NBER WORKING PAPER SERIES

GROWING LIKE INDIA:
THE UNEQUAL EFFECTS OF SERVICE-LED GROWTH

Tianyu Fan
Michael Peters
Fabrizio Zilibotti

Working Paper 28551
http://www.nber.org/papers/w28551

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2021, Revised July 2022

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 Sebastian Sotelo and Klaus Desmet for their thoughtful discussions at the NBER SI and the ASSA. We thank Daron Acemoglu, Treb Allen, David Atkin, Fabian Eckert, Reto Foellmi, Doug Gollin, Cormac O'Dea, Peter Klenow, Samuel Kortum, Ramon Marimon, Rachel Ngai, Richard Rogerson, María Sáez Martí, Todd Schoellman, Kjetil Storesletten, Danyang Xie, and seminar participants at the ASSA Meeting 2021, the Cowles Macro Conference, the Minneapolis FED, CREi, Dartmouth, HEC Lausanne, IIES, the NBER Summer Institute, 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Growing Like India: The Unequal Effects of Service-Led Growth
Tianyu Fan, Michael Peters, and Fabrizio Zilibotti
NBER Working Paper No. 28551
March 2021, Revised July 2022
JEL No. O1,O18,O4,O41

ABSTRACT

In many developing countries, structural transformation in recent years takes the form of a rapid rise of services with limited industrialization. Many such services are of a local nature and serve consumers (retail, restaurants, residential real estate, etc.). In this paper, we estimate the welfare effects of this pattern of development across the income ladder. We construct a spatial equilibrium model in which, over time, the expansion of consumer services is both a consequence (income effects) and a cause (productivity growth) of the development process. We estimate the model using Indian household data on sectoral employment and consumption expenditure. We find that productivity growth in non-tradable consumer services was an important driver of rising living standards, accounting for one-third of aggregate welfare gains. However, these gains were heavily skewed toward high-income households living in urban districts.

Tianyu Fan
Yale University
tianyu.fan@yale.edu

Michael Peters
Department of Economics
Yale University
28 Hillhouse Avenue
New Haven, CT 06511
and NBER
m.peters@yale.edu

Fabrizio Zilibotti
Department of Economics
Yale University
28 Hillhouse Avenue
New Haven, CT 06520
and NBER
fabrizio.zilibotti@yale.edu

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 and the employment share of agriculture is down to 42%.

In this paper we argue that productivity growth in the service sector played 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. In urban areas such as Delhi, Mumbai or Bangalore, more than 50% of the labor force is employed in such activities.

To quantify the welfare effects of productivity growth in the provision of CS, 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 restriction 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 appealing for non-tradable CS, where local price indices are not available and measurement issues about quality adjustments loom large.

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 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 bourgeoisie has benefited not only from the availability of better and cheaper goods but also—and to a far larger extent—from the growing local supply of CS that has changed the face of urban life. To the best of our knowledge, ours is the first paper that quantifies the unequal welfare impact of service-led growth for the development process.

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 are of a local nature. Together, these assumptions imply that the local CS sector caters to a particular subset of the income distribution. If, as we find, service-intensive products are luxuries, while goods with a low CS content are necessities, the CS sector provides non-tradable value-added whose main beneficiaries are the rich.

Non-homothetic preferences play a crucial role not only for the welfare analysis, but also in our estimation. The estimation of CS productivity is subject to an identification problem: An increase in the local employment share of CS could stem from local income growth coupled with nonhomothetic demand, and 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 demand for CS value-added.

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 has recently been popularized in the growth literature by Boppart [2014]. PIGL has two important properties, both of which feature prominently in our analysis. First, it allows 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 enable us 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 one derived from preferences over final products depends, in general, on the entire input-output matrix. We formally establish that—under PIGL preferences—the key parameter governing the income elasticity is common to the aggregate demand system for sectoral value-added and the final expenditure demand system at the individual level. This allows us to use micro-data on household expenditure to discipline the aggregate income elasticity for CS value-added.

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 about 400 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, to a large extent serve as inputs to the industrial sector and as such, their value-added can be shipped across locations.[1] 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 yields three 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 service-intensive products 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 since 1987. Third, and most importantly, the welfare impact of service-led growth was strikingly unequal. Productivity growth in CS was the main source of welfare gains for richer households, especially those living in urbanized districts. By contrast, for poorer households from rural districts, improvements in living standards hinged on productivity growth in agriculture.

In terms of positive implications, we also document that productivity growth in CS was a key driver of structural change. In particular, India's pattern of growth without industrialization is a direct consequence of service-led growth and agricultural employment would have declined far less in the absence of rising productivity in CS. Given that this reallocation from agriculture into the service sector is common in many developing countries today, we expect rising CS productivity to be a common feature of the development process.

We carry out the main analysis under a set of stark assumptions aimed to retain

[1] While the assumption that CS are consumed locally is stark, it 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.

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 the assumption that skills are perfect substitutes and assume, instead, that the labor services provided by people with different educational attainment are imperfect substitutes. Moreover, we allow skill intensity 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 main results and the broad picture are consistent and robust: productivity growth in CS played a key role for Indian growth since 1987, especially for the urban rich.

Related Literature: Our paper contributes to the literature on the structural transformation including, among many others, Kongsamut et al. [2001], Ngai and Pissarides [2007], Herrendorf et al. [2013, 2014, 2020], or Gollin et al. [2014].

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 [2019] argue that in more recent years ICT has been a major source of productivity growth. Their view is echoed by Eckert et al. [2020]. 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

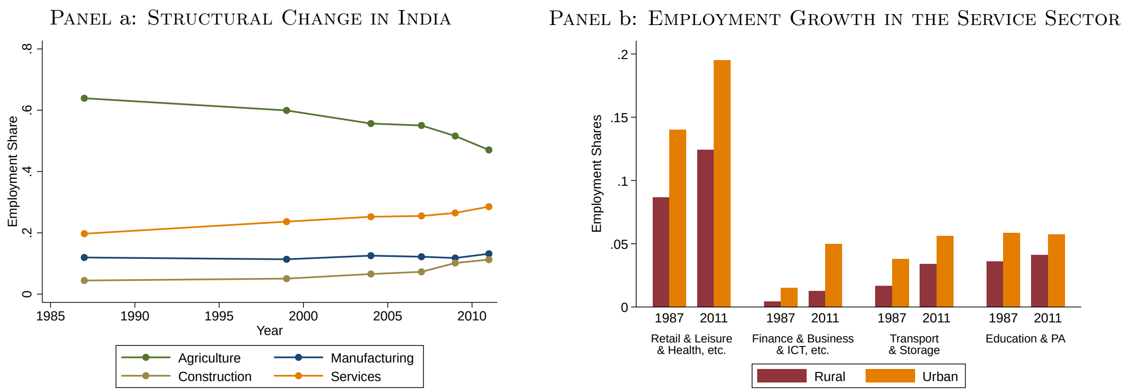

Figure 1: Structural Change in India: 1987–2011. The left panel shows the evolution of sectoral employment shares. The right panel shows employment shares for different service industries (see footnote 2 for details), separately for rural and urban districts. The figure is based on micro data from the NSS (see Section 4).

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].

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] use a class of generalized CES preferences related to 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.

2 Structural Change towards Services in India

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

The right panel of Figure 1 documents that a large part of the expansion in service employment originated in services that facilitate consumers' access to final consumption. We decompose the service sector into four subsectors.[2] The first group of service industries, which serve mostly consumers, employ almost 60% of all Indian service workers in 2011. Employment in these services grew significantly between 1987 and 2011. The second group of industries, which sells a significant part of their services to industrial firms, also grew significantly but only accounts for a tenth of service employment. For instance, ICT, notoriously a fast-growing industry in India, accounts for less than 1% of total employment in 2011. Transport services, which serve both consumers and industries, also grew. Finally, education and PA are mostly government-run activities and their employment share is constant over time.

Figure 1 also shows there are stark differences across local labor markets. We split India into rural and urban districts, broken down so that half of the workers belong to each type of district. Service activities are more prevalent in urban areas than in rural ones. While this is especially apparent in business-oriented services, non-tradable services catering to consumers are also substantially larger in cities and grew fast.

These patterns are representative for most of today's developing world. In Table 1 we report the change in sectoral employment shares and GDP per capita between 1991 and 2017 for a subset of countries and for the average of 27 developing countries. Most developing countries indeed grew like India and experienced falling agricultural employment without industrialization and fast employment growth in services.[3]

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

[3] In Appendix Section B-1, we report the results for each country. There, we also document that within the service sector, CS are much more prominent in poor countries.

Region	Change in ... empl. share (1991–2017)				GDP pc growth (1991–2017)
	Agricul.	Manufac.	Services	Constr.	
India	-0.22	0.01	0.13	0.09	320 %
Bangladesh	-0.29	0.03	0.21	0.06	170 %
Brazil	-0.19	-0.02	0.18	0.03	110 %
Kenya	-0.08	0.00	0.07	0.01	76 %
Philippines	-0.20	-0.02	0.18	0.04	100 %
Developing World	-0.18	0.00	0.15	0.04	157 %

Table 1: Growing Like India: 1991–2017. The table reports the change in sectoral employment shares and GDP per capita between 1991–2017 for selected 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.

3 Theory

We consider a model with R regions. Within each region there are three broad sectors of activity: 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 following important assumption about tradability: while food and goods are tradable across regions subject to iceberg costs, CS must be locally provided.[4] We assume that markets are frictionless and competitive.

In our benchmark model, we assume that the aggregate supply of labor is inelastically provided 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 all of these dimensions.

3.1 Technology

Each region r produces a measure one continuum of non-traded differentiated final products that enter consumers' utility. Each product is produced using the two tradable physical inputs—food and goods—and local CS workers.[5] For instance, a restaurant meal is a combination of tradable food and kitchen tools, and the 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

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

[4] 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.

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

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: $\sum_s \lambda_{ns} = 1$.[6]

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 content of CS (the retail store). A restaurant meal also requires food but has a larger CS content. Finally, personal services like haircuts or nanny services consist almost entirely of CS ($\lambda_{nCS} \approx 1$).

The 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

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

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

Nontradable CS vs tradable PS Equation (1) highlights the special role of the CS sector in our theory: its value-added is combined with that of tradable commodities to turn the latter into final products that local consumers can enjoy. We think of tradability as a critical difference between PS and CS. While CS value-added can only be consumed locally, the value-added from producer services like ICT or finance can be embodied in goods and is ultimately tradable.

When mapping the model to the data, we therefore include the value-added of PS in the industrial sector.[7] More formally, we let $H_{rGt} = H_{rMt} + H_{rPSt}$ in (3) include labor services provided in both the manufacturing and PS sector. We want to highlight that this specification does *not* restrict manufacturing and PS workers to

[6] 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.

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

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.[8] This allows cities like Delhi or Bangalore to have a comparative advantage in tradable services like finance or ICT, and they can export the value-added of such services to the rest of India (or, as part of our extension in Section 7, internationally).

3.2 Preferences and Demand System

A key aspect of our analysis is the demand system for the value-added created in the local CS sector. Following Boppart [2014] we assume consumers' preferences over the continuum of final products are in the PIGL class. PIGL preferences have three appealing properties for our purposes. First, they have simple and transparent aggregation properties that allow us to take a spatial demand system to the data. Second, they provide a simple mapping of preferences over final goods into preferences over value-added. Third, they allow us to derive analytic expressions for individual and aggregate welfare.

PIGL preferences do not have an explicit utility representation 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" highlights that this 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 $B(\mathbf{p}_r) = \exp\left(\int_{n=0}^{1} \beta_n \ln p_{rn} dn\right)$ 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 $\int_0^1 \kappa_n dn = 0$.

[8] 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, $A_{rGt} = \left(A_{rMt}^{\varsigma-1} + A_{rPSt}^{\varsigma-1}\right)^{1/(\varsigma-1)}$.

