mobility in India, we regard this to be a generous upper bound to how much mobility we can expect in response to counterfactual changes in productivity.

The results—reported in the last rows of Table 7—are both qualitatively and quantitatively similar to those in the baseline model (which the extended model encompasses as a particular case in which $\theta = 0$). We conclude that our results are robust to allowing quantitatively reasonable migration flows in response to counterfactual experiments.

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 local consumer services such as retail and restaurants accounts for the bulk of the decline in agricultural employment while industrial employment is often stagnant. This pattern of development raises two fundamental questions. First, can services be a source of productivity growth even at low levels of economic development? Second, if services are luxuries and must be enjoyed locally, how different are the welfare effects of service-led growth across different sectors of the population?

In this paper, we developed a methodology to answer these questions. 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.

³² Recall that \bar{w}^h describes the equivalent variation to achieve a given utility level.

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, productivity growth in sectors such as retail, hospitality, or transportation accounts for one third of the improvement in living standards between 1987 and 2011. Second, the welfare impact of service-led growth was strikingly unequal: it disproportionately benefited wealthy individuals in urban areas while leaving poor people almost unaffected. 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 find 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.

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, it would be interesting to know the extent to which other developing countries are growing "like India." 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 for the sustainability of growth. However, the distributional consequences of this type of growth could raise new concerns about inequality that remain invisible in aggregate statistics.

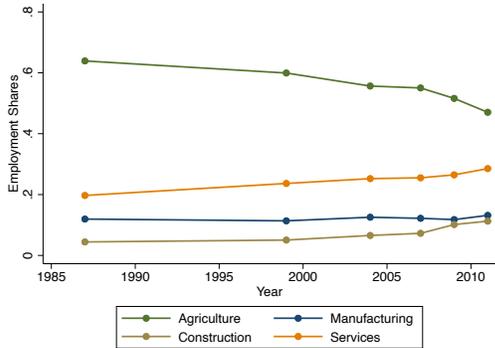
References

- 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.
- Burstein, A., M. Eichenbaum, and S. Rebelo: 2005, 'Large Devaluations and the Real Exchange Rate'. *Journal of Political Economy* **113**(4), 742–784.
- 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.
- 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.
- 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.
- 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.
- Singhari, S. and S. Madheswaran: 2016, 'The Changing Rates of Return to Education in India: Evidence from NSS Data'.

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



PANEL b: EMPLOYMENT GROWTH IN THE SERVICE SECTOR

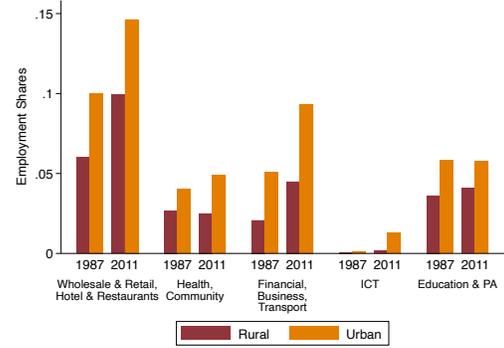


Figure 1: STRUCTURAL CHANGE IN INDIA: 1987–2011. The left panel shows the evolution of sectoral employment shares over time and is based on a standard ISIC classification. The right panel shows employment shares for different subsectors of the service sector independently for rural and urban localities. We rank districts by their urbanization rate and group them into rural and urban bins that account for roughly 50% of total employment.

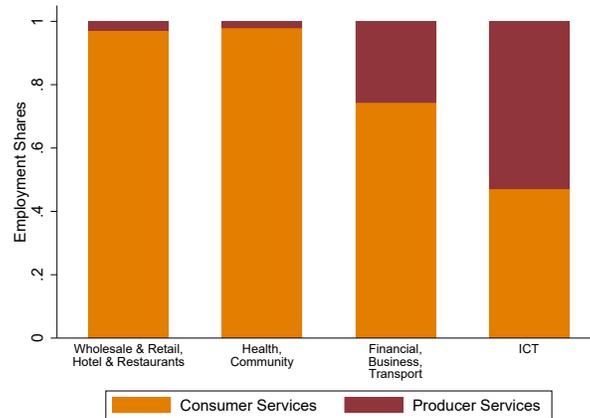


Figure 2: PRODUCER VS CONSUMER SERVICES IN DIFFERENT INDUSTRIES. The figure shows the share of producer and CS in 2011 in different industries within the service sector.

