Growing Like India:

The Unequal Effects of Service-Led Growth

Tianyu Fan, Michael Peters, and Fabrizio Zilibotti

Yale University*

November 2022

Abstract

In large part of today's developing countries, structural transformation takes the form of a rapid rise of services with limited industrialization. This development pattern raises some concerns because, allegedly, manufacturing is a more powerful source of sustained productivity growth than services. In this paper, we estimate sectoral productivity growth and its welfare consequences across Indian regions and households with heterogeneous income. We construct a spatial equilibrium model in which the expansion of the service sector is both a consequence, due to income effects, and a cause, due to productivity growth, of the development process. Our approach circumvents the notorious difficulties in measuring quality improvements in services. We estimate the model using Indian household data. We find that productivity growth in non-tradable consumer services such as retail, restaurants, or residential real estate, is an important driver of structural transformation and rising living standards. However, the welfare gains are heavily skewed toward high-income urban dwellers.

*This project stems from early discussions with Philippe Aghion and Robin Burgess. A previous version was circulated under the title: *Service-Led or Service-Biased Growth? Equilibrium Development Accounting across Indian Districts*. We are grateful to three referees and to Sebastian Sotelo and Klaus Desmet for their thoughtful discussions at the NBER SI and the ASSA. We thank Daron Acemoglu, Treb Allen, David Atkin, Johannes Boehm, Fabian Eckert, Reto Foellmi, Doug Gollin, Cormac O'Dea, Peter Klenow, Samuel Kortum, Ramon Marimon, Giuseppe Moscarini, Juan Pablo Nicolini, Rachel Ngai, Ezra Oberfield, Richard Rogerson, María Sáez Martí, Todd Schoellman, Kjetil Storesletten, Aleh Tsyvinski, Danyang Xie, and seminar participants at the ASSA Meeting 2021, the Cowles Macro Conference, the Minneapolis FED, CREi, Dartmouth, HEC Lausanne, IIES, Johns Hopkins SAIS, the NBER Summer Institute, New York University, MIT, Penn State, Peking University, RIDGE, STEG, SED, the Universitá di Bologna, the Universitá della Svizzera Italiana, the University of Sankt Gallen, the Workshop on Structural Transformation and Macroeconomic Dynamics, the World Bank, and Yale University. We thank Sarah Moon, Shengqi Ni, Pariroo Rattan, Haonan Ye, and Huihuang Zhu for excellent research assistance. Michael Peters and Fabrizio Zilibotti thank the "Minnesota Opportunity and Inclusive Growth Institute" for its generous hospitality.

1 Introduction

Urbanization and structural change are transforming the lives of hundreds of million of people across the globe. Consider India, the second most populous country: 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 (p.p.) and the employment share of agriculture is down to 42%.

In this paper we argue that productivity growth in the service sector plays a key role for this transformation and for the accompanying rise in living standards. We focus on non-traded services that serve final consumers, such as retail, restaurants, local transportation, or residential real estate. Such services, which we refer to as *consumer services* (CS), have seen a dramatic rise in employment and account for one third of aggregate employment in India, a share that increases to almost two thirds in urban districts such as Delhi or Mumbai.

To quantify the welfare effects of productivity growth in the provision of these services, we abandon the straightjacket of representative agent models and construct a multisectoral spatial equilibrium model in which people with heterogeneous income reside in different locations and consume different baskets of goods and services. We estimate the model using both micro and macro data. The estimation retrieves the spatial, sectoral, and time variation of productivity consistent with the equilibrium conditions of the theory. Our approach is in the wave of the development accounting literature: we recover the productivity distribution from the data and a set of restrictions imposed by the theory, but do not attempt to provide a theory of its determinants. An advantage of our methodology is that it does not rely on existing price indices of services. This is particularly important for non-tradable CS, where measurement issues about quality adjustments loom large.[1] Another advantage is that we use data on consumption rather than earnings, which would miss income from informal activities.

We use the estimated model to infer the heterogeneous welfare effects of structural change across both localities and the income distribution, building a bridge between

[1] Failure to account for quality changes can bias upwards price indexes for services over time. Suppose, for instance, that improvements in logistics reduce the cost of home-delivery making this service accessible to more consumers. Yet, shopping online is more expensive than traditional retail. Then, the average price paid by consumers for the service would grow. However, this reflects a convenience value for which consumers are willing to pay.

economic growth and economic development. By way of counterfactual analysis, we find that, while economic growth has improved living conditions in India across the board, the sources of welfare gains are diverse. In rural areas, poverty has fallen, mainly owing to productivity growth in agriculture. By contrast, the urban middle class has benefited not only from the availability of better and cheaper goods but also from the growing supply of local services that has changed the face of urban life. To the best of our knowledge, ours is the first paper that quantifies the unequal welfare effects of productivity growth in the service sector.

Our theory has two building blocks: (i) non-homothetic preferences, and (ii) the assumption that—while agricultural and industrial goods are traded across regions—CS must be provided locally. If, as we find, service-intensive products are "luxuries," these assumptions imply that the affluent urban residents are the main beneficiaries of service-led growth. Non-homothetic preferences also play a crucial role in our estimation procedure. The estimation of CS productivity is subject to an identification problem: An increase in local CS employment could stem from local demand (i.e., income growth originating from other sectors coupled with nonhomothetic preferences), but also from supply forces, namely, changing productivity of the local CS sector, which we refer to as *service-led* growth. The identification of their relative importance hinges on the income elasticity of the demand for different final goods.

To discipline this elasticity, we estimate households' Engel curves using detailed micro data on consumption expenditures. We parameterize preferences by an indirect utility function in the Price-Independent Generalized-Linear (PIGL) class. This preference class was introduced by Muellbauer [1976] and was recently revamped in the growth literature by Boppart [2014]. PIGL has two important properties. First, it features aggregation: 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, PIGL preferences enables one to seamlessly go back and forth between preferences defined over final expenditure and over sectoral value-added. This step is potentially treacherous, because, as shown by Herrendorf et al. [2013], the mapping between the parameters of the value-added demand system and the ones derived from preferences over final products depends, in general, on the entire input-output matrix. We formally establish that under our preferences the key parameter governing the income elasticity is common to both demand systems at the individual level. This allows us to use micro-data on household expenditure to estimate the elasticity of interest.

We apply our methodology to India, a fast-growing economy, with an average annual 4.2% growth rate during 1987–2011. Our estimation exploits individual geolocalized consumption and employment data, and we estimate sectoral productivity growth for 360 Indian districts. Our measurement of CS employment is consistent with the assumption that such activities are non-tradable in nature and that they contribute to households' local access to consumption goods (e.g., restaurants or retail shops) or directly enter their consumption basket (e.g., health or entertainment services). By contrast, producer services (PS) such as law services, ICT, or consulting activities, to a large extent serve as inputs to the industrial sector and as such, their value-added can be shipped across locations.[2] Leveraging this distinction, our estimation exploits a novel firm-level data set on firms in the service sector that reports whether firms sell to consumers or to other firms.

Our analysis delivers four main findings. First, at the spatial level, there are large sectoral productivity differences and the productivity gap between urban and rural districts is largest in CS. Thus, cities in India have a higher service employment share not only because their residents are richer but also because CS are provided more efficiently. Second, service-led growth played an important role for economic development. At the aggregate level, rising productivity of CS accounts for almost one-third of the increase in welfare. Third, and most importantly, service-led growth has yielded strikingly unequal welfare effects. Productivity growth in CS was the main source of welfare gains for richer households living in urbanized districts. By contrast, for poorer households from rural districts, living standards improved mostly owing to productivity growth in agriculture. Finally, productivity growth in CS was a key driver of structural change. Agricultural employment would have declined far less in the absence of rising productivity in CS.

We carry out the main analysis under a set of stark assumptions aimed to retain tractability and to focus on the main mechanism of the theory. In the second part of the paper, we relax three important assumptions. First, we consider an extension where India is an open economy with international trade flows calibrated to the data. In particular, we zoom in on the growing role of export of ICT services. Second, we relax

[2] The stark assumption that CS are consumed locally is in line with the findings of 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.

the assumption that skills are perfect substitutes and assume, instead, that labor inputs provided by people with different educational attainment are imperfect substitutes. Moreover, we allow skill intensities to vary across sectors (e.g., agriculture is less skill intensive), districts, and time (skill-biased technical change.) In this extension, changes in educational attainment are an engine of structural change and local comparative advantage. Finally, we allow for labor mobility across districts. While the quantitative results change to some extent in each extension, the broad picture is consistent and robust: service-led growth is a prominent feature of the Indian economy with major implications for the growth and distribution of welfare effects.

Related Literature: Our paper contributes to the literature on the structural transformation including, among others, Kongsamut et al. [2001], Ngai and Pissarides [2007], Herrendorf et al. [2013], Gollin et al. [2014], and García-Santana et al. [2021].

A recent literature focuses on the service sector. Buera and Kaboski [2012] emphasize the importance of skill-intensive services in the US since 1950. Hsieh and Rossi-Hansberg [2023] argue that in more recent years ICT has been a major source of productivity growth. Their view is echoed by Eckert et al. [2020]. Chatterjee and Giannone [2021], associate rising productivity in services with regional divergence. A few studies focus on services in the developing world, among them, Duarte and Restuccia [2010], who document large cross-country productivity differences, Gollin et al. [2015], who emphasize the relationship between urbanization and consumption of non-tradable services, and, most recently, Nayyar et al. [2021], who use cross-country data to highlight the promise of service-led growth in today's developing world. Desmet et al. [2015] and Dehejia and Panagariya [2016] study aspects of the development of the service sector in India, documenting its important role for cities. Atkin et al. [2018] study the welfare gains associated with the entry of global retail chains in Mexico, stemming from pro-competitive effects on the prices charged by domestic stores.

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 approach of Gancia et al. [2013], who exploit the restrictions imposed by an equilibrium model to identify sectoral productivity. We perform our accounting exercise in the context of a model with inter-regional trade linkages, commonly used in the economic geography literature; see, e.g., Redding and Rossi-Hansberg [2017] or Allen and Arkolakis [2014].

The link between spatial inequality and the process of structural change is a common ground with Budí-Ors and Pijoan-Mas [2022].

Non-homothetic preferences play a central role in our analysis. We are especially close to Boppart [2014] and Alder et al. [2022], who propose PIGL preferences to study the process of structural transformation. Eckert and Peters [2020] incorporate these preferences in a spatial model of structural change. Instead, Comin et al. [2021] and Matsuyama [2019] build on the class of generalized CES preferences postulated by Sato [1975]. In our paper, we use PIGL preferences because of their tractable aggregation properties. Our results on the unequal gains from service growth are reminiscent of Fajgelbaum and Khandelwal [2016], who measure the unequal gains from trade in a setting with non-homothetic preferences.

Road Map: The structure of the paper is as follows. Section 2 summarizes the key stylized facts of the growing role of services in India and the developing world. Section 3 lays out our theoretical framework. Sections 4 and 5 describe the data and our empirical methodology. Section 6 contains the main results on the unequal welfare effects of service-led growth. Section 7 contains the extensions of our analysis and a variety of robustness checks. Section 8 concludes. The Appendix contains details of the theoretical and empirical analysis. A Web Appendix available from the authors' webpages contains additional results.

2 Structural Change in the Developing World

Between 1987 and 2011, India experienced fast economic development: income per capita grew by a factor of three and the employment structure changed markedly. The upper left panel of Figure 1 highlights the pattern of growth without industrialization: the structural transformation is mostly an outflow out of agriculture and an inflow into services and construction whose employment shares increased, respectively, by nine and seven p.p. By contrast, manufacturing employment is stagnant. Today, the service sector accounts for about one-third of aggregate employment.

Large part of the expansion in service employment originated in activities that facilitate consumers' access to final consumption. The upper right panel of Figure 1 decomposes the service sector into four subsectors.[3] The first group serves mostly

[3] Using the official NIC classification, the four subsectors contain the following industries: (i) wholesale and retail trade; repair of motor vehicles and motorcycles; accommodation and food services; health and social work; arts and entertainment; other service activities; (ii) finance and insurance; ICT;

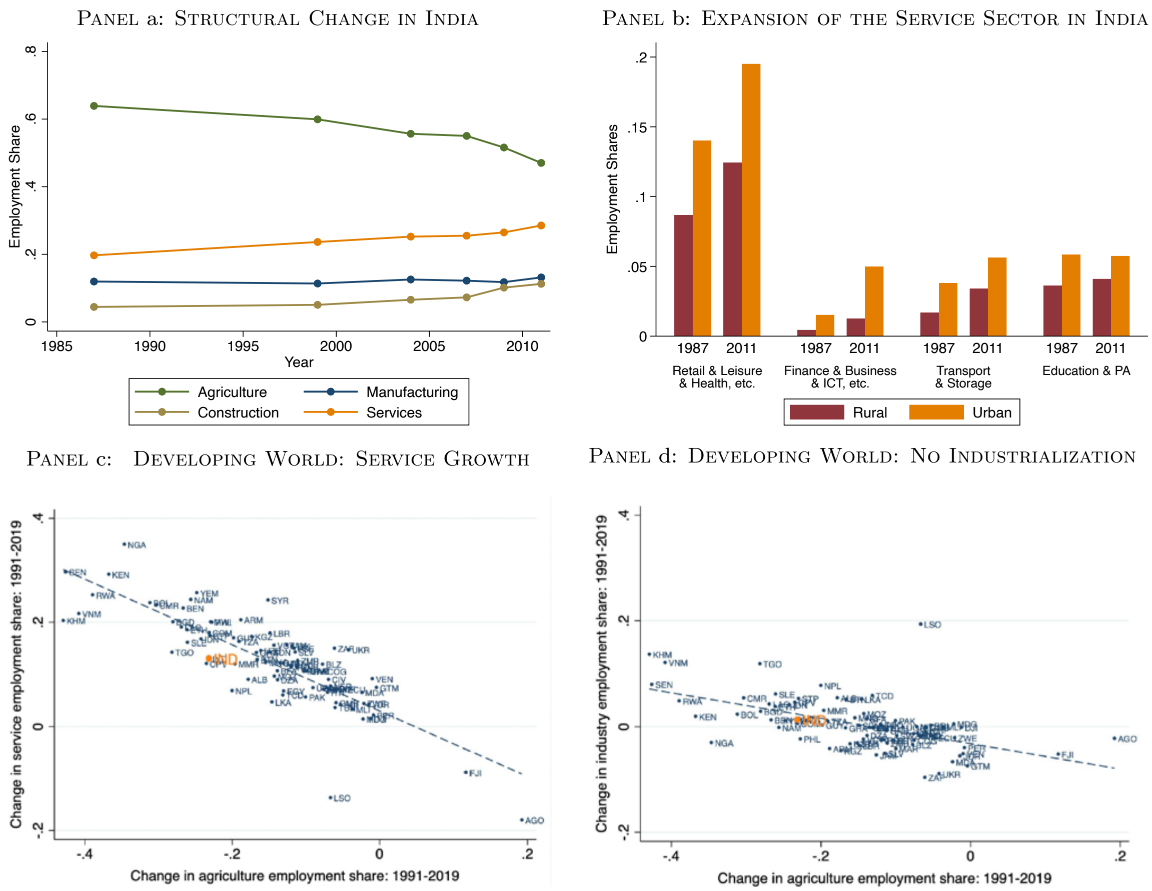

Figure 1: Structural Change towards Services in India and in the Developing World. The upper left panel shows the evolution of sectoral employment shares in India. The upper right panel shows employment shares for different service industries (see footnote 3 for details), separately for rural and urban districts. We split India into rural and urban districts, so that half of the population belongs to each type of district. The figure is based on micro data from the NSS (see Section 4). The lower left (lower right) panel shows the correlation between changes in agricultural employment shares and changes in service employment shares (manufacturing employment shares) between 1991 and 2019 for all non-OECD countries, that are poorer than China in 2019. We omit construction and utilities.

consumers. These service industries grew significantly since 1987 and employ almost 60% of all Indian service workers in 2011. The second group, which sells a significant part of their services to industrial firms, also grew substantially but only accounts for a tenth of service employment. For instance, ICT, a fast-growing industry, accounts for less than 1% of total employment in 2011. Transport services, which serve both consumers and industries, also expanded. Finally, the employment share of mostly government-run activities such as PA and education remained constant over time. Figure 1 also shows that all service activities are much more prevalent in urban areas.

This pattern of a decline in agriculture and an increase in services is by no means

real estate; professional, scientific, and technical activities; administrative and support services; publishing; (iii) transport and storage; and (iv) education and Public Administration (PA).

exceptional in today's developing world. In the lower panels of Figure 1 we display the cross-country relationship between the change in the employment share of agriculture and those of services (left panel) and manufacturing (right panel), respectively, between 1991 and 2019. To zoom in on the developing world, we include all non-OECD countries whose income per capita is below that of China in 2019. The left panel shows a strong negative relationship: a ten p.p. reduction in the agriculture share is matched on average by a 6.3 p.p. increase in the service share. The right panel shows that the relationship, albeit negative, is substantially weaker for the industrial sector: a ten p.p. reduction in the agriculture share is associated with a 2.4 p.p. increase in the manufacturing share. Crucially, the low speed of industrialization is *not* a mark of lacklustre development. If we restrict attention to the subsample of developing countries whose annual GDP grew by more than 4% per year, the employment share of agriculture declined by ca. 16 p.p, while the employment shares of industry and services grew, respectively, by 1 p.p. and 13 p.p. Figure 1 also shows that the typical developing country indeed grew *like India*: the observation for India, highlighted in orange, is not far from the regression line.[4] Nor is the predominance of CS relative to PS a prerogative of India. In Appendix Figure B-1, we show that the pattern of Panel b of Figure 1 is perfectly in line with the international evidence.

3 Theory

We consider a model with R regions and three broad sectors: agriculture (F for *food*), industry (G for *goods*), and CS. Consumers' preferences are defined over a continuum of final products, which combine the output of these three sectors. We make the important assumption that, while food and goods are tradable across regions subject to iceberg costs, CS must be locally provided.[5] Markets are frictionless and competitive.

We assume that labor is inelastically supplied in each region, that workers' human capital is perfectly substitutable across sectors, and that the economy is closed to international trade. In Section 7 we extend our model along each of these dimensions.

[4] Service-led growth is not a prerogative of the developing world. The main difference is that in richer nations the service sector mostly grows at the expense of manufacturing rather than agriculture. Even in a country like China, whose stellar growth has been led for decades by the industrial sector, the service sector has gained significant ground in the last ten years while the share of industrial employment is shrinking (Chen et al. [2022]).

[5] As we describe in detail below, we assume that the industrial sector employs both production workers and workers producing production services (PS). Because the value-added of, say, corporate lawyers and consultants is embodied in industrial goods, PS are ultimately tradable.

3.1 Technology

Each region produces a measure one continuum of non-traded differentiated final products using the two tradable goods—food and goods—and local CS workers as inputs.[6] For instance, a restaurant meal is a combination of food and kitchen tools, and of services provided by local cooks and waiters.

Formally, the production function for final good $n \in [0, 1]$ in region r at time t is

$$Y_{rnt} = \tilde{\lambda}_n x_{rFt}^{\lambda_{nF}} x_{rGt}^{\lambda_{nG}} \left(\mathcal{A}_{rnt} H_{rCSt}\right)^{\lambda_{nCS}}, \tag{1}$$

where x_{Ft}, x_{Gt}, denote the inputs of food and goods, H_{rCSt} is the number of efficiency units of labor delivering the CS allocated to the production of good n, and $\mathcal{A}_{rnt}$ reflects the productivity of providing CS for product n. We assume constant returns to scale: $\sum_s \lambda_{ns} = 1$.[7] The elasticities λ_{ns} determine the intensity of food, goods, and CS value-added in the production of product n. Intuitively, a home-cooked meal is a product with a large food content ($\lambda_{nF} \approx 1$) and a low CS content of (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 tradable food and industrial good are CES aggregates of regional varieties:

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

which are produced according to the linear productions functions

$$y_{rFt} = A_{rFt} H_{rFt} \quad \text{and} \quad y_{rGt} = A_{rGt} H_{rGt}, \tag{3}$$

where sectoral productivities A_{rst} can differ across regions. We refer to $\mathcal{A}_{rnt}$ in (1) as CS productivity even though it applies to all inputs. We show below that assuming CS must be supplied locally allows us to separately identify $\mathcal{A}_{rnt}$ from A_{rFt} and A_{rGt}.

Nontradable CS vs tradable PS: In our theory, tradability is the key difference between CS and PS. While CS value-added can only be consumed locally, the PS

---

[6] Burstein et al. [2005] also emphasize the non-tradable nature of CS and their large value-added share in final expenditure.

[7] The representation of technology in (1) is akin to the Cobb Douglas input-output structure commonly assumed in the production network literature—see, e.g., Acemoglu et al. [2012]. The scalar $\tilde{\lambda}_n \equiv \lambda_{nF}^{-\lambda_{nF}} \lambda_{nG}^{-\lambda_{nG}} \lambda_{nCS}^{-\lambda_{nCS}}$ is an inconsequential normalization to simplify expressions. We discuss below the importance of this functional form assumption.

value-added is embodied in goods and is ultimately tradable.

When mapping the model to the data, we include the value-added of PS in the industrial sector, namely, we let $H_{rGt} = H_{rMt} + H_{rPSt}$.[8] This specification does *not* restrict manufacturing and PS workers to being perfect substitutes. To see why, suppose industrial firms combine the inputs of manufacturing workers and PS to produce industrial goods using the technology $y_{rGt} = g_{rt}(H_{rMt}, H_{rPSt})$, where g_{rt} is a linearly homogeneous function. As long as firms maximize profits, the marginal products of H_{rMt} and H_{rPSt} are equalized and we can express aggregate output in the industrial sector in region r as $y_{rGt} = A_{rGt}H_{rGt}$, where high industrial productivity A_{rGt} can either stem from an advanced manufacturing production technology or an efficient provision of accounting and legal services to firms.[9] This allows cities like Delhi or Bangalore with a comparative advantage in tradable PS like finance or ICT to export the value-added of PS to the rest of India (and, in Section 7, even internationally).

3.2 Preferences and Demand System

Following Boppart [2014] we assume consumers' preferences over the continuum of final products are in the PIGL class. These preferences have two important properties. First, they admit aggregation, allowing us to take a spatial demand system to the data and perform welfare analysis. Second, they provide a simple mapping of preferences over final goods into preferences over value-added.

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

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

where e denotes total spending and $\mathbf{p}_r$ is the vector of prices in region r. The mnemonic FE is a reminder that the indirect utility function is defined over final expenditure and the prices of final products $n \in [0, 1]$. The functions $B(\mathbf{p})$ and $D(\mathbf{p})$ are restricted to be homogeneous of degree one and zero, respectively. We parametrize them as

[8] For simplicity, we restrict the value-added of PS to be embodied in industrial goods. According to the Indian Input-Output tables, the agricultural sector accounts for very little of intermediate input purchases from the service sector.

[9] Linear homogeneity allows us to write $y_{rGt} = g_{rt}(1 - s_{rPSt}, s_{rPSt})H_{rGt}$, where $s_{rPSt} = H_{rPSt}/H_{rGt}$. We can then write industrial TFP as $A_{rGt} \equiv \max_{s_{PS}} g_{rt}(1 - s_{PS}, s_{PS})$, that is, A_{rGt} is fully determined from the production function g_{rt}. For instance, suppose $g = \left[(A_{rMt}H_{rMt})^{(\varsigma-1)/\varsigma} + (A_{rPSt}H_{rPSt})^{(\varsigma-1)/\varsigma}\right]^{\varsigma/(\varsigma-1)}$. Then, $A_{rGt} = \left(A_{rMt}^{\varsigma-1} + A_{rPSt}^{\varsigma-1}\right)^{1/(\varsigma-1)}$.

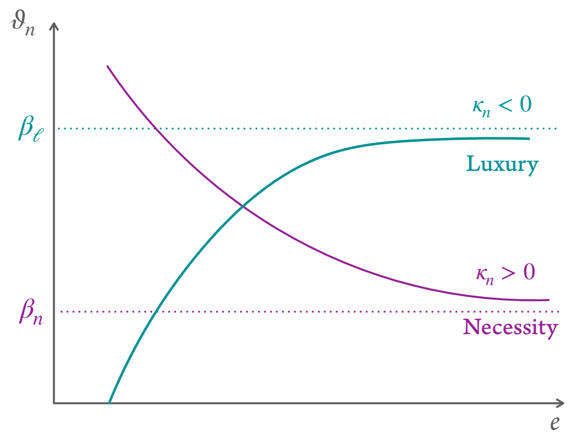

Figure 2: ENGEL CURVES. The figure shows the good-specific expenditure share as a function of income e (see (5)).

$B(\mathbf{p}_r) = \exp\left(\int_{n=0}^{1} \beta_n \ln p_{rn} dn\right)$ and $D(\mathbf{p}_r) = \left(\int_{n=0}^{1} \kappa_n \ln p_{rn} dn\right)$, where $\int_0^1 \beta_n dn = 1$ and $\int_0^1 \kappa_n dn = 0$[10]

By Roy's Identity, the expenditure share an individual with spending level e allocates to final good n is given by:

$$\vartheta_n^{FE}(e, \mathbf{p}_r) = \beta_n + \kappa_n \left(\frac{e}{\exp\left(\int_n \beta_n \ln p_{rn} dn\right)}\right)^{-\varepsilon}. \tag{5}$$

This expression highlights that the demand system is akin to a Cobb-Douglas specification (the term β_n) with a non-homothetic adjustment. In Figure 2, we depict the expenditure share as a function of expenditure. The expenditure share converges to β_n as income grows large. A good n is a luxury if $\kappa_n < 0$ (in which case β_n is approached from below) and a necessity if $\kappa_n > 0$ (in which case β_n is approached from above). Cobb-Douglas preferences are a special case when $\kappa_n = 0$. The slope of the Engel curves and the strength of income effects is governed by the parameter ε. This parameter—that we label the *Engel elasticity*—plays a key role in our analysis.

3.2.1 Final Expenditure and Value-Added

Equation (5) defines the expenditure shares over final products. For our purposes, it is essential to derive a demand system for the value-added produced by the three grand sectors F, G, and CS, because we estimate our model using data on sectoral employment. To derive this value-added demand system, note first that the prices of

[10] Our functional form for $D(\mathbf{p}_r)$ is more restrictive than the one in Boppart [2014]. In Section 7.3 we generalize the preference structure along the lines of his original contribution.

tradable goods are given by the usual CES price indices

$$P_{rst}^{1-\sigma} = \sum_{j=1}^{R} \tau_{rj}^{1-\sigma} A_{jst}^{\sigma-1} w_{jt}^{1-\sigma}, \quad \text{for } s \in \{F, G\}, \tag{6}$$

where $\tau_{rj} \geq 1$ is the iceberg cost of shipping variety j to region r. The price of final good n in region r is then given by $p_{rnt} = P_{rFt}^{\lambda_{nF}} P_{rGt}^{\lambda_{nG}} \left(\mathcal{A}_{rnt}^{-1} w_{rt}\right)^{\lambda_{nCS}}$, where w_{rt} denote the wage in region r. Plugging this expressions into the indirect utility function (4) yields a representation of consumers' preferences over sectoral value-added aggregates.

Proposition 1. *The value-added indirect utility function of consumers in region r is given by*

$$\mathcal{V}\left(e, \mathbf{P_{rt}}\right) = \frac{1}{\varepsilon} \left(\frac{e}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} P_{rCSt}^{\omega_{CS}}} \right)^{\varepsilon} - \sum_{s \in \{F,G,CS\}} \nu_s \ln P_{rst}, \tag{7}$$

where $\mathbf{P_{rt}} = (P_{rFt}, P_{rGt}, P_{rCSt})$, $P_{rCSt} \equiv A_{rCSt}^{-1} w_{rt}$, P_{rFt} *and* P_{rGt} *are given by* (6), *and*

$$\omega_s \equiv \int_n \lambda_{ns} \beta_n \, dn, \;\; \nu_s \equiv \int_n \lambda_{ns} \kappa_n \, dn, \;\; \text{and} \;\; \ln A_{rCSt} \equiv \int_n \frac{\beta_n \lambda_{nCS}}{\omega_{CS}} \ln \mathcal{A}_{rnt} \, dn. \tag{8}$$

The associated value-added expenditure shares are given by

$$\vartheta_{rst}\left(e, \mathbf{P_{rt}}\right) = \omega_s + \nu_s \left(\frac{e}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} P_{rCSt}^{\omega_{CS}}} \right)^{-\varepsilon}. \tag{9}$$

Proof. See Appendix A-1. □

Proposition 1 states three important properties of our theory. First, the indirect utility function defined over value-added also falls into the PIGL class and has the same functional form as the corresponding expressions over final products (4). In particular, the expenditure share over sectoral value added ϑ_{rst} in (9) features the same Engel elasticity as in (5). This result enables us to estimate ε from microdata for household expenditure shares on final products and then use it in the value-added demand system.

Second, the regional CS productivity index A_{rCS}, which is akin to the average CS productivity of all final products weighted by their CS content λ_{nCS} and their asymptotic spending share β_n, is a sufficient statistic for the local CS sector. Because preferences are nonhomothetic and CS are provided locally, productivity growth yields heterogeneous welfare effects. If goods with a high CS content are luxuries, productivity growth in CS is skewed toward the rich.[11] Moreover, CS productivity growth preeminently benefits local residents. Thus, if urban districts experience faster productivity growth, city dwellers are going to be the main beneficiaries of service-led growth. In contrast, the benefits from productivity growth in tradable sectors diffuse across locations through trade.

Third, the income elasticity of sectoral value-added depends on the *correlation* of the good-specific demand parameters κ_n with their factor intensities λ_{ns}. The expenditure share for sectoral value-added is rising in income if and only if $\nu_s < 0$, that is, if income-elastic *products* have a large *sectoral* input requirement. By contrast, if all goods were produced with equal factor proportions, or more generally if λ_{ns} were orthogonal to κ_n for all s, the demand for sectoral value-added would be homothetic even though the underlying demand for final products is nonhomothetic.

The closed-form expression of the mapping from the final-expenditure to the value-added demand system in Proposition 1 hinge on the assumption that the final good production function is Cobb-Douglas (cf. Equation (1)).[12] In the Web Appendix WA-1.1, we extend our analysis to a setting where (1) takes a CES form. In this case, we can still obtain an analytical characterization where the final-expenditure and value-added representations share the same Engel elasticty—i.e., we can derive the analogue of Equation (9). However, estimating the CES model would require additional data about the expenditure on individual final goods.

3.2.2 Heterogeneity and Aggregate Demand

Proposition 1 characterizes demand at the individual level. We now derive the aggregate demand system at the region level.

Suppose individuals differ in their human capital that determines the number of efficiency units of labor supplied to the market. Individual h's income is then given by $e^h_{rt} = q^h w_{rt}$, where q^h is the number of efficiency units of labor. Let $F_{rt}(q)$ denote

[11] In fact, the expenditure share $\vartheta_{CS}(e, \mathbf{P_{rt}})$ exactly measures the welfare exposure of a change in prices at the individual level. Formally, letting $e(\mathbf{P_{rt}}, V)$ denote the expenditure function associated with the utility level V given the price vector $\mathbf{P_{rt}}$, $\partial \ln e(\mathbf{P_{rt}}, V)/\partial \ln P_{rst} = \vartheta_{rst}(e, \mathbf{P_{rt}})$.

[12] More formally, using the expression for p_{rnt}, we can express $B(\mathbf{p}_r)$ as

$$\ln B(\mathbf{p}_r) = \int_n \beta_n \left(\ln P_{rFt}^{\lambda_{nF}} P_{rGt}^{\lambda_{nG}} P_{rCSt}^{\lambda_{nCS}} \right) = \sum_s \left(\int_n \beta_n \lambda_{ns} dn \right) \ln P_{rst} = \sum_s \omega_s \ln P_{rst},$$

that is, the price index B still has a constant price elasticity when we express it in terms of sectoral value added prices P_{rst}. In particular, the weight of *sectoral* prices, ω_s, reflects both the cost share λ and the expenditure share β, both of which are constant given the Cobb-Douglas assumptions.

the distribution function of q in region r at time t—which we empirically relate to the regional data on educational attainment.