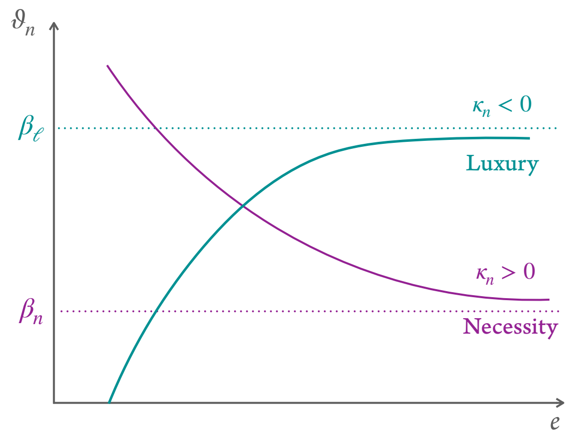

Figure 2: ENGEL CURVES. The figure shows the good-specific expenditure share as a function of income e (see (5)).

Roy's Identity implies that 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 with a non-homothetic adjustment. Moreover, spending e and prices $\mathbf{p}$ conveniently enter through a single summary statistic that resembles a notion of real income.

In Figure 2 we depict the expenditure share as a function of spending e. The parameter β_n determines the asymptotic expenditure share 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 emerge as 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 necessary 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 the price of final good n in region r is 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} is the wage per efficiency unit of human capital in region r. Equation (2) implies that the prices of 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. Plugging the expressions for p_{rnt} into (4) yields a representation of consumers' preferences over sectoral value-added aggregates.

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

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

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*

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

The associated expenditure shares over sectoral value-added are given by

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

Proof. See Appendix A-1. □

Proposition 1 plays an important role in our analysis. First, it establishes that the indirect utility function defined over value-added also falls in the PIGL class and has the same functional form as the corresponding expression over final products (4). Moreover, the expenditure share on sectoral value-added ϑ_{rst} in (9) features the same Engel elasticity ε as in (5). This property of PIGL preferences enables us to estimate ε from microdata for household expenditure shares on final products and then use it in the demand system defined over sectoral value-added.[9]

Proposition 1 also highlights that 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

[9] In Section A-1 in the Appendix we also derive the analogue of (9), if the production function for final products combines CS, food and goods in a CES fashion.

generally if λ_{ns} were orthogonal to κ_n for all s, the demand for sectoral value-added would be homothetic and independent of prices (i.e., Cobb Douglas) even though the underlying demand for final products was nonhomothetic.

Third, Proposition 1 shows that the regional CS productivity index A_{rCS} is akin to an average productivity of the technologies of all final products (weighted by their local CS content and their demand share) and acts as a sufficient statistic for the local CS sector. Because preferences are nonhomothetic and CS are provided locally, productivity growth has heterogeneous welfare effects. If goods with a high CS content have a high income elasticity, productivity growth in CS is skewed toward the rich.[10] Moreover, CS productivity growth mostly 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 are shared with other locations through trade.

3.2.2 Heterogeneity and Aggregate Demand

Proposition 1 characterizes demand at the individual level. We now derive the aggregate demand system. Suppose individuals differ in their human capital that they supply to all sectors. Individual h's income is then given by $e_{rt}^h = q^h w_{rt}$, where q^h is the number of efficiency units of labor. Let $F_{rt}(q)$ denote the distribution function of q in region r at time t—which we will 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

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

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

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

[10] 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 $\mathbf{P_{rt}}$, $\partial \ln e(\mathbf{P_{rt}}, V)/\partial \ln P_{rst} = \vartheta_{rst}(e, \mathbf{P_{rt}})$.

that of a consumer in region r who earns the average income $\mathbb{E}_{rt}[q]\, w_{rt}$ and has the inequality-adjusted preference parameter $\overline{\nu}_{rst}$ in (11). Crucially, the Engel elasticity parameter ε is the same as at the individual level.

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

$$\overline{\nu}_{rst} = \overline{\nu}_{s} = \frac{\zeta^{\varepsilon}(\zeta-1)^{1-\varepsilon}}{\zeta+\varepsilon-1}\nu_{s}. \tag{12}$$

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 $\mathbb{E}_{rt}[q]\, w_{rt} = \frac{\zeta}{\zeta-1}\underline{q}_{rt} w_{rt}$.

3.2.3 Welfare and Inequality

The aggregation properties of PIGL are especially handy for the welfare analysis, which is the core of our contribution. In particular, define the utilitarian welfare function at the regional level as $\mathcal{U}_{rt}(w_{rt}, \mathbf{P_{rt}}) \equiv \int \mathcal{V}(qw_{rt}, \mathbf{P_{rt}})\, dF_{rt}(q)$. Plugging in the indirect utility function in (7) yields

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

where $\nu_s^{\mathcal{U}} \equiv \nu_s \times ((\zeta-\varepsilon)(\zeta-(1-\varepsilon)))/(\zeta(\zeta-1))$. 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 another 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:*

$$w_{rt}H_{rCSt} = \left(\omega_{CS} + \overline{\nu}_{CS}\left(\frac{A_{rCSt}^{\omega_{CS}}\mathbb{E}_{rt}\left[q\right]w_{rt}^{1-\omega_{CS}}}{P_{rFt}^{\omega_F}P_{rGt}^{\omega_G}}\right)^{-\varepsilon}\right)w_{rt}H_{rt}, \quad (14)$$

where P_{rFt} and P_{rGt} are given by (6).

2. *Market clearing for tradable goods:*

$$w_{rt}H_{rst} = \sum_{j=1}^{R}\pi_{rsjt}\left(\omega_s + \overline{\nu}_s\left(\frac{A_{jCSt}^{\omega_{CS}}\mathbb{E}_{jt}\left[q\right]w_{jt}^{1-\omega_{CS}}}{P_{jFt}^{\omega_F}P_{jGt}^{\omega_G}}\right)^{-\varepsilon}\right)w_{jt}H_{jt}, \quad (15)$$

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

Proposition 2 characterizes the sectoral employment allocations and factor prices across space. The contrast between Equations (14) and (15) reflects the tradable nature of food and goods and the non-tradable nature of CS. The demand for CS value-added hinges both on local income and local CS productivity. For instance, the retail sector could be large in urban districts both because local consumers are, on average, more educated and richer, and 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 composite parameters $\overline{\nu}_s$ and ω_s and the aggregate CS index A_{rCSt}. They do not separately depend on the preference parameters defined over final consumption 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 employment characteristics and households' location of residence. We use data for 1987 and 2011. The NSS yields measures of average consumption (income) and sectoral employment shares at the district-year level. To measure income, we proxy earnings by average expenditure. We prefer this measure to direct information on wages to also capture income from informal employment.

Consistent with our theory, we measure employment shares in four sectors: agriculture, manufacturing, PS, and CS. For agriculture and manufacturing, we follow the ISIC classification. For services we exclude from our analysis some 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. However, in Section 7 we show that we also find an important role for the service sector when we merge construction with the industrial 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 and we use it to estimate the Engel elasticity ε. The EC covers all establishments engaged in the production or distribution of goods and services in India. It covers all sectors except crop production and plantation and collects information on each firm's location, industry, and employment. It contains approximately 24 million and 60 million establishments in 1990 and 2013, respectively. The relatively unexplored 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.[11] Finally, we use the ETD to measure the average relative price of agricultural goods.

Geography: To compare spatial units over time, we create a time-invariant definition of Indian districts.[12] Because the boundaries of several districts changed over time, we harmonized them using GIS software, relying on maps for the years 1987, 1991, 2001,

[11] 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 Section B-2.3 in the Appendix we compare the Service Survey with the EC and document that it is indeed representative of the distribution of firm size in India.

[12] Appendix B-3 describes in detail how we construct this crosswalk.

	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 2: 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.

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: A key aspect of our analysis is to split employment in the service sector into CS and PS. To do so, we combine information from the EC and the Service Survey. To capture the distinction in tradability, we assign firms to the CS sector if they sell to consumers and to the PS sector if they sell to other firms.[13] Ideally, we would base our analysis on firm-level input-output matrices. To the best of our knowledge, this information is not available in India. 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. We could thus, in principle, calculate the share of employment in every service industry-district cell distinguishing between firms selling to other firms and those serving consumers. However, the Service Survey contains too few firms to precisely estimate these employment shares in 360 districts and 6 subsectors. We therefore rely on the fact that the share of firms selling to other firms is highly correlated with firm size. Table 2 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 2 in the following way. First, we estimate the CS employment share by firm size for different service industries.[14] We then 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}

[13] We consider this a conservative choice because many firms might themselves sell to local consumers.
[14] We split the service sector into eight categories: "Retail and wholesale", "Hotels and restaurants", "Transport", "Finance", "Business services", "Health", "Community services", and "ICT".

	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 3: 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.

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 3 we report the resulting allocation of employment to CS. At the aggregate level, our procedure allocates 89% of service employment to CS rather than PS.[15] 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.

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. Given the sample size, splitting the destination of construction activities is possible only at the national, not the district, level. We refer the reader to Appendix B-4.3 for details.

In Section 7, we show that our results are robust to alternative measurement strategies. In particular, we report our results when we allocate ICT and business services entirely to PS and when we split service employment according to aggregate Input-Output-Tables. This reduces the share of CS employment to around 80% but leaves the qualitative results of our analysis unchanged.

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. At the sectoral level, we rely on earnings, which reflect differences in the use

[15] To corroborate our results, we also measured aggregate employment from the EC 2013. In the EC, industries such as 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. Note that even these sectors in part serve consumers as many lawyers (who are part of the business service industries) and banks sell their services to households.

of effective units of labor rather than bodies. To measure the distribution of human capital across districts, we follow the approach in the development accounting literature and leverage data on the regional distribution of schooling, together with an estimate of the Mincerian returns to schooling ρ (see Section 5.1 below).

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.

As we show in Appendix Table B-1, it is important to allow for human capital differences across years, sectors, and space. First, the level of schooling increased markedly between 1987 and 2011, which is itself a source of growth. Second, educational attainment differs vastly across sectors. That agriculture is the least skill-intensive industry and educational attainment is the highest in PS is not surprising. More interestingly, the CS sector also employs lots of educated individuals and is more skill-intensive than the manufacturing sector. Finally, there are large spatial differences whereby city dwellers are much more educated than the rural population. By explicitly measuring differences in human capital across time and space, we refrain from attributing these differences to changes in TFP.

5 Estimation: Identification and Results

We now turn to the estimation of the model. Our approach is in the tradition of development accounting; see, e.g., Caselli [2005], Hall and Jones [1999], and Gancia et al. [2013]). Whereas these studies infer productivity at the country level from an aggregate production function, we 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 parameters describing preferences 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 the 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 the expenditure for CS expands—as income rises. To estimate ε, we thus use the cross-sectional relationship between household income and expenditure shares on food.

In general, it is not legitimate to use expenditure data to infer structural parameters of the demand system for value-added aggregates. However, Proposition 1 establishes that, under PIGL preferences, the aggregate demand system for sectoral value-added and the individual demand system for final expenditure share the same elasticity parameter ε. With this in mind, let $\mathcal{F} \in [0, 1]$ denote the subset of the product space $[0, 1]$ that contains 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:

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

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 characteristics 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 4 reports the estimation results. We cluster standard errors at the region level to account for the correlation of spending shares through regional prices. The first column refers to a specification that, in addition to district fixed effects, only controls for whether the household lives in an urban or rural area (within districts), a full set of fixed effects for household size, and the number of workers within the house-

	Dependent variable: ln(food expenditure share)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\ln e$	-0.332***	-0.321***	-0.313***	-0.334***	-0.395***		
	(0.008)	(0.007)	(0.008)	(0.066)	(0.013)		
$\ln e\times$ 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%)		✓	✓	✓	✓	✓	✓
Addtl. Controls			✓	✓	✓	✓	✓
IV					✓		
N	101654	91495	91447	1129730	85923	91447	91447
R^2	0.476	0.425	0.437	0.635	0.307	0.446	0.439

Table 4: Income elasticity for food. The table shows the estimated coefficient ε of the regression (18). 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 fourth column. Standard errors, clustered at the district level or two-way clustered at the district and food item level (col 4), in parentheses. *** p<0.01, ** p<0.05, * p<0.1

hold. We estimate 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 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 grain from the Indian government. The estimated value for the elasticity is very similar to the estimates in columns (1) and (2).