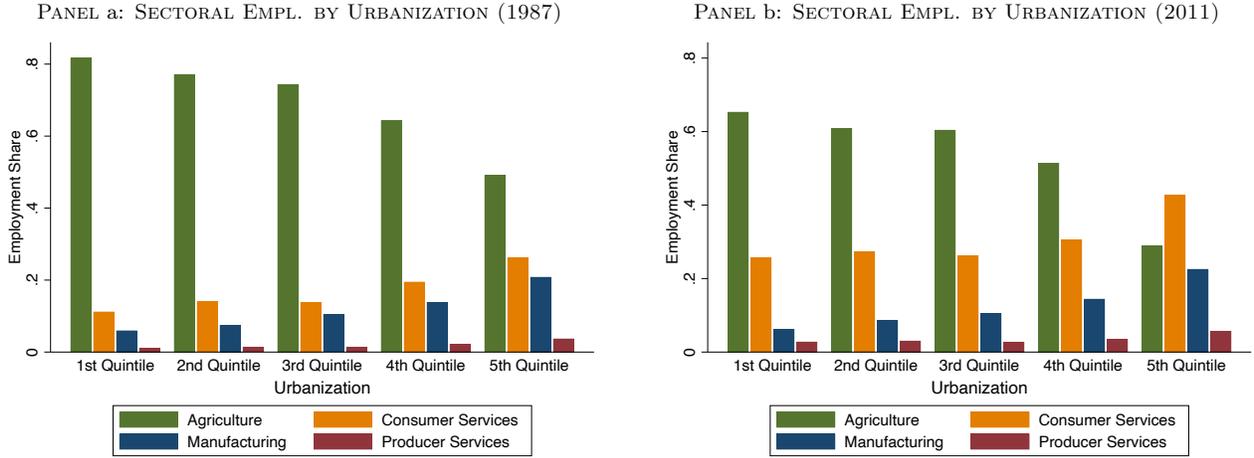


Figure 3: SECTORAL EMPLOYMENT OVER TIME AND SPACE. The figure plots the sectoral employment shares by urbanization quintile in 1987 and 2011.

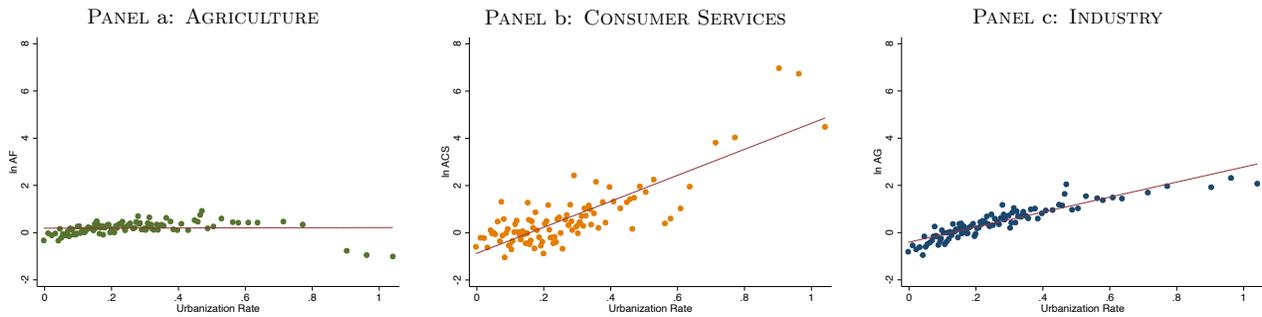
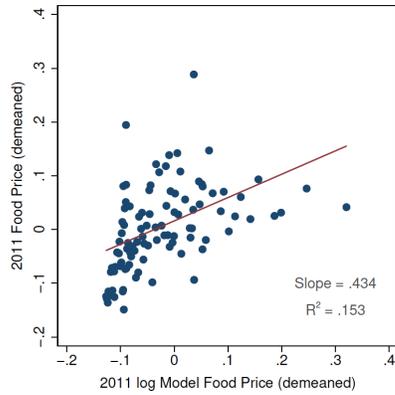


Figure 4: ESTIMATED SECTORAL PRODUCTIVITIES. The figure shows a bin scatter plot of the estimated sectoral labor productivities in agriculture, CS, and industry across urbanization-rate bins. Each plot is constructed by pooling the estimates for 1987 and 2011 after absorbing year effects.

PANEL a: FOOD PRICES: DATA VERSUS MODEL.



PANEL b: URBANIZATION AND FIXED EFFECTS OF CS SPENDING.

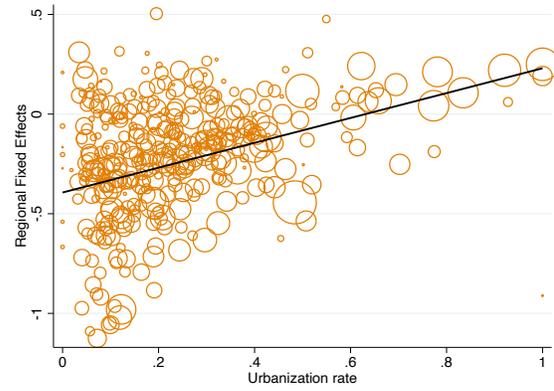


Figure 5: NON-TARGETED MOMENTS: FOOD PRICES AND CONSUMER SERVICE SPENDING. In the left panel we show a binscatter plot of regional log food prices in the model ($\ln p_{r,F}$) and the data (δ_r from (28)). In the right panel we display the correlation of the region fixed effect of a regression of log CS expenditure on individual income against the urbanization rate.