Because our analysis abstracts from savings and capital accumulation, income equals expenditure. Equation (5) thus implies that the *aggregate* spending share on value-added produced in sector s by consumers residing in region r is given by

$$\overline{\vartheta}_{rst} \equiv \frac{L_{rt}\int \vartheta_{rst}\left(qw_{rt}\right)qw_{rt}dF_{rt}\left(q\right)}{L_{rt}\int qw_{rt}dF_{rt}\left(q\right)} = \omega_s + \overline{\nu}_{rst}\left(\frac{A_{rCSt}^{\omega_{CS}}\mathbb{E}_{rt}\left[q\right]w_{rt}^{1-\omega_{CS}}}{P_{rFt}^{\omega_F}P_{rGt}^{\omega_G}}\right)^{-\varepsilon}, \quad (10)$$

$$\text{where} \qquad \overline{\nu}_{rst} \equiv \frac{\mathbb{E}_{rt}\left[q^{1-\varepsilon}\right]}{\mathbb{E}_{rt}\left[q\right]^{1-\varepsilon}}\nu_s, \quad (11)$$

having defined with slight abuse of notation the expectation operator $\mathbb{E}_{rt}[x] \equiv \mathbb{E}[x; F_{rt}(x)]$. Comparing (10) with (5) clarifes the sense in which PIGL allows for a representative household: the *aggregate* demand system in (10) is isomorphic to that of a consumer in region r who earns the average income $\mathbb{E}_{rt}\left[q\right]w_{rt}$ and has the inequality-adjusted preference parameter $\overline{\nu}_{rst}$ in (11). Crucially, the Engel elasticity parameter of the aggregate demand system (ε) is the same as at the individual level.

The inequality adjustment term $\mathbb{E}_{rt}\left[q^{1-\varepsilon}\right]/\mathbb{E}_{rt}\left[q\right]^{1-\varepsilon}$, depends, in general, on the distribution of efficiency units F_{rt}. The analysis further simplifies if we assume q follows a Pareto distribution with c.d.f. $F_{rt}\left(q\right) = 1 - \left(\underline{q}_{rt}/q\right)^{\zeta}$. In this case, Equation (11) boils down to

$$\overline{\nu}_{rst} = \overline{\nu}_s = \frac{\zeta^{\varepsilon}\left(\zeta-1\right)^{1-\varepsilon}}{\zeta+\varepsilon-1}\nu_s. \quad (12)$$

Thus, if income is Pareto distributed with a common tail parameter, $\overline{\nu}_s$ is the same for all regions and the adjustment relative to the micro parameter ν_s accounts for the income distribution (ζ) and the income elasticity (ε). Given $\overline{\nu}_s$, the distribution F_{rt} only enters through the average income term $\mathbb{E}_{rt}\left[q\right]w_{rt} = \frac{\zeta}{\zeta-1}\underline{q}_{rt}w_{rt}$.

3.2.3 Welfare and Inequality

The aggregation properties of PIGL are especially handy for the welfare analysis. To this aim, define the utilitarian welfare function at the regional level as $\mathcal{U}_{rt}\left(w_{rt}, \mathbf{P_{rt}}\right) \equiv \int \mathcal{V}\left(qw_{rt}, \mathbf{P_{rt}}\right)dF_{rt}\left(q\right)$. Plugging in the indirect utility function in (7) yields

$$\mathcal{U}_{rt}\left(w_{rt}, \mathbf{P_{rt}}\right) = \frac{\zeta^{1-\varepsilon}\left(\zeta-1\right)^{\varepsilon}}{\zeta-\varepsilon} \times \left(\frac{1}{\varepsilon}\left(\frac{\mathbb{E}_{rt}\left[q\right]w_{rt}}{P_{rFt}^{\omega_F}P_{rGt}^{\omega_G}P_{rCSt}^{\omega_{CS}}}\right)^{\varepsilon} - \sum_{s\in\{F,G,CS\}}\nu_s^{\mathcal{U}}\ln P_{rst}\right), \quad (13)$$

where $\nu_s^{\mathcal{U}} \equiv \overline{\nu}_s \times ((\zeta-\varepsilon)(\zeta-(1-\varepsilon)))/(\zeta(\zeta-1))$. Hence, utilitarian welfare is again a function in the PIGL class and is akin to the indirect utility of a representative agent with average income $\mathbb{E}_{rt}[q]\, w_{rt}$ and the inequality-adjusted taste parameter $\nu_s^{\mathcal{U}}$.

3.3 Equilibrium

Given the aggregate demand system defined in (10), we can now characterize the competitive equilibrium.

Proposition 2. *The sectoral labor allocations $\{H_{rFt}, H_{rGt}, H_{rCSt}\}_r$ and local wages $\{w_{rt}\}$ are determined by the following equilibrium conditions:*

1. *Market clearing for local CS:*
$$w_{rt}H_{rCSt} \;=\; \left(\omega_{CS} + \overline{\nu}_{CS}\left(\frac{A_{rCSt}^{\omega_{CS}}\mathbb{E}_{rt}[q]\, w_{rt}^{1-\omega_{CS}}}{P_{rFt}^{\omega_F}P_{rGt}^{\omega_G}}\right)^{-\varepsilon}\right) w_{rt}H_{rt}, \tag{14}$$
where P_{rFt} and P_{rGt} are given by (6).

2. *Market clearing for tradable goods:*
$$w_{rt}H_{rst} \;=\; \sum_{j=1}^{R}\pi_{rsjt}\left(\omega_s + \overline{\nu}_s\left(\frac{A_{jCSt}^{\omega_{CS}}\mathbb{E}_{jt}[q]\, w_{jt}^{1-\omega_{CS}}}{P_{jFt}^{\omega_F}P_{jGt}^{\omega_G}}\right)^{-\varepsilon}\right) w_{jt}H_{jt}, \tag{15}$$
where $s \in \{F, G\}$ and $\pi_{rsjt} = \tau_{rj}^{1-\sigma}A_{rst}^{\sigma-1}w_{rt}^{1-\sigma}/P_{jst}^{1-\sigma}$.

3. *Labor market clearing:* $\quad H_{rFt} + H_{rGt} + H_{rCSt} = H_{rt}.$

Proposition 2 characterizes the sectoral employment allocations and income distribution across space. The contrast between Equations (14) and (15) reflects the tradable nature of food and goods versus the non-tradable nature of CS. The demand for CS value-added hinges on both local income and local CS productivity. For instance, the retail sector could be large in urban districts either because local consumers are more educated and richer, or because more-efficient department store chains open branches in large cities. Instead, the demand for tradable goods originates from all localities.

The proposition also highlights that sectoral value-added and employment are fully determined by the parameters $\overline{\nu}_s$ and ω_s and the aggregate CS productivity index A_{rCSt}. They do not separately depend on the preference parameters defined over final goods $[\beta_n, \kappa_n]_{n=0}^1$, nor on the product-specific productivity $[\mathcal{A}_{rnt}]_{n=0}^1$. Similarly, the size of the industrial sector H_{rGt} only depends on A_{rGt}, and we do not need to impose more structure on how PS and manufacturing workers interact in production.

4 Empirical Analysis: Data and Measurement

Our analysis relies on five datasets: (i) the NSS Employment-Unemployment Schedule for the years 1987 and 2011 (the "NSS data"); (ii) the NSS Consumer-Expenditure Schedule for the same years; (iii) the Economic Census for the years 1990, and 2013 (the "EC"); (iv) a Special Survey of the Indian Service Sector for the year 2006 (the "Service Survey"); and (v) the Economic Transformation Database (ETD) provided by the Groningen Growth and Development Centre (GGDC)—see de Vries et al. [2021]. A more detailed description of these datasets is deferred to Appendix B-2.

The NSS is a household survey with detailed information on households' consumption, employment characteristics, and location of residence. We use this information to construct measures of average income and sectoral employment shares at the district-year level. We prefer to proxy income by consumption expenditure rather than relying on information on wages as the latter would miss income from informal employment.[13] Similarly, we explicitly include self-employed individuals, employees of household enterprises, and casual laborers.

Consistent with our theory, we measure employment shares in four sectors: agriculture, manufacturing, PS, and CS. For agriculture and manufacturing, we follow the NIC classification. For services, we exclude from our analysis subsectors in which the government plays a dominant role: public administration and defense, compulsory social security, education, and extraterritorial organizations and bodies. Finally, we merge construction and utilities with the service sector. Although the construction sector is often included in the industrial sector, the key distinction in our theory is tradability. Because construction and utilities are provided locally, we find it natural to merge them with services. In Section 7, we show that our main results do not hinge on this classification of the construction sector. Below we discuss in detail how we split service employment into CS and PS.

The NSS Consumer-Expenditure Schedule contains information on households' expenditure on different categories of final goods that we use to estimate the Engel elasticity ε. The EC covers all establishments engaged in the production or distribu-

[13] In Web Appendix WA-4.3, we document that average expenditure is strongly correlated with average wages and average income per capita at the district level.

tion of goods and services in India. It covers all sectors except crop production and plantation and collects information on each firm's location, industry, and employment. It contains approximately 24 million and 60 million establishments in 1990 and 2013, respectively. The Service Survey was conducted in 2006 and is representative of India's service sector. It covers almost 200,000 private enterprises subdivided into seven service industries.[14] Finally, we rely on ETD for measuring the average relative price of agricultural goods (while we do not use any published price index for services).

Geography: To compare spatial units over time, we create a time-invariant definition of Indian districts. Appendix B-3 describes in detail how we construct this crosswalk. 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 exclude two small districts that existed 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.

Consumer versus Producer Services: To split service employment into CS and PS, we combine information from the EC and the Service Survey. We aim to assign firms to the CS sector if they sell to consumers and to the PS sector if they sell to other firms. Ideally, we would use firm-level input-output matrices. To the best of our knowledge, this information is not available in India for the time period of our study. We therefore leverage micro data on the firms' downstream trading partners contained in the Service Survey, which reports whether a firm is selling mostly to consumers or to other firms. A limitation is that the Service Survey contains too few firms to precisely estimate the employment shares of firms selling to producers and to consumers in 360 districts and 6 subsectors. Therefore, we rely on the fact that the propensity to sell to other firms is highly correlated with firm size. Table 1 shows that larger firms are more likely to sell to firms. For example, only 6% of firms with three employees sell to other firms, while the share increases to 43% for firms with more than 50 employees.

We use the pattern in Table 1 in the following way. First, we estimate the CS

[14] These industries are (i) hotels and restaurants, (ii) transport, storage and communication, (iii) financial intermediation, (iv) real estate, renting and business activities, (v) education, (vi) health and social work, (vii) other personal service activities. In In Appendix B-2.3, we compare the Service Survey with the EC and document that it is indeed representative of the distribution of firm size in India.

	Firm size: Number of employees								
	1	2	3	4	5	6-10	11–20	21–50	51+
Share of PS firms	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.

employment share by firm size for different service industries.[15] Then, we use the *district*-specific size distribution from the EC to infer the aggregate CS employment share in district r. More formally, the CS employment share (relative to the total service sector) in subsector k in region r is given by $s_{rk}^{CS} = \sum_b \omega_{kb}^{CS} \ell_{kbr}$, where ω_{kb}^{CS} is the share of employment in firms selling to consumers 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. The spatial variation in CS employment thus stems from differences in (i) total service employment, (ii) the relative size of different service industries, and (iii) the distribution of firm size. In Appendix B-4.2, we describe this procedure in more detail.

In Table 2 we report the resulting allocation of employment to CS. At the aggregate level, our procedure allocates 89% of service employment to CS and 11% to PS. This allocation differs across subsectors. For instance, within the retail and restaurant sector, 97% of workers are employed by establishments catering to consumers. Instead, in the ICT sector, less than half of employment caters to consumers.[16]

In a similar vein, the construction sector serves both consumers (e.g., residential housing) and firms (e.g., business construction). To break these activities into PS and CS, we exploit information from the "Informal Non-Agricultural Enterprises Survey 1999–2000" (INAES) dataset, which covers the construction sector and also reports whether a firm sells to consumers or other firms. Because of sample size limitations, we can only split the destination of construction activities at the national, not the district, level. See Appendix B-4.3 for details.

In Section 7, we show that our results are robust to alternative measurement strategies. In particular, (i) we allocate ICT and business services entirely to PS, and (ii) we split PS and CS according to aggregate Input-Output-Tables. This reduces the share of CS employment to around 80%.

[15] We split the service sector into eight categories: "Retail and wholesale", "Hotels and restaurants", "Transport", "Finance", "Business services", "Health", "Community services", and "ICT".

[16] To corroborate our results, we also measured aggregate employment from the EC 2013. In the EC, wholesale, retail, restaurants, health, and community services account for 38% of total employment, which compares with approximately 6.5% for financial, business, and ICT services.

	Overall	In selected categories				Across space	
		Retail & Leisure & Health, etc.	Finance & Business, etc.	ICT	Transport & Storage	Urban	Rural
Share of CS	89	97	82	47	70	88	91

Table 2: Share of consumer service employment. The table reports the share of employment allocated to the CS sector. To aid readability we aggregate the service industries into four categories.

Human Capital Consistent with our theory, we measure each district's endowment of human capital $F_{rt}(q)$ and its distribution across sectors in terms of efficiency units of labor. 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, and measure the effect of each additional year by an estimated Mincerian returns to schooling ρ (see Section 5.1 below).

To measure the allocation of human capital to different sectors within each district, we use the observed distribution of earnings rather than a headcount of workers, because the former reflect differences in the use of effective units of labor. Appendix Table B-1, shows that educational attainment increased markedly between 1987 and 2011 and that is differs significantly across sectors, being the lowest in agriculture and the highest in PS. Interestingly, people working in the CS sector are on average more educated than those working in the industrial sector. There are also large spatial differences between more educated city dwellers and a less educated rural population. Measuring differences in educational attainment is important to separate the effect of human capital from that of changes in (disembodied) productivity.

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 from an aggregate production function, we rely on the equilibrium structure of our model and estimate the entire distribution of productivity $\{A_{rst}\}$ across sectors, space, and time.

The model has eight preference parameters and two parameters for the skill distribution: $\mathbf{\Omega} = \{(\varepsilon, \nu_{CS}, \nu_F, \nu_G, \omega_{CS}, \omega_F, \omega_G, \sigma), (\rho, \zeta)\}$. In addition, each region is characterized by a 3-tuple of regional productivity levels in agriculture, industry, and CS:

$\mathbf{A_{rt}} = \{A_{rFt}, A_{rGt}, A_{rCSt}\}$. Given the parameter vector $\mathbf{\Omega}$, there exists a one-to-one mapping from equilibrium skill prices $\{w_{rt}\}$ and sectoral employment allocations $\{H_{rst}\}$ to the underlying productivity fundamentals in $\mathbf{A_{rt}}$. 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 $\mathbf{A_{rt}}$ and its results.

5.1 Estimation of Structural Parameters

The Engel Elasticity: The elasticity ε ia a crucial parameter in our analysis. It determines how fast the expenditure on food shrinks—and, conversely, how fast it expands for CS—as income rises. To estimate ε, we use the cross-sectional relationship between household income and expenditure shares on food.

In general, it would not be legitimate to use expenditure data to infer structural parameters of the value-added demand system. However, Proposition 1 establishes that, under PIGL preferences, the demand system for sectoral value-added and the demand system for final expenditure have the same elasticity parameter ε. With this in mind, let $\mathcal{F} \in [0,1]$ denote the subset of the product space $[0,1]$ comprising all products classified as food items in the data. The spending share on food is then given by

$$\vartheta_{\mathcal{F}}^{FE}(e, \mathbf{p_r}) = \beta_{\mathcal{F}} + \kappa_{\mathcal{F}} \left(\frac{e}{\exp\left(\int_n \beta_n \ln p_{rn} dn\right)} \right)^{-\varepsilon}, \tag{16}$$

where $\beta_{\mathcal{F}} = \int_{n \in \mathcal{F}} \beta_n dn$ and $\kappa_{\mathcal{F}} = \int_{n \in \mathcal{F}} \kappa_n dn$. If the asymptotic expenditure share $\beta_{\mathcal{F}}$ is small—which is reasonable to assume for food items—Equation (16) yields a log-linear relationship between household income and expenditure shares:[17]

$$\ln \vartheta_{\mathcal{F}}^{FE}(e, \mathbf{p_r}) \approx \varepsilon \left(\int_n \beta_n \ln p_{rn} dn \right) - \varepsilon \times \ln e + \ln \kappa_{\mathcal{F}}. \tag{17}$$

We can then estimate ε from the linear regression

$$\ln \vartheta_{\mathcal{F}}^{h} = \delta_r + \varepsilon \times \ln e_h + x_h' \psi + u_{rh}, \tag{18}$$

where $\vartheta_{\mathcal{F}}^{h}$ denotes the food share of household h living in region r, e_h denotes total household spending, δ_r is a region fixed effects, and x_h is a set of household charac-

[17] The assumption that $\beta_{\mathcal{F}}$ is small is convenient but inconsequential. In Appendix C-1.1, we estimate ε from (16) without imposing this restriction. We find that $\beta_{\mathcal{F}} = 0$ is, in fact, the best estimate. Moreover, we also estimate ε for a range of value of $\beta_{\mathcal{F}}$ and find that they are very similar to the ones reported in Table 3. In column (3) of Table 3 we report the estimate for $\beta_{\mathcal{F}} = 0.05$.

	Food exp. share								Pooled data	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\ln e$	-0.332***	-0.321***	-0.369***	-0.313***	-0.334***	-0.395***			-0.230***	-0.392***
	(0.008)	(0.007)	(0.009)	(0.008)	(0.007)	(0.013)			(0.007)	(0.023)
$\ln e \times$ below median							-0.218***			
							(0.010)			
$\ln e \times$ above median							-0.415***			
							(0.011)			
$\ln e \times$ low urbanization								-0.291***		
								(0.007)		
$\ln e \times$ high urbanization								-0.358***		
								(0.012)		
Trim (top & bottom 5%)		✓	✓	✓	✓	✓	✓	✓	✓	✓
$\beta_{\mathcal{F}} = 0.05$			✓			✓			✓	✓
Serv. Categories									✓	✓
Addtl. Controls				✓	✓	✓	✓	✓	✓	✓
IV						✓				✓
N	101650	91492	91488	91443	1129730	85919	91447	91447	182068	171190
R^2	0.476	0.425	0.417	0.437	0.635	0.197	0.446	0.439	0.822	0.032

Table 3: ESTIMATES OF THE ENGEL ELASTICITY ε. The table shows the estimated coefficient ε of the regression (18). In columns 1 - 8, the dependent variable is the income share spent by each household on a set of 17 items classified as "food". These are: beverages; cereals; cereal substitutes; dry fruit, edible oil; egg, fish and meat; fresh fruit; intoxicants; milk and milk products; pan; packaged processed food products; pulses and products; salt and sugar; served processed food; spices; tobacco; vegetables. In all specifications, we control for a (within-district) urban/rural dummy, a set of fixed effects for household size, and the number of workers within the household. All regressions include region fixed effects; region-food item fixed are included in the fifth column. In columns 6 and 10, we instrument expenditure with a set of occupation fixed effects. In columns 9 and 10 we consider a pooled regression, where the dependent variables are $\ln\left(\vartheta^h_{\mathcal{F}} - \beta_{\mathcal{F}}\right)$ for food items and $\ln\left(\beta_{\mathcal{S}} - \vartheta^h_{\mathcal{S}}\right)$ for service items. Standard errors, clustered at the district level, in parentheses. *** p<0.01, ** p<0.05, * p<0.1

teristics that could induce a correlation between total spending $\ln e_h$ and food shares. Comparing (18) with (17), it is apparent that the terms $\left(\int_n \beta_n \ln p_{rn} dn\right)$ and $\ln(\kappa_{\mathcal{F}})$ are absorbed in the region fixed effects δ_r.

Table 3 reports the estimation results. We cluster standard errors at the district level. The first column refers to a specification that, in addition to district fixed effects, only controls for whether the household lives in an urban or rural area within each district, a full set of fixed effects for household size, and the number of workers within the household. We obtain an elasticity of 0.33 that is precisely estimated. In column (2), we trim the top and bottom 5% income levels, since we suspect these observations can contain some misreporting. The estimated elasticity is barely affected. In column (3), we set $\beta_{\mathcal{F}} = 0.05$ to match the average expenditure share of food at home in the US (CEX)—a proxy for the asymptotic food share. In column (4), we introduce additional household-level controls. In particular, we control through the inclusion of the respective fixed effects for (i) whether the household is self-employed (in agriculture or non-agriculture), (ii) whether the household is a regular wage earner or a casual laborer (in agriculture or non-agriculture), (iii) the household's religion, (iv) the household's social group, and (v) whether the household is eligible to purchase subsidized food from the government.

In column (5), we run a regression in which the unit of observation is the expenditure share on each of the 17 food items rather than the average expenditure on food and we control for region-food item fixed effects.[18] This increases the number of observations from ca. 91,000 to over 1.1 million. Reassuringly, the estimated elasticity is almost identical to that in the previous columns.

In column (6), we present the results from an IV regression addressing concerns about measurement error and unobserved income shocks that could bias the estimate. We instrument total expenditure with a full set of three-digit occupation fixed effects.[19] These fixed effects strongly predict total expenditure (F-Stat=140 in the first-stage regression). The exclusion restriction is that occupations only affect spending shares on food through their effect on income. The IV estimate of 0.395 is larger than the OLS estimate.

In Figure 3 we show a binscatter plot of the data for log food expenditure shares versus log expenditure after absorbing district-food item fixed effects, that is, corresponding to specification (5). Consistent with our PIGL specification, the relationship is indeed approximately log-linear. However, careful scrutiny reveals some mild concavity suggesting a higher elasticity for high-income consumers. In column (7) of Table 3, we allow different elasticities for households above and below the median income. The estimated elasticity is somewhat larger for high-income households. In column (8), we study the extent to which the elasticity differs between rural and urban districts. We define all districts in the top quartile of the distribution of urbanization as urban. While urban locations have higher elasticities, the differences are moderate.

Even though estimating ε from the expenditure system for food is consistent with our theory, we can also use information for other expenditure categories. The expenditure survey contains information on spending on some consumer service categories, such as domestic servants, barber shops, or tailor services—see Appendix C-1.2). In columns (9) and (10) we pool the expenditure shares on these services with those on food items and estimate ε using both sources of variation. More formally, we estimate (18) using as dependent variable $\ln\left(\vartheta^h_{\mathcal{F}} - \beta_{\mathcal{F}}\right)$ for food items and $\ln\left(\beta_{\mathcal{S}} - \vartheta^h_{\mathcal{S}}\right)$ for services. Note that $\beta_{\mathcal{S}} > \vartheta^h_{\mathcal{S}}$ if services are luxuries.[20]

[18] More formally, we run the regression $\ln \vartheta^h_{jr} = \delta_{jr} + \varepsilon \times \ln e_h + x'_h\psi + u_{jrh}$, where j denotes one of the 17 food items and δ_{jr} is a region-food item fixed effect.

[19] The survey assigns the occupation of the highest earning member to the entire household.

[20] Equation (16) implies that $\ln\left(s^h_{nrt}\right) = \upsilon_n + \varepsilon \exp\left(\int_n \beta_n \ln p_{rn} dn\right) - \varepsilon \ln e^h$, where for a necessity,

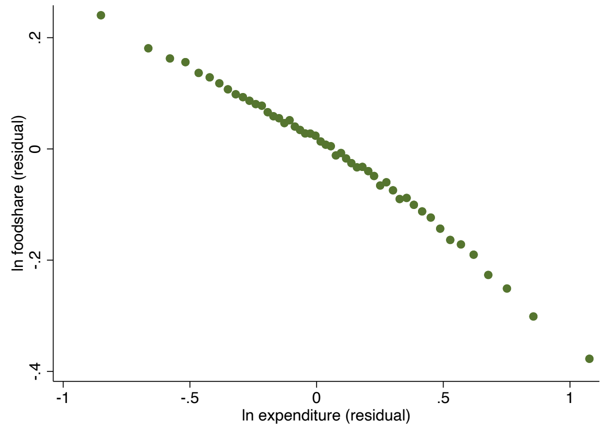

Figure 3: ENGEL CURVES IN INDIA. The figure shows a binscatter representation of the residual of a regression of the log expenditure share on food item j in region r on region-product fixed effects against the residual of a regression of the log income (total expenditure) on the same set of fixed effects. The slope coefficient of this plot yields the Engel elasticity. Cf. regression in column 5 of Table 3

We set $\beta_{\mathcal{S}}$ to match the expenditure share of the 99% quantile of the observed distribution in India. This yields $\beta_{\mathcal{S}} = 0.2$. For food items, we set $\beta_{\mathcal{F}} = 0.05$ as in column (3). In these regressions, we control for a full set of interactions of district-item fixed effects to account for price differences across both locations and type of final goods or services. While the OLS elasticity is smaller in column (9) than in column (4), the estimated coefficient in the IV regression of column (10) is almost identical to its analogue in column (5). We conclude that the results are robust to the inclusion of expenditure on some services.[21]

For our baseline analysis we take the Engel elasticity ε to be equal to the IV estimates of 0.395. Below we show that this is a conservative choice because the welfare gains attributed to CS productivity growth are decreasing in ε, implying that the effects we emphasize would be larger if we relied on the OLS rather than the IV estimates. Moreover, this estimate is closer to the estimates for rich households and urban location where concerns about non-measured subsistence food consumption might be less relevant.[22]

Other Preference Parameters: We estimate the six remaining parameters of the

$s^h_{nrt} = \vartheta^h_{nrt} - \beta_n$ and $\upsilon_n = \ln(\kappa_n)$ and, for a luxury, $s^h_{nrt} = \beta_n - \vartheta^h_{nrt}$ and $\upsilon_n = \ln(-\kappa_n)$.

[21] For our main specification, we rely exclusively on food expenditure data for two reasons. First, we believe they are more precisely measured. Second, we are guided by a precise prior on the asymptotic expenditure share. In Appendix C-1.2, we estimate ε using service expenditure alone. Reassuringly, the IV estimate of the Engel elasticity is not significantly different from that of column (6).

[22] If subsistence consumption accounts for a larger share of food consumption among the poor, richer households have higher market expenditures for food, biasing the Engel elasticity toward zero.

demand system, ω_s and $\overline{\nu}_s$, directly from the equilibrium conditions.[23] In Appendix A-2, we show that the market clearing conditions imply:

$$\sum_{r=1}^{R} w_{rt} H_{rFt} = \omega_F \sum_{r=1}^{R} w_{rt} H_{rt} + \overline{\nu}_F \sum_{r=1}^{R} \left(\omega_{CS} - \frac{H_{rCSt}}{H_{rt}} \right) w_{rt} H_{rt}. \tag{19}$$

Since these equations must hold for $t = 1987$ and $t = 2011$, they represent two moment conditions for the three parameters ω_F, ω_{CS} and $\bar{\nu}_F$. Note that these equations are independent of ε, trade costs, the elasticity of substitution σ, and the skill distribution. To attain identification, we exploit that ω_F pins down the asymptotic value-added share of the agricultural sector. In the US, the agricultural employment share (as well as its value-added share) is about 1%. Hence, we set $\omega_F = 0.01$ and use (19) for $t = 1987$ and $t = 2011$ to identify $\overline{\nu}_F$ and ω_{CS}.

As we show in Appendix A-2, $\overline{\nu}_{CS}$ is not separately identified from the level of productivity A_{rCSt}. Hence, without loss of generality, we normalize it to -1. The remaining parameters ω_G and $\overline{\nu}_G$ are pinned down by the homogeneity restrictions of the indirect utility function. Finally, we externally calibrate the trade elasticity σ and set it to five, a consensus estimate in the literature.

In the first panel of Table 4 we report the resulting estimates. The implied 70% asymptotic value-added share of CS, ω_{CS}, is reasonable.[24] For instance, the value-added share of the service sector in the US (that is not a targeted moment and includes PS and CS) has averaged 77% throughout the last decade. The asymptotic value-added share of the good producing sector (that includes both manufacturing and PS) is 30%. Moreover, $\overline{\nu}_G = -0.258$, which implies that industrial goods are also luxuries.

Skill Parameters ζ and ρ: To link observable schooling s_i to unobservable human capital q_i, we assume that $q_i = \exp(\rho s_i) \times \upsilon_i$, where s_i denotes the number of years of education, ρ is the annual return to schooling, and υ_i is an idiosyncratic shock, which we assume to be iid and which satisfies $\mathbb{E}[\upsilon_i] = 1$. Log earnings of individual i in region r at time t, y_{irt}, are thus given by a standard Mincerian regression $\ln y_{irt} = \ln w_{rt} + \rho s_i + \ln \upsilon_i$ and we can estimate ρ from the within-region variation between

[23] The market-level demand system depends on the aggregate preference parameters $\overline{\nu}_s$ which are related to the primitive micro-level preference parameters ν_s via (11). Identifying ν_s is only required to quantify the welfare consequences of service-led growth, not to estimate the model.

[24] Our model implies that the regional CS income share cannot exceed ω_{CS}. For $\omega_{CS} = 0.696$, four small districts violate the constraint. In these cases, we topcode the share of CS and split the excess proportionally between the other two sectors. In practice, this issue is inconsequential because these districts account for a mere 0.15% and 0.23% of Indian valued added in 1987 and 2011, respectively.

Parameter		Target	Value
Preference parameters	ε	Engel elasticity	0.395
	ω_F	Agricultural spending share US	0.01
	ω_{CS}	Equation (19), $t \in \{1987, 2011\}$	0.696
	ω_G	Implied by $\sum_s \omega_s = 1$	0.294
	$\overline{\nu}_F$	Equation (19), $t \in \{1987, 2011\}$	1.258
	$\overline{\nu}_{CS}$	Normalization	-1
	$\overline{\nu}_G$	Implied by $\sum_s \overline{\nu}_s = 0$	-0.258
	σ	Set exogenously	5
Skill parameters	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.

earnings and education. 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 et al. [2016]. In Appendix C-7, we show that our results are robust to using a higher return to education. Given the estimate of ρ, we then calculate the average amount of human capital per region as $\mathbb{E}_{rt}[q] = \sum_s \exp(\rho \times s)\ell_r(s)$, where $\ell_r(s)$ denotes the share of people in region r with s years of education.

To estimate the tail parameter of the skill distribution ζ, recall that the distribution of income in region r is given by $G_r(y) = 1 - \left(\underline{q}_r w_r / y\right)^{\zeta}$. We therefore estimate ζ from the tail of the income distribution within-regions. This procedure yields an estimate of $\zeta \approx 3$ (see Appendix C-2). With this estimate at hand, we can also compute the lower bound $\underline{q}_{rt}$ form $\mathbb{E}_{rt}[q_i] = \frac{\zeta}{\zeta-1}\underline{q}_{rt}$.

Trade costs τ: To calibrate the matrix of trade costs, we follow the literature on gravity equations and parameterize trade costs as a power function of distance, that is, $\tau_{rk} = \bar{\tau} d_{rk}^{\varrho}$, where d_{rk} is the geographic distance between districts r and k (Head and Mayer [2014]). We calibrate ϱ to match a distance elasticity of trade flows of -1.35 as reported in Monte et al. [2018].

5.2 Estimation of Productivity Fundamentals $\mathbf{A_t}$

In this section, we summarize the methodology to estimate $\mathbf{A_t}$, referring the reader to Appendix A-2 for details. 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 $\mathbf{A_t}$.

Consider first the identification of A_{rCSt}. The CS market clearing condition (14) implies that, for each region r, the local CS employment share is given by

$$\frac{H_{rCSt}}{H_{rt}} = \omega_{CS} + \overline{\nu}_{CS} \times \left(\underbrace{P_{rFt}^{-\omega_F} P_{rGt}^{-\omega_G}}_{\text{Prices}} \times \underbrace{\mathbb{E}_{rt}\left[q\right]}_{\text{Skills}} \times \underbrace{w_{rt}^{1-\omega_{CS}}}_{\text{Wages}} \times \underbrace{A_{rCSt}^{\omega_{CS}}}_{\text{Productivity}} \right)^{-\varepsilon}, \tag{20}$$

where $\bar{\nu}_{CS} < 0$ and $H_{rCSt}/H_{rt} < \omega_{CS}$, since CS are luxuries. Equation (20) highlights the role of income (through wages and the local supply of skills) and productivity in determining the size of the CS sector. Inverting the relationship yields a unique solution for A_{rCSt} as a function of observables. Given demand as determined by human capital, wages, and local prices of goods, the CS productivity A_{rCSt} is increasing in the observed employment share H_{rCSt}/H_{rt}. Conversely, given the employment share H_{rCSt}/H_{rt}, A_{rCSt} is decreasing in the local demand factors.

This structural decomposition of the observed variation in CS employment shares into income effects and service-led growth is a key step of our equilibrium accounting methodology. We emphasize that the estimates of A_{rCSt} do not rely on any published CS price index—an important advantage given the notorious measurement difficulties. We not even use data for the local prices of goods in (20) as these can be inferred from the respective equilibrium conditions. The only published price index we target is the average price of food relative to industrial goods.

The procedure to estimate productivity in the tradable sectors is different. Equation (15) implies relative productivity across two locations in sector s is given by (see Appendix A-2 for the derivation)

$$\frac{A_{rs}}{A_{js}} = \left(\frac{H_{rs}}{H_{js}}\right)^{\frac{1}{\sigma-1}} \times \left(\frac{w_r}{w_j}\right)^{\frac{\sigma}{\sigma-1}} \times \left(\frac{\sum_{d=1}^{R} \tau_{rd}^{1-\sigma} P_{dst}^{\sigma-1} \overline{\vartheta}_{dst} w_{dt} H_{dt}}{\sum_{d=1}^{R} \tau_{jd}^{1-\sigma} P_{dst}^{\sigma-1} \overline{\vartheta}_{dst} w_{dt} H_{dt}}\right)^{\frac{1}{1-\sigma}}. \tag{21}$$

Relative productivity A_{rs}/A_{js} is determined by three factors: relative employment shares H_{rs}/H_{js}, relative wages w_r/w_j, and relative demand as summarized by producer market access. A large employment share (holding wages fixed) and high wages (holding the employment share fixed) indicate that the location provides its goods at low prices. The market access term captures the correction associated with geography: *ceteris paribus*, the employment share in tradable goods is larger in districts that are close to centers of demand.