In column (4), 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.[16] 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.

[16] More formally, in column 4 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.

In column (5), we present the results from an IV regression addressing concerns about measurement error that could bias the estimated Engel elasticity. We instrument total expenditure with a full set of three-digit occupation fixed effects.[17] 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 through their effect on income. The IV estimate of 0.395 is indeed 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 (4). 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 (6) of Table 4, we investigate this issue more formally by allowing different elasticities for households above and below the median income. The estimated elasticity is indeed somewhat larger for high-income households.[18] Finally, in column (7) we analyze the extent to which the elasticity differs between rural and urban localities.[19] While urban locations have higher elasticities, the differences are quantitatively small.

For our baseline analysis we take the Engel elasticity ε to be equal to the IV estimate 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 are larger if we rely 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.[20] In Section 7, we show that our results are qualitatively robust to all estimates reported in Table 4 and we also report the results of an alternative calibration of ε that does not rely on expenditure data.

Other Preference Parameters: We estimate the six remaining parameters of the demand system, ω_s and $\overline{\nu}_s$, directly from the equilibrium conditions.[21] In Appendix

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

[18] Allowing for the elasticity to vary with income in the theory would take us outside of the class of PIGL preferences. We leave to future research generalizations in this direction. In this paper, we maintain the assumption of a constant ε and study a range of values in Section 7.

[19] We define urban locations as the ones in the top quartile of the distribution of urbanization.

[20] Suppose subsistence consumption accounts for a large share of food consumption among poor rural households and is unmeasured. As these households get richer, they would purchase a higher share of their food consumption in the market, biasing the Engel elasticity towards zero.

[21] The market-level demand system depends on the aggregate preference parameters $\overline{\nu}_s$ which are

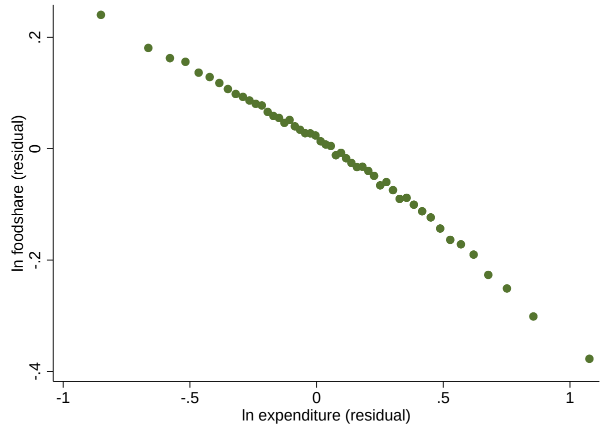

Figure 3: Engel curves in India. The figure shows a binscatter representation of the residual of a regression of the log expenditure share on food item j in region r on region-product fixed effects against the residual of a regression of the log income (total expenditure) on the same set of fixed effects. The slope coefficient of this plot yields the Engel elasticity. Cf. regression in column 5 of Table 4

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 achieve 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 5 we report the resulting estimates. We view the implied 70% asymptotic expenditure share on CS as reasonable.[22] For instance, the value-added

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.

[22] 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

Parameter		Target	Value
Preference parameters	ε	Engel elasticity	0.395
	ω_F	Agricultural spending share US	0.01
	ω_{CS}	Agricultural employment share 2011	0.696
	ω_G	Implied from $\sum_s \omega_s = 1$	0.294
	$\overline{\nu}_F$	Agricultural Employment share 1987	1.258
	$\overline{\nu}_{CS}$	Normalization	-1
	$\overline{\nu}_G$	Implied from $\sum_s \overline{\nu}_s = 0$	-0.258
	σ	Set exogenously	5
Skill parameters	ρ	Mincerian schooling returns	0.056
	ζ	Earnings distribution within regions	3

Table 5: STRUCTURAL PARAMETERS. The table summarizes the estimated structural parameters. The details of the estimation are discussed in the text.

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 production and PS) is 30%. Moreover, $\overline{\nu}_G = -0.258$, which implies that industrial goods are also luxury goods, although their income elasticity is smaller than for CS.

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 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].[23] Given this 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 tail of the income distribution within-regions. This procedure yields an estimate of $\zeta \approx 3$ (see Appendix C-1). With this estimate at hand, we can also compute the lower

proportionally between the other two sectors. In practice, this issue is inconsequential because these districts account altogether for a mere 0.15% and 0.23% of the total valued added of India in 1987 and 2011, respectively.

[23] In Section 7, we show that our results are robust to using a higher return to education.

bound $\underline{q}_{rt}$ form $\mathbb{E}_{rt}[q_i] = \frac{\zeta}{\zeta-1}\underline{q}_{rt}$.

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

Given the structural parameter vector $\mathbf{\Omega}$, data on local wages and sectoral employment allocations, as well as time-series data on relative prices and aggregate income, the equilibrium conditions uniquely identify a set of local productivity fundamentals $\mathbf{A_t}$. We summarize the methodology to estimate $\mathbf{A_t}$ in this section, referring the interested reader to Appendix A-2 for details.

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

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

Equation (20) highlights the role of income effects and service-led productivity growth. CS employment depends on the local supply of skills ($\mathbb{E}_{rt}\left[q\right]$), local wages ($w_{rt}$), the local prices of tradable goods (P_{rFt} and P_{rGt}), and local productivity (A_{rCSt}). Controlling for the level of human capital, local prices, and the equilibrium wage, CS productivity A_{rCSt} is increasing in the observed employment share H_{rCSt}/H_{rt}.[24] Conversely, holding the employment share H_{rCSt}/H_{rt} constant, CS productivity A_{rCSt} is decreasing in both human capital and factor prices. From (20) we can uniquely solve for A_{rCSt}.

This structural decomposition of the observed variation in CS employment shares into the part that is service led (i.e., A_{rCSt}) versus the part that is driven by income effects (i.e., $\mathbb{E}_{rt}\left[q\right] w_{rt}^{1-\omega_{CS}}$ and $P_{rFt}^{\omega_F} P_{rGt}^{\omega_G}$) is a key step of our equilibrium accounting methodology. Note in particular that our estimates of A_{rCSt} do not rely on data on price deflators in the CS sector. We view this as an advantage given the notorious difficulty in measuring the price of non-tradable services.

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

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

[24] Recall that CS are a luxury, that is, $\nu_{CS} < 0$ and $\frac{H_{rCSt}}{H_{rt}} < \omega_{CS}$.

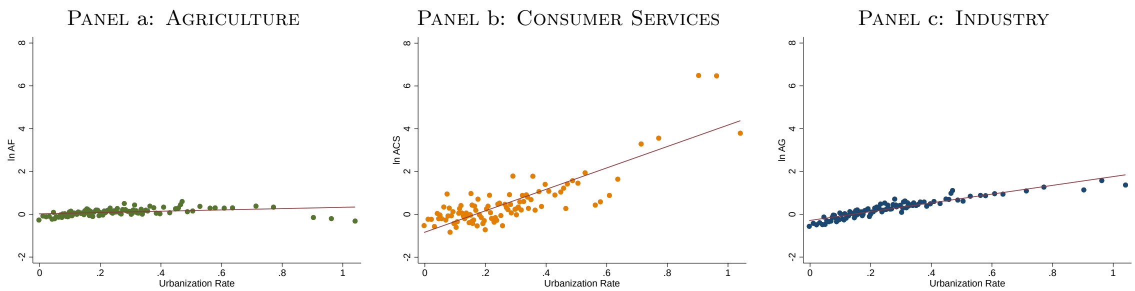

Figure 4: Estimated Sectoral Productivities. The figure shows a binscatter plot of the estimated sectoral labor productivities in agriculture, CS, and industry across urbanization-rate bins. Each plot is constructed by pooling the estimates for 1987 and 2011 after absorbing year effects.

Relative productivity A_{rs}/A_{js} is driven 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 is a correction term that summarizes the possibility that a location can have a high employment share in tradable goods not because of high efficiency but rather because it is close to centers of demand.

Equations (20) and (21) determine the distribution of sectoral productivity across locations. 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. As we discuss in more detail in the Appendix, we target two aggregate moments to achieve identification. First, we target a 2.8 growth factor for real income per person, which matches real GDP per capita growth according to the World Bank (WDI). Second, we target the change in the price of agricultural goods relative to industrial goods as reported in the ETD.[25] 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}.

Results Figure 4 summarizes the cross-sectional pattern of our productivity estimates by displaying a binscatter plot of A_{rst} as a function of the urbanization rate. The relationship between productivity and urbanization is increasing for CS (Panel

[25] We measure real GDP in terms of the numeraire industrial good. Because of nonhomothetic preferences, we cannot define a standard aggregate consumption price index. All sectoral average prices discussed in this section are constructed as weighted averages across Indian districts in a specific year, using the districts' expenditure shares as weights.

[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 1.42.

	Sectoral Productivity Growth					
	10th	25th	50th	75th	90th	Aggregate
Consumer Services (g_{rCS})	-1.3	0.4	2.6	6.3	11.0	4.6
Agriculture (g_{rF})	0.4	1.2	1.9	2.7	3.3	2.1
Industry (g_{rG})	1.8	2.6	3.4	4.3	5.1	3.8

Table 6: 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 expenditure-weighted average in 2011.

(b)) and in the industrial sector (Panel (c)). For agriculture, the relationship is relatively flat and slightly hump-shaped. The declining portion corresponding to districts with an urbanization rate above 50% likely reflects the scarcity of land (a factor of production from which we abstract) in urban areas.

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 of an abundance of human capital; it also reflects high CS productivity. Among the tradable goods, productivity is significantly more dispersed in the industrial than in the agricultural sector. To understand why, note a district's relative productivity reflects its sectoral earning share relative to its skill price (see (21)). The "compressed" productivity distribution in agriculture reflects the observation that wages are negatively correlated with the employment share of agriculture across districts. By contrast, wages are positively correlated with the employment share of industry, implying a wider productivity dispersion.

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 6. Two facts are salient. First and foremost, productivity in the CS sector grew in the vast majority of districts. In the median region, CS productivity grew by 2.6% annually between 1987 and 2011. This is less than productivity growth in the industrial sector and higher than in agriculture. Second, productivity growth was unequal across space, particularly so in the CS sector. 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 4.6% and hence exceeded productivity growth in the two tradable sectors.[27]

[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 the Appendix, where we also report

Change in Price Index		Annual Productivity Growth				
Agricult.	Retail & Hospitality	Agricult.	Manufac.	Mining	Finance & Business	Retail & Hospitality
52%	24%	2.6%	5.3%	4.2%	4.1%	4.2%

Table 7: ETD Price Deflators and Productivity Growt. The table reports changes in sectoral price deflators (relative to manufacturing) and 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.

In Section C-3 in the Appendix we analyze the cross-sectional variation in productivity growth in more detail. In particular, we show that productivity growth is positively correlated with the urbanization rate in 1987. This correlation is also the reason why the expenditure-weighed average of productivity growth exceeds the growth experience of the median locality.

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-4.

Relative Price Indexes: Our methodology allows us to recover disaggregated productivity estimates for the entirety of India. While we are not aware of alternative productivity growth estimates at the sector-region level, the ETD provides estimates of aggregate price deflators and nationwide productivity growth for 12 aggregate sectors. In terms of our theory, their sector "Trade, restaurants, and hotels" is closest to our notion of the CS sector.