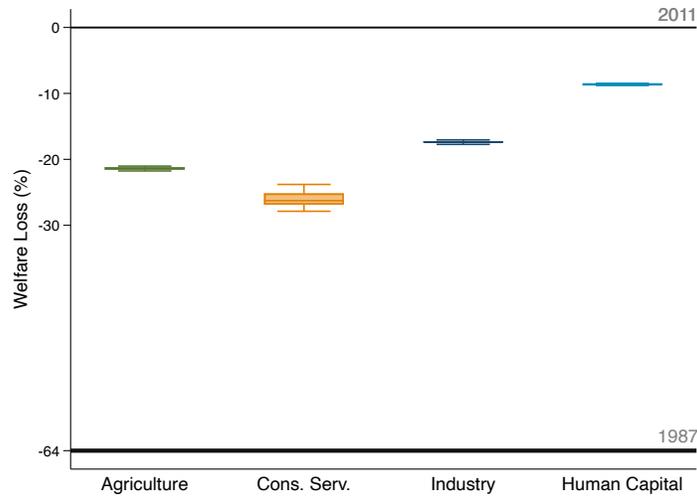


Figure 6: AGGREGATE WELFARE EFFECTS. The figure displays the percentage welfare losses ΔW associated with counterfactually setting productivity in agriculture, CS, and industry, as well as the level of human capital, to their respective levels in 1987 in all Indian districts. We compute the distribution 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.

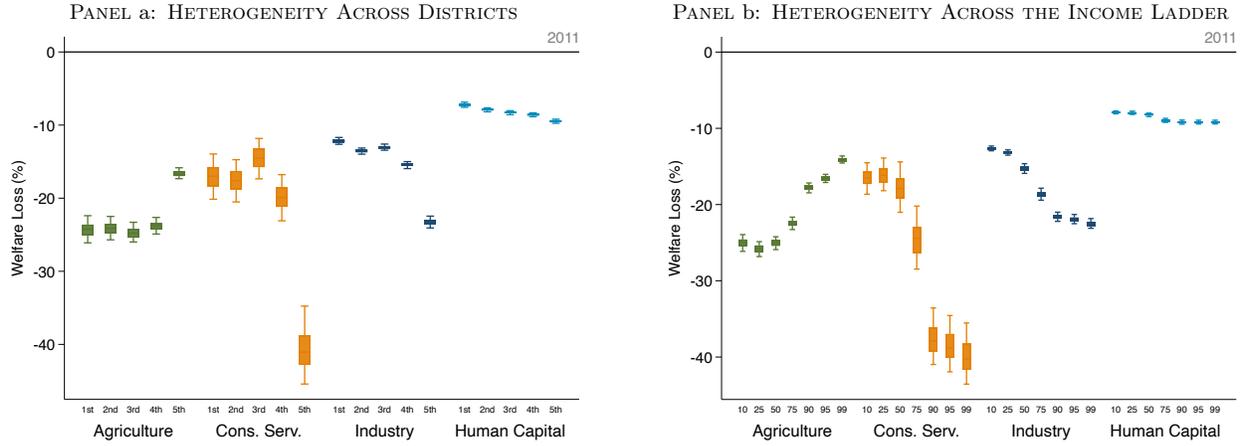


Figure 7: THE HETEROGENEOUS WELFARE IMPACT OF SERVICE-LED GROWTH. The figure displays the average percentage welfare losses associated with counterfactually setting productivity in agriculture, CS, and industry, as well as human capital, at the respective 1987 level, broken down by urbanization quintile 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