Equations (20)–(21) determine the distribution of sectoral productivity across lo-

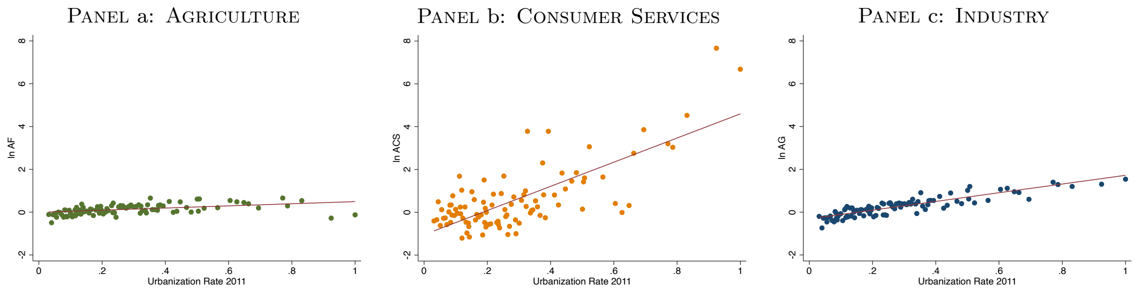

Figure 4: Estimated Sectoral Productivities. The figure shows a binscatter plot of the (demeaned) estimated sectoral labor productivities in agriculture, CS, and industry across urbanization-rate bins for 2011.

cations. To determine the level, we must pin down the average productivity growth for each sector between 1987 and 2011, which then determines the sectoral aggregate price levels. To this aim, we target two moments—see Appendix A-2. First, we target a 4.2% annual growth rate for real income per capita, which matches real GDP per capita growth according to the World Bank (WDI) using the industrial good as the numeraire.[25]. Second, we target the change in the price of agricultural goods relative to industrial goods as reported in the ETD. Empirically, agricultural prices rose by a factor of 1.52 relative to prices in the industrial sector.[26] Given these moments, our model identifies all productivity levels A_{rst} in both years of interest.

Results Figure 4 summarizes the cross-sectional pattern of our estimated productivities by way of a binscatter plot displaying $\ln A_{rs2011}$ as a function of the urbanization rate in 2011. Productivity in CS and industry is increasing with urbanization—see panels b and c. For agriculture, the relationship is flatter and slightly hump-shaped. The declining portion among more-urbanized districts likely reflects the scarcity of land (a factor of production from which we abstract) in urban areas.

Interestingly, both the productivity dispersion and its correlation with urbanization is strongest in the CS sector. Hence, the large employment share of CS in urban locations is not a mere consequence of high wages or an abundance of human capital; it also reflects high CS productivity. Among the tradable goods, productivity is significantly more dispersed in industry than in agriculture. To understand why, note that in Equation (21) a district's relative productivity reflects its sectoral earning share relative to its skill price. The "compressed" productivity distribution in agriculture re-

[25] We take GDP in terms of industrial goods as our measure of real GDP, because industrial goods are tradable. When we compute real GDP using a chained Fisher index, we get a growth rate of 4.6%.

[26] The ETD data covers the time period between 1990 and 2010. The ETD's precursor (the 10-sector database by the GGDC) ranges from 1987 to 2011 and reports a price change of 1.52.

	Sectoral Productivity Growth					
	10th	25th	50th	75th	90th	Aggregate
Consumer Services (g_{rCS})	-1.3	0.4	2.6	6.3	10.9	4.0
Agriculture (g_{rF})	0.3	1.1	1.8	2.7	3.3	2.0
Industry (g_{rG})	1.8	2.7	3.5	4.4	5.1	3.6

Table 5: Regional distribution of sectoral productivity growth. The table reports moments of the distribution of sectoral productivity growth. 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. Column 6 reports the population-weighted average in 2011.

flects the observation that wages are negatively correlated with the employment share of agriculture across districts. By contrast, wages are positively correlated with the employment share of industry, implying a wider productivity dispersion.

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, which we summarize in Table 5. Two patterns are salient. First and foremost, CS productivity grew in the vast majority of districts. In the median region, it grew by 2.6% annually between 1987 and 2011—less than productivity growth in the industrial sector and more than in agriculture. Second, productivity growth was unequal across space. In CS, the top 10% of locations experienced productivity growth exceeding 11%. When we aggregate regional productivity growth, we find that productivity growth in CS was about 4% and hence exceeded productivity growth in the two tradable sectors.[27]

In Appendix C-4 we show that local productivity growth is positively correlated with the urbanization rate in 1987. This correlation is also the reason why the population-weighed average of productivity growth exceeds the growth experience of the median locality. There we also show that estimated distribution of productivity growth is robust to the different values of ε reported in Table 3.

5.3 Nontargeted Moments

In this section we compare our model to a variety of nontargeted moments. We summarize here the main findings, while deferring details to Appendix C-5.

Nationwide Sectoral Productivity Growth: Our methodology allows us to recover sectoral productivity estimates for the entirety of India. While we are not aware

[27] To account for measurement error, we winsorize the top and bottom 3% of the estimated distribution of productivity growth in CS. The details are discussed in Appendix C-6, where we also report robustness results for these choices.

Annual Productivity Growth				
Agricult.	Manufac.	Mining	Finance & Business	Trade, Restaurants & Hotels
2.6%	5.3%	4.2%	4.1%	4.2%

Table 6: Annual Productivity Growth (ETD). The table reports aggregate productivity growth estimates from the ETD published by the GGDC for the period 1990–2010. Productivity is measured as real value added per worker.

of alternative estimates at the sector-region level, the ETD provides estimates of nationwide productivity growth for 12 aggregate sectors. In terms of our theory, their sector "Trade, restaurants, and hotels" is closest to our notion of CS.

In Table 6, we report annual sectoral productivity growth according to the ETD. These estimates are for value-added per worker growth and thus not directly comparable to ours. Nevertheless, they confirm the important role of the service sector for Indian growth. Productivity growth in the Indian retail sector was 4.2%. Productivity growth in manufacturing was higher, growth in the agricultural sector was appreciably slower. These figures are broadly in line with our findings reported in Table 5.[28]

Elasticities of Substitution and Income Elasticities: Given our estimated preference parameters, we can calculate the elasticities of substitution and the income elasticities. For the class of PIGL preferences, neither of them are structural parameters but vary with relative prices and total expenditure. In Appendix A-3, we show that the Allen-Uzawa elasticity of substitution between sectors s and k is given by

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

while the spending elasticity is given by $\frac{\partial \ln \vartheta_s e}{\partial \ln e} = 1 - \varepsilon \frac{\vartheta_s - \omega_s}{\vartheta_s}$.

In Table 7 we report the elasticities of substitution and the sectoral spending elasticities. Our estimates imply that CS and industrial goods are complements, with an elasticity of substitution ranging between 0.5 and 0.9, that agricultural and CS value-added are substitutes with an elasticity between 1.2 and 1.7, and that agricultural and industrial output are also substitutes, but with a smaller elasticity.

We find these results economically plausible. As the (quality-adjusted) price of CS-

[28] We also constructed pseudo-price indexes for CS in our model by averaging P_{CSrt} across districts, for $t \in \{1987, 2011\}$, using district population in 2011 as weights. Relative to the industrial price index, this CS pseudo-price index grows by 38% between 1987 and 2011. The ETD data reports an increase in the relative price of "Trade, Restaurants and Hotels" by 24%.

Urbanization	Elasticities of substitution			Spending elasticities		
quantile	Agr. & CS	Ind. & CS	Agr. & Ind.	Agr.	CS	Ind.
1 (Rural)	1.7	0.5	1.2	0.6	1.7	1.2
5 (Urban)	1.2	0.9	1.1	0.6	1.2	1.1

Table 7: ELASTICITIES OF SUBSTITUTION AND INCOME ELASTICITIES. The table reports the average elasticities of substitution between the respective pairs of sectoral output and the average income elasticities. Rural (urban) locations are defined as being in the lowest (highest) urbanization quantile.

intensive restaurants declines, individuals substitute away from home-cooked meals, making agricultural and CS value-added substitutes. Similarly, falling prices of industrial value-added might increase the spending share on CS value-added if consumers' reallocate their spending to products that heavily rely on CS. These results are also broadly in line with the existing literature. A number of papers document evidence of complementarity between goods and services either in two-sector models or in three-sector models where all elasticities are forced to be identical—see Herrendorf et al. [2014], Comin et al. [2021], and Duernecker et al. [2017]. Given the small size of the agricultural sector in the US, this is consistent with our finding of industrial goods and services to be complements. In terms of spending elasticities, we estimate CS and industrial goods to be luxuries and agricultural output to be a necessity. This is consistent with Comin et al. [2021] who report spending elasticities of 0.57, 1.15 and 1.29 for Tanzania.

Local Food Prices: Finally, our estimated model predicts local food prices that can be compared with the data inferred from the expenditure survey. In Appendix C-5 we show that these prices are strongly correlated across districts.

6 The Unequal Effects of Service-Led Growth

We now turn to our two main questions: How important was productivity growth in the service sector to rising living standards? And how skewed were these benefits across geography and the income distribution?

To quantify the welfare effects of CS growth, we compute counterfactual equilibria where we set CS productivity growth since 1987 to zero in all districts. The resulting changes in wages and employment allocations thus reflect the productivity growth in CS, holding constant productivity growth in tradable sectors and taking general equilibrium effects into account. We repeat the same exercise for productivity growth in agriculture and industry, respectively.

As in Baqaee and Burstein [2023], we measure welfare changes in terms of equivalent variations relative to the *status quo* in 2011. In other words, we calculate what share of its 2011 income a household residing in region r endowed with human capital q would be willing to forego to avoid the change of prices and wages associated with a counterfactual return of productivity in sector $s \in \{F, G, CS\}$ to the 1987 level in all Indian districts. More formally, let $x_r = (w_r, \mathbf{P}_r)$ and $\hat{x}_r = \left(\hat{w}_r, \hat{\mathbf{P}}_r\right)$ denote prices and wages in region r in 2011 and in a counterfactual scenario, respectively. Let $\varpi^q\left(\hat{x}_r|x_r\right)$ denote the percentage change in income an individual with skill level q facing prices and wages x_r requires, to achieve the same level of utility as under $\hat{x}_r$. For instance, if $\varpi^q = -20\%$, the consumer would be indifferent between giving up 20% of her 2011 income and the counterfactual allocation. Using the indirect utility function $\mathcal{V}$ given in (4), $\varpi^q\left(\hat{x}_r|x_r\right)$ is implicitly defined by

$$\mathcal{V}(qw_r(1+\varpi^q\left(\hat{x}_r|x_r\right)), \mathbf{P}_r) \equiv \mathcal{V}(q\hat{w}_r, \hat{\mathbf{P}}_r). \tag{22}$$

In Appendix A-4, we derive an analytical expression for $\varpi^q\left(\hat{x}_r|x_r\right)$.

6.1 Results: Sources of Welfare Growth in India

To highlight the unequal effects of service-led growth, we first zoom in on three districts with different characteristics. Then, we consider various levels of aggregation.

Three Indian Districts: Consider three selected districts: Bangalore, Chengalpattu, and Bankura. Bangalore is a fast-growing large urban district. Chengalpattu is a dynamic industrial district in Tamil Nadu that includes the southern suburbs of the megacity of Chennai.[29] Bankura is a rural district in West Bengal, which is mostly reliant on agriculture. Table 8 provides some descriptive statistics for these districts. Household income is significantly higher in Bangalore and Chengalpattu. Both the patterns of sectoral specialization and the estimated productivity growth are markedly diverse. In 2011 the employment share of CS is about 56% in Bangalore, 51% in Chengalpattu, and 28% in Bankura. There are large differences in CS productivity growth ranging from 2.4% in Bankura to 11% in Bangalore. Industrial productivity growth is high in both Chengalpattu and Bangalore, consistent with the boom of manufacturing

[29] We use the border of Chengalpattu in 1987. This district was split into Kancheepuram and Thiruvallur between 1991 and 2001. A district of Chengalpattu has then be reunified in 2019.

District	Urban	Population	Avg.	Emp. Share (%)			Prod. Growth (%)		
	Share		Income	Agr.	Ind.	CS	Agr.	Ind.	CS
Bangalore	0.77	10.6	3781	8	36	56	3.5	6.0	10.8
Chengalpattu	0.67	8.1	2807	12	37	51	2.8	5.0	8.6
Bankura	0.07	3.0	1597	64	7	28	1.5	2.1	2.4

Table 8: THREE INDIAN DISTRICTS. The table reports descriptive economic and demographic statistics in 2011 for the selected districts discussed in the text. The productivity growth are our estimates.

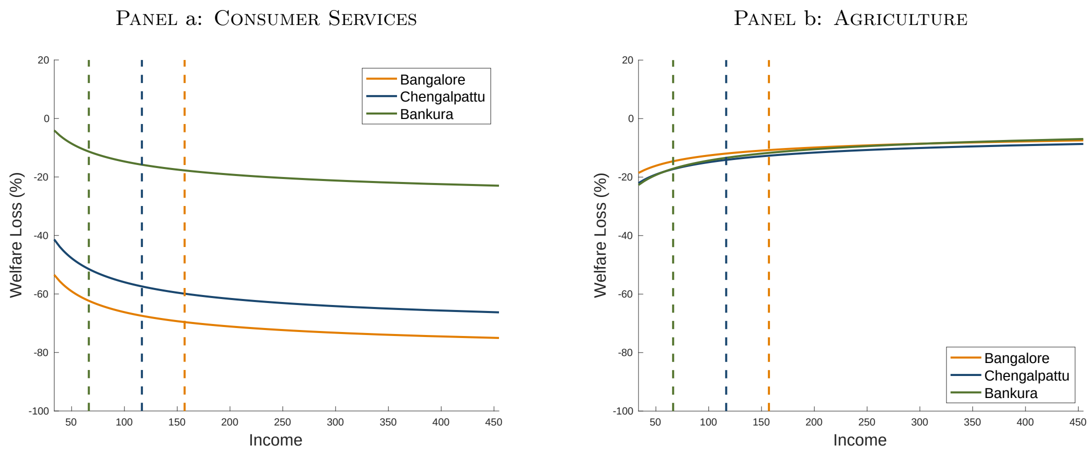

Figure 5: COUNTERFACTUAL WELFARE CHANGES. The figure displays the average percentage welfare losses associated with counterfactually setting productivity in CS (left panel) and agriculture (right panel) to their 1987 level in all Indian districts for households with different income levels living in Bangalore, Bankura, and Chengalpattu. The median income of Indian households is normalized to 100. The dashed lines indicate the average income in each district.

activity in the Chennai area and the ICT development in Bangalore. Productivity growth is lower in all sectors in Bankura.

In the left panel of Figure 5, we display the welfare effects of resetting CS productivity for the entirety of India to its 1987 level. We depict these effects separately for the three districts as a function of household income and indicate local average incomes with dashed vertical lines. The welfare effects of service-led growth vary significantly across space and income ladder. In rural Bankura, gains are small, especially for very poor households. The reason is twofold. First, the expenditure share on CS is low. Second, CS productivity growth is much lower than in Chengalpattu and Bangalore. Within each district, the gains from service-led growth are increasing in income. Even in Bankura, the equivalent variation for rich households exceeds 20% of their 2011 income. For the richest household in Bangalore, the corresponding figure is ca. 70%.

For comparison, in the right panel we depict the equivalent variations of agricultural productivity growth. For very poor households in Bankura, the equivalent variation

Panel a: Inequality Across Districts

Panel b: Inequality Across the Income Ladder

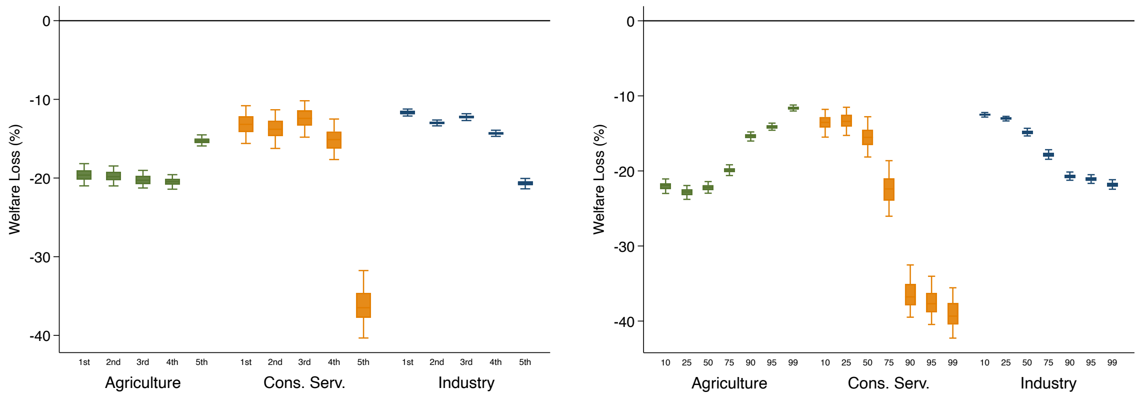

Figure 6: The Unequal Effects of Service-Led Growth. The figure displays the average percentage welfare losses (using district population as weights) associated with counterfactually setting productivity in agriculture, CS, and industry, to the respective 1987 level, broken down by urbanization quintile in 2011 (Panel a and by the 10th, 25th, 50th, 75th, 90th, 95th, and 99th percentiles of the income distribution in 2011 (Panel b). We compute the distribution of such welfare losses using a nonparametric 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% quantiles of the bootstrap distribution.

is 22%. The bulk of welfare gain accrues from agricultural productivity growth in the entirety of India (rather than from changes in local productivity), which reduces food prices overall. The diffusion of the effects of productivity growth in agriculture via trade explains why the spatial differences are small.

Average Effects: To draw more general lessons, we now compute average welfare effects at different levels of aggregation. In Figure 6 (left panel) we depict the population-weighted average equivalent variation in different urbanization quintiles.

Because the effects are based on an estimated model, they entail sampling uncertainty. To quantify the extent of this uncertainty, we estimate the distribution of the welfare effects using a nonparametric bootstrap procedure (Horowitz [2019])—see Appendix WA-5 for details. In Figure 6 we report these distributions as a boxplot. Each box shows the 25%–75% quantiles of the distribution of 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% quantiles, respectively.

The benefits of agricultural productivity growth are larger in rural than in urbanized districts. For households in the four lowest quintile of urbanization, the average equivalent variation is about 20%. For the top quintile of urbanization, it drops to 15%. By contrast, the gains from productivity growth in CS and industry are skewed toward urban locations. The average equivalent variation for CS is a staggering 37% for the most urbanized quintile.

The right panel of Figure 6 focuses on income distribution, showing the welfare effects at the 10th, 25th, 50th, 75th, 90th, 95th, and 99th percentiles. As expected, the benefits of productivity growth in CS and industry are increasing in income, whereas the pattern is opposite for agriculture. In the case of CS, the equivalent variation for the top decile of the income distribution is comparable to that for the top quintile of urbanization. Interestingly, for households below the median income, the welfare effects of productivity growth in CS and in the industrial sector are roughly equal and are smaller than those from agriculture. For the top 25%, service-led growth was quantitatively much more important.

Finally, in the left panel of Figure 7 we report the population-weighted average equivalent variation across all Indian districts. On average, Indians would have been willing to sacrifice 21% of their income in lieu of giving up the observed productivity growth originating in the CS sector. 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 increase in economic well-being. Productivity growth in agriculture and industry were also important sources of welfare improvement, albeit smaller than CS.

In summary, productivity growth in CS played an important role for economic development in India. In urban areas and for rich households, growth in CS was the dominant source of rising living standards. By contrast, technical progress in agriculture was the most important source of welfare gains for below-median households.

6.2 Structural Change

Figure 1 shows that growth without industrialization is a salient feature of the recent economic development in India and in the developing world. In this section, we show that productivity growth in the *CS* was an important engine of this process.

Structural Change in the Theory: We start by highlighting some analytical properties of the model. We first show how changes in prices affect sectoral spending shares. Differentiating Equation (10) for any two sectors s and k yields

$$\frac{\partial \ln \overline{\vartheta}_{rst}}{\partial \ln P_{rkt}} = \varepsilon \omega_k \frac{\overline{\vartheta}_{rst} - \omega_s}{\overline{\vartheta}_{rst}}. \tag{23}$$

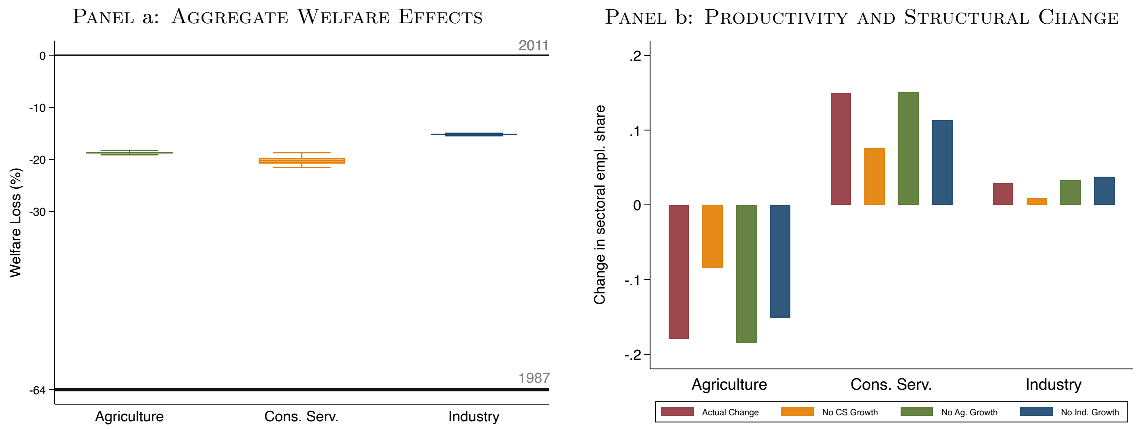

Figure 7: Aggregate Welfare Effects and Structural Change. In the left panel we show the analogue of Figure 6 with welfare effects aggregated up to the nationwide level. In the right panel we show changes in sectoral employment (in efficiency units). We depict the actual change in India (red bars) and the counterfactual results in the absence of productivity growth in the CS sector (orange bars), agriculture (green bars), and industrial sector (blue bars).

Because F is a necessity, whereas G and CS are luxuries, $\overline{\vartheta}_{rFt} > \omega_F$, whereas $\overline{\vartheta}_{rGt} < \omega_G$, and $\overline{\vartheta}_{rCSt} < \omega_{CS}$. Thus, productivity growth in *any* sector increases the expenditure share on goods and CS and decreases the expenditure share on food. In the case of CS, $\overline{\vartheta}_{rCSt} = H_{rCSt}/H_{rt}$. Thus, productivity growth in *any* sector increases the employment share of CS. However, the elasticity in (23) depends on the sectoral origin of productivity growth. In particular, $\frac{\partial \ln \vartheta_{rCSt}/\partial \ln P_{rCSt}}{\partial \ln \vartheta_{rCSt}/\partial \ln P_{rFt}} = \frac{\omega_{CS}}{\omega_F}$, which is a large number in our calibration, implying that a given productivity growth rate in CS causes more structural change than an equal productivity growth rate in agriculture.

To gauge the magnitude of such differences, consider the Indian economy in 1987. A 10% increase in A_{rCS} in all districts would trigger changes in the employment shares of F, G, and CS of, respectively, -1, 0.2, and 0.8 p.p. Note that this split, whereby 80% of the decline in agriculture gets absorbed in the service sector, is quantitatively in line with the experience of most developing countries, documented in Figure 1. Instead, a 10% increase in A_{rF} in all districts would yield much smaller changes of -0.015, 0.003, and 0.012 p.p.[30] While productivity growth in agriculture causes structural change, its quantitative effects are small in our calibration.

Structural Change in the Estimated Model: We now consider the productivity changes we inferred from the calibrated model. The right panel of Figure 7 shows

[30] These employment effects, which include general equilibrium effects, track closely those predicted by Equation (23). According to the latter, a 10% increase in A_{rCS} yields changes in the expenditure shares of F, G, and CS of, respectively, -1.1, 0.2, and 0.9 p.p. A 10% increase in A_{rF} yields corresponding changes of -0.018, 0.004, and 0.14 p.p.

the sectoral reallocation between 1987 and 2011. All figures are in effective units of labor. In contrast to the welfare analysis, sampling variation plays a minor role for these results and we do not include the standard errors to improve readability. The red (leftmost) bars show the actual data for India: agricultural employment declined by 18 p.p. and CS increased by 15 p.p. The industrial sector, which contains PS, only saw an increase by 3 p.p. The remaining three bars depict the counterfactual change in the sectoral employment shares when we shut down (one at a time) productivity growth in CS, agriculture, and industry, respectively. Productivity growth in CS (orange bars) is responsible for the lion's share of the structural transformation. In its absence, the agricultural employment share would have only declined by about 9 p.p. (instead of 18 p.p.) and the rise in CS employment would have only been 8 p.p (instead of 15 p.p.). Hence, our theory does attribute an important role to service-biased growth, whereby labor reallocates from agriculture to services even in the absence of CS productivity growth. However, quantitatively, this mechanism explains only half of the observed structural transformation in India.

The Role of Agriculture: In line with Equation (23), the figure shows very modest effects of productivity growth in agriculture on structural change (green bars). Somewhat surprisingly, it appears to have marginally *increased* employment in agriculture and slowed down employment growth in industry and CS. This result runs opposite to the experiment above in which we increased agricultural productivity uniformly across all districts. To understand why, consider a different experiment where all districts are initially identical and agricultural productivity increases in the single region $\tilde{r}$. Region $\tilde{r}$ now has a comparative advantage in food production that pushes up employment in that sector. Also, the ensuing increase in local wages deters growth in CS. In this case, a high *relative* productivity in agriculture slows down local structural change. This prediction is in line with the findings of recent empirical papers of Asher et al. [2022], who study the long-run impacts of irrigation canals on structural change in India, and Kelly et al. [2022], who document a negative effect of local agricultural productivity on the onset of the British Industrial Revolution.[31] Returning to Figure 7, asymmetric agricultural productivity growth (as those implied by our estimates) yield in general ambiguous net average effects.

[31] Concerning the nationwide effect of agricultural productivity growth, Gollin et al. [2021] find that the intensity of adoption of high-yielding crop varieties (the Green Revolution) sped up the decline of agriculture. Moscona [2020] finds opposite effects at the regional level: districts that saw faster agricultural productivity growth increased their employment share, as implied by our model.

Taking Stock: Service-led growth explains most of India's structural transformation between 1987 and 2011. Without productivity growth in CS, India would still be a much more rural economy today.

7 Robustness

In this section, we discuss the robustness of the welfare analysis. In Section 7.1, we study the sensitivity of our results to changes in structural parameters, most notably, the Engel elasticity ε. In Section 7.2, we revisit some measurement choices concerning the split between CS and PS. In Section 7.3, we generalize our preference structure. In Section 7.4, we study various generalizations of the model (open economy, skill heterogeneity, spatial mobility). For each experiment, we report equivalent variations associated with productivity growth in CS in Table 9. We defer the corresponding results for agricultural and industrial productivity growth to Appendix C-7.

7.1 Sensitivity to Structural Parameters

The Engel elasticity ε is the most important parameter in our theory. The effect of CS productivity is decreasing in ε, because a high elasticity attributes a large share of employment growth in the CS sector to income effects.

For our analysis, we relied on the IV estimate of $\varepsilon = 0.395$ (column 6 in Table 3). In the first row of Table 9, we present an alternative calibration based on the elasticity estimated for the sample of high-income households, $\varepsilon = 0.415$, which is the largest elasticity in Table 3. The effects are marginally smaller but very similar to the baseline results. In Appendix C-7, we allow ε to be larger in more-urbanized than in less-urbanized districts, according to the estimates of column 8 in Table 3. We find only a marginal reduction in the inequality of welfare effects across districts. In the second row, we set $\varepsilon = 0.32$, the OLS estimate of the Engel elasticity. This change reduces the income effects and magnifies the importance of service-led growth, especially in cities.[32] In summary, our main results are robust to the entire range of ε estimated in Table 3.[33]

[32] Boppart [2014] estimates Engel elasticities for the US from CEX and PSID. His estimates range between 0.22 and 0.29. Because his preference specification is slightly different (e.g. he abstracts from agriculture), his elasticities are not directly comparable with ours. Nevertheless, the results in Table 9 indicate that lower income elasticities would magnify the welfare effects associated with CS growth.

[33] We also considered a calibration where we do not estimate ε but calibrate it targeting the aggregate productivity growth of the Indian retail sector (4.2%) according to ETD (see Table 6). This yields $\varepsilon = 0.384$, which is smaller than our baseline estimate. The resulting welfare gains are slightly larger.

	Aggregate Effects	by Urbanization Quintiles		by Income Quantiles		
		1st	5th	10th	50th	90th
Baseline	-20.5	-13.1	-36.7	-13.6	-14.5	-37.6
	Alternative calibrations of ε (Section 7.1)					
$\varepsilon = 0.415$ (High Income Households)	-19.4	-12.2	-35.3	-12.6	-13.5	-36.2
$\varepsilon = 0.321$ (OLS estimate)	-25.2	-17.0	-42.7	-17.8	-19.0	-43.4
	Alternative measurement choices (Section 7.2)					
Allocate PS share based on WIOD	-18.9	-13.4	-31.4	-13.9	-14.0	-32.0
Allocate ICT & Business to PS	-17.3	-15.1	-24.8	-14.1	-12.4	-25.2
Allocate Construction to Industry	-11.1	-1.6	-30.5	-2.9	-4.4	-25.1
	Alternative modeling assumptions (Section 7.4)					
Open economy	-17.8	-11.8	-31.5	-12.6	-12.2	-31.6
Imperfect skill substitution	-18.4	-10.0	-33.4	-10.0	-10.7	-34.8
Spatial labor mobility	-18.6	-13.4	-30.1	-	-	-

Table 9: THE IMPORTANCE OF SERVICE-LED GROWTH—ROBUSTNESS. The table reports a summary of the robustness tests described in the main text. The numbers indicate percentage equivalent variations associated with setting the 2011 productivity level in the CS sector to the corresponding 1987 level in all Indian districts.

In Appendix C-7, we also discuss the sensitivity of our results to changes in other parameters: the asymptotic food share ω_F, the tail of the skill distribution ζ, the educational return ρ, and the elasticity of substitution across local varieties σ (all other parameters are either point-identified in our theory or pinned down by normalization.) The effects of these changes are quantitatively small and do not affect our conclusions.

7.2 Measurement: The PS-CS Split

Our split of service employment into PS and CS reported in Table 2 hinges on whether service firms sell mostly to firms or consumers. Our data-driven approach could underestimate the PS sector if firms report sales to small firms as sales to individuals. To address this concern, we consider two alternative classifications.

First, we use aggregate Input-Output-Tables from the WIOD to measure the share of service output that is used as an intermediate input in the industrial and agricultural sector. In India this number is ca. 20%. Thus, we increase the relative size of the PS sector so that it accounts for 20% of value-added on the service sector altogether. This procedure implies that we assign 18% rather than 11% of service employment to PS.

Second, we treat business services and ICT as only producing tradable services and allocate them entirely to PS, while retaining our baseline approach for the remaining service industries. We view this as a generous upper bound as in reality a sizeable portion of services in finance or law are sold to consumers (e.g., divorce lawyers). In this case, PS accounts for 22% service employment. Because business services are especially salient in cities, this measurement choice reduces the role of CS in urban areas relative to the first exercise based on WOID.

The results are shown in rows 3 and 4 of Table 9. As expected, attributing a larger share of the expanding service sector to PS reduces the estimated productivity growth in CS. Accordingly, the average welfare effect of service-led growth declines by 1.6 and 3.2 p.p., respectively, but remains large. At the spatial level, welfare inequality shrinks, but overall CS productivity growth continues to benefit mostly the urban dwellers.

Finally, we turn attention to the construction sector. In our analysis, we merge construction with the CS sector because of its non-tradable nature. However, the traditional classification treats it as part of the industrial sector. Although our classification is more logical in the framework of our theory, we report the result of following the traditional classification in row 5 of Table 9. The reclassification increases the average welfare effect of productivity growth in the industrial sector, which goes up to 20%—see Appendix Table C-7. We still find CS to be an important contributor to aggregate welfare growth, however the magnitude is appreciably smaller. Interestingly, the welfare effects of service-led growth are even more skewed in favor of urban districts than in our baseline estimate, because the construction sector is relatively more salient in rural areas. The smaller aggregate welfare effect is therefore mostly driven by rural districts, where construction accounts for the bulk of non-tradable activities. By contrast, service-led growth in urban locations is not primarily driven by construction.