In the first two columns of Table 7 we report the changes in the sectoral price indices for agriculture and the retail sector relative to the manufacturing sector between 1990 and 2010. Agricultural prices rose by 52% and prices in the CS sector rose by 24%. While the relative price growth of agricultural products is a target of our estimation, and hence matched by construction, the relative price growth of CS is not. When we compute an aggregate relative price index for CS in our model as the expenditure share-weighted average of regional prices, we find an increase of 29%, which is close to the ETD figure.

Nationwide Sectoral Productivity Growth: In the remaining columns of Table 7, we report annual sectoral productivity growth according to the ETD. These estimates

robustness results for these choices (see Section C-5).

Urbanization	Elasticities of substitution			Income elasticities		
quantile	Agr. & CS	Ind. & CS	Agr. & Ind.	Agr.	CS	Ind.
1 (Rural)	1.6	0.6	1.2	0.7	1.6	1.2
5 (Urban)	1.2	0.9	1.1	0.7	1.2	1.1

Table 8: 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.

are based on value-added per worker growth and hence not directly comparable to our estimates of physical productivity. Nevertheless, they confirm the important role of the service sector for Indian growth. Productivity growth in the Indian retail sector was 4.2% and hence very similar to productivity growth in business services. Productivity growth in manufacturing was higher, growth in the agricultural sector was appreciably slower. These figures are qualitatively in line with our findings, although our model predicts slightly faster growth in the service sector—see Table 6. In Section 7 we present an alternative calibration strategy that explicitly targets the aggregate productivity growth estimates reported in Table 7.

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.[28] In Table 8 we report the Allen-Uzawa elasticities of substitution between each pair of sectoral value-added and the respective spending elasticities. Our estimates imply that CS and industrial goods are complements, with an elasticity of substitution ranging between 0.6 and 0.9, that agricultural and CS value-added are substitutes with an elasticity between 1.2 and 1.6, and that agricultural and industrial value-added are also substitutes, but with a smaller elasticity.

We find these results economically plausible. As the (quality-adjusted) price of CS-intensive restaurants declines, individuals might 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 also heavily rely on CS. They are

[28] As we show in Section A-3 in the Appendix, the elasticity of substitution between sectors s and k is given by $EOS_{sk} = 1 - \varepsilon \frac{(\vartheta_s - \omega_s)(\vartheta_k - \omega_k)}{\vartheta_s \vartheta_k}$ and the spending elasticity is given by $\frac{\partial \ln \vartheta_s e}{\partial \ln e} = 1 - \varepsilon \frac{\vartheta_s - \omega_s}{\vartheta_s}$.

also qualitatively consistent with the common wisdom in the literature. Duernecker et al. [2017] estimate services and goods to be complements. Comin et al. [2021] use US data to estimate a common elasticity of substitution and find evidence of complementarity. 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. Herrendorf et al. [2014] also estimate sectoral value-added to be complementary, even though their preferred estimates are closer to the case of Leontief preferences.

In terms of spending elasticities, we estimate CS to be a strong luxury, industrial goods to be a weak luxury and agricultural output to be a necessity, which is again consistent with the results in Comin et al. [2021]. Quantitatively, they for example find that the spending elasticities for Tanzania are 0.57, 1.15 and 1.29. They also show that spending elasticities of CS and industrial goods are lower in richer countries (which is consistent with our results across urbanization quantiles). Moreover, as in Boppart [2014], Comin et al. [2021], and Alder et al. [2022], Table 8 also shows that our model generates sustained income effects, with agricultural value-added being a necessity even in urban districts where individuals are rich.

As an additional analysis of nontargeted moments, in Appendix Section C-4 we estimate the relationship between income and the spending share on CS-intensive goods. Specifically, we estimate the same specification as in (18) except that we use households' expenditure share on services as the dependent variable. These expenditures include, for example, domestic servants, barber shops, or tailor services. We also add entertainment expenses such as movie theaters or club fees. Consistent with our results on sectoral value-added, we find that CS-intensive goods are luxuries.

We also use these expenditure share data to validate our estimates of regional CS productivity A_{rCS}. Our theory implies that, conditional on total expenditure, CS shares should be large in regions where prices are low, i.e., where A_{rCS} is large relative to local wages. Consistent with this prediction, our estimates of A_{rCS} are positively correlated with the estimated regional fixed effects of the CS expenditure system.

Local Price Levels: Finally, our estimated model predicts local prices that can be compared with the data. Because the expenditure survey reports both total expenditure and the total quantity bought for a variety of food items, we can compute average prices. In Section C-4 in the Appendix we show that these prices are strongly correlated with the prices of agricultural goods in our model even though we do not use this information in our calibration.

6 The Unequal Effects of Service-Led Growth

We now turn to our two main questions of interest: How important was productivity growth in the service sector to rising living standards? And how skewed were these benefits across different socioeconomic groups?

To quantify the welfare effects of CS growth, we compute counterfactual equilibria where we set CS productivity growth to zero in all districts. The resulting changes in wages and employment allocations thus reflect the effect of CS productivity growth, holding constant productivity growth in all other sectors and taking general equilibrium effects into account. Our model then allows us to compute the welfare effects for consumers and how these effects vary across space and the income distribution.

As in Baqaee and Burstein [2021] 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 $\varpi^q(\hat{x}_r|x_r)$ 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 $\hat{x}_r$. Hence, 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), $\varpi^q(\hat{x}_r|x_r)$ is implicitly defined by

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

In Appendix A-4, we derive an analytical expression for $\varpi^q(\hat{x}_r|x_r)$. In particular, we show that ϖ only depends on the level of human capital q if preferences are non-homothetic. Following a similar procedure, and exploiting the aggregation properties of PIGL preferences, we also calculate percentage equivalent variations for the utilitarian aggregation at the regional and national level.

6.1 Results: Sources of Welfare Growth in India

To highlight the inherent inequality of service-led growth, we first zoom in on three districts with different characteristics. Then, we summarize our findings at different

District	Urban Share	Population (millions)	Avg. Income	Emp. Share (%) Agr.	Ind.	CS	Prod. Growth (%) Agr.	Ind.	CS
Bangalore	0.77	10.6	3781	8	36	56	3.5	5.7	10.9
Chengalpattu	0.67	8.1	2807	12	37	51	2.9	4.7	8.8
Bankura	0.07	3.0	1597	64	7	28	1.5	1.9	2.5

Table 9: 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.

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] Finally, Bankura is a rural district in West Bengal, which is mostly dependent on agriculture and representative of rural India.

Table 9 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 different. In 2011 the employment share of CS is about 56% in Bangalore, 51% in Chengalpattu, and 28% in Bankura. Chengalpattu is the most industrial among the three. There are large differences in productivity growth between the local CS sectors—from 2.5% growth in Bankura to 11% in Bangalore. Industrial productivity growth is high in both Chengalpattu and Bangalore, consistent with the boom of manufacturing activity in the Chennai area and the ICT development in Bangalore. Productivity growth is lower across the board in Bankura, reflecting the general trend of urban-biased growth in India.

In the left panel of Figure 5 we display the welfare effects of resetting CS productivity for the entirety of India to the level of 1987. We depict these effects separately for the three districts as a function of household income and highlight average incomes with dashed vertical lines. The figure shows that the welfare effects of service-led growth vary dramatically across both space and the income ladder. In a rural location like Bankura, gains are small, especially for very poor households. The first reason is that the provision of local CS—and, hence, the average expenditure share—is low. Second, CS productivity grew significantly less than in Chengalpattu and Bangalore,

[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.

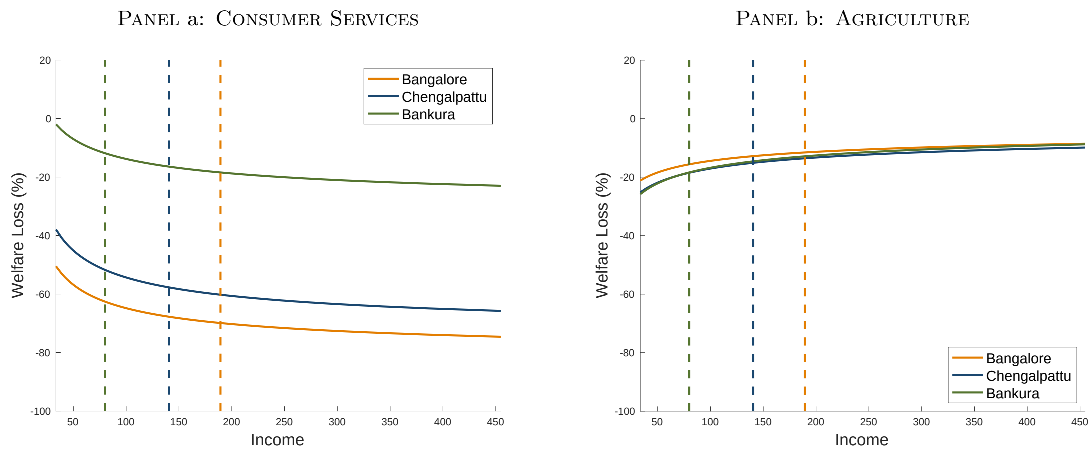

Figure 5: Counterfactual Welfare Changes. The figure displays the average percentage welfare losses associated with counterfactually setting productivity in CS (left panel) and agriculture (right panel) to their 1987 level in all Indian districts for households with different income levels living in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

where the welfare effects of service-led growth are much larger.

The figure also highlights that the benefits of service-led growth are increasing in income because CS are luxuries. Even in Bankura, the equivalent variation for rich households exceeds 20% of their 2011 income. For the very rich in Bangalore, the equivalent variation due to CS growth is closer to 70% of their income.

As a comparison, in the left panel, we depict the welfare effects of agricultural productivity growth. Very poor households in Bankura would rather sacrifice 30% of their 2011 income than experience the productivity setback in agriculture. In this case, the benefits for the poor mainly accrue from agricultural productivity growth in the whole of India, which reduces food prices overall. This diffusion of productivity growth via trade also explains why the spatial differences are small.

India-Wide Effects: To gather more general lessons, we now compute the welfare impacts for the entirety of India. In the left panel of Figure 6 we depict the (income-weighted) average welfare effects aggregated at the level of quintiles of urbanization.

These welfare implications stem from our estimated model and are therefore associated with sampling uncertainty.[30] To quantify the extent of this uncertainty, we thus estimate the distribution of the welfare effects using a nonparametric bootstrap procedure (Horowitz [2019])—see Section OA-4 in the Appendix for details. In Figure

[30] Intuitively, because the underlying micro data is a sample of individuals, the measured sectoral employment shares are random variables and so are our estimated structural parameters. Given the accounting nature of our analysis, our estimates of productivity fundamentals $\mathbf{A_t}$ (and, in turn, the counterfactual exercises of shutting down productivity growth) inherit this sampling uncertainty.

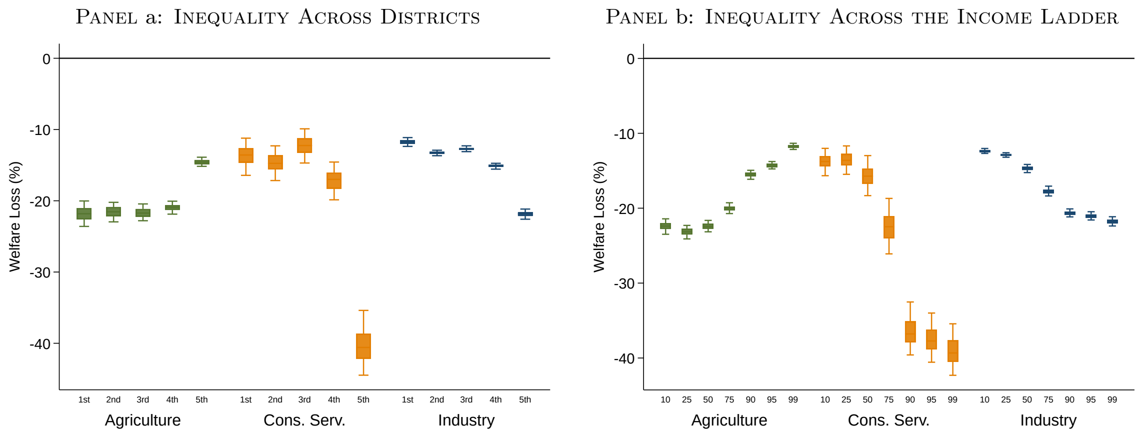

Figure 6: The Unequal Effects of Service-Led Growth. The figure displays the average percentage welfare losses, ϖ^q, associated with counterfactually setting productivity in agriculture, CS, and industry, to the respective 1987 level, broken down by urbanization quintile in 2011 (Panel (a)) and by the 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.