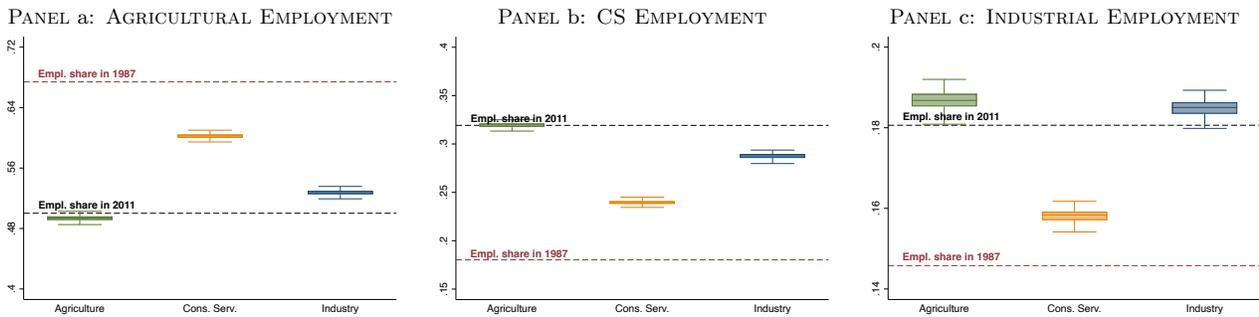
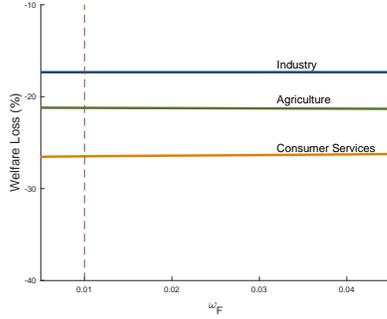
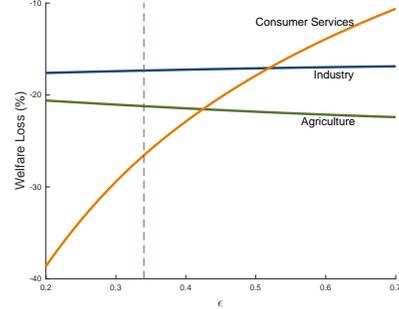


Figure 8: SECTORAL PRODUCTIVITY GROWTH AND STRUCTURAL CHANGE. Each panel shows the counterfactual employment share in the respective sector when we set productivity in agriculture, CS, and industry to their 1987 level. The dashed horizontal lines show employment in 1987 and 2011, for reference. We compute 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.

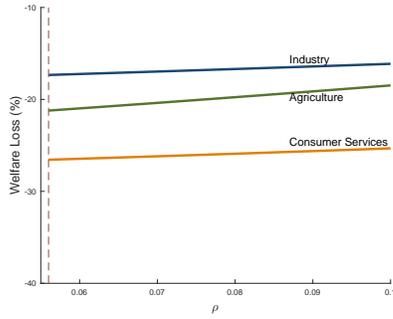
PANEL a: LONG-RUN SHARE OF AGRICULTURE ω_F



PANEL b: INCOME ELASTICITY ϵ



PANEL c: RETURN TO EDUCATION ρ



PANEL d: SKILL DISTRIBUTION ζ

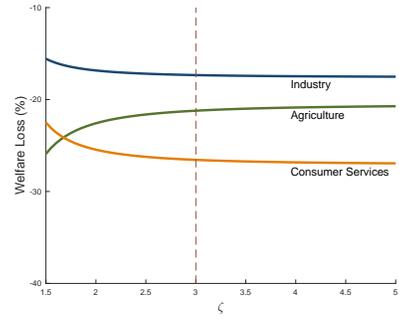


Figure 9: ROBUSTNESS ANALYSIS. Panels (a), (b), (c), and (d) show the aggregate welfare effects as a function of the preference parameters ω_F , ϵ , the Mincerian rates of return to education ρ , and the tail parameter of the skill distribution ζ . The vertical dashed line corresponds to the parameter value in our benchmark analysis.

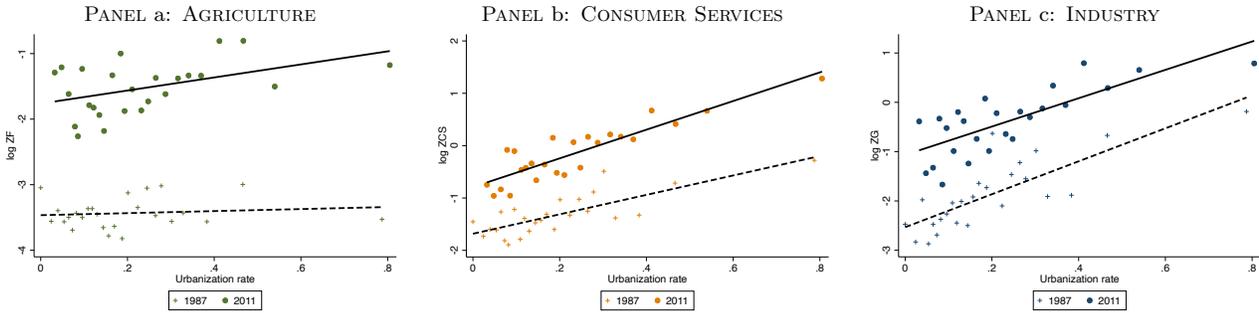


Figure 10: SKILL BIAS OF TECHNOLOGY. The figure shows a binned scatter plot of Z_{rF} , Z_{rCS} , and Z_{rG} as a function of the urbanization rate in 1987 (dashed line and “+” markers) and 2011 (solid line dots).