7.3 Generalized PIGL Preferences

In our analysis, we parameterized the indirect utility function by setting $D(\mathbf{P_r}) = \sum_s \nu_s \ln P_{rs}$ in the value-added representation. In this section, we generalize the approach of Boppart [2014] to a three-sector environment. We assume a CES function:[34]

$$D(\mathbf{P_r}) = \frac{1}{\gamma}\left(\left(\sum_{s\in\{F,G,CS\}} P_{rs}^{\nu_s}\right)^{\gamma} - 1\right), \tag{24}$$

[34] This specification preserves the isomporphism between the expenditure and the value-added approach—see Web Appendix WA-2.1. Here, for simplicity, we write the value-added functions.

where $\sum_s \nu_s = 0$. The associated expenditure share is given by

$$\vartheta_{rst}\left(e, \mathbf{P_r}\right) = \omega_s + \nu_s \left(\frac{e}{\prod_{j\in\{F,G,CS\}} P_{rj}^{\omega_j+\gamma\nu_j/\varepsilon}}\right)^{-\varepsilon}. \tag{25}$$

This new specification flexibly adjusts the weights of the pseudo-price index by a term that depends on the new parameter γ that is set to zero in our baseline specification. In other words, the parameter γ affects the strength of relative price effects. In particular, Equation (23) now reads:

$$\frac{\partial \ln \overline{\vartheta}_{rst}}{\partial \ln P_{rk}} = \left(\gamma\nu_k + \varepsilon\omega_k\right) \times \frac{\vartheta_{rst} - \omega_s}{\vartheta_{rst}}. \tag{26}$$

Here, the relationship between prices and expenditure shares hinges on the sign of $\gamma\nu_k + \varepsilon\omega_k$. If $\gamma = \gamma^* \equiv -\varepsilon\frac{\omega_{CS}}{\nu_{CS}} > 0$, the CS employment share is independent of A_{rCSr}, preventing the identification of the local CS productivities from local employment data. If $\gamma > \gamma^*$, a fall in P_{rCSt} reduces $\overline{\vartheta}_{rst}$ and H_{rCSt}/H_{rt}. Web Appendix Figure WA-1 illustrates the paradoxical implications of this calibration of γ for the Indian economy. In the cross-section, the model associates high local employment shares in CS with low productivity in CS. Over time, it attributes growing employment shares in CS to negative productivity growth in CS. Cities like Mumbai, Delhi or Bangalore would have *lower* productivity in the CS sector against the intuitive argument that cities attract large and more efficient retailers or health providers. What's more, estimated productivity *growth* in CS is negative in a large number of districts (and on average), more so in urban districts where the CS employment share grew the most. We find this topsy-turvy pattern implausible and, hence, restrict attention to the range $\gamma < \gamma^*$.

To see how γ affects the welfare effects of service-led growth, consider again the three districts of Bangalore, Chengalpattu, and Bankura. The left panel of Figure 8 shows how the welfare effects associated with the estimated productivity growth in CS over the period 1997-2011 vary as functions of γ.[35] The special case of $\gamma = 0$ corresponds to our baseline analysis. The welfare effects are increasing in γ. As $\gamma \to \gamma^*$, the model requires larger and larger variations in CS prices (hence, productivities) to rationalize the observed variation in employment shares. Over time, it requires a larger productivity growth in CS, which magnifies the welfare effects. Note that, while

[35] For a given value of γ, we can identify our model with preferences based on (24) from exactly the same moments as our baseline model. We always recalibrate all other parameters when varying γ.

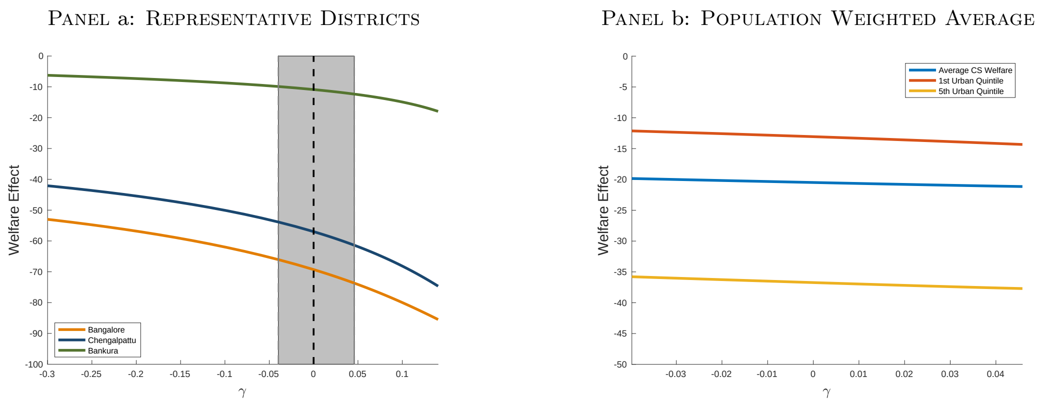

Figure 8: Welfare Effects of CS Productivity Growth as a function of γ. In the left panel we depict the welfare effect of CS growth as a function of γ in three representative districts of India. We depict the range of γ where $0 \leq EOS_{CS,G} \leq 1$ for 90% of districts as the shaded areas. In the right panel we show the aggregate welfare effect and the welfare effect for the 1st and 5th quantile of the urbanization rate.

the welfare effects become unboundedly large as $\gamma \to \gamma^*$, they decline only slowly in the range of negative γ. Appendix Figure WA-2 shows how changes in γ affect the distribution of productivity growth in CS in the region where $\gamma < \gamma^*$. Increasing γ raises both the average and the spread of productivity growth.

We can further compress the range of plausible γ's by considering the implied Allen-Uzawa elasticities of substitution between G and CS (see Web Appendix WA-2.3 for details). As we have shown in Table 7, our baseline analysis of $\gamma = 0$ implies that $EOS_{CS,G} \in (0, 1)$, i.e. goods and CS are closer complements than under Cobb-Douglas preferences. This is in line with the estimates in the literature and with Boppart [2014]. If we focus on the range of γ such that $EOS_{CS,G} \in (0, 1)$ in at least 90% of the Indian districts, we obtain $\gamma \in [-0.05, 0.04]$. This range is also consistent with the aggregate rate of CS productivity growth. If we calibrate γ to match the rate of 4.2% as reported in Table 6, we find $\gamma = 0.02$.

In the left panel of Figure 8, we highlight this range as the shaded area. In the right panel, we zoom in on that range and depict the population weighted-average welfare effect at both the aggregate level and for different urbanization quintiles. The welfare effects are quantitatively very similar to our baseline estimates.

7.4 Other Generalizations of the Theory

In this section, we outline three generalizations of the theory that we present more formally in Appendix A-5 and Web Appendix WA-3.

Open Economy: Our main analysis treats India as a closed economy. However, international trade, in particular exports of ICT services, has become increasingly important. To incorporate these dimensions, we extend our model to allow for international trade. We assume households, both in India and in the rest of the world, consume differentiated industrial goods sourced from many countries. To capture India's comparative advantage in ICT, we assume India is an ICT exporter and exports the entirety of its ICT value-added. We calibrate the parameters so as to generate trade flows like in the data. The results of the counterfactual analysis are shown in Table 9. International trade—especially, recognizing the tradable nature of ICT services—mildly reduces the welfare effect of productivity growth in CS, especially in cities, which (as shown in Table 1) saw the fastest increase in ICT employment. Nevertheless, CS continue to play an important role for aggregate growth and for urban areas in particular.

Imperfect Substitution and Skill Bias in Technology: Our main analysis assumes that all workers' efficiency units are perfect substitutes. We generalized our model assuming workers with different educational attainments are imperfect substitutes. Because agricultural workers have, on average, lower educational attainment, an increase in the skill endowment could be responsible for the reallocation of workers from agriculture to CS (see, e.g., Porzio et al. [2022] or Schoellman and Hendricks [2023]). By ignoring skill-based specialization, our Ricardian model could exaggerate the importance of productivity in determining the expansion of the service sector.

We postulate two skill groups and define workers to be skilled if they have completed secondary school. We assume the production functions to be of the CES form

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

where H^{+} and H^{-} denote high- and low-skilled workers, respectively. Note that the technology admits differences in both Hicks-neutral TFP and skill bias (Z_{rst}) across sector-districts and time. We calibrate the elasticity of substitution between high- and low-skilled workers to 1.8, a standard estimate in the literature. The results in Table 9 show that the quantitative role for the CS sector is very similar to the one of our baseline calibration. If anything, the unequal effects across the income ladder are more pronounced because skilled individuals are more likely to work in the CS sector.

This extension yields two additional findings. 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].

Spatial Mobility: In our baseline model, we assumed people to be spatially immobile. However, a counterfactual decline in CS productivity could prompt people to move out of cities. Labor mobility could then work as a form of insurance, thereby reducing the equivalent variation associated with CS productivity growth. To gauge the quantitative importance of labor mobility, we re-estimate our model in the presence of an endogenous location choice, which we model as a discrete choice, where individuals receive idiosyncratic preference shocks and locations differ in amenities.

Allowing for an endogenous location choice does not affect the estimation of the parameters nor the productivities. The spatial distribution of amenities rationalizes the observed population distribution as an equilibrium outcome, given the equilibrium prices. To calculate the welfare effects in the presence of spatial mobility, we first set local amenities so that the spatial equilibrium matches the distribution of the Indian population across districts in 2011. Next, we sample one million fictitious households and associate each of them with a vector of realizations of the geographic preference shock (one per district). Then, we counterfactually reset the CS productivity distribution to the 1987 level, allowing people to optimally relocate to their preferred district. Finally, we calculate the equivalent variation for each sampled household.

In the last row of Table 9, we report the results of an experiment in which we calibrate the elasticity of labor mobility to match the observed migration flows in India.[36] As expected, labor mobility lowers the equivalent variation of productivity growth in CS, but the difference is moderate—from an average 21.6% to 18.6%. The effect is somewhat more conspicuous for households that chose to reside in urban areas in the baseline economy of 2011. Intuitively, resetting CS productivity to the 1987 level reduces the economic appeal of urban areas. The option to migrate allows some households to partially offset the economic losses by moving to districts that better suit their geographic preferences. Altogether, the results do not alter the broad picture.

[36] See Appendix A-5.3 for details of the model and its calibration. We do not report the results by income, because individuals draw their human capital after moving.

8 Conclusion

Service-led growth is a widespread feature of the contemporary world. The classic argument of Baumol [1967] suggests that this trend—if solely driven by income effects—could lead to economic stagnation. This view has been recently echoed by Rodrik [2016] who expresses concern for the premature deindustrializatioon of many developing countries. In this paper, we develop a novel methodology to determine the importance of productivity growth in services an engine of growth and structural transformation. The methodology lends itself to a quantitative analysis of both the aggregate welfare effects of service-led growth and its distributional consequences.

Our application to India delivers two main results. First, productivity growth in consumer services was both fast and important for welfare, accounting for one-third of the improvement in living standards between 1987 and 2011. However, service-led growth in India was also unequal: it disproportionally benefited the urban middle-class and affluent dwellers while being far less consequential for poor people. The reasons is twofold: service productivity grew particularly fast in urban areas and richer consumers spend more on service-intensive goods owing to nonhomothetic preferences. While our analysis suggests that low employment growth in the manufacturing sector could be less of a threat to the sustainability of future growth than economists previously thought, it also raises novel concerns about inequality that remain invisible in aggregate statistics.

Our approach has some important limitations. Among them, two are particularly pressing. First, our accounting perspective takes productivity as exogenous. Understanding the drivers of sectoral productivity over time and across space is a question of first-order importance, especially for policy guidance. Second, we did not consider the effects of service-led growth on other dimensions of inequality such as gender disparity. For instance, Rendall [2013] argues that the growth of services tends to reduce gender inequality by decreasing the labor demand for physical strength. We plan to study these aspects in future research.

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi: 2012, 'The Network Origins of Aggregate Fluctuations'. *Econometrica* **80**(5), 1977–2016.

Acemoglu, D. and F. Zilibotti: 2001, 'Productivity Differences'. *Quarterly Journal of Economics* **116**(2), 563–606.

Alder, S., T. Boppart, and A. Müller: 2022, 'A Theory of Structural Change That Can Fit the Data'. *American Economic Journal: Macroeconomics* **14**(2), 160–206.

Allen, T. and C. Arkolakis: 2014, 'Trade and the Topography of the Spatial Economy'. *Quarterly Journal of Economics* **129**(3), 1085–1140.

Asher, S., A. Campion, D. Gollin, and P. Novosad: 2022, 'The Long-run Development Impacts of Agricultural Productivity Gains: Evidence from Irrigation Canals in India'. Mimeo, Johns Hopkins.

Atkin, D., B. Faber, and M. Gonzalez-Navarro: 2018, 'Retail Globalization and Household Welfare: Evidence from Mexico'. *Journal of Political Economy* **126**(1), 1–73.

Baqaee, D. and A. T. Burstein: 2023, 'Welfare and output with income effects and taste shocks'. *The Quarterly Journal of Economics*. forthcoming.

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.

Budí-Ors, T. and J. Pijoan-Mas: 2022, 'Macroeconomic Development, Rural Exodus, and Uneven Industrialization'. Mimeo, CEMFI.

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., pp. 679–741.

Chatterjee, S. and E. Giannone: 2021, 'Unequal Global Convergence'. Working Paper.

Chen, X., G. Pei, Z. M. Song, and F. Zilibotti: 2022, 'Tertiarization Like China'. *NBER Working Paper* (w30272).

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

de Vries, G., L. Arfelt, D. Drees, M. Godemann, C. Hamilton, B. Jessen-Thiesen, A. I. Kaya, H. Kruse, E. Mensah, and P. Woltjer: 2021, 'The Economic Transformation Database (ETD): Content, Sources, and Methods'. WIDER Technical Note 2/2021.

Dehejia, R. and A. Panagariya: 2016, 'The link between manufacturing growth and accelerated services growth in India'. *Economic Dev. 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.

Duernecker, G., B. Herrendorf, and A. Valentinyi: 2017, 'Structural change within the service sector and the future of Baumol's disease'. *Available at SSRN 3082293*.

Eckert, F., S. Ganapati, and C. Walsh: 2020, 'Skilled Scalable Services: The New Urban Bias in Economic Growth'. Minneapolis Reserve Bank of Minneapolis, Institute WP 25.

Eckert, F. and M. Peters: 2020, 'Spatial Structural Change'. Mimeo, Yale University.

Fajgelbaum, P. D. and A. K. Khandelwal: 2016, 'Measuring the unequal gains from trade'. *The Quarterly Journal of Economics* **131**(3), 1113–1180.

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 (10th World Congress of the Econometric Society)*. Cambridge Univ. Press, pp. 373–418.

García-Santana, M., J. Pijoan-Mas, and L. Villacorta: 2021, 'Investment Demand and Structural Change'. *Econometrica* **89**(6), 2751–2785.

Gervais, A. and J. B. Jensen: 2019, 'The tradability of services: Geographic concentration and trade costs'. *Journal of International Economics* **118**, 331–350.

Gollin, D., C. W. Hansen, and A. M. Wingender: 2021, 'Two blades of grass: The impact of the green revolution'. *Journal of Political Economy* **129**(8), 2344–2384.

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 Ouput Per Worker Than Others?'. *Quarterly Journal of Economics* **114**(1), 83–116.

Head, K. and T. Mayer: 2014, 'Gravity equations: Workhorse, toolkit, and cookbook'. In: *Handbook of international economics*, Vol. 4. Elsevier, pp. 131–195.

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.

Horowitz, J. L.: 2019, 'Bootstrap methods in econometrics'. *Annual Review of Econ.* **11**, 193–224.

Hsieh, C.-T. and E. Rossi-Hansberg: 2023, 'The Industrial Revolution in Services'. *Journal of Political Economy Macroeconomics*. forthcoming.

Kelly, M., J. Mokyr, and C. O'Grada: 2022, 'The Mechanics of the Industrial Revolution'. *Journal of Political Economy* **120**.

Kongsamut, P., S. Rebelo, and D. Xie: 2001, 'Beyond balanced growth'. *Review of Economic Studies* **68**(4), 869–882.

Matsuyama, K.: 2019, 'Engel's Law in the Global Economy: Demand-Induced Patterns of Structural Change, Innovation, and Trade'. *Econometrica* **87**(2), 497–528.

Monte, F., S. J. Redding, and E. Rossi-Hansberg: 2018, 'Commuting, migration, and local employment elasticities'. *American Economic Review* **108**(12), 3855–90.

Moscona, J.: 2020, 'Agricultural Development and Structural Change, Within and Across Countries'. Working Paper.

Muellbauer, J.: 1976, 'Community Preferences and the Representative Consumer'. *Econometrica* **44**(5), 979–999.

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

Ngai, L. R. and 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: 2022, 'The human side of structural transformation'. *American Economic Review* **112**(8), 2774–2814.

Redding, S. J. and E. Rossi-Hansberg: 2017, 'Quantitative Spatial Economics'. *Annual Review of Econ.* **9**(1), 21–58.

Rendall, M.: 2013, 'Structural Change in Developing Countries: Has it Decreased Gender Inequality?'. *World Development* **45**, 1–16.

Rodrik, D.: 2016, 'Premature deindustrialization'. *Journal of economic growth* **21**(1), 1–33.

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

Schoellman, T. and L. Hendricks: 2023, 'Skilled Labor Productivity and Cross-Country Income Differences'. *American Economic Journal: Macroeconomics*. forthcoming.

Simonovska, I. and M. E. Waugh: 2014, 'The elasticity of trade: Estimates and evidence'. *Journal of international Economics* **92**(1), 34–50.

Singhari, S., S. Madheswaran, et al.: 2016, *The changing Rates of Return to Education in India: Evidence from NSS data*. Institute for Social and Economic Change Bangalore.

SUPPLEMENTARY APPENDIX

APPENDIX A: THEORY

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

A-1 Proof of Proposition 1

To derive expression (7), note that the definition of p_{rnt} implies that

$$\int_n \beta_n \ln p_{rnt} dn = \ln P_{rFt} \int_n \beta_n \lambda_{nF} dn + \ln P_{rGt} \int_n \beta_n \lambda_{nG} dn + \ln w_{rt} \int_n \beta_n \lambda_{nCS} dn - \int_n \beta_n \lambda_{nCS} \ln \mathcal{A}_{rnt} dn.$$

Using the definitions of ω_s and A_{rCSt}, we obtain $\int_n \beta_n \ln p_{rnt} dn = \omega_F \ln P_{rFt} + \omega_G \ln P_{rGt} + \omega_{CS} \ln \left(A_{rCSt}^{-1} w_{rt}\right)$. Similarly, $\int_n \kappa_n \ln p_{rnt} dn = \nu_F \ln P_{rFt} + \nu_G \ln P_{rGt} + \nu_{CS} \ln \left(A_{rCSt}^{-1} w_{rt}\right)$, where ν_s is defined in (8). Substituting these expression in (4), and recalling that $P_{rGt} = \left(A_{rCSt}^{-1} w_{rt}\right)$, yields the expression for $\mathcal{V}^{FE}(e, \mathbf{p}_{rt})$ in (7).

To derive expression (9), note that sector s receives a share λ_{ns} of total revenue of good n. Hence,

$$\vartheta(e, \mathbf{P}_{rt}) = \frac{\int \lambda_{ns} e \vartheta_n^{FE}(e, \mathbf{P}_{rt}) dn}{e} = \omega_s + \nu_s \left(\frac{e}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} \left(A_{rCSt}^{-1} w_{rt}\right)^{\omega_{CS}}} \right)^{-\varepsilon},$$

which is the expression in (9). In Web Appendix WA-1.1, we extend this analysis to the case of a CES production function for final goods.

A-2 Estimation of Parameters and Productivity (Sections 5.1 and 5.2)

In this section we describe in more detail how we estimate the productivity fundamentals $\{A_{rst}\}$ and the structural parameters ω_{CS} and ν_F. Consider a single time period. Given regional data on educational attainment and sector-region data on earnings, we calculate $\{[w_r]_r, H_{rF}, H_{rG}, H_{rCS}\}_r$ in a model-consistent way. Human capital in location r is given by $H_{rt} = L_{rt} \sum_e exp(\rho \times e) \ell_{rt}(e)$, where ρ is the return to education, and $\ell_{rt}(e)$ denotes the share of people in region r with e years of education at time t. Then, labor supply is given by

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

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

Step 1: Estimate ω_{CS} and ν_F. The two structural parameters are jointly identified from aggregate market clearing conditions. The local market clearing Equations (14)–(15), imply the two aggregate resources constraints for tradable goods $s = F, G$:

$$\sum_{r=1}^{R} w_{rt} H_{rst} = \sum_{r=1}^{R} \sum_{j=1}^{R} \pi_{rsjt} \left(\omega_s + \nu_s \left(\frac{A_{jCSt}^{\omega_{CS}} \mathbb{E}_{jt}[q] w_{jt}^{1-\omega_{CS}}}{P_{jFt}^{\omega_F} P_{jGt}^{\omega_G}} \right)^{-\varepsilon} \right) w_{jt} H_{jt} \quad \text{(A-1)}$$

One of the constraints is redundant due to Walras' Law. We can substitute the local market clearing condition for CS (14) into the aggregate resources constraint for agriculture to obtain

$$\sum_{r=1}^{R} w_{rt} H_{rFt} = \omega_F \sum_{r=1}^{R} w_{rt} H_{rt} - \frac{\nu_F}{\nu_{CS}} \sum_{r=1}^{R} \left(\omega_{CS} - \frac{H_{rCSt}}{H_{rt}} \right) w_{rt} H_{rt}. \quad \text{(A-2)}$$

Given data on $\{w_r, H_{rs}\}$, (A-2) yields a single equation in three unknowns: ω_F, $\frac{\nu_F}{\nu_{CS}}$, and ω_{CS}. We externally calibrate ω_F. Also, it is clear from the set of CS market clearing conditions in (14) that ν_{CS} is not separately identified from the average CS productivity *level* A_{*CSt}. As such a level is not interesting for us, it is legitimate to normalize $\nu_{CS} = -1$. Conditional on a choice for ω_F, we can then use (A-2) in 1987 and 2011 to uniquely pin down ω_{CS} and ν_F.

Step 2: Estimate the local price vector $\{p_{rFt}, p_{rGt}, p_{rCSt}\}_r$. Given the structural parameters, there is a unique local price vector that rationalizes all market clearing conditions (14)–(15). We set the average level of the price of goods as the numeraire: $\left(\sum_r (p_{rGt})^{1-\sigma}\right)^{\frac{1}{1-\sigma}} = 1$.

Using the trade shares $\pi_{rsjt} = \tau_{rj}^{1-\sigma} A_{rst}^{\sigma-1} w_{rt}^{1-\sigma} / P_{jst}^{1-\sigma}$, we can write the market clearing condition for tradable goods (15), as

$$w_{rt} H_{rst} = A_{rst}^{\sigma-1} w_{rt}^{1-\sigma} \left(\sum_{j=1}^{R} \tau_{rj}^{1-\sigma} P_{jst}^{\sigma-1} \bar{\vartheta}_{jst} w_{jt} H_{jt} \right), \quad \text{for } s \in \{F, G\}.$$

Rearranging terms yields

$$A_{rst} = w_{rt}^{\frac{\sigma}{\sigma-1}} H_{rst}^{\frac{1}{\sigma-1}} \left(\sum_{j=1}^{R} \tau_{rj}^{1-\sigma} P_{jst}^{\sigma-1} \bar{\vartheta}_{jst} w_{jt} H_{jt} \right)^{\frac{1}{1-\sigma}}, \quad \text{for } s \in \{F, G\},$$

which is equation (21) in the main text.

None of our results depends on the level of food prices in 1987. We pin down the change in aggregate food prices relative to goods prices between 1987–2011 by targeting the published data analogue P_{FGt}^{Data}:

$$\sum_{r=1}^{R} \frac{w_{rt} H_{rt}}{\sum_{j=1}^{R} w_{jt} H_{jt}} \times \frac{P_{rFt}}{P_{rGt}} = P_{FGt}^{\text{Data}}. \tag{A-3}$$

We compute the equilibrium price vector as the fixed point of these conditions.

Step 3: Determine the level of the nominal wage. We proxy income by expenditure. The NSS data on expenditure is reported in rupees. Given the price vector computed in Step 2, we thus set the level of earnings to match a given growth of the real GDP per capita. Since we use final goods as the numeraire, we take real GDP per capita to be denominated in goods.

Step 4: Estimate $\{A_{rst}\}_r$ Given the nominal wage and the local price vector, sectoral productivity is simply given by $A_{rst} = w_{rt}/p_{rst}$.

A-3 The Elasticity of Substitution (Section 5.3)

In this section we derive the elasticity of substitution implied by the theory. For simplicity we suppress the region and time subscripts and denote sectoral prices by P_s. The Allen-Uzawa elasticity of substitution between sectoral output s and k is given by $EOS_{sk} \equiv \frac{\frac{\partial^2 e(P,V)}{\partial P_s \partial P_k} e(P,V)}{\frac{\partial e(P,V)}{\partial P_s} \frac{\partial e(P,V)}{\partial P_k}}$. The expenditure function is given by

$$e(P,V) = \left(V + \sum_s \nu_s \ln P_s \right)^{1/\varepsilon} \varepsilon^{1/\varepsilon} \prod_{s \in \{F,G,CS\}} P_s^{\omega_s}.$$

In the Web Appendix WA-1.2, we prove that

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

A-4 The Equivalent Variation (Section 6)

To measure welfare changes, we calculate equivalent variations (EV) relative to the 2011 *status quo*. Consider the indirect utility of an individual in r with human capital q:

$$\mathcal{V}(q w_r, \mathbf{P_r}) = \frac{1}{\varepsilon} \left(\frac{q w_r}{\prod_s P_{rs}^{\omega_s}} \right)^{\varepsilon} - \sum_s \nu_s \ln P_{rs}. \tag{A-4}$$

We implicitly define the EV for an individual with skills q, $\varpi^q(\hat{x}_r|x_r)$ implicitly by

$$\mathcal{V}\left(qw_r(1+\varpi^q(\hat{x}_r|x_r)),\mathbf{P_r}\right)\equiv\mathcal{V}\left(q\hat{w}_r,\hat{\mathbf{P}}_\mathbf{r}\right), \tag{A-5}$$

where $x_r \equiv (w_r, \mathbf{P_r})$. Hence, ϖ_r^q is the percentage 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.

Using equations (A-4) and (A-5) we can solve for $\varpi^q(\hat{x}_r|x_r)$ as

$$1+\varpi^q(\hat{x}_r|x_r)=\prod_s\left(\frac{\hat{w}_r/\hat{P}_{rs}}{w_r/P_{rs}}\right)^{\omega_s}\times\left(1-\left(\sum_s\nu_s\ln\left(\frac{\hat{P}_{rs}}{P_{rs}}\right)\right)\varepsilon\left(\frac{q\hat{w}_r}{\prod_s\hat{P}_{rs}^{\omega_s}}\right)^{-\varepsilon}\right)^{1/\varepsilon} \tag{A-6}$$

The EV comprises two parts. The first part, $\prod_s\left((\hat{w}_r/\hat{P}_{rs})/(w_rP_{rs})\right)^{\omega_s}$, is akin to the usual change in real wage. This would be the entire EV if preferences were homothetic, i.e., if $\nu_s=0$. The second part captures the unequal effects of productivity growth under nonhomothetic preferences.

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 change of *regional* spending power $\overline{\varpi}_r(\hat{x}_r|x_r)$ the representative agent in district r facing prices P_r would require to attain indifference. As before $\overline{\varpi}_r(\hat{x}_r|x_r)$ is implicitly defined by

$$\mathcal{U}\left(\mathbb{E}_r[q]w_r(1+\overline{\varpi}_r(\hat{x}_r|x_r)),\mathbf{P}_r\right)=\mathcal{U}(\mathbb{E}_r[q]\hat{w}_r,\hat{\mathbf{P}}_r), \tag{A-7}$$

where $\mathcal{U}$ is defined in (13). One can show that $\overline{\varpi}_r(\hat{x}_r|x_r)$ satisfies an expression similar to the one given in (A-6). As a measure of aggregate welfare, we report the average EV using district population as weights:

$$\overline{\varpi}=\sum_r\overline{\varpi}_r\frac{\mathbb{E}_r[q]L_{r2011}}{\sum_r\mathbb{E}_r[q]L_{r2011}}.$$

This is a purely statistical measure that does not rest on an aggregation result.

A-5 Generalizations of Theory (Section 7.4)

In this section we discuss in more extended terms the extensions discussed in Section 7.4. Further technical are available in Web Appendix WA-3.

A-5.1 Open Economy

In this section, we describe the environment and calibration strategy of the open-economy extension. We defer the technical analysis to Web Appendix WA-3.

We assume households, 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%). In terms of exports, according to the World Bank, the export of goods and merchandise increased from 11.3 billion (4.1% of GDP) in 1987 to 302.9 billion (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].

A-5.2 Imperfect Substitution and Skill Bias in Technology

In this section, we describe the environment and calibration strategy of the Imperfect Substitution and Skill Bias in Technology extension. We defer the technical analysis to Web Appendix WA-3.

In this extension, workers with different educational attainments are imperfect substitutes in production. Table WA-2shows that 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. [2022] or Schoellman and Hendricks [2023]). By ignoring such skill-based specialization, our Ricardian model could potentially exaggerate the importance of technology for the development of the service sector.

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:

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

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. 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.[1] 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}. Because sectoral productivity is now determined by two parameters, we set both A_{rs} and Z_{rs} to the respective 1987 level when running counterfactuals.

This extension also allows us to uncover additional facts about the skill bias in technology. 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.

A-5.3 Spatially Mobile Workers

In this section, we describe the environment and calibration strategy of the Spatially Mobile Workers extension. The model is in the vein of economic geography models à la Redding and Rossi-Hansberg [2017], where individuals' migration decisions are modelled as discrete choice problem, with individuals receiving idiosyncratic preference shocks and locations differ in a scalar amenity. Specifically, we assume that individuals make their location choices prior to knowing their particular skill realisation q and draw q from region-specific skill distribution $F_{rt}(q)$. Letting $v_{rt}(q)$ denote the utility of an individual with skills q in region r at time t, we value of settling in location r is given

$$V^i_{rt} = \mathcal{B}_{rt} \int v_{rt}(q) dF_{rt}(q) u^i_{rt}, \tag{A-8}$$

[1] 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.

$\mathcal{B}_{rt}$ is a location amenity, and u_{rt}^i is an idiosyncratic preference shock for location r, which we assume to be Frechet-distributed; $P\left(u_{rt}^i \leq u\right) = e^{-u^{-\eta}}$. The share of people locating in region r at time t is thus given by

$$L_{rt} = \frac{\left(\mathcal{B}_{rt} \int v_{rt}(q) dF_{rt}(q)\right)^{\eta}}{\sum_j \left(\mathcal{B}_{jt} \int v_{jt}(q) dF_{jt}(q)\right)^{\eta}} L. \tag{A-9}$$

In Web Appendix WA-3.3, we formally lay out the model and characterize its equilibrium. In particular, we discuss how we cardinalize consumers' expected consumption utility $\int v_{rt}(q) dF_{rt}(q)$ using the equivalent variation ϖ_{rt} to measure location amenities $\mathcal{B}_{rt}$ and idiosyncratic preferences u_{rt}^i in monetary terms. We also show that all our estimates of both structural parameters and sectoral productivities are exactly the same as in the model with immobile labor, because we can use (A-9) to rationalize the observed population distribution through an appropriate choice of amenities $\mathcal{B}_{rt}$.

To perform counterfactuals, we need an estimate of the spatial labor supply elasticity η, 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. We therefore discipline this elasticity by ensuring that in a counterfactual where we set productivity to their 1987 level in all sectors, the amount of spatial reallocation is as high as what occurred in India between 1987 and 2011. While we think of this choice as an upper bound on the elasticity of spatial supply, we also tested the robustness of our results to higher-elasticity scenarios.

With our calibrated model at hand, we then compute the welfare impact of service-led growth in the presence of spatial mobility in the following way. Using the equilibrium conditions laid out in Proposition 2 together with the spatial labor supply equation (A-9), we can compute equilibrium wages and prices for any change in local productivity. Given these wages and prices we then simulate the optimal migration behavior of 1m individuals, given their initial realization of idiosyncratic preference shocks, u_{rt}^i. The counterfactual welfare change for an individual i that located in region r in 2011 but moved to location j after the counterfactual productivity change is then given by $V_{jCF}^i / V_{r2011}^i - 1$, where V_{rt}^i is given in (A-8). In Table 9 in the main text we report the population weighted average either at the national level of by urbanization quantile. Note that in the absence of mobility individuals from r have a counterfactual utility of V_{rCF}^i, which exactly coincides with our baseline results, given that we cardinalized the location value v_{rt} in monetary terms.