6 we report these distributions as a boxplot. Each box shows the 25%–75% quantiles of the distribution of aggregate welfare gains. The line within the box indicates the median and the two vertical lines on the top and the bottom indicate the 5% and 95% quantiles, respectively.

Unsurprisingly, the benefits of agricultural productivity growth are skewed toward rural areas. On average, households in the lowest quintile of urbanization are prepared to sacrifice 21% of their 2011 income to avoid going back to the 1987 productivity level in agriculture. The equivalent variation declines sharply in the top quintile, where productivity growth in agriculture is only worth 14%. By contrast, the benefits from productivity growth in CS and the industrial sector are skewed toward urban locations. This pattern is most pronounced for the CS sector, whose productivity growth is worth 41% of 2011 income for the most urbanized quintile. Our estimates of the distributions of these welfare gains make the urban-rural split of India also statistically precise. While we cannot reject that the welfare consequences of sectoral productivity growth are the same across the lower four quintiles, the nature of growth in the top urban quintile seems to have been qualitatively different: there, welfare gains were mostly service led, while the benefits from agricultural productivity growth were modest.

In the right panel of Figure 6 we focus on inequality across people and decompose the welfare effects across the Indian income distribution. We focus on the 10th, 25th, 50th, 75th, 90th, 95th, and 99th percentiles. As expected, the benefits of productivity growth in CS and, to a lesser extent, industry are sharply increasing in income, whereas the opposite is true for agriculture. Interestingly, the welfare change for the top 99% attributable to CS productivity growth is comparable to the top quintile of the urbanization distribution because not all rich people live in cities. Furthermore, we uncover significant differences in the sources of welfare growth between the top 25% and the bottom 75% of the population. For the bottom 75%, the welfare effects of productivity growth in agriculture, CS, and the industrial sector are roughly of the same size. For the top 25%, service-led growth was quantitatively much more important.

Finally, in the left panel of Figure 7 we aggregate welfare effects up to the nationwide level. Even at the aggregate level, a substantial part of welfare gains were service led. On average, the Indian population would have been willing to reduce its income in 2011 by 24% in lieu of giving up the observed productivity growth originating in the CS sector. Furthermore, with 90% probability, the welfare gains of service-led growth are between 22% and 27%. 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.

Figure 7 also shows that agricultural productivity was an important source of welfare improvement. The salience of agriculture is hardly surprising given its large employment share in India. The smaller welfare effects of productivity growth in the industrial sector is perhaps more surprising. The equivalent variation amounts to 17% and is very precisely estimated.

In summary, productivity growth in non-tradable services catering to consumers played an important role for economic development in India. However, the incidence of such productivity advances was heavily skewed. In urban areas and for rich households, growth in CS was the dominant factor of rising living standards. By contrast, technical progress in agriculture was the main source of welfare gains for the poor, living in rural areas.[31]

6.2 Structural Change: Growth Without Industrialization

A key aspect of the process of economic development in India and, more generally, in today's developing world (see Table 1) is a decline of agriculture that is not accompanied by significant industrialization. In this section, we show that productivity growth

[31] These findings are consistent with Chatterjee and Giannone [2021], who show that rising productivity in services is associated with regional divergence.

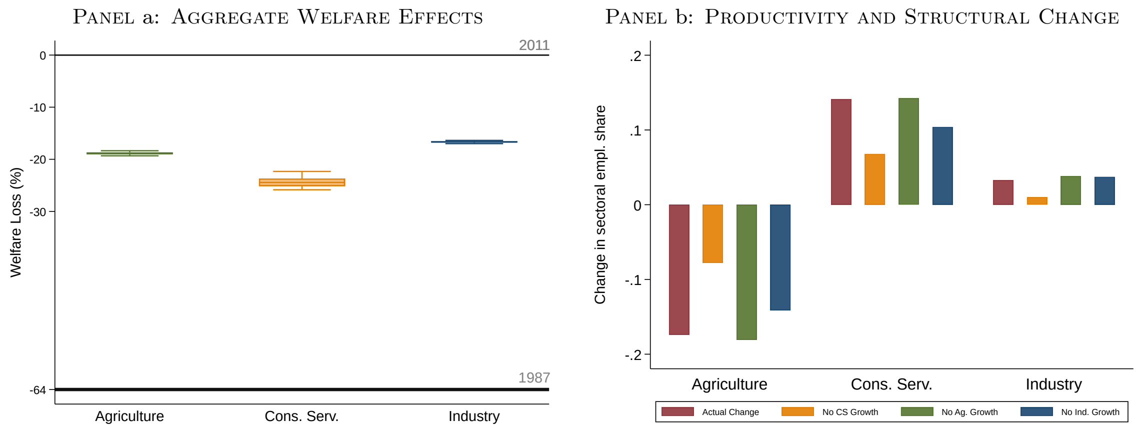

Figure 7: Aggregate Welfare Effects. In the left panel we show the analogue of Figure 6 with welfare effects aggregated up to the nationwide Indian level. In the right panel we show the change in sectoral employment. We depict the actual change in India (red bars) and the counterfactual results in the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

in the CS sector plays an important role for this pattern of sectoral transformation.

The right panel of Figure 7 depicts the extent of sectoral reallocation between 1987 and 2011. All figures are in effective units of labor.[32] The red bars show the actual Indian experience: agricultural employment declined by 17% and CS increased by 14%. The industrial sector, which contains PS, only saw an increase by 3%.

The remaining bars depict the sectoral reallocation implied by counterfactuals in which we shut down productivity growth in one sector at a time. The figure shows vividly that productivity growth in CS (orange bars) is responsible for the lion's share of the structural transformation. Absent productivity growth in CS, the agricultural employment share would have only declined by about 8% and the rise in CS employment would have been 7%. Note that the income effects associated with productivity growth originating in other sectors and from human capital accumulation played an important role: CS employment would have grown even without any CS productivity growth. Yet, tertiarization would have been far less spectacular than observed in the data.

The reason productivity growth in CS led to a marked decline in agricultural employment is that, without it, Indian consumers would be poorer and CS would be more expensive. Both forces would foster the demand for agricultural goods. The income effect would increase agricultural demand because food is a necessity. The substitution effect would reinforce this effect because food and CS are (mild) substitutes.

The green bars show that productivity growth in agriculture had only very modest

[32] 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.

effects on structural change. If anything, it marginally *increased* employment in agriculture and slowed employment growth in industry and CS. This result is in line with the findings of Asher et al. [2022] and Kelly et al. [2022], who document a negative effect of agricultural productivity on local industrialization in India and Britain during the Industrial Revolution. Similarly, productivity growth in the industrial sector did play the role of a pull factor, but in itself accounts for a small part of the observed path of structural change in India.

To sum up, service-led growth explains the lion's share of India's structural transformation between 1987 and 2011. Not only would Indian consumers be substantially worse off in welfare terms, but India would still be a much more rural economy today.

7 Robustness

In this section, we discuss the robustness of the welfare effects reported in Figures 6 and 7. First, in Section 7.1, we study the sensitivity of our results to changes in the structural parameters, in particular the Engel elasticity ε. Next, in Section 7.2, we revisit some measurement choices concerning the split between CS and PS. Finally, in Section 7.3, we study generalizations of the model. In particular, we extend our model to an open economy setting, we consider a production structure where skills are imperfectly substitutable, and we allow for workers to be spatially mobile.

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 main analysis, we estimated ε using micro data on Engel curves. In this section, we present an alternative calibration that does not rely on consumption expenditure at all. In particular, we re-estimate our model and calibrate ε by targeting the aggregate productivity growth of the Indian retail sector, that is 4.2%, as reported in the ETD (see Table 7). Doing so yields an estimate of $\varepsilon = 0.426$, which is slightly larger then our baseline estimate of $\varepsilon = 0.395$.

In the first panel of Table 10, we report the implied welfare effects of CS productivity growth according to this alternative calibration.[33] For parsimony, we only report

[33] In Table 10 we only report the welfare effects of productivity growth in the CS sector. For completeness we report the analogue of Table 10 for agricultural and industrial productivity growth in

	Aggregate Effects	by Urbanization Quintiles		by Income Quantiles		
		1st	5th	10th	50th	90th
Baseline	-24.7	-13.6	-41.1	-13.7	-14.7	-37.7
	Alternative calibrations of ε (Section 7.1)					
$\varepsilon = 0.426$ (Match CS Prod. Growth in ETD)	-23.1	-12.4	-39.0	-12.5	-13.3	-35.7
$\varepsilon = 0.321$ (OLS estimator)	-29.6	-17.6	-47.1	-18.0	-19.2	-43.4
	Alternative measurement choices (Section 7.2)					
Allocate PS share based on WIOD	-22.1	-13.9	-34.4	-14.2	-14.3	-32.1
Allocate ICT & Business to PS	-19.1	-15.6	-25.8	-14.4	-12.7	-25.2
Allocate Construction to Industry	-15.3	-0.6	-34.0	-3.0	-4.6	-25.2
	Alternative modeling assumptions (Section 7.3)					
Open economy	-21.1	-11.9	-34.9	-12.5	-12.2	-31.6
Imperfect skill substitution	-22.2	-10.6	-36.8	-10.2	-10.9	-34.9
Spatial labor mobility	-25.6	-14.8	-41.0	-15.1	-16.0	-38.0

Table 10: 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.

the effects at the aggregate level, for the top and bottom urbanization quintiles and for three quantiles of the income distribution. The results are quite similar to our baseline results. At the aggregate level, the welfare effect of service-led growth declines by 1.6 percentage points but remains larger than that of the two tradable sectors. Furthermore, the unequal effects of service-led growth are as large as in our baseline analysis.

In the third row, we report the results of setting $\varepsilon = 0.32$, the OLS estimate of the Engel elasticity. This change reduces the income effects and hence amplifies the role of service-led growth relative to our baseline results. Given that the highest estimate of ε in Table 4 was 0.42, changes in ε that are in line with either the micro or the macro data would not alter the main picture, in spite of somewhat affecting the quantitative results.[34]

Because the main focus of our analysis is on the *unequal* welfare effects, we further explore how far *heterogeneity* in Engel elasticities could affect our conclusions. In our theory, it is in principle possible to assume different elasticities across districts, at least as long as people are not spatially mobile. In the data, the Engel elasticity is larger

Section C-6 in the Appendix.

[34] We also calculated the ε for the aggregate welfare effect of growth in CS to be as large as in the agricultural (industrial) sector. The values for ε are 0.51 (0.57) and hence larger than what either the micro or macro data suggests.

in more-urbanized districts. This has two opposite implications on welfare: on the one hand, more of the expansion of CS activity is due to income effects, yielding a lower estimated productivity growth in CS in cities. On the other hand, households' welfare is more sensitive to productivity growth in CS. As long as productivity growth in services continues to be higher in urban districts, this points to higher spatial inequality. In Panel A of Appendix Figure C-3 we provide a tentative evaluation of these effects by assuming a higher (lower) elasticity in Bangalore (Bankura). While this yields a mild reduction in inequality, the quantitative effect is small. The conclusion that service-led growth is skewed toward urban districts is therefore robust.