APPENDIX B: DATA AND MEASUREMENT

In this section, we extend the discussion of empirical issues in Sections 2 and 4.

B-1 International Evidence

In Figure 1 we showed that most of service employment in India is concentrated in sectors that serve consumers. Figure B-1 shows that this pattern is not a prerogative of India. India is in line with the international pattern, conditional on its GDP per capita.

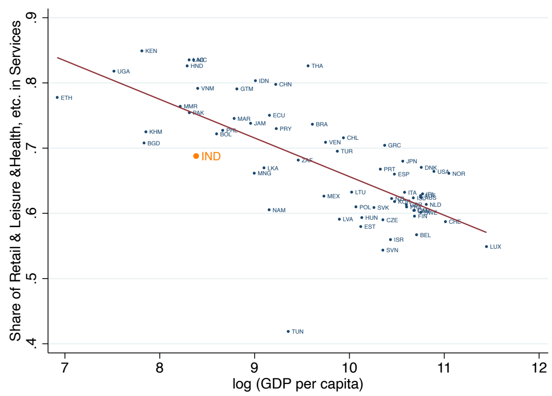

Figure B-1: The Composition of Services and Economic Development. The figure shows a cross-country scatter plot. On the vertical axis, it plots the share of "Retail & Leisure & Health" (the first group of service industries in panel b of Figure 1) in total service employment excluding Education & PA in 2010. On the horizontal axis, it plots the GDP per capita. The data are from the International Labor Office, which uses the ISIC classification.

B-2 Data Sources

In this section we describe the five datasets we use.

B-2.1 National Sample Survey (NSS)

The National Sample Survey (NSS) is a representative survey that has been conducted by the government of India to collect socioeconomic data at the household level since 1950. Each round of the survey consists of several schedules that cover different topics like consumer expenditure, employment and unemployment, participation in education, etc. We focus on the *consumer expenditure module* and the *employment and unemployment module* and use data from rounds 43, 55, 60, 64, 66, and 68 of NSS, which span the years 1987 to 2011. The survey covers the entirety of India except for a few regions due to unfavorable field conditions.[2] In 1987 (2011), our data comprises about 126,000 (101,000) households and 650,000 (455,000) individuals.

We use the "employment and unemployment" module to measure sectoral employment shares and total earnings. An individual is defined as being employed if his/her usual principal activity is one of the following: (i) worked in household enterprises (self-employed); (ii) worked as a helper in household enterprises; (iii) worked as a regular salaried/wage employee; (iv) worked as a casual wage labor in public works; (v) worked as casual wage labour in other types of work. We describe the details of our sectoral employment classification in Section B-4 below.

We proxy household income by total expenditure. More specifically, we measure total household expenditure and divide it by household labor force (number of household members with age above 15 and below 65). We winsorize the expenditure data at 98th percentiles to reduce measurement error.

As we describe in more detail in Section B-2.5, the NSS provides two measures of expenditure. The so-called uniform reference period (URP) measure simply measures total expenditure as expenditure within the last 30 days. The mixed reference period (MRP) measure asks respondents for the total expenditure within the last year for a subset of durable goods to account for the lumpiness of purchases. As a measure of total spending we thus prefer the MRP classification. For the year 2011, the MRP measure is directly contained in the employment module. For the year 1987, the employment module only contains the URP measure. To have a consistent measure in both years, we thus merge the 1987 expenditure module with the 1987 employment module at the household level and compute the MRP measure directly from the data on detailed spending categories. In practice, this choice is inconsequential because the URP measure and MRP measure are highly correlated across space.

We estimate human capital using information on educational attainment and Mincerian returns—see Section 4. In Table B-1, we report the resulting distribution of human capital across time, space and sectors of production. In Web Appendix Table WA-2 we report the same composition when we classify PS and CS workers according to the NIC classification.

B-2.2 Economic Census

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

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

	Less than primary	Primary, upper primary, and middle	Secondary	More than secondary
		Aggregate Economy (1987–2011)		
1987	66.78%	22.03%	7.99%	3.19%
2011	40.32%	30.10%	18.79%	10.79%
		By Sector (2011)		
Agriculture	53.72%	29.23%	14.45%	2.60%
Manufacturing	32.63%	35.31%	20.68%	11.39%
CS	29.85%	32.24%	23.40%	14.51%
PS	28.04%	30.13%	22.03%	19.81%
		By Urbanization (2011)		
Rural	46.97%	30.00%	16.30%	6.84%
Urban	33.69%	30.30%	21.27%	14.73%

Table B-1: EDUCATIONAL ATTAINMENT. The table shows the distribution of the educational attainment over time (Panel A), by sector of employment (Panel B) and across space (Panel C). The breakdown of rural and urban districts is chosen so that approximately half of the population live in rural districts and half live in urban districts.

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

In Table B-2 we report some summary statistics of the EC in various years. In the most recent year, 2013, the EC has information on almost 60 million firms. The majority of them is very small: they employ on average around two employees, and 55% of them have a single employee. The share of firms with more than 100 employees is 0.06%.

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

Table B-2: THE ECONOMIC CENSUS: SUMMARY STATISTICS. The table reports the number of firms, total employment, average employment, and the share of firms with one, less than five, and more than 100 employees.

B-2.3 Service Sector in India: 2006–2007

The Service Sector in India (2006–2007) dataset is part of an integrated survey by the NSSO (National Sample Survey Organisation) in its 63rd round. In the 57th round (2001–2002), the dataset was called Unorganized Service Sector. With the inclusion of the financial sector and large firms, the dataset was renamed as Service Sector in India and is designed to be representative of India's service sector. In Table B-3 we compare this Service Survey with the Economic Census for a variety of subsectors within the service sector. Table B-3 shows that the service survey is consistent with the EC, that is, average firm size and the share of firms with less than five employees are quite comparable in most subsectors.

The Service Survey covers a broad range of service sectors, including hotels and restaurants (Section H of NIC 04); transport, storage and communication (I); financial intermediation (J); real estate, renting and business activities (K); education (M); health and social work (N); and other community, social and personal service activities (O). Excluded are the following subsectors: railways transportation; air transport; pipeline transport; monetary intermediation (central banks, commercial banks, etc); trade unions; government and public sector enterprises; and firms that appeared in the Annual Survey of Industries frame (ASI 2004–2005). In terms of geography, the survey covers the whole of the Indian Union except for four districts and some remote villages.[3] The survey was conducted in a total number of 5,573 villages and 7,698 urban blocks. A total of 190,282 enterprises were ultimately surveyed.

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

NIC2004	Sector	Number of firms		Average employment		Less than 5 employees	
		EC	Service Survey	EC	Service Survey	EC	Service Survey
55	Hotels and restaurants	1499101	30744	2.52	2.49	90%	91%
60	Land transport; transport via pipelines	1317904	41065	1.67	1.24	97%	99%
61	Water transport	7914	174	4.35	1.92	90%	98%
63	Transport activities; travel agencies	188474	2101	3.40	3.33	86%	85%
64	Post and telecommunications	723119	22885	2.06	1.41	96%	99%
65–67	Financial intermediation	293489	16331	5.61	3.81	69%	82%
70	Real estate activities	70128	3648	2.18	1.64	93%	96%
71	Renting of machinery and household goods	365246	5387	2.00	1.77	94%	97%
72	Computer and related activities	66414	1060	6.01	13.45	83%	86%
73	Research and development	2097	5	16.66	4.58	66%	89%
74	Other business activities	519696	10610	2.81	1.92	90%	95%
85	Health and social work	783644	11930	3.39	1.99	88%	95%
91	Activities of membership organizations	1002996	2837	1.82	1.32	94%	98%
92	Recreational, cultural, and sporting activities	222061	2698	2.95	2.91	85%	82%
93	Other service activities	1419685	26132	1.74	1.54	97%	99%

Table B-3: Economic Census and Service Survey. The table reports statistics about the number of firms and their employment from the Economic Census 2005 and Service Survey 2006.

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

B-2.4 INAES 1999–2000

The Informal Non-Agricultural Enterprises Survey (INAES) is part of the 55th survey round of the NSSO. It covers all informal enterprises in the non-agricultural sector of the economy, excluding those engaged in mining, quarrying and electricity, gas and water supply.[4] The survey provides information on operational characteristics, expenses, value added, fixed asset, loans, and factor income. For our analysis we use two pieces of information: the number of employees and whether the main customer is another firm or a household. We use this dataset to allocate employment in the construction sector to either consumer or producer services.

B-2.5 Household Expenditure Survey

The regressions in Table 3 are based on individual expenditure data from the National Sample Survey, Round 68, Schedule 1.0. The dataset contains detailed information on a large set of spending categories. In Table B-4, we report the categories we use in this paper.

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

Table B-4: Broad classification of NSS expenditure survey. The table reports the classification of broad expenditure items in the Expenditure Survey.

We classify the categories 1–17 as food. We also use the spending categories 20 and 24 on services in the pooled regressions of columns 9 and 10 in Table 3. In Web Appendix WA-4.2, we report the more detailed breakdown of consumer services across subcategories.

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

Spending on category c is measured as spending within a particular reference period. For all categories, subjects report total spending during the last 30 days. For durable goods as well as medical and educational spending (i.e., categories 29–34), the subjects additionally report total spending in the last year. This second concept of expenditure aims to account for the lumpiness of purchases. For this group we therefore take 1/12 of annual spending as our measure of monthly expenditure. We measure total spending as the sum of all spending across all categories to calculate the spending share on food and consumer services. In Web Appendix WA-4.2, we report a set of descriptive statistics on the cross-sectional distribution of spending, food shares and CS shares.

In the regressions of Table 3, we control for additional household-level covariates. These include the total size of the household and the number of members of age 15–65. We also control for additional household demographics such as:

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

- the household's religion—Hinduism, Islam, Christianity, Sikhism, Jainism, Buddhism, Zoroastrianism, or other;

- the household's social group—scheduled tribe, scheduled case, backward class, and other;

- whether the household is eligible to receive a rationing card.

B-3 Geography: Harmonizing Regional Borders

In this section, we describe the procedure we used to harmonize the geographical boundaries in order to construct a consistent panel of districts. The borders of numerous Indian districts have changed between 1987 and 2011. The left panel of Figure B-2 plots the districts' boundaries in 2001 and 2011. The purple line represents the boundaries in 2001, and the red line represents the boundaries in 2011.

The most common type of redistricting is a *partition* in which one district has been separated into several districts in subsequent years. The second type is a *border move* in which the shared border between two districts has been changed. The third is a *merge* in which two districts were merged into a single district.

To attain a consistent geography, we take a region to be the smallest area that covers a single district or a set of districts with consistent borders over time. In the case of a partition, the region is constructed as the district in the pre-partition year. In the case of a border move, we construct the union of two districts. The right panel of Figure B-2 shows the official Indian districts in 2011 (dashed red lines) and our

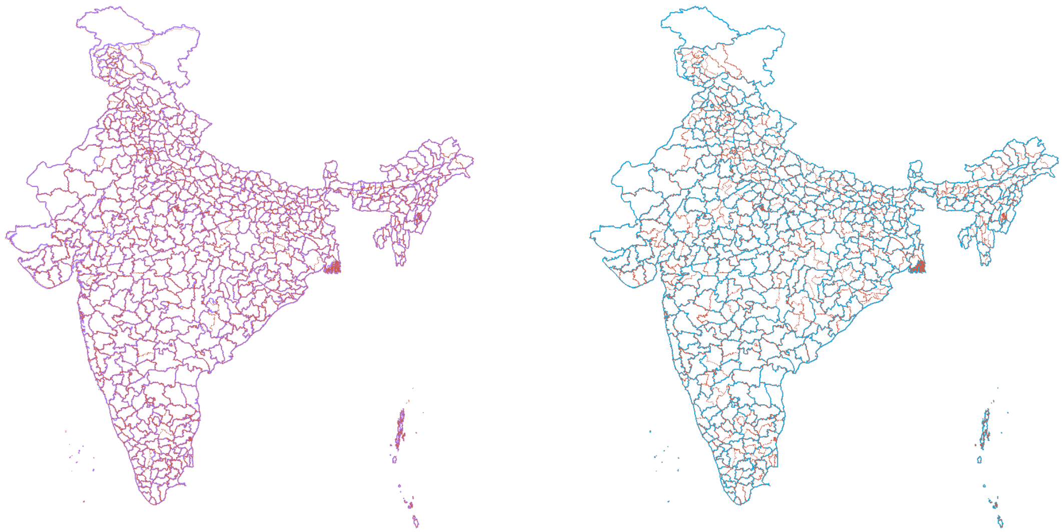

Figure B-2: District Borders in India 1987–2011. The left figure plots the districts' boundaries in 2011 and 2011. The purple line represents the boundaries in 2001 and the dashed red line represents the boundaries in 2011. The right figure shows the official Indian districts in the year 2011 (dashed red lines) and the time-invariant geographical units we construct (solid blue lines) upon which our analysis is based.

modified districts (solid blue lines). We exclude from the analysis two small districts that existed in 2011 but not in 1987. We also exclude districts with less than 50 observations because the small sample would yield imprecise estimates of the sectoral employment shares.

B-4 Classification of Industries

We distinguish four sectors: agriculture, manufacturing, consumer services and producer services. To map these categories to the data, we first construct in Section B-4.1 six broad industries. Then, in Section B-4.2, we assign employment in services and construction to CS and PS, respectively.

B-4.1 Broad Industry Classification

We classify economic activities into six industries: (i) Agriculture, (ii) Manufacturing, (iii) Construction and Utilities, (iv) Services, (v) Information and Communications Technology (ICT) and (vi) Public Administration and Education. The classification relies on the official National Industrial Classification (NIC). Because the NIC system changes over time, we construct a concordance table between 2-digit industries of different versions of the NIC based on official documents and detailed sector descriptions. This concordance system allows us to compare sectoral employment patterns over time. We report the classification in Web Appendix Tables WA-7 and WA-8.

B-4.2 Attributing Employment to CS and PS

We separate CS and PS using the Service Survey (see Section B-2.3), which reports the identity of the main *buyer* of a given firm. We refer to firms that mainly sell to other firms as PS firms and firms that mainly sell to consumers as CS firms.

Ideally, we would calculate the employment share of PS firms in each subsector of the service sectors and in each region. Unfortunately, the sample size of the Service Survey is not sufficiently large to estimate these averages precisely. We therefore generate the regional variation in employment shares by using regional variation in the firm-size distribution and differences in the employment share of PS firms by firm size. Empirically, within each subsector, large firms are much more likely to sell to firms. In Web Appendix Figure WA-5, we plot the employment share of PS firms as a function of firm size in the data. We show in Web Appendix Table WA-9 that the same pattern is present within 2- and 3-digit industries. We operationalize our procedure as follows:

1. For each 2-digit subsector k within the service sector listed in Table WA-7 and size bin b we calculate the employment share of PS firms as

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

   Here, f denotes a firm, $1\left\{f \in PS\right\}$ is an indicator that takes the value 1 if firm f is a PS firm and l_f denotes firm employment. In practice we take three size bins, namely "1 or 2 employees," "3–20 employees," and "more than 20" employees. We weigh observations with the sampling weights provided in the Service Survey.[5]

2. We then use the Economic Census (see Section B-2.2) and calculate the share of employment of firms in size bin b in subsector k in region r as $\ell_{kbr} = \frac{\sum_{f\in(k,b,r)} l_f}{\sum_{f\in(k,r)} l_f}$.

3. We then combine these two objects to calculate the share of employment of PS firms in region r in subsector k as $s_{rk}^{PS} = \sum_b \ell_{kbr}\omega_{kb}^{PS}$.

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

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

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

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

Five service subsectors are not covered by the Service Survey. For firms in publishing and air transport, we assign all employment to PS; for firms in retail trade (except motor vehicle and the repair of personal goods), we assign all employment to CS; and for firms in wholesale trade and firms engaged in the sale and repair of motor vehicles, we use the average PS share from the subsectors for which we have the required information.

B-4.3 Construction and Utilities

We merge employment in construction and utilities with services. To separate CS from PS, we follow a similar strategy as for the service industries. We use the INAES 1999-2000 discussed in Section B-2.4.

From the description of the NIC, some subsectors are clearly for public purposes. We therefore classify 5-digit level industries within the construction sector into public and private and drop all subsectors that we classify as public. These account for roughly 9.2% of total construction employment. See Table WA-10 in Web Appendix WA-4.2 for a detailed classification.

For all subsectors attributed to the private sector, we estimate the CS and PS share based on the information in the INAES. The survey has information on firms in the construction sector and reports the identity of the main buyer of the firm. In particular, we observe in the data whether the firm sells to (i) the government, (ii) a cooperative or marketing society, (iii) a private enterprise, (iv) a contractor or intermediary, (v) a private individual, or (vi) others. We associate all firms that answer (ii), (iii), or (iv) with PS firms and all firms that answer (v) with CS firms. We then calculate the PS share of a given private subsector as total PS employment relative to total CS and PS employment in the respective subsector, that is, for subsector k we calculate the PS share as $\omega_k^{PS} = \frac{\sum_{f \in k} 1\{f \in PS\} l_f}{\sum_{f \in k} l_f}$, where l_f denotes firm employment, and $1\{f \in PS\}$ is an indicator for whether firm f is a PS firm.

In Table B-5, we report the relative employment shares of public employment (as classified in Table WA-10), CS, and PS in the construction sector as a whole. The share of public employment is around 10%. Among the private subsectors, 12.9% of employment is associated with the provision of producer services.

B-5 Urbanization and Spatial Structural Change

In Figure B-3, we show the structural transformation in India across time and space. We focus on urbanization as our measure of spatial heterogeneity.[6] This is a mere

[6] The urbanization rate is the share of the population living in urban areas according to the definition of the NSS. The NSS defines an urban location in the following way: (i) all locations with a municipality, corporation or cantonment and locations defined as a town area, (ii) all other locations that satisfy the following criteria: (a) a minimum population of 5,000, (b) at least 75% of the male population

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

Table B-5: Composition of the construction sector. The table shows the relative employment shares of PS, CS, and public employment in the construction sector in different years. We associate public employment to sectors classified as "public" in Table WA-10. The classification of employment in the private subsectors to CS and PS is explained in the main text. The last row reports the relative employment share of PS within the private subsectors.

descriptive device, because there is a a strong positive correlation between urbanization and expenditure per capita in the NSS data in 2011. Figure B-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 larger when one considers earnings instead of employment, because earnings are higher in service industries and in cities.

Panel a: Sectoral Empl. by Urbanization (1987) Panel b: Sectoral Empl. by Urbanization (2011)

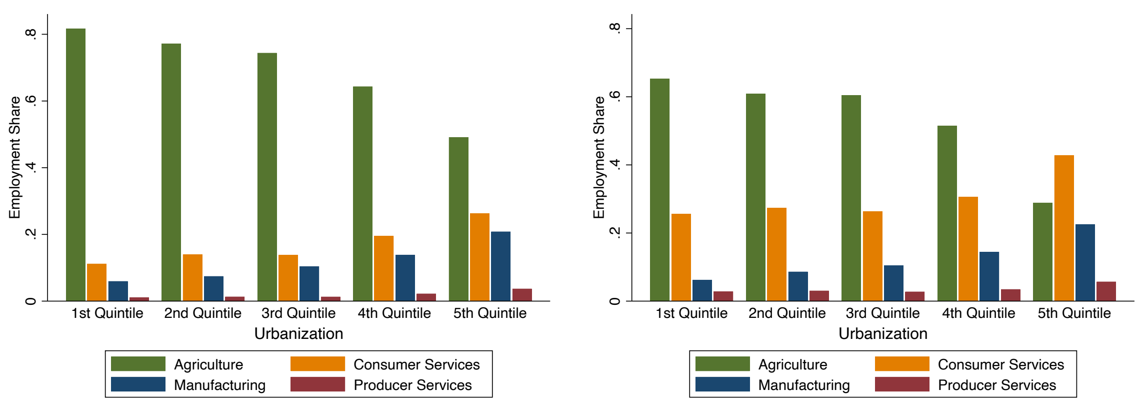

Figure B-3: Sectoral Employment over Time and Space. The figure plots the sectoral employment shares by urbanization quintile in 1987 and 2011.

is employed outside of agriculture, and (c) a density of population of at least 1,000 per square mile.

APPENDIX C: ESTIMATION

In this section, we discuss details of the estimation.

C-1 Estimating the Engel elasticity ε

C-1.1 Nonlinear estimation

In Section 5.1 we estimate the Engel elasticity ε under the assumption that the asymptotic expenditure on food is small. This allowed us to estimate ε from log-linear regression of food shares and total expenditure. In this section we estimate the ε without this assumption and focus directly on the non-linear expression for food expenditure shares given in equation (16).

Equation (16) implies that the log food share satisfies the equation

$$\ln\left(\vartheta_{\mathcal{F}}^{FE}\left(e,\mathbf{p_r}\right)-\beta_{\mathcal{F}}\right)=\ln\left(\kappa_{\mathcal{F}}\exp\left(\int_n\beta_n\ln p_{rn}dn\right)^{-\varepsilon}\right)-\varepsilon\ln e. \tag{C-1}$$

We can thus consider the empirical regression

$$\ln\left(\vartheta_{\mathcal{F}}^{h}-\beta_{\mathcal{F}}\right)=\delta_r+\varepsilon\times\ln e_h+x_h'\psi+u_{rh}, \tag{C-2}$$

where $\vartheta_{\mathcal{F}}^{h}$ denotes the food share of household h living in region r, e_h denotes total household spending, δ_r is a region fixed effects, and x_h is a set of household characteristics. We now use (C-2) to estimate both $\beta_{\mathcal{F}}$ and ε without restricting $\beta_{\mathcal{F}}=0$. We stress that we do not use the estimate of $\beta_{\mathcal{F}}$ in our analysis. $\beta_{\mathcal{F}}$ is the *final good* expenditure share on food which is part of the final consumption vector, while our structural estimation relies on preference parameters of the value added demand system. Hence, the value of $\beta_{\mathcal{F}}$ only matters insofar as it affects the estimate of ε. Also, focusing on the transformed dependent variable $\ln\left(\vartheta_{\mathcal{F}}^{h}-\beta_{\mathcal{F}}\right)$ is computationally convenient, because we can estimate (C-2) as a linear regression. This is makes it easy to control for the regional fixed effects δ_r.

In Table C-1, we report the results. We focus on the specification with household controls of column 2 (for the OLS) and column 6 (for the IV) of Table 3 in the main text. The table shows the estimates of ε and the associated R^2 for different choices of $\beta_{\mathcal{F}}$. In panel A, we report the OLS estimates; in panel B, we report the IV estimates. The first column is the case of $\beta_{\mathcal{F}}=0$, which is our baseline estimate.

Two results emerge. First, the estimate of ε is not sensitive to $\beta_{\mathcal{F}}$ in a range where the asymptotic expenditure on food item does not exceed 6% (the expenditure share on food items in the US is 5%). Second, a comparison of the R^2 shows that the specification with $\beta_{\mathcal{F}}=0$ delivers the best fit to the data, even though the difference across columns is small.

	Dependent variable: ln(food expenditure share)						
β_F	0	0.01	0.02	0.03	0.04	0.05	0.06
			Panel A: OLS estimates				
$\ln e$	-0.319***	-0.327***	-0.336***	-0.345***	-0.355***	-0.366***	-0.378***
	(0.007)	(0.008)	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)
N	91474	91474	91474	91474	91474	91474	91474
R^2	0.4283	0.4278	0.4273	0.4266	0.4258	0.4247	0.4233
			Panel B: IV estimates				
$\ln e$	-0.395***	-0.405***	-0.416***	-0.427***	-0.439***	-0.452***	-0.466***
	(0.013)	(0.014)	(0.014)	(0.014)	(0.015)	(0.015)	(0.016)
N	85916	85916	85916	85916	85916	85916	85916
R^2	0.3099	0.3097	0.3095	0.3093	0.3089	0.3084	0.3076

Table C-1: Income elasticity for food: Non-linear estimation. The table shows the estimated coefficient ε of the regression (C-2) for different choices of $\beta_{\mathcal{F}}$. All variables are defined as in Table 3. For all regressions we trim the top and bottom 5% of the income distribution and we control for region fixed effects, a (within-district) urban/rural dummy, a set of fixed effects for household size, and the number of workers within the household. In panel A we report the OLS estimates. In panel B we report the IV estimates. Standard errors are clustered at the district level. In all specifications we consider a balanced sample excluding individuals whose food expenditure is below 6%. The results in the unbalanced sample including all individuals are almost identical.

C-1.2 Consumer Service Expenditure Regression

In columns 9 and 10 of Table 3 we pool data on food shares and data on service expenditure shares. To measure service expenditures, we follow the official classification of the NSS expenditure module. As seen in Tables WA-3 and WA-4), these expenditures include, for example, domestic servants, barber shops, or tailor services. We also add entertainment expenses such as movie theaters or club fees.

In the left panel of Figure C-1, we plot the cross-sectional distribution of service expenditure shares in our data. The figure shows that the variation is sizable and most consumers India spend between 0 and 15% of their income on consumer services. The 99% quantile of the distribution, shown as the solid line, is 0.2.

It is useful to recall that, since CS spending is a luxury, our theory implies that $\kappa_{\mathcal{S}} < 0$ and that the asymptotic expenditure share $\beta_{\mathcal{S}}$ exceeds the observed spending share $\vartheta^h_{\mathcal{S}rt}$ for all households. Equation (16) thus implies that

$$\ln\left(\beta_{\mathcal{S}} - \vartheta^{FE}_{\mathcal{S}}\left(e, \mathbf{p}_r\right)\right) = \ln\kappa_{\mathcal{S}} + \varepsilon \ln\left(\exp\left(\int_n \beta_n \ln p_{rn} dn\right)\right) - \varepsilon \ln e. \qquad \text{(C-3)}$$

Hence, the relationship between $\vartheta^{FE}_{\mathcal{S}}\left(e, \mathbf{p}_r\right)$ and total expenditure e is positive; the relationship between $\ln\left(\beta_{\mathcal{S}} - \vartheta^{FE}_{\mathcal{S}}\left(e, \mathbf{p}_r\right)\right)$ and $\ln e$ is negative and in fact log-linear with a slope coefficients of ε.

To identify ε from a regression based on (C-3), we need to estimate $\beta_{\mathcal{S}}$. Because $\beta_{\mathcal{S}}$ is the asymptotic expenditure share, we take it to be the 99% quantile of the

Panel a: Distribution of CS Spending Shares.

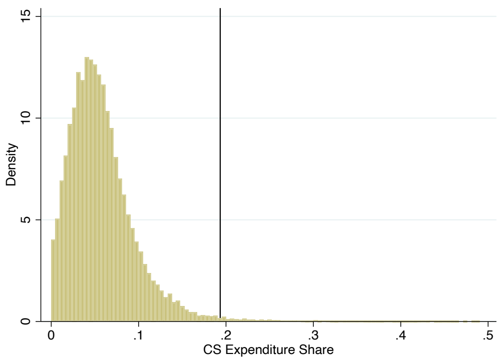

Panel b: CS Spending Share and Income.

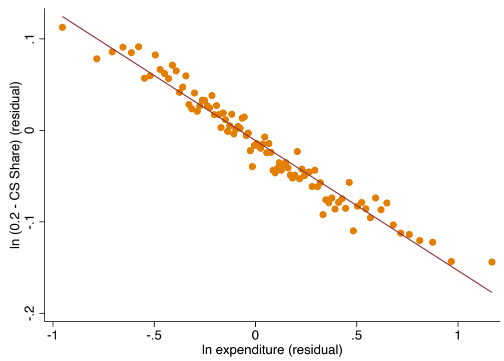

Figure C-1: Consumer Service Spending. In the left panel we display the cross-sectional distribution of spending share on services. In the right panel we display a binscatter plot of the relationship between (the log of) total expenditure and (the log of) the differences between the actual expenditure share on consumer services and and the asymptotic expenditure share 0.2, that is $\ln\left(0.2 - \vartheta^h_{rCSt}\right)$.

expenditure share distribution in India, which turns out to be 0.2. This value is shown as the solid line in the left panel of Figure C-1. Given this value for $\beta_{\mathcal{S}}$, we estimate ε from the same regression as in our baseline analysis contained in the main text, that is

$$\ln\left(\beta_{\mathcal{S}} - \vartheta^h_{\mathcal{S}}\right) = \delta_r + \varepsilon \times \ln e_h + x'_h\psi + u_{rh}, \tag{C-4}$$

where the region fixed effect δ_r absorbs the constant $\kappa_{\mathcal{S}}$ and the vector of regional prices.

The results are reported in Table C-2. The first two columns contain different specifications of estimating (C-4) via OLS. The implied elasticity is negative but smaller than what we estimate for the specification based on food expenditure. In the last two columns, we report the IV specification, where - as in the baseline - we instrument total expenditure e with full set occupation fixed effects. Doing so increases the the elasticity substantially and we now estimate a value of around 0.3, which is still slightly lower but in the same ballpark of the IV estimate based on food expenditure.

Finally, in the right panel of Figure C-1 we display the relationship between (the log of) household expenditure and the adjusted expenditure share graphically. While the relationship shows more noise relative to the specification based on the food expenditure shown in Figure 3, it is again approximately linear.

C-2 Estimating the Shape of the Human Capital Distribution (ζ)

We estimate the tail parameter of the distribution of efficiency units ζ from the distribution of income. Our model implies that total income and expenditure of individual h is given by $e^h_{rt} = q^h w_{rt}$, where q follows a Pareto distribution $f_{rt}(q) = \zeta \underline{q}^{\zeta}_{rt} q^{-(\zeta+1)}$.

	Dependent variable: ln(0.2 - CS Exp Share)			
	(1)	(2)	(3)	(4)
$\ln e$	-0.115***	-0.097***	-0.263***	-0.328***
	(0.010)	(0.010)	(0.023)	(0.039)
Trim (top & bottom 5%)	✓	✓	✓	✓
Addtl. Controls		✓		✓
IV			✓	✓
N	90672	90625	85312	85269
R^2	0.132	0.138	0.027	0.003

Table C-2: Income elasticity for Cons. Serv. Standard errors, clustered at the district level in parentheses. *** p<0.01, ** p<0.05, * p<0.1

This implies that

$$\ln\left(f_{rt}(q)\right) = \ln(\zeta \underline{q}_{rt}^{\zeta}) - (\zeta + 1)\ln(q). \tag{C-5}$$

We estimate ζ from a regression of the (log of the) upper tail density on log efficiency units that we calculate as $q^h_{rt} = \frac{e^h_{rt}}{w_{rt}}$. In Table C-3, we report the estimated ζ based on (C-5). We report both the estimate based on the full sample (column 1) and the estimates by urbanization quintile (columns 2–6). We also report our estimates based on two measures of earnings: total expenditures per capita (as in our main analysis) and total income, which is also reported in the NSS data.

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

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

Table C-3: Identification of ζ. The table reports the estimate of ζ based on (C-5). In the first columns we report the estimates for the years 1987 and 2011. In the remaining columns we perform our estimation separately for different quantiles of the urbanization distribution.

C-3 The Relative Price of Agricultural Goods

Our estimation uses the relative price of agricultural goods (relative to manufacturing goods) to identify the relative productivity in the agricultural sector (relative to manufacturing). The Ministry of Planning and Program Implementation (MOSPI) of the Government of India reports value added by 2-digit sectors at current prices and constant prices from 1950–2013[7] We then construct the sectoral price index as the ration between sectoral value added in current prices relative to constant prices. We normalize both price indexes in the year 2005 to unity. We then calculate the relative price of agricultural products as $p_t^{AM} = p_t^A / p_t^M$. To check the validity of our results, we also use two additional data sources to calculate the relative price. The first is the GGDC 10-Sector Database[8], which provides long-run data on sectoral productivity performance in Africa, Asia, and Latin America. This dataset reports the annual series of value added at current national prices and value added at constant 2005 national prices. We follow the same procedures to calculate the relative price.

The second is the Wholesale Price Index (WPI) from the Office of the Economic Advisor.[9] The WPI tracks ex-factory prices for manufactured products and market prices for agricultural commodities.[10] Again, we use the same method to calculate the relative prices, and normalize the relative price in the year 2005 to 1.

In Figure C-2 we plot the relative price of agricultural goods to manufacturing goods. Since the pattern from the different data sources is very similar, we use the results based on GGDC in our analysis.

C-4 Estimates of CS Productivity Growth

In Section 5.2 we showed: (i) CS productivity is systematically higher in urbanized locations (see Figure 4), and (ii) productivity growth is spatially dispersed (see Table 5). In this section we provide more details on the correlates of our estimates of CS productivity growth and how they depend on the demand system we use.

Consider first Table C-4, where we regress sectoral productivity growth in region r, that is, $\ln A_{rs2011} - \ln A_{rs1987}$, on the 1987 urbanization rate in region r. Urban locations experienced higher productivity growth, especially in CS and the Industrial Sector (which, recall, includes some business services). Recall that the information

[7] Data are available at `http://www.mospi.gov.in/data`. See "Summary of macro economic aggregates at current prices, 1950–51 to 2013–14" and "Summary of macro economic aggregates at constant(2004–05) prices, 1950–51 to 2013–14."