Next, we consider heterogeneity in ε across income levels. Incorporating a variable income elasticity in our theory would be complicated and is beyond the scope of the paper. To gauge a sense of its quantitative effect, we run the following experiment. 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 6 in Table 4. Panel B of Appendix Figure C-3 displays the results, focusing again on the districts of Bangalore and Bankura. We find that welfare inequality is now *larger* than in the benchmark case in which all agents have the same preferences. 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 Section C-6 in the Appendix, we also study the sensitivity of our results to changes in other parameters. We focus on 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 because all other parameters are either point-identified in our theory or pinned down by normalization. The effects of changing these parameters are quantitatively very small and do not affect our conclusions in any way.

7.2 Measurement: The PS-CS Split

Our split of service employment into PS and CS reported in Table 3 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 used 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 roughly 20%. Therefore, we increase the relative size of the PS sector so that it accounts for 20% of value-added in the service sector at the aggregate level. 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 arguably sold to consumers and non-tradeable in nature. 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 4 and 5 of Table 10. As expected, the importance of productivity growth in CS decreases when we attribute a larger share of the expanding service sector to PS. At the aggregate level, the welfare effect of service-led growth declines by 2.6 and 5.6 percentage points, respectively, but is still sizable. At the spatial level, even though there is now less inequality than in our baseline analysis, CS productivity growth still mostly benefits the urban rich.

Finally, we turn 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 construction as part of the industrial sector. Although we regard our classification as more logical in the framework of our theory, we report the result of following the traditional classification in row 6 of Table 10. This reclassification increases the average welfare effect of productivity growth in the industrial sector, which goes up to 21%—see Appendix Table Table C-5. Nevertheless, we still find CS to be an important contributor to aggregate welfare growth. Interestingly, we see a large effect on the spatial heterogeneity because the construction sector is relatively more salient in rural areas. Thus, the welfare effects of service-led growth are even more skewed in favor of urban districts than in our baseline estimate. While for the most rural districts, construction accounts for the bulk of non-tradable activities, service-led growth in urban locations is not primarily driven from construction.

7.3 Generalizations of the Theory

In this section, we outline three generalizations of the theory that we present more formally in Appendix C-6.

Open Economy: Our main analysis treats India as a closed economy. However, international trade, in particular exports of ICT services, has become increasingly important. We incorporate these dimensions in an extension. 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 services, we assume India is an ICT exporter and exports the entirety of 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 10. International trade—especially, recognizing the tradable nature of ICT services—mildly reduces the welfare effect of productivity growth in CS, especially in cities, which saw the fastest increase in ICT employment. Nevertheless, CS continue to play an important role for aggregate growth and for urban areas in particular.

Imperfect Substitution and Skill Bias in Technology: Our analysis assumes that workers endowed with different efficiency units are perfect substitutes. In an extension, we generalize our model by assuming workers with different educational attainments are imperfect substitutes. Because agricultural workers have, on average, lower educational attainment than those employed in service industries, an increase in the skill endowment could be responsible for the reallocation of workers from agriculture to CS (see, e.g., Porzio et al. [2020] or Schoellman and Hendricks [2020]). By ignoring such skill-based specialization, our Ricardian model could exaggerate the importance of technology for the development of the service sector.

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:

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

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 calibrate the elasticity of substitution between high- and low-skilled workers to 1.8, a standard estimate in the literature. The results in Table 10 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 also uncovers two interesting 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].

Spatial Mobility: In our baseline model, we assumed people to be spatially immobile. However, a decline in CS productivity could lead people to move out of cities. 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 problem, where individuals receive idiosyncratic preference shocks and locations differ in amenities.[35]

We lay out and solve the model in Appendix OA-2.3. Here, we summarize the main ideas. We assume that individuals are free to locate in the region of their choosing. 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. Workers are subject to idiosyncratic preference shocks for each location which are Frechet-distributed. In the counterfactual exercises, we allow workers to migrate to their preferred location in the counterfactual environment.

Allowing for an endogenous location choice does not affect the estimation of the parameters nor of the productivities. Intuitively, given the observed population, we can estimate the model exactly as in our baseline analysis. The spatial distribution of amenities rationalizes then the observed population distribution as an equilibrium outcome. However, the spatial labor supply elasticity affects the counterfactuals. We calibrate such elasticity so that, holding local amenities fixed, resetting the productivity of CS in 2011 to the 1987 level in all districts triggers a spatial reallocation of the same magnitude as the total migration flow observed in India between 1987 and 2011. The

[35] Budí-Ors and Pijoan-Mas [2022] show that the rural exodus played a key role for the structural transformation in Spain.

results, which we report in the last row of Table 10, are very similar to those in the baseline model. We conclude that allowing for quantitatively reasonable migration responses to changes in the economic environment does not alter our results.

8 Conclusion

Tertiarization is well underway not only in mature but also in many developing economies. In a classic contribution, Baumol [1967], expressed the concern that this trend—driven by income effects—might lead to stagnation in labor productivity. In this paper, we develop a novel methodology to determine the importance of different sectors as 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 accounts for one-third of the improvement in living standards between 1987 and 2011. Second, the welfare impact of service-led growth is strikingly unequal: it disproportionally benefited wealthy individuals in urban areas while leaving poor people almost unaffected. The reasons are that service productivity grew particularly fast in urban areas and that richer consumers care more about the consumption of services owing to nonhomothetic preferences.

Our approach has several limitations that we hope to overcome in future research. Two are particularly important. First, our accounting approach takes CS 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, one should explore the extent to which these patterns are similar in other countries. While many developing countries are growing like India at the aggregate level, the patterns of spatial development might be different. While our analysis suggests that low employment growth in the manufacturing sector might be less concerning than previously thought, it also raises new concerns about inequality that remain invisible in aggregate statistics.

References

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

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: 2021, 'Welfare and output with income effects and taste shocks'. *NBER Working Paper* (w28754).

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.

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.

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., 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.

Herrendorf, B., R. Rogerson, and A. Valentinyi: 2013, 'Two Perspectives on Preferences and Structural Transformation'. *American Economic Review* **103**(7), 2752–89.

Herrendorf, B., R. Rogerson, and A. Valentinyi: 2014, 'Growth and Structural Transformation'. In: *Handbook of economic growth*, Vol. 2. Elsevier, pp. 855–941.

Herrendorf, B., R. Rogerson, and A. Valentinyi: 2020, 'Structural Change in Investment and Consumption—A Unified Analysis'. *Review of Economic Studies* **88**(3), 1311–1346.

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

Hsieh, C.-T. and E. Rossi-Hansberg: 2019, 'The Industrial Revolution in Services'. NBER Working Paper No. 25968.

Kelly, M., J. Mokyr, and C. O'Grada: 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.

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: 2020, 'The Human Side of Structural Transformation'. CEPR DP15110.

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

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

Schoellman, T. and L. Hendricks: 2020, 'Skilled Labor Productivity and Cross-Country Income Differences'. Mimeo, Federal Reserve of Minneapolis.

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

Singhari, S., 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 Appendices A, B, C are for online publication.

ONLINE APPENDIX A: THEORY

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

A-1 Proof of Proposition 1

To derive the expression in (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 get

$$\int_n \beta_n \ln p_{rnt} dn = \omega_F \ln P_{rFt} + \omega_G \ln P_{rGt} + \omega_{CS} \ln \left(A_{rCSt}^{-1} w_{rt} \right).$$

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 the final-good indirect utility function $\mathcal{V}^{FE}$ in (4) yields

$$\mathcal{V}^{FE}\left(e, \mathbf{p}_{rt}\right) = \frac{1}{\varepsilon} \left(\frac{e}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} \left(A_{rCSt}^{-1} w_{rt} \right)^{\omega_{CS}}} \right)^{\varepsilon} - \nu_F \ln p_{rFt} - \nu_G \ln p_{rGt} - \nu_{CS} \ln \left(A_{rCSt}^{-1} w_{rt} \right),$$

which is the expression in (7).

To derive the expenditure share over sectoral value added, $\vartheta_{rst}\left(e, \mathbf{P}_{rt}\right)$ in (9), note that sector s receives a share λ_{ns} of total revenue of good n. Hence, given a spending level e and prices $\mathbf{P}_{rt}$, sector s receives a share

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

which is the expression in (9). In Section OA-1.1 in the Online Appendix we extend this analysis to the case of CES preferences.

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

In this section we describe the details of our strategy to estimate the productivity fundamentals $\{A_{rst}\}$ and two structural parameters, ω_{CS} and ν_F. Consider a single time period. Given data on educational attainment (by region) and earnings (by sector-region), we can calculate $\{[w_r]_r, H_{rF}, H_{rG}, H_{rCS}\}_r$ in a model-consistent way. The supply of 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. We then calculate sectoral labor supply as

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

where w_i is the wage of individual i (in region r at time 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 demand parameters ω_{CS} and ν_F The two structural parameters are jointly identified from aggregate market clearing conditions. The local market clearing Equations (14) to (15), imply the two aggregate resources constraints for tradable goods $s = F, G$:

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

One of the aggregate resources 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 arrive at

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

Given data on $\{w_r, H_{rs}\}$, this is, for a given year, a single equation in three unknowns: ω_F, $\frac{\nu_F}{\nu_{CS}}$, and ω_{CS}. From the CS market clearing condition (14), it is apparent that ν_{CS} is not separately identified from the *level* of productivity in the consumer service sector, A_{rCSt}. Hence, under the assumption that consumer services are a luxury, we can wlog normalize $\nu_{CS} = -1$. For a given choice of ω_F we can therefore use (A-2) in 1987 and 2011 to uniquely solve for 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 from (14) to (15). We set the average level of the price of goods as the numeraire, i.e. $\left(\sum_r (p_{rGt})^{1-\sigma}\right)^{\frac{1}{1-\sigma}} = 1$. In addition, one can show that all our results are insensitive to the level of food prices in 1987. Finally, we target the change in aggregate food prices (relative to goods prices)

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

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

Step 3: Determine the level of the nominal wage The NSS data on expenditure (our measure of income) is reported in rupees. Given the vector of prices computed in Step 2, we thus chose the level of earnings to match a given growth of the real GDP per capita. In our model, we use final goods as the numeraire, and thus 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 implied elasticity of substitution. For notational 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

$$EOS_{sk} = \frac{\frac{\partial^2 e(P,V)}{\partial P_s \partial P_k} e\left(P, V\right)}{\frac{\partial e(P,V)}{\partial P_s} \frac{\partial e(P,V)}{\partial P_k}}.$$

The expenditure function is given by

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

As we show in Section OA-1.2, this implies that

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

A-4 Derivation of Equivalent Variation (Section 6)

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

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

We define the equivalent variation for an individual with skills q, $\varpi^q\left(\hat{x}_r|x_r\right)$ implicitly by[1]

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

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. If, for example, $\varpi_r^q = -20\%$, the consumer would be indifferent between giving up 20% of her 2011 income and a counterfactual allocation in which productivity in a particular sector is reset to the 1987 level.

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

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

The expression in (A-6) highlights that the equivalent variation consists of two parts. The first part, $\prod_s \left((\hat{w}_r/\hat{P}_{rs})/(w_r P_{rs})\right)^{\omega_s}$, is akin to the usual change in the real wage. This is the only part that is present if preferences are homothetic, that is if $\nu_s = 0$. Because this change is common across all consumers within a location, this aspect of sectoral productivity growth is necessarily equal. The second part captures the presence of nonhomothetic preferences and induces unequal effects of productivity growth. Consider, for example, a decline in CS prices, that is, $\ln \hat{P}_{rs}/P_{rs} < 0$. Because CS are luxuries, $\sum_s \nu_s \ln\left(\hat{P}_{rs}/P_{rs}\right) > 0$ so that rich individuals for whom $\left(q\hat{w}_r/\prod_s \hat{P}_{rs}^{\omega_s}\right)^{-\varepsilon}$ is small, have a higher willingness to pay for lower CS prices.