[8] The data are available at `https://www.rug.nl/ggdc/productivity/10-sector`

[9] The data are available at `https://eaindustry.nic.in/`

[10] One issue with this is that the base year (and the basket of goods) changes during different time periods. Two series are relevant to our research. The first one is the series with the base year 1993, which is available from 1994 though 2009. The second one is the series with the base year 2004, which is available from 2005 though 2016.

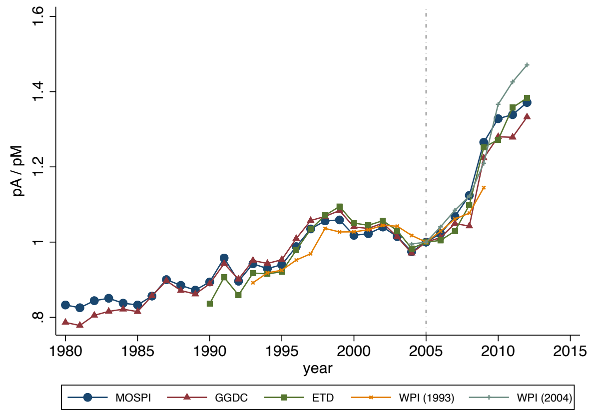

Figure C-2: RELATIVE PRICE OF AGRICULTURAL TO MANUFACTURING GOODS. The figure shows the relative prices of agricultural products from the different sources mentioned in the main text. "MOSPI" refers to the data from the Indian Government that is used in our analysis. "GGDC" stems from the GGDC 10-Sector Database. "ETD" is the new revised version of GGDC database. "WPI (1993)" and "WPI (2004)" are based on the Wholesale Price Index with a 1993 base year and a 2004 base year respectively.

on urbanization is not used on our estimation. Hence, cities not only have higher CS productivity in levels but also experiences faster growth.[11]

In Figure C-3 we show the extent to which our productivity estimates depends on our estimated demand system. Specifically, we depict the distribution of CS productivity *growth*, $\ln A_{rCS2011} - \ln A_{rCS1987}$, as a function of the Engel elasticity ε. We consider five value of this elasticity that span the range of estimated based on our results in Table 3: our baseline estimate (0.395, column 6), the estimate for high-income households (0.415, column 7), the estimate for urban locations (0.358, column 8), the OLS estimate (0.321, column 2), and the estimate based on food and service expenditure (0.23, column 9), which is the smallest estimate in our analysis. Figure C-3 shows that the estimated distribution of growth rates is quite stable. For the smallest ε of 0.23, the dispersion is slightly larger, reflecting the fact that local employment shares depend on $A_{ArCSt}^{\omega_{CS}\varepsilon}$ (see (20)). Because the importance of service-led growth is decreasing in ε, we focus our robustness analysis on the range where $\varepsilon > 0.3$.

C-5 Non-targeted Moments: Additional Results

As we mention in the main text, we can use the data from the expenditure survey to validate our estimates agricultural productivity and hence food prices. The expenditure survey reports both total expenditure and the total quantity bought for a variety of food items. We thus compute the price of product n in region r, p_{nr}, as the ratio between total expenditure and total quantity and then run the regression

$$\ln p_{nr} = \delta_r + \delta_n + u_{nr}, \tag{C-6}$$

[11] We also ran the regressions in Table C-4 based on the 2011 urbanization rate. The positive correlation between productivity growth and urbanization is if anything stronger.

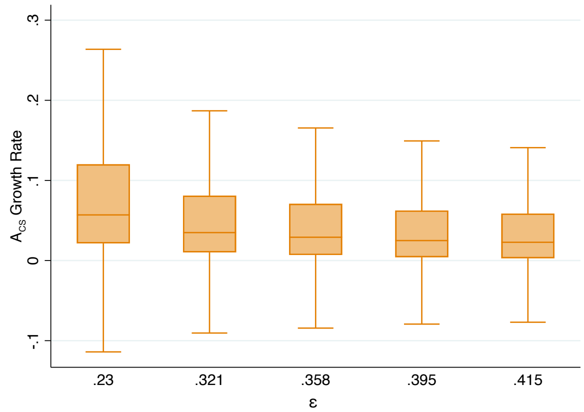

Figure C-3: CS Productivity Growth and the Engel elasticity ε. The figure shows the cross-sectional distribution CS productivity growth rate, $\frac{\ln A_{rCS2011} - \ln A_{rCS1987}}{2011-1987}$, as a function of ε. We always display a boxplot that indicates the median, the interquartile range and the upper and lower adjacent values.

	Productivity Growth		
	Agriculture	Industry	Cons. Serv.
1987 urbanization	0.239**	0.437***	2.529***
	(0.079)	(0.086)	(0.457)
Weight (1987 Pop)	✓	✓	✓
N	360	360	360
R^2	0.025	0.067	0.079

Table C-4: Productivity Growth and Urbanization. The table reports the results of univariate regressions of sectoral productivity growth, $\ln(\frac{A_{rs2011}}{A_{rs1987}})$, on the urbanization rate in 1987. We weigh all regressions by the population size in 1987.

where δ_r and δ_n are region and product fixed effects. The estimated fixed effect $\hat{\delta}_r$ thus describes the average food price in region r.

In Figure C-4 we show the correlation between the estimated $\hat{\delta}_r$ and the regional price of agricultural goods in the model, that is $\ln p_{rFt}$. 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).

C-6 Outliers in Quantitative Analysis

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

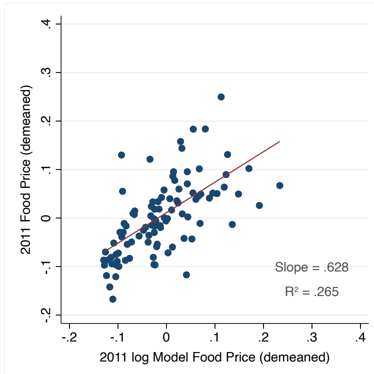

Figure C-4: Food Prices: Model vs Data. The figure shows a binscatter plot of regional log food prices in the data ($\hat{\delta}_r$ from (C-6)) and the model ($\ln p_{rF}$)

In Table C-5, we report the upper and lower quantiles of the regional distribution of welfare changes for the different counterfactuals. Consider for example the agricultural sector. If agricultural productivity had not grown since 1987, the most adversely affected region would have seen its welfare decline by 56% in terms of an equivalent variation. Conversely, some regions would have seen their welfare increase. The last row of Table C-5 shows that some regions would have seen very large gains if CS productivity had not grown. These are regions where CS productivity *declined* between 1978 and 2011. As explained above, this pattern is entirely driven by a few districts being close to the theoretical threshold of ω_{CS}. For comparison, in the last row we report the estimated distribution of the welfare effects in our baseline analysis, where we truncate the productivity growth distribution at the top and bottom 3%. This has large effects on the welfare effects in the right tail of the distribution.

	Regional Welfare Changes (%)									
	Min	1%	2%	3%	5%	95%	97%	98%	99%	Max
Agriculture	-56.0	-45.3	-43.2	-42.1	-39.7	3.7	7.3	13.6	17.7	48.1
Industry	-34.0	-28.4	-27.0	-25.8	-24.2	-5.9	-2.5	-2.3	-0.9	28.3
Cons. Serv.	-98.9	-96.2	-89.9	-86.5	-77.3	19.1	44.8	136.6	317.6	1491.7
Cons. Serv. (Baseline)	-93.8	-93.0	-88.0	-85.9	-77.0	19.0	36.6	40.2	70.5	91.1

Table C-5: Distribution of Welfare Losses. The table reports the lower and upper percentiles of the regional distributions of sectoral welfare losses.

These extreme values at the bottom of the regional productivity growth distribution have aggregate effects. For our baseline analysis we trim the top and bottom 3% of the productivity growth distribution and set regional productivity growth in such regions to the 3% and 97% quantile respectively. In Table C-6 we report the change in aggregate in the absence of CS productivity growth as a function of this trimming cutoff. Without any trimming, the aggregate effect is -18.0%. Once such outliers are truncated, we recover our baseline results of a welfare loss of about -20.5%. In the last row of Table C-6 we report the aggregate employment share of the affected districts. The changes in the aggregate effects of CS growth are not driven by few large districts but by a small number of small districts with very large changes in CS productivity.

	Trimming Cutoff					
	No Trimming	1%	2%	3%	4%	5%
Wefare Loss	-18.0%	-19.4%	-20.1%	-20.5%	-20.8%	-20.9%
Employment Share	0	0.5%	1.9%	3.2%	5.4%	8.0%

Table C-6: Welfare Losses with Different Trimming Cutoffs. The table reports the aggregate welfare effects of productivity growth in the CS sector for different trimming rules. A trimming cutoff $x\%$ means that we set the $x\%$ highest and lowest productivity growth rates to $1 - x\%$ and $x\%$ respectively.

C-7 Details of Robustness Analysis (Section 7)

In Figure C-5 we report the results of our analysis discussed in Section 7, where we allow for heterogeneity in the Engel elasticity ε. In the left panel of Figure C-5, we assume our baseline estimate of $\varepsilon = 0.395$ in Bangalore and $\varepsilon = 0.29$ in rural Bankura as suggested by column 8 of Table 3. Doing so yields a mild reduction in spatial inequality but the quantitative effect is small.

In the right panel, we allow for heterogeneous ε across the income ladder. In particular, we estimate productivity growth in CS based on the benchmark Engel elasticity of 0.395. Then, we consider (a zero measure of) households with income above and below median with elasticities of 0.415 and 0.218, respectively, corresponding to the estimates of column 7 in Table 3. The right panel of Figure C-5 highlights that this *amplifies* the differential welfare impact of service-led growth between rich and poor households. The reason is intuitive: rich agents consume more and care more about the provision of CS. This suggests that a model with increasing Engel elasticities by income is likely to deliver even more unequal welfare effects of service-led growth.

In the main text, we focused on the robustness of our results with respect to the Engel elasticity. Here we report our results for ω_F and ζ. We always recalibrate the entire model, when changing one of the parameters.

We summarize our results in Figure C-6, where we plot the implied impact of sectoral productivity growth as a function of the respective parameters. In the left panel of we report for completeness the effect of ε. As discussed in the main text, for the impact of service-led growth to become small, one would need to believe in an estimate of the Engel elasticity, which is much larger then suggested by both the micro data on Engel curves and the macro data on productivity growth.

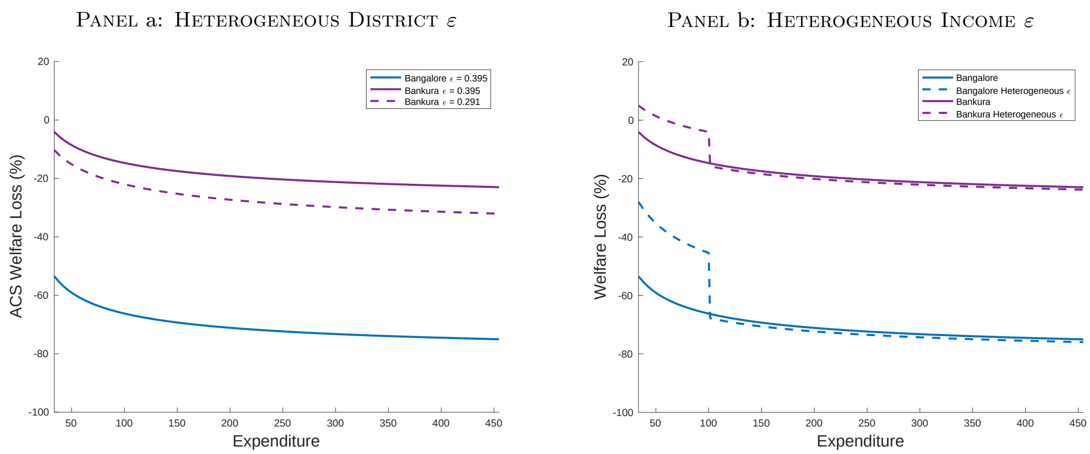

Figure C-5: Heterogeneous Engel Elasticities. In the left panel we allow for heterogeneous ε across locations. We assume that ε of individuals in Bangalore (Bankura) is 0.395 (0.291), which is in line with the results reported in Table 3. In the right panel we allow for different ε across individuals. In line with Table 3, we assume that individuals above (below) the median income have ε of 0.415 (0.218).

In the middle panel we focus on ω_F, which we calibrate to 1% so as to match the value added 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 ω_F. Panel (b) of Figure C-6 shows that the implied welfare impact of productivity growth in the CS sector is, if anything, slightly larger the higher ω_F. Our choice of $\omega_F = 0.01$ is therefore conservative.

Finally, in panel (c) of Figure C-6 we show the effect of the tail of the skill distribution ζ. Note that this only changes the mapping from the "aggregate" demand parameter $\bar{\nu}_s$ to the micro parameter ν_s. All our productivity estimates are independent of ζ. Figure C-6 shows that the higher ζ, the higher the importance of CS growth relative to agricultural productivity. This reflects the importance of nonhomothetic demand. The smaller ζ, the higher 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 C-6 shows this intuition is borne out but that the effects are quantitatively moderate.

We also analyzed the effect of the skill return ρ. Our estimate of 5.6% is on the lower end of typical Mincerian regressions. 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. Our results are essentially insensitive to this parameter. Similarly, our results are virtually unchanged for different values of the elasticity of substitution σ.

In Table C-7 we report the analogue to Table 9, that is the welfare effects of agricultural and industrial productivity growth. Table C-7 shows that our baseline results are not significantly affected by either the alternative modelling assumptions or the alternative measurement choices.

Panel a: Engel Elasticity ε — Panel b: Long-Run Food Share ω_F — Panel c: Skill distribution ζ

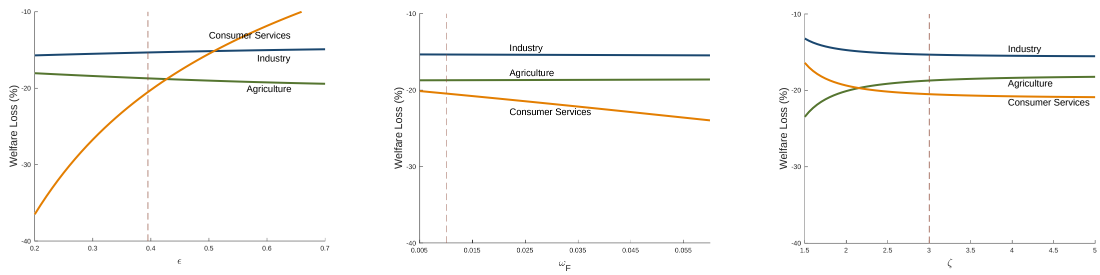

Figure C-6: Robustness Analysis. Panels (a), (b), and (c) show the aggregate welfare effects as a function of the preference parameters ε, ω_F, and the tail parameter of the skill distribution ζ. The vertical dashed line corresponds to the parameter value in our benchmark analysis.

	Agriculture					**Industry**				
	Aggregate Effects	**by Urbanization Quintiles**		**by Income Quantiles**		**Aggregate Effects**	**by Urbanization Quintiles**		**by Income Quantiles**	
		1st	**5th**	**10th**	**90th**		**1st**	**5th**	**10th**	**90th**
Baseline	**-18.7**	**-19.6**	**-15.3**	**-21.8**	**-15.1**	**-15.3**	**-11.8**	**-20.8**	**-12.5**	**-20.7**
						Alternative calibrations of ε (Section 7.2)				
$\varepsilon = 0.415$ **(High Income Households)**	**-18.8**	**-19.7**	**-15.3**	**-22.0**	**-15.1**	**-15.3**	**-11.8**	**-20.8**	**-12.4**	**-20.7**
$\varepsilon = 0.321$ **(OLS estimator)**	**-18.5**	**-19.4**	**-15.1**	**-21.3**	**-15.2**	**-15.5**	**-12.0**	**-20.8**	**-12.8**	**-20.7**
						Alternative measurement choices (Section 7.2)				
Allocate PS share based on WIOD	**-18.5**	**-19.4**	**-15.6**	**-21.4**	**-15.5**	**-17.1**	**-13.0**	**-23.6**	**-13.7**	**-23.6**
Allocate ICT & Business to PS	**-18.8**	**-19.8**	**-16.0**	**-21.7**	**-15.8**	**-16.3**	**-12.2**	**-23.0**	**-12.7**	**-22.9**
Allocate Construction to Industry	**-18.8**	**-21.6**	**-12.6**	**-23.0**	**-13.8**	**-19.4**	**-11.9**	**-30.9**	**-13.3**	**-29.5**
						Alternative modeling assumptions (Section 7.4)				
Open economy	**-18.8**	**-19.6**	**-15.5**	**-21.8**	**-15.5**	**-17.6**	**-14.2**	**-22.8**	**-14.9**	**-22.5**
Imperfect skill substitution	**-23.4**	**-25.8**	**-18.4**	**-25.1**	**-20.2**	**-14.7**	**-10.5**	**-21.3**	**-9.9**	**-22.6**
Spatial labor mobility	**-18.2**	**-18.9**	**-15.2**			**-15.3**	**-12.0**	**-20.2**		

Table C-7: The importance of service-led growth—Robustness.

Web Appendix for "Service-Led or Service-Biased Growth? Equilibrium Development Accounting Across Indian Districts."

by Tianyu Fan, Michael Peters, and Fabrizio Zilibotti

November 13, 2022

- *Not for Publication Unless Requested* -

WA-1 Additional theoretical results

WA-1.1 CES Production Function for Final Goods

In this section we generalize the results of Section A-1 in the Appendix to the case, where the production of final goods combines tradable goods and local CS in a CES way. Specifically, suppose that

$$y_n = \left(\lambda_{nF} x_F^{\frac{\varsigma-1}{\varsigma}} + \lambda_{nG} x_G^{\frac{\varsigma-1}{\varsigma}} + \lambda_{nCS} \left(\mathcal{A}_{rnt} H_{nCS}\right)^{\frac{\varsigma-1}{\varsigma}}\right)^{\frac{\varsigma}{\varsigma-1}}, \tag{WA-1}$$

where the parameters λ_{ns} are sectoral weights, which are specific to good n. The good-specific price index is then given by

$$p_{rnt} = \left(\lambda_{nF}^{\varsigma} P_{rFt}^{1-\varsigma} + \lambda_{nG}^{\varsigma} P_{rGt}^{1-\varsigma} + \lambda_{nCS}^{\varsigma} \left(\mathcal{A}_{rnt}^{-1} w_{rt}\right)^{1-\varsigma}\right)^{\frac{1}{1-\varsigma}}.$$

Similarly, the cost shares of food, industrial goods, and CS for final good n are given by

$$\chi_{rnt}^{F} = \lambda_{nF}^{\varsigma} \left(\frac{P_{rFt}}{p_{rnt}}\right)^{1-\varsigma} \quad \text{and} \quad \chi_{rnt}^{G} = \lambda_{nF}^{\varsigma} \left(\frac{P_{rGt}}{p_{rnt}}\right)^{1-\varsigma} \quad \text{and} \quad \chi_{rnt}^{CS} = \lambda_{nF}^{\varsigma} \left(\frac{\mathcal{A}_{rnt}^{-1} w_{rt}}{p_{rnt}}\right)^{1-\varsigma}. \tag{WA-2}$$

This implies that

$$\int_n \kappa_n \ln p_{rnt} dn = \int_n \ln \left(\lambda_{nF}^{\varsigma} P_{rFt}^{1-\varsigma} + \lambda_{nG}^{\varsigma} P_{rGt}^{1-\varsigma} + \lambda_{nCS}^{\varsigma} \left(\mathcal{A}_{rnt}^{-1} w_{rt}\right)^{1-\varsigma}\right)^{\frac{\kappa_n}{1-\varsigma}} dn$$

and

$$\exp\left(\int_n \beta_n \ln p_{rnt} dn\right) = \exp\left(\int_n \ln \left(\lambda_{nF}^{\varsigma} P_{rFt}^{1-\varsigma} + \lambda_{nG}^{\varsigma} P_{rGt}^{1-\varsigma} + \lambda_{nCS}^{\varsigma} \left(\mathcal{A}_{rnt}^{-1} w_{rt}\right)^{1-\varsigma}\right)^{\frac{\beta_n}{1-\varsigma}} dn\right).$$

The indirect utility function (in terms of sectoral value added) can thus be written as

$$V\left(e,\mathbf{P}_{rt}\right)=\frac{1}{\varepsilon}\left(\frac{e}{B\left(\mathbf{P}_{rt}\right)}\right)^{\varepsilon}-D\left(\mathbf{P}_{rt}\right),$$

where

$$\begin{aligned} B\left(\mathbf{P}_{rt}\right) &= \exp\left(\int_n \ln\left(\lambda_{nF}^{\varsigma}P_{rFt}^{1-\varsigma}+\lambda_{nG}^{\varsigma}P_{rGt}^{1-\varsigma}+\lambda_{nCS}^{\varsigma}\left(\mathcal{A}_{rnt}^{-1}w_{rt}\right)^{1-\varsigma}\right)^{\frac{\beta_n}{1-\varsigma}}dn\right)\\ D\left(\mathbf{P}_{rt}\right) &= \int_n \ln\left(\lambda_{nF}^{\varsigma}P_{rFt}^{1-\varsigma}+\lambda_{nG}^{\varsigma}P_{rGt}^{1-\varsigma}+\lambda_{nCS}^{\varsigma}\left(\mathcal{A}_{rnt}^{-1}w_{rt}\right)^{1-\varsigma}\right)^{\frac{\kappa_n}{1-\varsigma}}dn. \end{aligned}$$

The resulting expenditure shares on sectoral value added are then again given by $\vartheta_{rst}=-\frac{\partial V(e,\mathbf{P}_{rt})}{\partial P_{rst}}P_{rst}/\frac{\partial(e,\mathbf{P}_{rt})}{\partial e}e$. The expressions above imply

$$\vartheta_{rst}=\int_n \beta_n\chi_{rnt}^{s}\left(\mathbf{P}_{rt}\right)dn+\left(\int_n \kappa_n\chi_{rnt}^{s}\left(\mathbf{P}_{rt}\right)dn\right)\left(\frac{e}{B\left(\mathbf{P}_{rt}\right)}\right)^{-\varepsilon}, \tag{WA-3}$$

where $\chi_{rnt}^{s}\left(\mathbf{P}_{rt}\right)$ are the sectoral cost shares for good n given in (WA-2). The notation $\chi_{rnt}^{s}\left(\mathbf{P}_{rt}\right)$ stresses that these shares depend on the regional prices of tradable goods and CS. Equation (WA-3) is a direct generalization of the Cobb-Douglas structure considered in the main text. There, the spending shares $\chi_{rnt}^{s}\left(\mathbf{P}_{rt}\right)$ are constant and given by $\varsigma_{rnt}^{s}\left(\mathbf{P}_{rt}\right)=\lambda_{ns}$. In this more general formulation, the value added demand system still falls in the PIGL class (and has the same Engel elasticity ε as the final good demand system), but the other parameters depend on regional prices. In particular, (WA-3) can be written as

$$\vartheta_{rst}=\omega_{rst}+\nu_{rst}\left(\frac{e}{B\left(\mathbf{P}_{rt}\right)}\right)^{-\varepsilon}, \tag{WA-4}$$

where $\omega_{rst}\equiv\int_n\beta_n\chi_{rnt}^{s}\left(\mathbf{P}_{rt}\right)dn$ and $\nu_{rst}\equiv\int_n\kappa_n\chi_{rnt}^{s}\left(\mathbf{P}_{rt}\right)dn$. This is exactly the same representation as in our baseline analysis, except that ω_{rst} and ν_{rst} are no longer constant. Note, however, that it is still the case that $\sum_s\omega_{rst}=1$ and $\sum_s\nu_{rst}=0$ as required.

Equation WA-3 clarifies which aspects of our analysis hinge on the assumption of the final good production function (WA-1) to take the Cobb-Douglas form.

First, note that our strategy to estimate the Engel elasticity ε is still valid. Equation WA-3 implies that the expenditure share on food is given by

$$\vartheta_{\mathcal{F}}^{FE}=\int_{n\in\mathcal{F}}\beta_n\chi_{rnt}^{s}\left(\mathbf{P}_{rt}\right)dn+\left(\int_{n\in\mathcal{F}}\kappa_n\chi_{rnt}^{s}\left(\mathbf{P}_{rt}\right)dn\right)\left(\frac{e}{B\left(\mathbf{P}_{rt}\right)}\right)^{-\varepsilon}. \tag{WA-5}$$

If the asymptotic expenditure share on food is small, i.e., $\int_{n\in\mathcal{F}}\beta_n\chi_{rnt}^{s}\left(\mathbf{P}_{rt}\right)dn\approx 0$, (WA-5) shows that a cross-sectional regression of log food shares on log expenditure still identifies ε, because $\int_{n\in\mathcal{F}}\kappa_n\chi^s_{rnt}\left(\mathbf{P}_{rt}\right)$ is common across individuals within a location and hence absorbed in the region fixed effect.

Second, to calibrate our model in this more general case, we would require additional data. In addition to the elasticity of substitution ς of the production function (WA-1), we would need to know the good-specific sectoral weights $\{\lambda_{nF},\lambda_{nG},\lambda_{nCS}\}_n$, the asymptotic good-specific spending shares $\{\beta_n\}_n$ and the good-specific homotheticity parameters $\{\kappa_n\}_n$. The sectoral weights $\{\lambda_{ns}\}_n$ are needed to compute the good-specific sectoral cost shares χ^F_{rnt} given a set of sectoral prices $\mathbf{P}_{rt}$; see (WA-2). Given χ^s_{rnt}, one then needs $\{\beta_n\}_n$ and $\{\kappa_n\}_n$ to compute the demand shifters ω_{rst} and ν_{rst} in (WA-4). Given this additional information, our estimation procedure applies directly to this more general case. However, it would require data on cost shares and consumer demand at the disaggregated good level, which is not available in our context. For the case of a Cobb-Douglas production function, Proposition 1 shows that this information is *not* needed because the aggregate demand system only depends on the two sufficient statistics $\omega_s\equiv\int_{n=0}^1\lambda_{ns}\beta_n dn$ and $\nu_s\equiv\int_{n=0}^1\lambda_{ns}\kappa_n dn$, which we can directly estimate from aggregate data.

WA-1.2 Elasticity of Substitution

In this section we derive the expression for the elasticity of substitution given in A-3. Recall that the expenditure function is given by

$$e\left(P,V\right) = \left(V+\sum_s \nu_s\ln P_s\right)^{1/\varepsilon}\varepsilon^{1/\varepsilon}\prod_{s\in\{F,G,CS\}}P_s^{\omega_s}.$$

Then,

$$\begin{aligned}\frac{\partial e\left(P,V\right)}{\partial P_s} &= \left(V+\sum_s \nu_s\ln P_s\right)^{1/\varepsilon}\varepsilon^{1/\varepsilon}\prod_{s\in\{F,G,CS\}}P_s^{\omega_s}\left(\frac{\frac{1}{\varepsilon}\nu_s}{V+\sum_s\nu_s\ln P_s}+\omega_s\right)\frac{1}{P_s}\\ &= e\left(P,V\right)\left(\frac{\frac{1}{\varepsilon}\nu_s}{V+\sum_s\nu_s\ln P_s}+\omega_s\right)\frac{1}{P_s},\end{aligned}$$

and

$$\frac{\partial^2 e(P,V)}{\partial P_s \partial P_k} = \frac{\partial e(P,V)}{\partial p_k}\left(\frac{\frac{1}{\varepsilon}\nu_s}{V+\sum_s \nu_s \ln P_s}+\omega_s\right)\frac{1}{P_s} - e(P,V)\frac{\frac{1}{P_s}\frac{1}{\varepsilon}\nu_s\nu_k\frac{1}{P_k}}{\left(V+\sum_s \nu_s \ln P_s\right)^2}$$
$$= e(P,V)\frac{1}{P_k}\frac{1}{P_s}\left\{\left(\frac{\frac{1}{\varepsilon}\nu_k}{V+\sum_s \nu_s \ln P_s}+\omega_k\right)\left(\frac{\frac{1}{\varepsilon}\nu_s}{V+\sum_s \nu_s \ln P_s}+\omega_s\right)\right\}$$
$$- e(P,V)\frac{1}{P_k}\frac{1}{P_s}\varepsilon\frac{\frac{1}{\varepsilon}\nu_s\frac{1}{\varepsilon}\nu_k}{\left(V+\sum_s \nu_s \ln P_s\right)^2}$$

Now note that

$$\begin{aligned}\frac{\frac{1}{\varepsilon}\nu_k}{V+\sum_s \nu_s \ln P_s}+\omega_k &= \nu_k\frac{1}{\varepsilon}\left(V+\sum_s \nu_s \ln p_s\right)^{-1}+\omega_k \\ &= \nu_k\left(\frac{e(P,V)}{\prod_{s\in\{F,G,CS\}} P_s^{\omega_s}}\right)^{-\varepsilon}+\omega_k=\vartheta_k.\end{aligned}$$

Hence,

$$\begin{aligned}\frac{\partial e(P,V)}{\partial P_s} &= e(P,V)\,\vartheta_s\frac{1}{P_s} \\ \frac{\partial^2 e(P,V)}{\partial P_s \partial P_k} &= e(P,V)\frac{1}{P_k}\frac{1}{P_s}\left\{\vartheta_k\vartheta_s-\varepsilon\frac{\frac{1}{\varepsilon}\nu_s}{V+\sum_s \nu_s \ln P_s}\frac{\frac{1}{\varepsilon}\nu_k}{V+\sum_s \nu_s \ln P_s}\right\} \\ &= e(P,V)\frac{1}{P_k}\frac{1}{P_s}\left\{\vartheta_k\vartheta_s-\varepsilon(\vartheta_s-\omega_s)(\vartheta_k-\omega_k)\right\}.\end{aligned}$$

This implies that

$$\begin{aligned}EOS_{sk} &= \frac{e(P,V)\frac{1}{P_k}\frac{1}{P_s}\left\{\vartheta_k\vartheta_s-\varepsilon(\vartheta_s-\omega_s)(\vartheta_k-\omega_k)\right\}e(P,V)}{e(P,V)\,\vartheta_s\frac{1}{P_s}e(P,V)\,\vartheta_k\frac{1}{P_k}} \\ &= 1-\varepsilon\frac{(\vartheta_s-\omega_s)(\vartheta_k-\omega_k)}{\vartheta_s\vartheta_k}.\end{aligned}$$

WA-2 PIGL generalization (Section 7.3)

WA-2.1 Details for the general PIGL specification

Let the indirect utility function be

$$V^{FE}\left(e,[p_i]_{i=0}^{1}\right)=\frac{1}{\varepsilon}\left(\frac{e}{B(p)}\right)^{\varepsilon}-D(p) \tag{WA-6}$$

where

$$B(p) = \exp\left(\int_0^1 \beta_n \ln p_n dn\right) \text{ with } \int_n \beta_n dn = 1$$

and

$$D(p;\gamma) = \frac{1}{\gamma}\left[\left(\exp\left(\int_n \kappa_n \ln p_n dn\right)\right)^{\gamma} - 1\right] \text{ with } \int_n \kappa_n dn = 0.$$

Note that $\lim_{\gamma\to 0} D(p;\gamma) = \int_n \kappa_n \ln p_n dn$ as in the baseline model. The expenditure share of an individual with spending e on good n is then given by

$$\vartheta_n^{FE}(e) = -\frac{\frac{\partial V}{\partial p_s} p_s}{\frac{\partial V}{\partial e} e} = \beta_n + \kappa_n \left(\exp\left(\int_n \kappa_n \ln p_n dn\right)\right)^{\gamma} \left(\frac{e}{\exp\left(\int_0^1 \beta_n \ln p_n dn\right)}\right)^{-\varepsilon}.$$

To derive the expenditure shares on *sectoral* value added, note that

$$p_{rnt} = P_{rFt}^{\lambda_{nF}} P_{rGt}^{\lambda_{nG}} \left(\mathcal{A}_{rnt}^{-1} w_{rt}\right)^{\lambda_{nCS}}$$

Hence,

$$\vartheta_{rst}(e) = \omega_s + \nu_s \left(P_{rFt}^{\nu_F} P_{rGt}^{\nu_G} \left(A_{rCSt}^{-1} w_{rt}\right)^{\nu_{CS}}\right)^{\gamma} \left(\frac{e}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} \left(A_{rCSt}^{-1} w_{rt}\right)^{\omega_{CS}}}\right)^{-\varepsilon} \quad \text{(WA-7)}$$

where ω_s, ν_s and A_{rCSt} are defined as in Proposition 1. The aggregate expenditure share on sectoral value added in region r is then given by

$$\begin{aligned}
\overline{\vartheta}_{rst} &= \omega_s + \nu_s \left(P_{rFt}^{\nu_F} P_{rGt}^{\nu_G} \left(A_{rCSt}^{-1} w_{rt}\right)^{\nu_{CS}}\right)^{\gamma} \left(\frac{1}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} \left(A_{rCSt}^{-1} w_{rt}\right)^{\omega_{CS}}}\right)^{-\varepsilon} \frac{\int (q w_{rt})^{1-\varepsilon} dF_{rt}(q)}{\int (q w_{rt}) dF_{rt}(q)} \\
&= \omega_s + \frac{E_{rt}\left[q^{1-\varepsilon}\right]}{E_{rt}[q]^{1-\varepsilon}} \nu_s \left(P_{rFt}^{\nu_F} P_{rGt}^{\nu_G} \left(A_{rCSt}^{-1} w_{rt}\right)^{\nu_{CS}}\right)^{\gamma} \left(\frac{w_{rt} E_{rt}[q]}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} \left(A_{rCSt}^{-1} w_{rt}\right)^{\omega_{CS}}}\right)^{-\varepsilon},
\end{aligned}$$

where, as before, $e = q w_{rt}$ with $q \sim F_{rt}(q)$. Under our distributional assumptions on F_{rt}, $\frac{E_{rt}\left[q^{1-\varepsilon}\right]}{E_{rt}[q]^{1-\varepsilon}} = \frac{\zeta^{\varepsilon}(\zeta-1)^{1-\varepsilon}}{\zeta+\varepsilon-1}$, and we can express $\overline{\vartheta}_{rst}$ as

$$\overline{\vartheta}_{rst} = \omega_s + \overline{\nu}_s \left(P_{rFt}^{\omega_F \varepsilon + \gamma \nu_F} P_{rGt}^{\omega_G \varepsilon + \gamma \nu_G}\right) \left(\frac{w_{rt}}{A_{rCSt}}\right)^{\omega_{CS}\varepsilon + \nu_{CS}} \left(w_{rt} E_{rt}[q]\right)^{-\varepsilon}, \quad \text{(WA-8)}$$

with $\overline{\nu}_s = \frac{\zeta^{\varepsilon}(\zeta-1)^{1-\varepsilon}}{\zeta+\varepsilon-1} \nu_s$.

The indirect utility function over value added associated with (WA-6) is given by

$$V(e, [P_{rst}]) = \frac{1}{\varepsilon}\left(\frac{e}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} \left(A_{rCSt}^{-1} w_{rt}\right)^{\omega_{CS}}}\right)^{\varepsilon} - \frac{1}{\gamma}\left(P_{rFt}^{\gamma\nu_F} P_{rGt}^{\gamma\nu_G} \left(A_{rCSt}^{-1} w_{rt}\right)^{\gamma\nu_{CS}} - 1\right). \quad \text{(WA-9)}$$

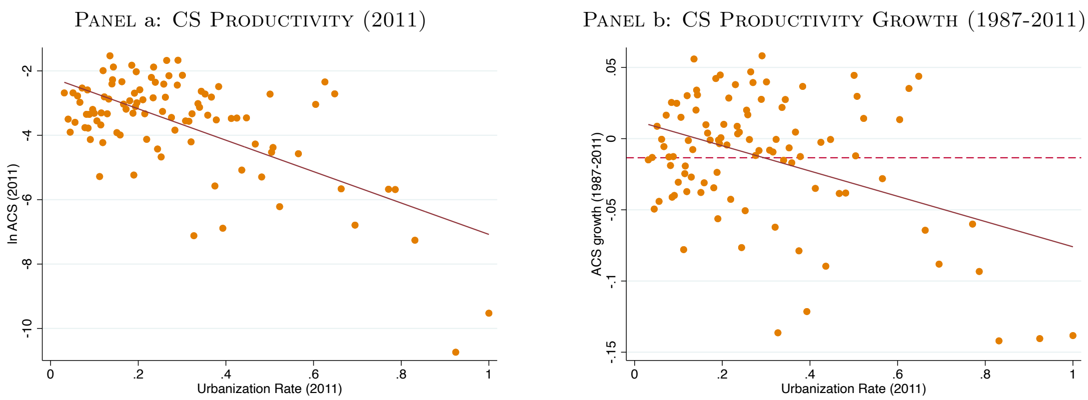

Figure WA-1: $\gamma > \gamma^*$: Implications for CS Productivity. In the left (right) panel we depict the correlation between the estimated CS productivity in 2011 (CS productivity growth rate between 1987 and 2011) and the urbanization rate in 2011. These estimates stem from a calibration of our model with $\gamma = 0.5$.

Given (WA-9), we can also compute the certainty equivalent of a counterfactual change of prices. As before, define the certainty equivalent ϖ of a counterfactual allocation $\left(\hat{w}_{rt}, \hat{P}_{rt}\right)$ given the current allocation (w_{rt}, P_{rt}) as

$$V\left(qw_{rt}\left(1+\varpi\right), \left[P_{rst}\right]\right) \equiv V\left(q\hat{w}_{rt}, \left[\hat{P}_{rst}\right]\right).$$

Using (WA-9), we can solve for ϖ as

$$1+\varpi = \prod_s \left(\frac{\hat{w}_{rt}/\hat{P}_{rst}}{w_{rt}/P_{rst}}\right)^{\omega_s} \times \left\{1 - \varepsilon\left(\frac{q\hat{w}_{rt}}{\prod_s \hat{P}_{rst}^{\omega_s}}\right)^{-\varepsilon} \frac{1}{\gamma}\left(\left(\prod_s \left(\frac{\hat{P}_{rst}}{P_{rst}}\right)^{\nu_s}\right)^{\gamma} - 1\right)\left(\prod_s P_{rst}^{\nu_s}\right)^{\gamma}\right\}^{1/\varepsilon}.$$

It can be shown that this expression reduces to the expression in (A-6) if $\gamma \to 0$.

WA-2.2 Implications for CS productivity

In Section 7.3 we discussed the paradoxical implications of a parametrization that involves $\gamma > \gamma^* - \varepsilon\frac{\omega_{CS}}{\nu_{CS}}$. In Figure WA-1 we display these predictions graphically. These figures stem from a calibration of our model, which imposes $\gamma = 0.5$ and is otherwise calibrated to the same moments as our baseline model.

In the left panel, we show the cross-sectional correlation between the urbanization rate in 2011 and $\ln A_{rCS2011}$. As highlighted in the text, there is a strong negative correlation, that is, cities have *low* productivity in the provision of consumer services. In the right panel we focus on productivity growth in CS, i.e., $\ln A_{rCS2011} - \ln A_{rCS1987}$. Again, the correlation with the urbanization rate is negative. Moreover, the average productivity growth rate, indicated as the dashed line, is negative. These implications not only strike us as non-sensible but they are also at odds with empirical estimates of aggregate productivity growth that point towards positive growth in the services sector; see Table 6.

In Figure WA-2 we display the distribution of the estimates productivity growth

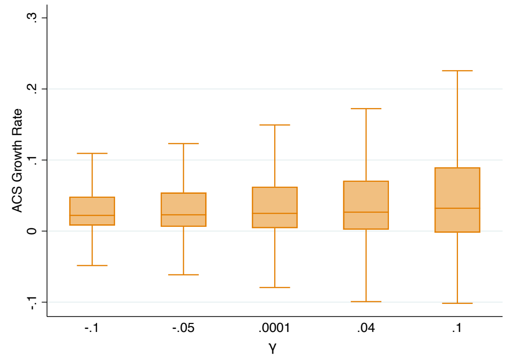

Figure WA-2: Estimated CS Productivity Growth: The Role of γ. The figure shows a boxplot of the cross-sectional distribution of estimated CS productivity growth rate, $\frac{\ln A_{rCS2011}-\ln A_{rCS1987}}{2011-1987}$, as a function γ. The solid line within the box shows the 50th percentile. The box shows the interquartile range. The lines at the lower and upper end show the upper and lower adjacent values. The upper (lower) adjacent value is defined as the 75% (25%) quantile plus (minus) 1.5 times the interquartile range.

rate in the CS sector, $\frac{\ln A_{rCS2011}-\ln A_{rCS1987}}{2011-1987}$, as a function of γ (for the range where 90% of regions have an $ESO_{CS,G}$ between 0 and 1. The distribution fans out for high level of γ as we approach γ^*. However, the average rate of CS productivity growth is relatively constant.

WA-2.3 The Elasticity of Substitution

We now derive the Allen-Uzawa elasticity of substitution for the preference specification in (WA-6). The Allen-Uzawa elasticity of substitution between sectors s and k is defined by

$$EOS_{sk} \equiv \frac{\frac{\partial^2 e(p,V)}{\partial p_s \partial p_k} e\left(p, V\right)}{\frac{\partial e(p,V)}{\partial p_s}\frac{\partial e(p,V)}{\partial p_k}}. \tag{WA-10}$$

The expenditure function associated with the indirect utility function of value added in (WA-9) is given by

$$e\left(p, V\right) \quad = \quad \left(V + \frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j} - \frac{1}{\gamma}\right)^{1/\varepsilon} \varepsilon^{1/\varepsilon}\prod_s p_s^{\omega_s}.$$

We now derive the different components of EOS_{sk}^{AU} as defined in (WA-10).

1. The partial elasticity of the expenditure function is given by

$$
\begin{aligned}
&= \frac{\partial e\left(p,V\right)}{\partial p_s} \\
&= \varepsilon^{1/\varepsilon}\prod_s p_s^{\omega_s}\left[\frac{1}{\varepsilon}\left(V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}\right)^{\frac{1}{\varepsilon}-1}\tilde{\nu}_s\prod_j p_j^{\gamma\tilde{\nu}_j}+\left(V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}\right)^{1/\varepsilon}\omega_s\right]\frac{1}{p_s} \\
&= \varepsilon^{1/\varepsilon}\left(\prod_s p_s^{\omega_s}\right)\left(V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}\right)^{1/\varepsilon}\left[\frac{1}{\varepsilon}\frac{\nu_s\prod_j p_j^{\gamma\nu_j}}{V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}}+\omega_s\right]\frac{1}{p_s} \\
&= e\left(p,V\right)\left[\frac{1}{\varepsilon}\frac{\nu_s\prod_j p_j^{\gamma\nu_j}}{V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}}+\omega_s\right]\frac{1}{p_s}. \qquad \text{(WA-11)}
\end{aligned}
$$

Now note that

$$\frac{1}{\varepsilon}\frac{\nu_s\prod_j p_j^{\gamma\nu_j}}{V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}}+\omega_s=\vartheta_s. \qquad \text{(WA-12)}$$

Substituting (WA-12) in (WA-11) yields

$$\frac{\partial e\left(p,V\right)}{\partial p_s}=e\left(p,V\right)\vartheta_s\frac{1}{p_s}. \qquad \text{(WA-13)}$$

2. The cross-partial elasticity of the expenditure function is given by

$$
\begin{aligned}
&= \frac{\partial^2 e\left(p,V\right)}{\partial p_s\partial p_k} \\
&= \frac{1}{p_s}\left(\frac{\partial e\left(p,V\right)}{\partial p_k}\left[\frac{1}{\varepsilon}\frac{\tilde{\nu}_s\prod_j p_j^{\gamma\nu_j}}{V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}}+\omega_s\right]+e\left(p,V\right)\frac{1}{\varepsilon}\frac{1}{p_k}\left(\frac{\nu_s\nu_k\gamma\prod_j p_j^{\gamma\nu_j}\left(V-\frac{1}{\gamma}\right)}{\left(V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}\right)^2}\right)\right) \\
&= \frac{1}{p_s}\frac{1}{p_k}e\left(p,V\right)\left[\frac{1}{\varepsilon}\frac{\nu_k\prod_j p_j^{\gamma\nu_j}}{V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}}+\omega_k\right]\left[\frac{1}{\varepsilon}\frac{\nu_s\prod_j p_j^{\gamma\nu_j}}{V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}}+\omega_s\right]+ \\
&\quad \frac{1}{p_s}\frac{1}{p_k}e\left(p,V\right)\frac{1}{\varepsilon}\left(\frac{\nu_s\nu_k\gamma\prod_j p_j^{\gamma\nu_j}\left(V-\frac{1}{\gamma}\right)}{\left(V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}\right)^2}\right)
\end{aligned}
$$

Using (WA-12) we get

$$\frac{\partial^2 e\left(p,V\right)}{\partial p_s\partial p_k}=\frac{1}{p_s}\frac{1}{p_k}e\left(p,V\right)\left(\vartheta_k\vartheta_s+\frac{1}{\varepsilon}\left(\frac{\nu_s\nu_k\gamma\prod_j p_j^{\gamma\nu_j}\left(V-\frac{1}{\gamma}\right)}{\left(V+\frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j}-\frac{1}{\gamma}\right)^2}\right)\right)$$

Furthermore, note that

$$\frac{\nu_s \nu_k \gamma \prod_j p_j^{\gamma\nu_j} \left(V - \frac{1}{\gamma}\right)}{\left(V + \frac{1}{\gamma}\prod_j p_j^{\gamma\nu_j} - \frac{1}{\gamma}\right)^2} = (\vartheta_k - \omega_k)(\vartheta_s - \omega_s)\,\varepsilon\left(\gamma \frac{\left(\frac{e}{\prod_s p_s^{\omega_s}}\right)^{\varepsilon}}{\prod_j p_j^{\gamma\nu_j}} - \varepsilon\right) \tag{WA-14}$$

and that (using (WA-7))

$$\frac{\left(\frac{e}{\prod_s p_s^{\omega_s}}\right)^{\varepsilon}}{\prod_j p_j^{\gamma\nu_j}} = \frac{\nu_s}{\vartheta_s - \omega_s} \text{ for all } s. \tag{WA-15}$$

Hence,

$$\frac{\partial^2 e(p,V)}{\partial p_s \partial p_k} = \frac{1}{p_s}\frac{1}{p_k} e(p,V)\left(\vartheta_k\vartheta_s + (\vartheta_k - \omega_k)(\vartheta_s - \omega_s)\left(\frac{\gamma\nu_s}{\vartheta_s - \omega_s} - \varepsilon\right)\right).$$

We can thus compute the Allen-Uzawa elasticity as

$$EOS_{sk} = 1 + \left(\frac{\gamma\nu_s}{\vartheta_s - \omega_s} - \varepsilon\right)\frac{(\vartheta_k - \omega_k)(\vartheta_s - \omega_s)}{\vartheta_s\vartheta_k}.$$

WA-3 Generalizations of Theory: Formal Details

In this section we provide additional formal details for the extension of our theory discussed in Sections 7.4 in the main text and A-5 in the Appendix.

WA-3.1 Open economy

In this model we present the formal analysis for the open economy extension.

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

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

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

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

$$P_G = \left(p_{G,D}^{1-\eta} + \varphi^{\eta} p_{G,ROW}^{1-\eta}\right)^{\frac{1}{1-\eta}}. \quad \text{(WA-16)}$$

The expenditure share on Indian goods is $\frac{p_{G,D}C_{G,D}}{P_G C_G} = \left(\frac{P_{G,D}}{P_G}\right)^{1-\eta}$. Combining this expression with Equation (WA-16) yields the expenditure shares

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

For simplicity we subsume trade costs in the relative price of foreign goods and assume there are no intra-country shipment costs for exporting goods. We do, however, still assume (as in the baseline model) that there are intra-country trade costs for domestically consumed food and goods.

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

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

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

Consider next the exported ICT services.[12] We assume that the ROW buys a bundle of regional varieties of ICT services

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

where y_{rICTt} denotes the quantity of services produced in region r and exported to the rest of the world. ICT services are produced in region r according to the production

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

function $y_{rICTt} = A_{rICTt} H_{rt}$. Hence, the price of ICT services is given by

$$p_{ICT} = \left(\sum_r p_{rICT}^{1-\sigma} \right)^{\frac{1}{1-\sigma}} = \left(\sum_r \left(\frac{w_r}{A_{rICT}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}.$$

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

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

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

We do allow for the international trade cost; however, it is not separately identified from the foreign demand shifter in our estimation. In addition, there is no ICT exporting cost.

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

1. Market clearing for agricultural goods:

$$w_{rt} H_{rFt} = \sum_{j=1}^{R} \pi_{rFjt} \left(\omega_F + \nu_F \left(\frac{A_{jCSt}^{\omega_{CS}} \mathbb{E}_{jt}[q]\, w_{jt}^{1-\omega_{CS}}}{P_{jFt}^{\omega_F} \left(P_{rGt}^{Agg} \right)^{\omega_G}} \right)^{-\varepsilon} \right) w_{jt} H_{jt}$$

   where $\pi_{rFot} = \tau_{ro}^{1-\sigma} A_{oFt}^{\sigma-1} w_{ot}^{1-\sigma} / P_{rFt}^{1-\sigma}$

2. Market clearing for manufacturing goods:

$$w_{rt} H_{rFt} = \sum_{j=1}^{R} \pi_{rGjt} \frac{P_{jGt}^{1-\eta}}{\left(P_{jGt}^{Agg} \right)^{1-\eta}} \left(\omega_G + \nu_G \left(\frac{A_{jCSt}^{\omega_{CS}} \mathbb{E}_{jt}[q] w_{jt}^{1-\omega_{CS}}}{P_{jFt}^{\omega_F} \left(P_{jGt}^{Agg} \right)^{\omega_G}} \right)^{-\varepsilon} \right) w_{jt} H_{jt}$$
$$+ \left(\frac{w_{rt}^{1-\sigma} A_{rGt}^{\sigma-1}}{\sum_{j=1}^{R} w_{jt}^{1-\sigma} A_{jGt}^{\sigma-1}} \right) \times \left(\sum_j w_{jt}^{1-\sigma} A_{jGt}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{Gt}$$

   where $\left(P_{jGt}^{Agg} \right)^{1-\eta} = P_{jGt}^{1-\eta} + \varphi^{\eta} p_{G,ROW,t}^{1-\eta}$ and $\pi_{rGot} = \tau_{ro}^{1-\sigma} A_{oGt}^{\sigma-1} w_{ot}^{1-\sigma} / P_{rGt}^{1-\sigma}$

3. Market clearing for local CS:

$$w_{rt} H_{rCSt} = \left(\omega_{CS} + \nu_{CS} \left(\frac{A_{rCSt}^{\omega_{CS}} \mathbb{E}_{rt}[q]\, w_{rt}^{1-\omega_{CS}}}{P_{rFt}^{\omega_F} \left(P_{rGt}^{Agg} \right)^{\omega_G}} \right)^{-\varepsilon} \right) w_{rt} H_{rt}$$

4. Market clearing for local ICT services:

$$w_{rt} H_{rICTt} = \left(\frac{w_{rt}^{1-\sigma} A_{rICTt}^{\sigma-1}}{\sum_{j=1}^{R} w_{jt}^{1-\sigma} A_{jICTt}^{\sigma-1}} \right) \times \underbrace{\left(\sum_{j} w_{jt}^{1-\sigma} A_{jICTt}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICTt}}_{\text{ICT exports}}$$

5. Labor market clearing:

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

6. Balanced Trade:

$$\underbrace{\left(\left(\sum_{j} w_{jt}^{1-\sigma} A_{jGt}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{Gt} + \left(\sum_{j} w_{jt}^{1-\sigma} A_{jICTt}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICTt} \right)}_{\text{Exports}} = \underbrace{\sum_{j=1}^{R} \frac{\left(\omega_G + \nu_G \left(\frac{A_{jCSt}^{\omega_{CS}} \mathbb{E}_{jt}[q] w_{jt}^{1-\omega_{CS}}}{P_{jFt}^{\omega_F} \left(P_{jGt}^{Agg} \right)^{\omega_G}} \right)^{-\varepsilon} \right) w_{jt} H_{jt}}{\varphi^{-\eta} \left(\frac{P_{rGt}}{p_{G,ROW,t}} \right)^{1-\eta} + 1}}_{\text{Imports}}$$

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

$$p_{G,IND} = \left(\sum_{r} \left(\frac{w_{rt}}{A_{rGt}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} = 1.$$

Given the productivities $\{A_{rFt}, A_{rGt}, A_{rCSt}, A_{rICTt}\}_r$, the population distribution $\{H_{rt}\}_r$, the demand shifters of the foreign sector $(\Upsilon_{ICTt}, \Upsilon_{Gt})$ and the other preference parameters of the model, we can calculate

$$\{x_t, \{w_{rt}, H_{rFt}, H_{rGt}, H_{rCSt}, H_{rICTt}\}_r\}.$$

Identification of Productivity Fundamentals For the economy with trade we need to identify the following additional objects:

$$\left\{ [A_{rICTt}]_{r=1}^{R}, \Upsilon_{Gt}, \Upsilon_{ICTt} \right\}.$$

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

1. Relative ICT payments across localities for ICT exports:

$$\frac{w_{rt} H_{rICTt}}{w_{jt} H_{jICTt}} = \frac{w_{rt}^{1-\sigma} A_{rICTt}^{\sigma-1}}{w_{jt}^{1-\sigma} A_{jICTt}^{\sigma-1}}.$$

These are $R-1$ equations to determine A_{rICTt} up to scale, that is,

$$A_{rICTt} = A_{ICTt} a_{rICTt} \text{ with } \sum_r a_{rICTt}^{\sigma-1} = 1$$

yields

$$a_{rICTt} = \left(\frac{H_{rICTt} w_r^{\sigma}}{\sum_j H_{jICTt} w_{jt}^{\sigma}} \right)^{\frac{1}{\sigma-1}}.$$

Because the level of ICT productivity A_{ICTt} is not separately identified from the aggregate demand shifter Υ_{ICTt}, without loss of generality we can set $A_{ICTt} = 1$.[13]

2. To identify Υ_{ICT} we use that

$$\begin{aligned} \sum_r w_r H_{rICTt} &= \sum_r \left(\frac{w_{rt}^{1-\sigma} A_{rICTt}^{\sigma-1}}{\sum_{j=1}^{R} w_{jt}^{1-\sigma} A_{jICTt}^{\sigma-1}} \right) \left(\sum_j w_{jt}^{1-\sigma} A_{jICTt}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICTt} \\ &= \left(\sum_j w_j^{1-\sigma} a_{jICTt}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICTt}. \end{aligned} \tag{WA-17}$$

The right hand-side is total value added of the ICT sector, which we can calculate directly in the data. Given that w_{jt} and a_{jICTt} are observed, we can calculate Υ_{ICTt}.

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

$$\text{Total value added in manufacturing} = \sum_r w_{rt} H_{rGt}$$

and

$$\text{Total value added of exports} = \left(\sum_j w_{jt}^{1-\sigma} A_{jGt}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{Gt}.$$

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

$$M_{1t} = \frac{\left(\sum_j w_{jt}^{1-\sigma} A_{jGt}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{Gt}}{\sum_r w_{rt} H_{rGt}} = \frac{P_{G,IND}^{1-\eta} \Upsilon_{Gt}}{\sum_r w_{rt} H_{rGt}} = \frac{\Upsilon_{Gt}}{\sum_r w_{rt} H_{rGt}}. \tag{WA-18}$$

[13] Note that the equilibrium condition for ICT exports implies that

$$w_{rt} H_{rICTt} = \left(\frac{w_{rt}^{1-\sigma} A_{rICTt}^{\sigma-1}}{\sum_j^R w_{jt}^{1-\sigma} A_{jICTt}^{\sigma-1}} \right) \left(\sum_j w_{jt}^{1-\sigma} A_{jICTt}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICTt} = \left(\frac{w_{rt}^{1-\sigma} a_{rICTt}^{\sigma-1}}{\sum_j^R w_{jt}^{1-\sigma} a_{jICTt}^{\sigma-1}} \right) \left(\sum_j w_{jt}^{1-\sigma} a_{jICTt}^{\sigma-1} \right)^{\frac{1-\eta}{1-\sigma}} A_{ICTt}^{\eta-1} \Upsilon_{ICTt}.$$

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

Therefore, for a given moment of the export share of manufacturing M_{1t} and data on $\{w_{jt}, H_{jGt}\}_j$ we can solve for Υ_{Gt}.

WA-3.2 Imperfect Skill Substitution

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

$$w_{rt}^{-} H_{rst}^{-} + w_{rt}^{+} H_{rst}^{+} \quad = \quad \sum_{j=1}^{R} \pi_{rsjt} \times \vartheta_s \left(\underline{q}_{jt}^{-} w_{jt}^{-}, \underline{q}_{jt}^{+} w_{jt}^{+}, s_{jt}^{Y,-}, \mathbf{p_{jt}} \right) \overline{w}_{rt} L_{rt},$$

where $\overline{w}_{rt}$ denotes average income, $\pi_{rsot} = \tau_{ro}^{1-\sigma} p_{rst}^{1-\sigma} / P_{rst}^{1-\sigma}$, and $P_{rst}^{1-\sigma} = \sum_o \tau_{ro}^{1-\sigma} p_{ost}^{1-\sigma}$. The market clearing condition for non-tradable CS implies

$$w_{rt}^{-} H_{rCSt}^{-} + w_{rt}^{+} H_{rCSt}^{+} = \vartheta_{CS} \left(\underline{q}_{rt}^{-} w_{rt}^{-}, \underline{q}_{rt}^{+} w_{rt}^{+}, s_{rt}^{Y,-}, \mathbf{p_{rt}} \right) \overline{w}_{rt} L_{rt}. \tag{WA-19}$$

Finally, labor market clearing implies

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

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

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

1. We classify individuals into high and low skill workers by their years of schooling. We assume workers with at least secondary schooling are high-skill workers.

2. As in our baseline model, we assume a Mincerian return $\rho = 5.6\%$ per year of schooling within skill groups. This allows us to measure the aggregate skill supplies H_{rt}^{-} and H_{rt}^{+} for each region.

3. As in our baseline model, we use the observed sectoral earnings shares by skill group to measure sectoral labor supplies. Specifically, for each skill group $j = \{-, +\}$ and sector s, we calculate

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

where w_i is the wage of individual i.

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

These data are sufficient to uniquely solve for $\{A_{rst}, Z_{rst}\}$ and to perform the counterfactual analysis reported in Section 7.4.

WA-3.3 Spatial Mobility

Model Setting In this section, we describe how we incorporate spatial labor mobility into the baseline model. We assume that individuals are free to locate in the region of their choosing. Given the long-run focus of our analysis, we assume that individuals learn their productivity q after settling in region r. This productivity is drawn from the location-specific distribution $F_{rt}(q)$. Intuitively, by settling in location r, individuals have access to the local schooling system and they take this form of local human capital accumulation into account when making their location choice.

Formally, we assume that the utility of individual i to settle in location r at time t given the wage vector $\hat{w}_{rt}$ and the price vector $\hat{\mathbf{P}}_{\mathbf{rst}}$ is given by

$$V_{rt}^{i} \equiv \mathcal{B}_{rt} E_{rt}[q] w_{rt} \left(1 + \overline{\varpi}_{rt}\left(\hat{w}_{rt}, \hat{\mathbf{P}}_{\mathbf{rst}} | w_{rt}, \mathbf{P}_{\mathbf{rst}}\right)\right) u_{rt}^{i},$$

where $\overline{\varpi}_{rt}$ is the equivalent variation, $w_{rt}, \mathbf{P}_{\mathbf{rst}}$ are the wages and prices in the calibrated equilibrium in 2011, $\mathcal{B}_{rt}$ is a location amenity, and u_{rt}^i is an idiosyncratic preference shock for location r.[14] By cardinalizing consumers' spatial preferences with $\overline{\varpi}_{rt}$, we measure spatial amenities $\mathcal{B}$ and u_r in money terms. As a result, the overall utility of a location in the original equilibrium is simply $\mathcal{U}_{rt}^i = \mathcal{B}_{rt} E_{rt}[q] w_{rt} u_{rt}^i$.

We assume that workers' idiosyncratic preference shocks for each location u_{rt}^i are Frechet-distributed with parameter η, that is, $P\left(u_{rt}^i \leq u\right) = e^{-u^{-\eta}}$. Under these assumptions, one can show that the spatial allocation of labor is given by

$$L_{rt} = \frac{\left(v_{rt}\mathcal{B}_{rt}\right)^{\eta}}{\sum_j \left(v_{jt}\mathcal{B}_{jt}\right)^{\eta}} L. \tag{WA-20}$$

[14] Note that individuals evaluate locations based on the average money-metric utility $\overline{\varpi}_{rt}$, because they do not know their specific human capital realization q when making their location choice.

where $v_{rt} \equiv E_{rt}[q]w_{rt}\left(1+\overline{\varpi}_{rt}\left(\hat{w}_{rt},\hat{\mathbf{P}}_{\mathbf{rst}}|w_{rt},\mathbf{P}_{\mathbf{rst}}\right)\right)$ denotes the systematic part of regional utility. Holding $\sum_j (v_{jt}\mathcal{B}_{jt})^{\eta}$ constant, the partial elasticity with respect to the money-metric utility is given by η.[15] Note that η is not equal to the empirically estimated labor supply elasticity with respect to local wages due to the presence of non-homothetic preferences.

Estimation Allowing for spatial mobility requires us to estimate additional parameters. First, we need to estimate the level of exogenous amenities $\mathcal{B}_{rt}$. Second, we need the labor supply elasticity η.

Using the set of Equations (WA-20), we can identify $\mathcal{B}_{rt}$ given the observed allocation of labor and wages and given an estimate of η. Hence, we cannot separately identify η without additional information. However, given η we can estimate $\mathcal{B}_{rt}$ to rationalize the population distribution given the observed wages and emplomyment allocation.

Because we are mainly interested in understanding how the option of labor mobility affects our welfare counterfactuals, we discipline η by their implied migration response. For our main exercise we chose η so that the cross-sectional standard deviation of employment growth induced by setting productivity in all sectors to their 1987 level is the same as the one observed in the data between 1987 and 2011. More specifically, let $\hat{L}_r$ denote number of people in region r in the counterfactual equilibrium where local amenities are given by $\mathcal{B}_{r2011}$ but productivities take their 1987 value, i.e. A_{rs1987}. Let $\hat{\ell}_r = \hat{L}_r / \sum_r \hat{L}_r$ and $\ell_r = L_r / \sum_r L_r$ denote the respective population shares. The cross-sectional standard deviation of population share changes is then given by

$$\Sigma \equiv sd(\hat{\ell}_r - \ell_r).$$

We then choose η such that Σ coincides with the observed counterpart between 1987 and 2011, that is $sd(\ell_{r1987} - \ell_{r2011})$. This implies that $\eta = 0.68$. To generate twice the standard deviation, we would require $\eta = 2.34$.

The Welfare Effect of Service-led Growth in the Presence of Mobility We compute the welfare effect of service-led growth in the presence of spatial mobility in the following way: Given the elasticity η we first estimate the vector of local amenities in 2011, $\mathcal{B}_{r2011}$, to rationalize the observed population distribution given wages and sectoral employment shares.

We then set the vector of productivities in the CS sector to their level in 1987, $A_{rCS1987}$, and solve for the counterfactual level of wages $\hat{w}_r$ and prices $\hat{\mathbf{P}}_{\mathbf{rst}}$ using the

[15] It is also possible to explicitly allow for congestion externalities, where local amenities depend on the size of the population. If, for example, amenities were given by $\mathcal{B}_{rt} = B_{rt}L_{rt}^{-\delta}$ with B_{rt} being a time-varying, exogenous district characteristic, the parameter δ would parameterize the strength of local congestion through housing prices or the reduced availability of public goods. In our setup without moving costs, δ plays a very similar role to η as they both affect the aggregate labor supply.

equilibrium conditions stated in Proposition 2, together with the labor supply equation (WA-20).

Given the new equilibrium wages and prices, we estimate the average welfare losses. To so, we simulate the optimal behavior of 1m individuals. More specifically, consider an individual i that draws a vector of idiosyncratic location tastes $\hat{u}^i = \{\hat{u}^i_r\}_{r=1}^R$ from $F(\hat{u}^i_r) = e^{-(\hat{u}^i_r)^{-\eta}}$. All draws are independent across locations. Given $\hat{u}^i$, the utility for individual i to move to location r in the observed equilibrium in 2011 is given by

$$V^i_{r2011} = \mathcal{B}_{r2011} E_{r2011}[q] \, w_{r2011} \hat{u}^i_r, \tag{WA-21}$$

and the actual utility of individual i is given by

$$V^i_{2011} = \max_r \{V^i_{r2011}\}. \tag{WA-22}$$

In the counterfactual equilibrium, the utility of individual i to settle in location j is given by

$$V^i_{jCF} = \mathcal{B}_{j2011} E_{j2011}[q] \, w_{j2011} \left(1 + \varpi_j \left(w_{CF}, P_{CF} | w_{2011}, P_{2011}\right)\right) \hat{u}^i_j. \tag{WA-23}$$

Equation (WA-23) highlights that the counterfactual utility, V^i_{jCF}, consists of (i) the location amenity $\mathcal{B}_{j2011}$, which does not change; (ii) the expected skill level $E_{j2011}[q]$ at the destination j, given the actual distribution of human capital in 2011; (iii) the equivalent wage of working and consuming in j given the counterfactual wage and prices, $w_{j2011}\left(1 + \varpi_j \left(w_{CF}, P_{CF} | w_{2011}, P_{2011}\right)\right)$; and (iv) person i's idiosyncratic preference, $\hat{u}^i_j$, which also determined the initial location choice (WA-21). Hence, we assume that people keep their initial location preference, $\hat{u}^i_j$, when contemplating to change locations. Individuals that moved to Delhi because of a high location preference $\hat{u}^i_{Delhi}$ are likely to stay in Delhi.