In a similar vain, 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 $\overline{\varpi}_r\left(\hat{x}_r|x_r\right)$ the representative agent in district r facing prices P_r would require to attain indifference. As before $\overline{\varpi}_r\left(\hat{x}_r|x_r\right)$ is implicitly defined by

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

where $\mathcal{U}$ is defined in (13). One can show that $\overline{\varpi}_r\left(\hat{x}_r|x_r\right)$ satisfies an expression similar to the one given in (A-6).

To arrive at an aggregate level of welfare changes, we calculate the equivalent variation at the national level by averaging the local income variations using regional

[1] Recall that we defined $x_r = (w_r, \mathbf{P_r})$.

income shares as weights:

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

A-5 Generalizations of Theory (Section 7.3)

In this section we provide further details on the extension of our theory discussed in Section 7.3 in the main text. Additional details, including all formal derivations, are available in Section OA-2 in the Online Appendix.

A-5.1 Open Economy

Thus far, we have treated India as a closed economy. However, international trade, in particular exports of ICT services, has become increasingly important for India. In this section, we extend our model to an open-economy environment. For brevity, we only summarize the main features of the extended model. The technical analysis can be found in Online Appendix OA-2.

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

To calibrate this model, we need information on the revenue of ICT services, the exports and imports of goods, and an estimate of the trade elasticity. We measure ICT revenue from the income share of ICT workers. We classify as ICT service workers all those employed in the following service industries: (i) telecommunications, (ii) computer programming, (iii) consultancy and related activities software publishing, and (iv) information-service activities. In our NSS data, these activities constitute 0.72% of total employment in 2011 (in 1987, it was a less than 0.1%). ICT workers earn, on average, higher wages than other workers. When one considers the earning share, they account for 1.56% of total earnings in 2011 (in 1987, it was 0.11%). Given the small size of the ICT sector in 1987, we assume it was zero and target the earnings share in 2011. In terms of exports, according to the World Bank, the export of goods and merchandise increased from 11.3 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 our model, we allow for individual heterogeneity in human capital but maintain that workers endowed with different efficiency units are perfect substitutes for one another. In this section, we generalize our model by assuming workers with different educational attainments are imperfect substitutes in production (see Section OA-2 in the Online Appendix for details). As we showed in Table OA-2, agricultural workers have, on average, lower educational attainment than those employed in service industries. Thus, an increase in the skill endowment could be responsible for the reallocation of workers from agriculture to CS (see, e.g., Porzio et al. [2020] or Schoellman and Hendricks [2020]). By ignoring such skill-based specialization, our Ricardian model could exaggerate the importance of technology for the development of the service sector.

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.[2] 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

[2] 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.

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 our baseline model, workers are exogenously assigned to regions. In the counterfactual analysis, we assumed people to be spatially immobile. However, people could decide to leave urban areas in response to sector-region productivity changes. To gauge the quantitative importance of labor mobility, we re-estimate our model in the presence of a migration choice. Formally, as is standard in models of economic geography (see, e.g., Redding and Rossi-Hansberg [2017]), we model migration as a discrete choice problem, where individuals receive idiosyncratic preference shocks and locations differ in a scalar amenity.

In Section OA-2.3 in the Appendix we discuss the solution of this model in more detail. We first show that all our estimates of both structural parameters and sectoral productivities are the same as in the model with immobile labor. Intuitively, given the observed population, we can estimate the model exactly as in our baseline analysis. We can then residually estimate the spatial distribution of amenities $\mathcal{B}_{rt}$ to rationalize the observed population distribution as an equilibrium outcome.

To perform counterfactuals, we need an estimate of the spatial labor supply elasticity, 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. We also tested the robustness of the results to higher-elasticity scenarios.

ONLINE APPENDIX B: DATA AND MEASUREMENT

In this section, we discuss details of the data and measurement issues discussed in Section 4.

B-1 International Evidence

In Table 1 in the main text, we documented the absence of industrialization for a variety of developing countries using the data from International Labor Organization (ILO). In Table 1 we focused on a selected subset of countries and the average of 27 developing countries.[3] In Table OA-1 in the Online Appendix we also report Table 1 for each country separately.

In Figure 1 in the main text, we showed that service employment in India is to a large extent concentrated in consumer rather then producer services. Again, India is the norm rather then the exception. In Figure B-1 we display the correlation between the employment share in consumer service relative to the entire service sector and income per capita in 2010. Poor countries have a much higher share of their service workforce in consumer services then rich countries and India does not appear to be an outlier - if anything, producer services are slightly more important in India.

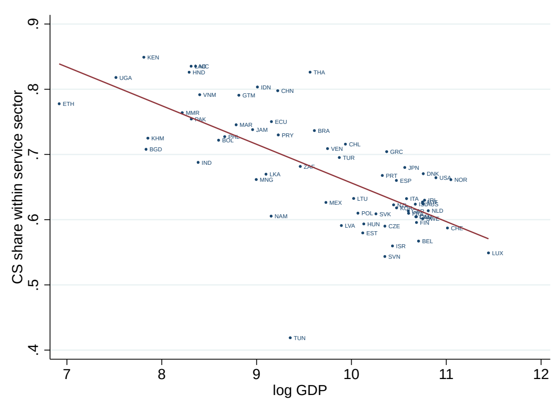

Figure B-1: Producer vs. consumer services and economic development. The figure shows the share of CS relative to the entire service sector in 2010 as a function of income per capita.

[3] We classify the following countries as developing countries: Bangladesh, Bolivia, Brazil, China, Ecuador, Guatemala, Honduras, Indonesia, India, Jamaica, Kenya, Cambodia, Lao People's DR, Sri Lanka, Morocco, Myanmar, Mongolia, Namibia, Nicaragua, Pakistan, Philippines, Paraguay, Thailand, Tunisia, Uganda, Viet Nam, South Africa

B-2 Data Sources

Our analysis relies on five datasets: (i) The National Sample Survey (NSS), (ii) The Economic Census (EC), (iii) The Service Sector in India: 2006-2007, (iv) The Informal Non-Agricultural Enterprises Survey 1999-2000 (INAES), and (v) the Household Expenditure survey. In this section we describe these datasets in detail.

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.[4] 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.

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

As we describe in more detail in Section B-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. For all years except 1987, expenditure is reported using the MRP classification. To make the results comparable across years, we merge the expenditure module (described in Section B-2.5) with the employment module in 1987 at the household level and use the MRP measure contained in the expenditure module. In practice, this choice is inconsequential because

[4] 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.

the different measures are highly correlated.

To measure human capital, we utilize information on educational attainment. We classify individuals' education into four levels: (i) less than primary; (ii) primary, upper primary, and middle; (iii) secondary; (iv) more than secondary. We then associate different years of schooling to each category to estimate annual returns. Building on the official classification in India, we attribute 0, 3, 6, and 9 years respectively.

In Table B-1, we report the distribution of human capital across time, space and sectors of production. First, educational attainment grew substantially between 1987 and 2011. Second, there are large different across space and individuals in cities have much higher levels of education. Third, there are large sectoral differences across sectors. Workers in agriculture have the lowest level of education, the PS sector has the highest level of education. The CS sector is more educated than the manufacturing sector.[5]

	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.

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.

[5] In Table OA-2 in the Online Appendix we report the same composition when we classify PS and CS workers according to the NIC classification.

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.[6] 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	0.90	0.98
63	Transport activities; travel agencies	188474	2101	3.40	3.33	86%	85%
64	Post and telecommunications	723119	22885	2.06	1.41	96%	99%
65–67	Financial intermediation	293489	16331	5.61	3.81	69%	82%
70	Real estate activities	70128	3648	2.18	1.64	93%	96%
71	Renting of machinery and household goods	365246	5387	2.00	1.77	94%	97%
72	Computer and related activities	66414	1060	6.01	13.45	83%	86%
73	Research and development	2097	5	16.66	4.58	66%	89%
74	Other business activities	519696	10610	2.81	1.92	90%	95%
85	Health and social work	783644	11930	3.39	1.99	88%	95%
91	Activities of membership organizations	1002996	2837	1.82	1.32	94%	98%
92	Recreational, cultural, and sporting activities	222061	2698	2.95	2.91	85%	82%
93	Other service activities	1419685	26132	1.74	1.54	97%	99%

Table B-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.

B-2.4 Informal Non-Agricultural Enterprises Survey 1999–2000 (INAES)

We use this dataset to allocate employment in the construction sector to either consumer or producer services. The Informal Non-Agricultural Enterprises Survey is part of the 55th survey round of the NSSO. It covers all informal enterprises in the non-agricultural sector of the economy, excluding those engaged in mining, quarrying and electricity, gas and water supply.[7] The Informal Non-Agricultural Enterprises Survey collects 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.

[6] 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.

[7] 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).

B-2.5 Household Expenditure Survey

To estimate the expenditure elasticity ε we rely on data on consumer expenditure. This data is contained in the National Sample Survey, Round 68, Schedule 1.0. The dataset reports detailed information on a large set of spending categories. In Table B-4 we report the broad classifications. The data also contains a finer allocation of spending within each category. For the purpose of this paper, we rely only on the classification in Table B-4.

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 consumers' spending on food as categories 1–17. We classify spending on consumer services as all spending in the consumer service category (category 24) and entertainment (category 20). In Section OA-3.2 in the Online Appendix we report the more detailed breakdown of consumer services across subcategories.

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 Section OA-3.2 in the Online Appendix we report a set of descriptive statistics on the cross-sectional distribution of spending, food shares and CS shares.

For our regression analysis reported in Table 4, we control for additional household-level covariates. We control for the size of the household and the number of (potential) workers in the household, which we define as all individuals between ages 15 and 65. We also control for additional household demographics, namely

- the type of the household, which for rural areas is one of (i) self-employed in agriculture, (ii) self-employed in non-agriculture, (iii) regular wage/salary earner,

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

- 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.

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

B-3 Geography: Harmonizing Regional Borders

In this section we describe our procedure to harmonize the geographical boundaries to construct a consistent panel of time-invariant localities. This need arises because the borders of numerous Indian districts have changed between 1987 and 2011. This is seen in the left panel of Figure B-2 that 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 regional re-districting is a *partition* in which one district has been separated into several districts in the 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 carry out the analysis on a panel of districts with a consistent geography, we construct regions that have consistent borders in 1987 and 2011. To keep the number of regions as large as possible, a region is always the smallest area that covers a single district or a set of districts with consistent borders over time. For instance, in the case of a partition, the region is constructed as the district in the pre-partition year. In the case of a border move, a region is constructed as the union of two districts. We construct a regional map with consistent borders from 1987 to 2011. The right panel of Figure B-2 shows the official Indian districts in the year 2011 (dashed red lines) and the time-invariant geographical units, that we for simplicity also refer to as districts (solid blue lines). We exclude from the analysis two small districts that did not exist in 1987 but did in 2011. Furthermore, because our methodology requires us to calculate sectoral employment shares at the district level, we exclude districts with less than 50 observations as these do not allow us to credibly estimate such shares.