Now consider an individual i that settled in location r in the original equilibrium and in j in the counterfactual. The utility change of individual i is given by

$$\varpi^i_{r,MOB} \equiv \frac{V^i_{jCF}}{V^i_{r2011}} - 1, \tag{WA-24}$$

where the subscript MOB stresses that $\varpi^i_{r,MOB}$ takes the option value of moving into account. Also note that V^i_{jCF} and V^i_{r2011} are already cardinalized in monetary terms so that $\varpi^i_{r,MOB}$ already has the interpretation of an equivalent variation, taking into account the potential changes in location amenities encapsulated in $\mathcal{B}_{j2011}$ and $\hat{u}^i_j$. Using (WA-21), (WA-23), and (WA-24), we can express $\varpi^i_{r,MOB}$ as

$$1 + \varpi^i_{r,MOB} = \frac{V^i_{rCF}}{V^i_{r2011}} \frac{V^i_{jCF}}{V^i_{rCF}} = \underbrace{\left(1 + \varpi_r \left(w_{CF}, P_{CF} | w_{2011}, P_{2011}\right)\right)}_{\text{EV of staryers}} \times \underbrace{\frac{V^i_{jCF}}{V^i_{rCF}}}_{\text{Insurance}} \tag{WA-25}$$

Hence, the overall welfare effect is the product of equivalent variation of stayers and the term V^i_{jCF}/V^i_{rCF}, which captures that the option of spatial mobility offers insurance: if the situation in location r deteriorates too much, you can move to j. Note that by virtue of individual i moving from r to j, $V^i_{jCF} \geq V^i_{rCF}$. This implies that

$$\overline{\varpi}^i_{r,MOB} \geq \overline{\varpi}_r \left(w_{CF}, P_{CF} | w_{2011}, P_{2011}\right), \tag{WA-26}$$

i.e. the welfare loss of falling CS productivity will necessarily be smaller once the option of spatial mobility is taken into account.

Given the simulated migration choices for N individuals, we compute the aggregate welfare effect as

$$\overline{\varpi}_{AGG,MOB} = \frac{1}{N} \sum_{i=1}^{N} \overline{\varpi}^i_{r,MOB}. \tag{WA-27}$$

Similarly, the welfare effect of individuals who sorted into region r in the initial equilibrium is given by

$$\overline{\varpi}_{r,MOB} = \frac{1}{N_r} \sum_{i=1}^{N_r} \overline{\varpi}^i_{r,MOB}, \tag{WA-28}$$

where N_r denotes the number of individuals in region r. Because we simulate the initial distribution using the observed factor prices and calibrated location amenities, this distribution coincides with the actual region population distribution in 2011. In Table 9 in the main text we report $\overline{\varpi}_{AGG,MOB}$ in column 1 and $\overline{\varpi}_{r,MOB}$, aggregated by urbanization quintiles, in columns 2 and 3. Because we assume that individuals redraw their human capital after moving, the welfare effects by income quantile are not well-defined.

In our main analysis we showed that cities were the main beneficiaries of service-led growth, both because they experienced particularly fast productivity growth in CS and their residents are, on average, richer. This implies that cities should, on average, lose residents if CS productivity is reset to the level in 1987. In the left panel of Figure WA-3 we report the implications of our model. There is indeed a strong negative relationship and cities are predicted to experience very large negative population growth. In the right panel, we depict the actual change in the local population between 1987 and 2011 by the urbanization rate in 2011. For ease of comparison with the counterfactual results shown in the left panel, we plot $L_{r,1987}/L_{r,2011} - 1$, i.e. how much *smaller* the location was in 1987 relative to its population in 2011. The figure shows that population growth was very unbalanced and that the cities in 2011 experienced a dramatic rise in their population.

Figure WA-3 shows that the extent of mobility induced by service-led growth was of a similar magnitude than what is observed in the data. Recall that we calibrated

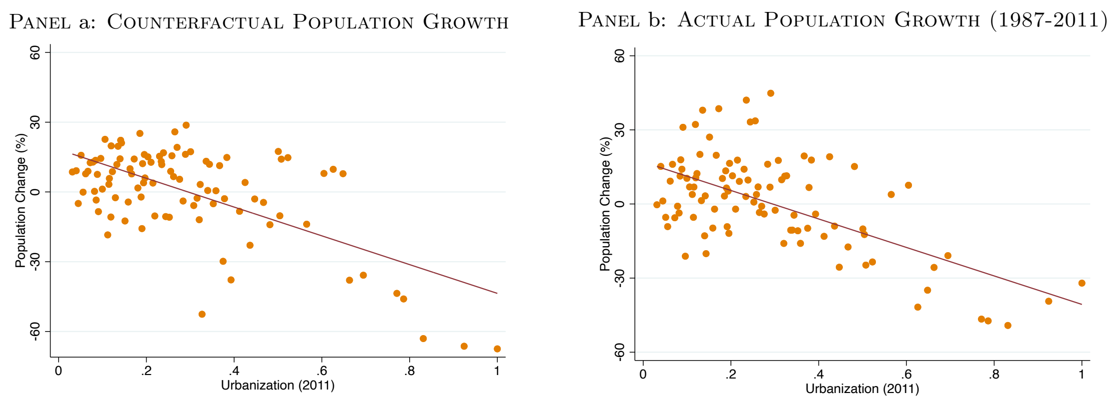

Figure WA-3: Local Population Growth. In the left panel we depict the correlation between the urbanization rate in 2011 and implied population growth in response to resetting CS productivity to its level in 1987. In the right panel we depict the correlation between the urbanization rate in 2011 and change in the local population between 1987 and 2011, $L_{r,1987}/L_{r,2011} - 1$.

η to match the cross-sectional standard deviation of local population changes shown in the right panel of Figure WA-3 if all productivities had been set back to their 1987 level. The left panel such suggests that changes in service productivity account for a large share of this dispersion, which is not entire surprising given their non-tradable nature. Importantly, the implied population changes in the left panel are arguably a very generous upper bound. Because higher mobility reduces the welfare losses due to technological regress, our calibration provides a conservative estimate of the the gain from service-led growth in the presence of spatial mobility.

WA-4 Additional empirical results

WA-4.1 Growth Without Industrialization: Country-Specific Results

In Table WA-1 we report the change in sectoral employment shares and income per capita for 27 developing countries. While there are, of course, idiosyncratic differences across countries, the broad pattern of "growth without industrialization" is observed in most of the developing world.

WA-4.2 Data

In this section, we report additional details on the data, described in Section B-2 in the Appendix.

In Table B-1, we reported the distribution of human capital across time, space and sectors of production. In Table WA-2 we report the same composition when we classify PS and CS workers according to the NIC classification, that is, we allocate workers in wholesale, retail, hotel, restaurants, health, and community services to CS, and workers in financial and business services, transport, and ICT to PS. This classification increases

Region	Change in ... empl. share (1991–2017)				GDP pc Growth	Region	Change in ... empl. share (1991–2017)				GDP pc Growth
	Agricul.	Manufac.	Services	Constr.	(1991–2017)		Agricul.	Manufac.	Services	Constr.	(1991-2017)
India	-0.22	0.01	0.13	0.09	320						
Bangladesh	-0.29	0.03	0.21	0.06	170	Bolivia	-0.15	-0.02	0.13	0.05	239
Brazil	-0.19	-0.02	0.18	0.03	110	China	-0.40	-0.06	0.37	0.08	433
Ecuador	-0.09	-0.03	0.09	0.03	82	Guatemala	0.17	-0.11	-0.03	-0.02	92
Honduras	-0.12	-0.01	0.12	0.00	71	Indonesia	-0.24	0.04	0.16	0.04	189
Jamaica	-0.09	-0.07	0.15	0.01	69	Kenya	-0.08	-0.00	0.07	0.01	76
Cambodia	-0.55	0.16	0.30	0.09	212	Lao People's DR	-0.24	0.04	0.17	0.03	452
Sri Lanka	-0.17	-0.02	0.17	0.02	285	Morocco	-0.04	-0.04	0.08	-0.00	52
Myanmar	-0.28	0.02	0.22	0.04	509	Mongolia	-0.18	-0.00	0.12	0.06	313
Namibia	-0.33	-0.01	0.28	0.06	97	Nicaragua	-0.03	-0.03	0.04	0.02	70
Pakistan	-0.07	0.03	0.03	0.01	71	Philippines	-0.20	-0.02	0.18	0.04	100
Paraguay	-0.11	-0.02	0.10	0.03	149	Thailand	-0.27	0.03	0.22	0.02	190
Tunisia	-0.15	-0.04	0.16	0.04	73	Uganda	-0.06	-0.02	0.06	0.01	119
Viet Nam	-0.33	0.11	0.16	0.07	371	South Africa	-0.13	-0.06	0.15	0.04	43
Developing World	-0.18	-0.00	0.15	0.04	157						

Table WA-1: Growing Like India: 1991–2017. The table reports the change in sectoral employment shares and GDP per capita between 1991–2017 for 27 countries. The employment data comes from the ILO. The data on GDP comes from the Penn World Tables. In the last column we report the averages across 27 developing countries.

the skill content of workers in CS and PS, mostly because it implies that construction workers are not assigned as service workers. However, qualitatively, it is still the case that PS and CS workers are more educated than workers in the manufacturing sector or in agriculture.

	Less than primary	Primary, upper primary, and middle	Secondary	More than secondary
		Aggregate Economy (1987 – 2011)		
1987	66.79%	22.03%	7.99%	3.19%
2011	40.32%	30.10%	18.79%	10.79%
		By Sector (2011)		
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%
		By Urbanization (2011)		
Rural	46.97%	30.00%	16.30%	6.84%
Urban	33.69%	30.30%	21.27%	14.73%

Table WA-2: Educational Attainment. The table shows the distribution of 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.

In Table B-4 in the Appendix we reported the different broad spending categories of the Expenditure Survey. In Tables WA-3 and WA-4 we report the more detailed classification of the consumer service (category 24) and entertainment spending (category 20) categories.

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

Table WA-3: Expenditure items within consumer services. This table reports the detailed expenditure items within the category consumer services (category 24 in Table B-4)

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

Table WA-4: Expenditure items within entertainment. This table reports the detailed expenditure items within the category entertainment (category 20 in Table B-4)

In Table WA-5 we report a selected set of summary statistics for the main variables of interest. In total we have expenditure data for slightly more than 100,000 households. In the first two rows we show the distribution of household expenditure for the case of measuring durable spending at the monthly frequency (the uniform reference period *URP*) and at the annual frequency (the mixed reference period *MRP*). Table WA-5 shows that the dispersion in spending is much higher for the URP case, especially in the right tail. We therefore use the MRP measure as our measure of total expenditure.

Table WA-5 also reports a set of statistics for the distribution of food shares and consumer service spending shares. The full distribution is shown in Figure WA-4. Through the lens of our theory, this dispersion is generated through heterogeneity in income and relative prices.

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

Table WA-5: NSS expenditure survey—Summary statistics. The table reports selected summary statistics from the NSS expenditure survey.

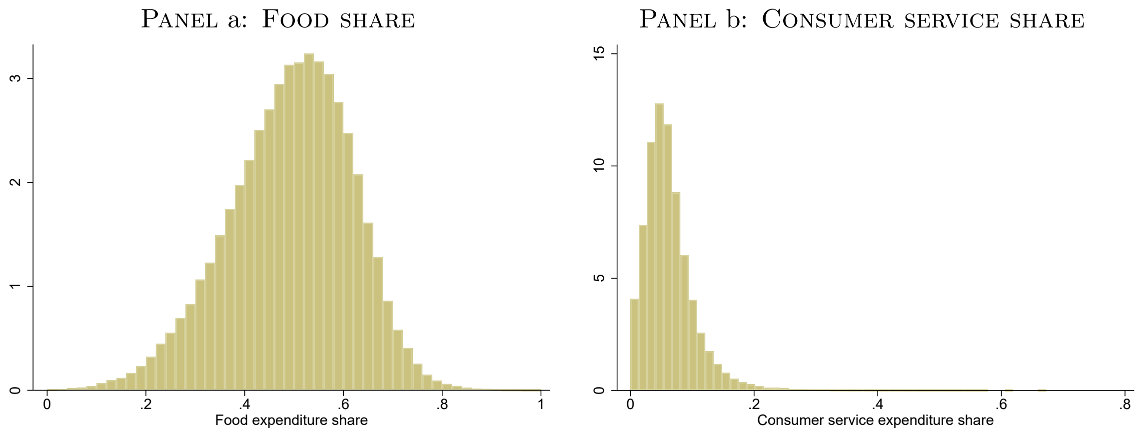

Figure WA-4: Distributon of food and consumer service expenditure shares. The figure shows the unconditional distribution of the expenditure shares for food (left panel) and consumer services (right panel).

For our estimation of the Engel elasticity ε, we ran a specification for the expenditure share on individuals food items. In Table WA-6 we report the cumulative expenditure share on the top ten food varieties in the expenditure survey.

1987	Cumulative Share	**2011**	Cumulative Share
Rice	18.2	Cereal: s.t.	9.1
Milk (liquid)	29.0	Fuel and light: s.t.	16.9
Atta	37.3	Milk & milk products	24.7
Fire-wood and chips	41.9	Milk: liquid (litre)	31.7
Sugar (crystal)	44.7	Rice: o.s.	36.4
Mustard oil	47.2	Vegetables: s.t.	40.2
Ground nut oil	49.5	Edible oil: s.t.	43.3
Arhar (tur)	51.6	Egg, fish & meat: s.t.	46.2
Cooked meals	53.3	Served processed food: s.t.	49.1
Potato	54.9	Wheat/atta: o.s.	51.9

Table WA-6: NSS expenditure survey: expenditure shares of the ten most important food varieties. The table reports the cumulative expenditure shares on the ten most important food categories.

In Table WA-7 we report the offial NIC classification of India and how we aggregate the different subsectors in the six sectors Agriculture, Manufacturing, Construction and Utilities, Services, Information and Communications Technology (ICT) and Public Administration and Education.

In Table WA-8 we summarize our concordance between the different NIC classifications in 1987, 1998 & 2004 and 2008. To ensure comparability over time, we harmonize the sectoral classification at the 2008 level.

To classify employment into PS and CS employment, we rely on the fact that large firms are more likely to sell to firms as opposed to consumers. In Figure WA-5 we

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

Table WA-7: Industrial Classification. The table reports the industrial classifications into six broad sectors.

depict the employment share of PS firms as a function of firm size in the raw data. Among small firms, more than 95% of firms mostly sell to consumers. Among firms with more than 50 employees, almost half of firms sell mostly to other firms.

In Table WA-9 we show that the same pattern is present within 2- and 3-digit industries regardless of whether we use sampling weights. In particular, we regress a dummy variable for whether the firm sells mainly to other firms on different firm size dummies. The coefficients are generally positive and increasing.

To assign construction employment to PS and CS, we first classify industries within construction at the 5-digit level into public and private firms. In Table WA-10 we report our classification. We drop all public subsectors from our analysis. These account for roughly 9.2% of employment in the construction sector.

WA-4.3 Expenditure, Wages and Income per capita

In our main analysis we measure district-level income by average consumption expenditures. In doing so, we hope to better capture the informal sector. However, this meausure is strongly correlated with average wages and independent estimates of GDP pc at the district level. In the left panel of Figure WA-6 we plot the correlation between expenditure per capita and average wages in 2011 as a binscatter plot. In the right panel we perform the same exercise with GDP pc.[16] This data is only available in 2005, hence we report the correlation with average expenditure in the NSS survey

[16] We thank Johannes Boehm and Ezra Oberfield for sharing their data with us.

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

Table WA-8: Concordance between 2-digit industry Classes. The table reports the classification of NIC codes in different years to the broad sectoral categories of Table WA-7.

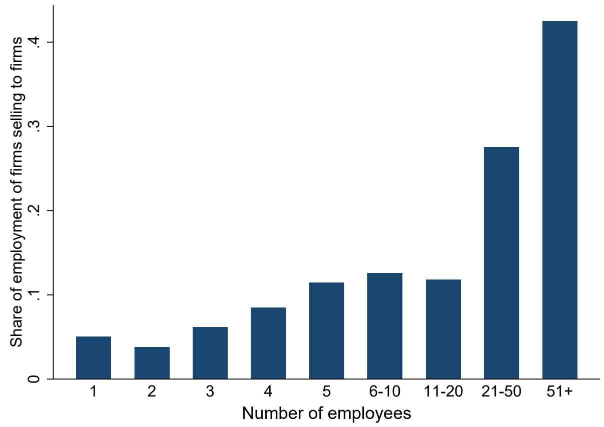

Figure WA-5: PRODUCER SERVICE SHARE BY FIRM SIZE. The figure shows the share of service firms whose main customers are other firms (as opposed to private individuals) with a breakdown by firm size.

of 2004. Figure WA-6 shows that expenditure per capita is strongly correlated with other measures of income per capita.

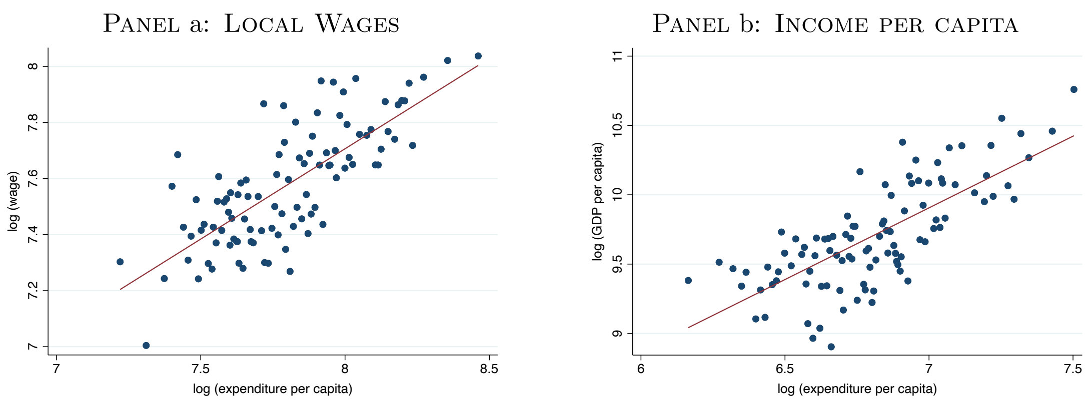

Figure WA-6: EXPENDITURE, WAGES AND GDP PC. In the left panel we show the correlation between expenditure per capita and average wages in 2011 across districts. In the right panel, we show the correlation with GDP per capita in 2005, the only year where this information is available.

WA-4.4 Urbanization and Aggregate Growth

In Figure WA-7 we report the time-series change in the urbanization rate (panel a) and in income per capita (panel b). The urbanization rate is the share of the population living in urban areas according to the definition of the NSS. The NSS defines an urban location in the following way: (i) all locations with a Municipality, Corporation or Cantonment and locations defined as a town area, (ii) all other locations that satisfy the following criteria: (a) a minimum population of 5000, (b) at least 75 percent of the

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

Table WA-9: Corporate Customers and Firm Size. Columns 1 and 2 (3 and 4) control for 2 (3) digit industry fixed effects. Columns 2 and 4 weigh each observation by the sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

male population are employed outside of agriculture, and (c) a density of population of at least 1000 per square mile. This share increased from around 22% in 1987 to 29% in 2010. Income per capita, shown in the right panel, stems from the World Bank. Between 1987 and 2010, income per capita increased by a factor of almost 3.

Urbanization and Income per Capita

For some of our analysis we choose urbanization as our measure of spatial heterogeneity. We do so as a descriptive device and interpret urbanization as a broad proxy for regional economic development. Figure WA-8 shows that there is a strong positive correlation between urbanization and expenditure per capita in the NSS data in 2011.

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

Table WA-10: Classification of the construction sector. The table reports how we classify different subsectors in the construction sector as either public or private sectors.

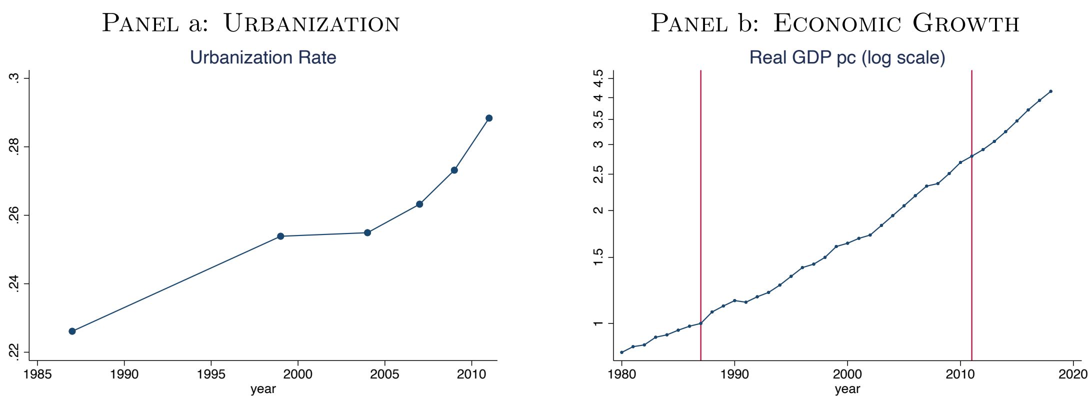

Figure WA-7: Economic Growth in India: 1987 - 2011. This figure shows the evolution of the urbanization rate (Panel a) and income per capita (Panel b). The urbanization rate is the share of population living in urban areas according to the definition of the NSS. Income per capita stems from World Bank.

Spatial Structural Change: Sectoral Income Shares

In Figure B-3 in the main text we report sectoral employment shares as a function of the urbanization rate. In Figure WA-9 we report sectoral income shares by urbanization quintiles in 1987 (Panel a) and in 2011 (Panel b). If anything, the patters we describe in Figure B-3 are more pronounced because earnings are higher in service industries and in cities.

WA-5 The Bootstrap Procedure

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

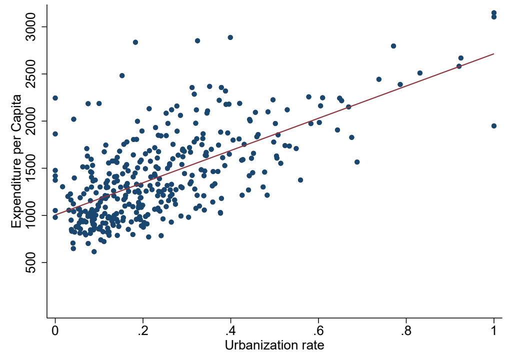

Figure WA-8: Expenditure per capita vs. Urbanization. The figure shows a scatter plot of the average expenditure per capita in the NSS data across district-level urbanization rates in 2011.

Panel a: Sectoral Income by Urbanization (1987)

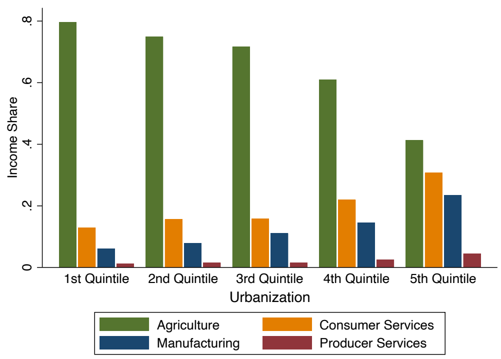

Panel b: Sectoral Income by Urbanization (2011)

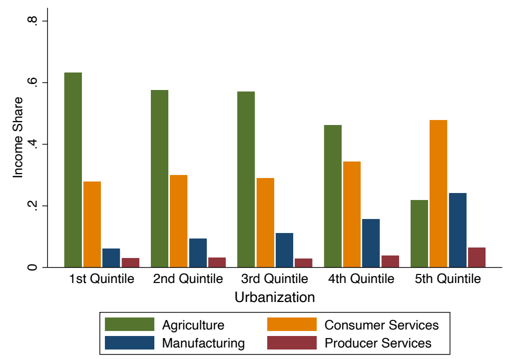

Figure WA-9: Spatial Structural Change in India. The figure plots plots the sectoral income shares by urbanization quintile in 1987 and 2011.

1. From the underlying micro data of the NSS, we draw households randomly with replacement and we sample, within each district, the same number of households as the current dataset.[17]

2. Given this bootstrap sample, we recalculate all statistics used in our accounting procedure, that is, sectoral employment shares, sectoral income shares, and the supply of human capital at the district level.

3. We then redo our entire analysis on this bootstrap sample:

   (a) We re-estimate the structural parameters that rely on this data, that is, the income elasticity ε (by targeting the estimated income elasticity of the

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

expenditure of food reported in Table 3) and the preference parameters ν_F and ω_{CS} (as explained in Section 5),

(b) We re-estimate the productivity fundamentals $\mathbf{A_t}$, and

(c) We calculate our counterfactuals by setting sectoral productivity growth between 1987 and 2011 to zero.

4. This procedure provides us with alternative estimates of the welfare effects and the impact on the structural transformation. Let $\Delta\varpi_r^{q(b)}$, $\Delta\overline{\varpi}_r^{(b)}$ and $\Delta\overline{\varpi}^{(b)}$ denote the individual, regional, and aggregate welfare impact from bootstrap iteration b. Similarly, let $L_{s2011}^{CF_F,(b)}$, $L_{s2011}^{CF_CS,(b)}$ and $L_{s2011}^{CF_I,(b)}$ denote counterfactual employment share in sector s in bootstrap iteration (b) in 2011 if productivity in agriculture (F), CS, and Industry (I) had not grown since 1987. We always use the same choices to treat outliers as in our baseline analysis (see Section C-6).

5. We replicate this procedure B times and hence arrive at the vector

$$\left\{\Delta\varpi_r^{q(b)}, \Delta\overline{\varpi}_r^{(b)}, \Delta\overline{\varpi}^{(b)}, L_{s2011}^{CF_F,(b)}, L_{s2011}^{CF_CS,(b)}, L_{s2011}^{CF_I,(b)}\right\}_{b=1}^{B}. \qquad \text{(WA-29)}$$

In practice we take $B = 200$.

6. From WA-29 we can estimate the distribution of the statistics of interest. For example, the τth quantile of the distribution of aggregate welfare gains, $m_{\Delta\overline{\varpi}}^{\tau}$, can be estimated from the empirical distribution

$$\frac{1}{B}\sum_{b=1}^{B} 1\left[\Delta\overline{\varpi}^{(b)} \leq m_{\Delta\overline{\varpi}}^{\tau}\right] \leq \tau.$$

The quantiles for the other objects of interest are calculated similarly.

7. In the box plots in Figures 6 and 7 we plot the 5%, 25%, 50%, 75% and 95% quantiles of the respective distribution.

Note that, for simplicity, this procedure only captures the sampling variation stemming from the NSS micro data. Hence, we do not, for example, resample firms in the Economic Census or the firm survey to re-estimate the relative weights of PS versus CS employment within the different subsectors of the service sector (see Section B-4).

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

In Figure WA-11 we show the estimated distribution of the welfare losses depicted in Figures 6 and 7. We show the losses attributable to productivity growth in agriculture

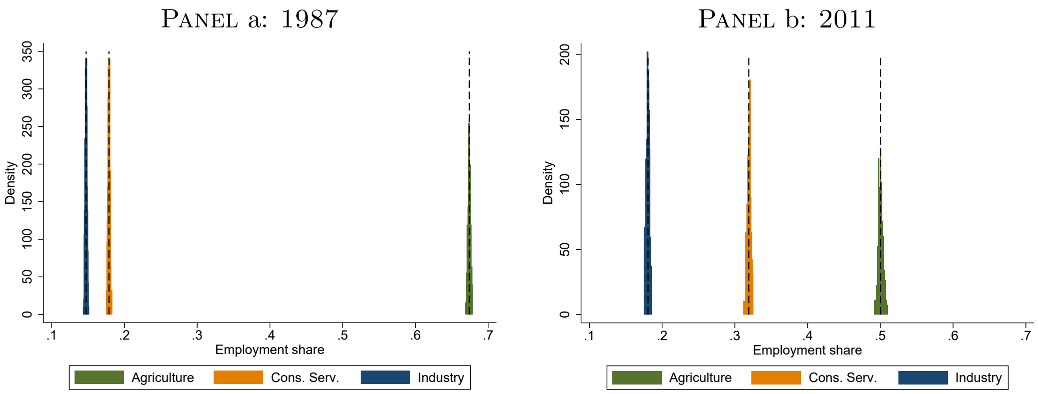

Figure WA-10: Bootstrap Distribution of Aggregate Employment Shares. The figure shows the bootstrap distribution of the aggregate sectoral employment share in 1987 (left panel) and 2011 (right panel). The vertical dashed line corresponds to the empirically observed value.

(Panel a), in CS (Panel b), and in the industrial sector (Panel c). For each case we depict the aggregate welfare losses and the losses for the first and fifth urbanization quintile on the left and for different quantiles of the income distribution on the right. The distributions are well-behaved and do not seem to be driven by extreme outliers.

Panel a: No Productivity Growth in Agriculture

Aggregate Welfare Losses

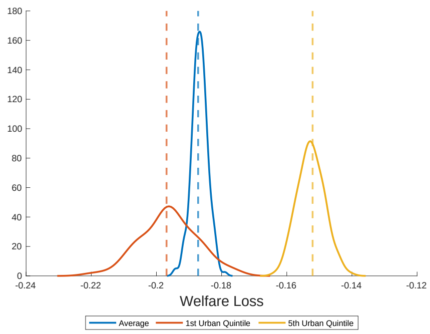

Individual Welfare Losses

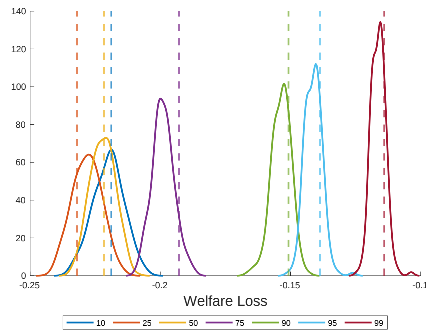

Panel b: No Productivity Growth in CS

Aggregate Welfare Losses

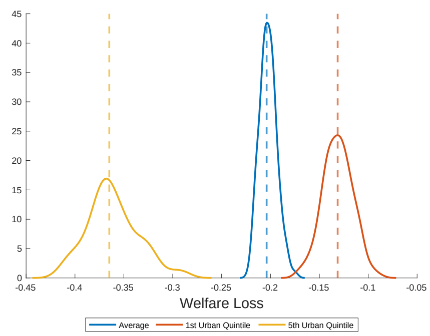

Individual Welfare Losses

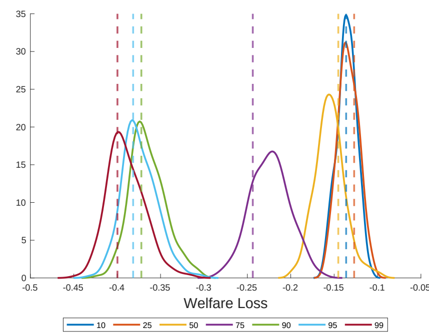

Panel c: No Productivity Growth in the industrial sector

Aggregate Welfare Losses

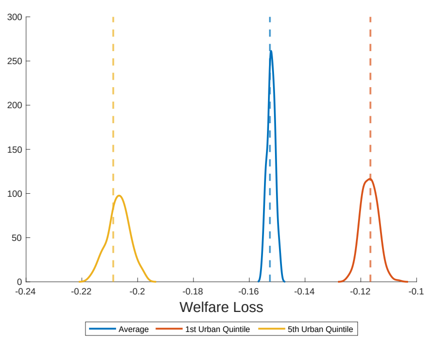

Individual Welfare Losses

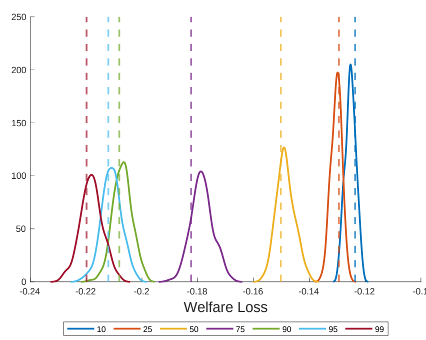

Figure WA-11: Bootstrap Distribution of Welfare Losses. The figure shows the bootstrap distribution of the welfare losses when we counterfactually set sectoral productivity in 2011 to its level in 1987. In panel (a) we shut down productivity growth in agriculture, in panel (b) we shut down productivity growth in CS and in panel (c) we shut down productivity growth in the industrial sector. Within each panel, on the left we show the aggregate welfare losses and the losses for the first and fifth urbanization quintile. On the right we show the losses for the different quantiles of the income distribution.