B-4 Classification of Industries

At the heart of our analysis is the sectoral composition of regional employment. In our theory we distinguish between four sectors: agriculture, manufacturing, consumer

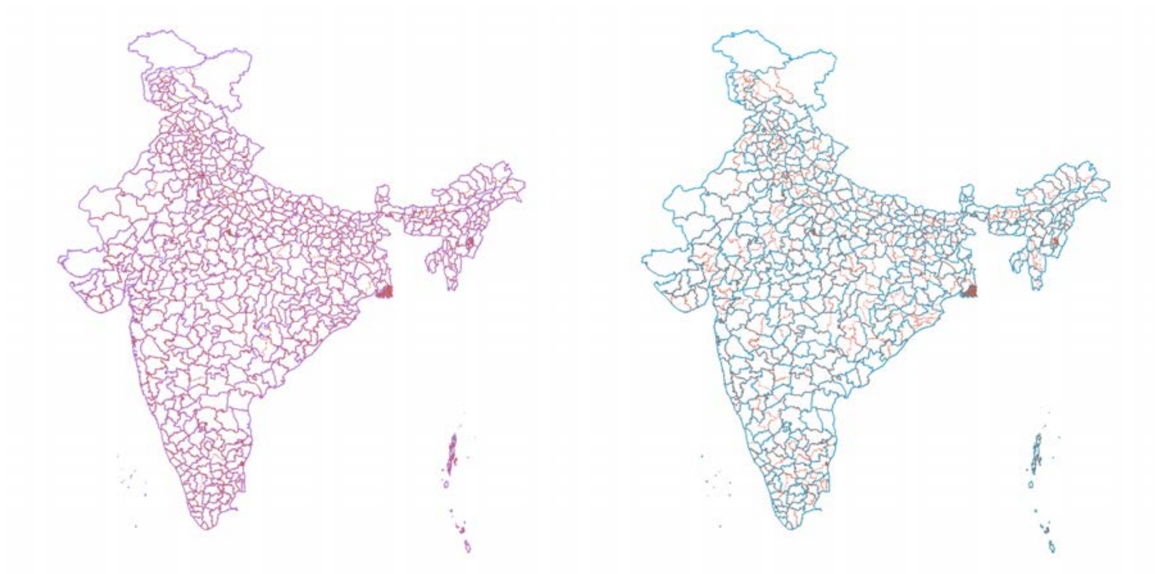

Figure B-2: District Borders in India 1987–2011. The left figure plots the districts' boundaries in 2011 and 2011. The purple line represents the boundaries in 2001 and the dashed red line represents the boundaries in 2011. The right figure shows the official Indian districts in the year 2011 (dashed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

services and producer services. To map these concepts to sectors in the data, we first construct six broad industries (see Section B-4.1). In a second step we then attribute employment in services and construction to consumer and producer services respectively (see Section B-4.2).

B-4.1 Broad Industry Classification

We divide 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. To do so we rely on India's official classification system, the National Industrial Classification (NIC). Because the NIC classification system changes over time, we construct a concordance table between 2-digit industries of different versions of the NIC based on official NIC documents and detailed sector descriptions. This concordance system allows us to compare sectoral employment patterns over time. We report the classification of industries in Tables OA-7 and OA-8 in Section OA-3.2 in the Online Appendix.

B-4.2 Attributing Employment to CS and PS

Our theory highlights the difference between PS, which are inputs in the production of goods, and CS, which are bought directly by consumers. To attain a systematic classification, we rely on the Service Survey (see Section B-2.3) that reports the identify of the main *buyer* of a given firm. We therefore refer to firms that sell to other firms as PS firms and firms that sell to consumers as CS firms.

Ideally, we would calculate the employment share of PS firms in each subsector of the service sectors and in each region. Given the large number of regions and subsectors, the sample size of the Service Survey is not sufficiently large to estimate these averages precisely.

We therefore generate the regional variation in employment shares by using regional variation in the firm-size distribution and differences in the employment share of PS firms by firm size. Empirically, large firms are – within their subsector – much more likely to sell to firms. In Section OA-3.2 in the Online Appendix, we depict the employment share of PS firms as a function of firm size in the raw data (see Figure OA-2) and show that the same pattern is present within 2- and 3-digit industries (see Table OA-9).

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

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

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

Here, f denotes a firm, $1\{f \in PS\}$ is an indicator that takes the value 1 if firm f is a PS firm and l_f denotes firm employment. In practice we take three size bins, namely "1 or 2 employees," "3–20 employees," and "more than 20" employees. We always weigh observations with the sampling weights provided in the Service Survey.[8]

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 NSS.

[8] 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 subsectors within the service sector 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 vehciles 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, where have the required information.

B-4.3 Construction and Utilities

As explained in the text we also attribute employment in construction and utilities to either CS or PS. We follow a similar strategy as for the service sector. To do so, we use the Informal Non-Agricultural Enterprises Survey 1999-2000 (see Section B-2.4).

From the description of the National Industry Classification, some subsectors are clearly for public purposes. We therefore classify 5-digit level industries within the construction sector into public and private and drop for our analysis all subsectors that we classify as public. These account for roughly 9.2% of total construction employment. See Table OA-10 in Section OA-3.2 in the Online Appendix for the detailed classification.

For all subsectors attributed to the private sector, we estimate the CS and PS share based on the information in the Informal Non-Agricultural Enterprises Survey. The survey has information on firms in the construction sector and reports the 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 OA-10), CS, and PS in the construction sector as a whole. The share of public employment is around 10% with a slight bump in 2009, presumably a consequence of the financial crisis in 2008. Among the private subsectors, 12.9% of employment is associated with the provision of producer services.

B-5 Urbanization and Spatial Structural Change

In Figure B-3 we quantify the structural transformation in India across both time and space. We focus on urbanization as our measure of spatial heterogeneity.[9] This

[9] 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,

	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 OA-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.

is a mere descriptive device, becuase 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 even starker when one looks at 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)

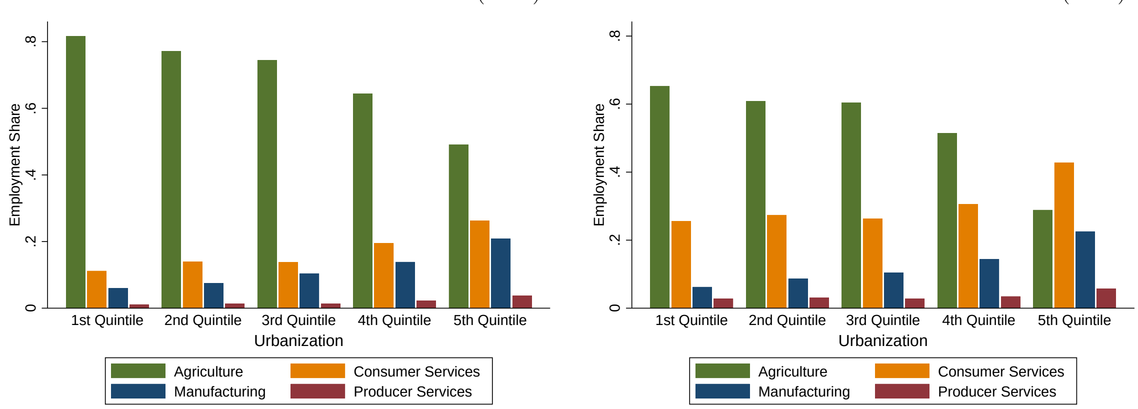

Figure B-3: Sectoral Employment over Time and Space. The figure plots the sectoral employment shares by urbanization quintile in 1987 and 2011.

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

ONLINE APPENDIX C: ESTIMATION

In this section, we report additional details of the estimation.

C-1 Estimating the Shape of the Human Capital Distribution (ζ)

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

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

Hence, we estimate ζ from a regression of the (log of the) upper tail density on log efficiency units that we calculate as $q_{rt}^h = \frac{e_{rt}^h}{w_{rt}}$. In Table C-1 we report the estimated ζ based on (C-1). 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 (which we use as our earnings measure for our main analysis) and total income, which is also reported in the NSS data.

Table C-1 contains two results. First, the estimated tail parameter for the aggregate economy is slightly below 3, is stable across years, and does not depend on the exact measure of earnings. Second, the estimated tail parameter is declining in the urbanization rate indicating that urban locations have higher inequality. Our estimates also indicate that inequality was lower in 2011 than in 1987. For our quantitative model, we set ζ to an average value of 3. In Section 7 we show that our results are robust to a variety of choices for ζ. Hence, for simplicity, we abstract from the heterogeneity in ζ across urbanization quantiles.

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

Table C-1: IDENTIFICATION OF ζ. The table reports the estimate of ζ based on (C-1). 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-2 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[10] 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[11], 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.[12] The WPI tracks ex-factory prices for manufactured products and market prices for agricultural commodities.[13] 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-1 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 MOPSI in our analysis.

C-3 Productivity Growth and Urbanization

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 6). In Table C-2 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).

[10] 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."

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

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

[13] 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.

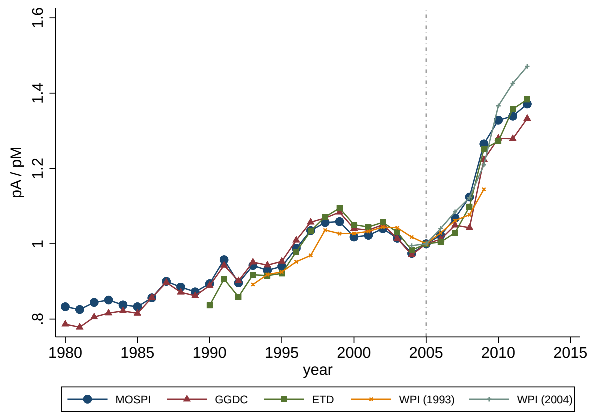

Figure C-1: Relative price of agricultural to manufacturing goods. The figure shows the relative prices of agricultural products from the different sources mentioned in the main text. "MOSPI" refers to the data from the Indian Government that is used in our analysis. "GGDC" stems from the GGDC 10-Sector 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.

	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-2: 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.

C-4 Non-targeted Moments: Additional Results

In our main analysis we have used data on food shares to estimate the Engel elasticity ε. Alternatively, we could have used data on the expenditure share of CS. We prefer food expenditures for two reasons. First, expenditure on food items is likely to be better measured. Second, the log-linear specification in (18) only recovers a consistent estimate of ε if the asymptotic expenditure β_n is small. While this is plausible for the case of food, the asymptotic spending share on CS intensive products has to be positive if such goods are luxuries.

Nevertheless, our model is consistent with two important features of the data on consumer service spending. First, we run—in the model and in the data—the same specification as in (18) except that we use households' expenditure share on CS as the dependent variable. We follow the official classification of the NSS expenditure module to assign expenditures to CS. As seen in Tables OA-3 and OA-4), these expenditures

Panel a: A_{rCS} and Fixed Effects of CS Spending.

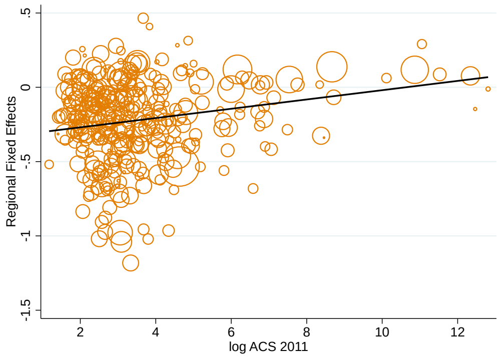

Panel b: Food prices—Data versus Model.

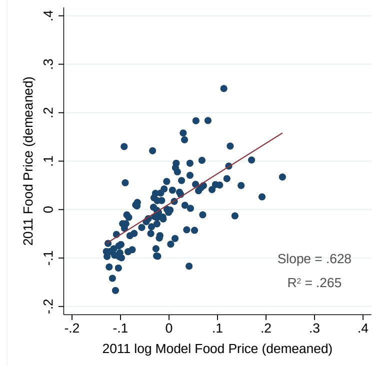

Figure C-2: Regional Variation in Consumer Spending. In the left panel we display the correlation of the region fixed effect of a regression of log CS expenditure on individual income against our estimates of consumer service productivity. The right panel shows a binscatter plot of regional log food prices in the data ($\hat{\delta}_r$ from (C-2)) and the model ($\ln p_{rF}$)

include, for example, domestic servants, barber shops, or tailor services. We also add entertainment expenses such as movie theaters or club fees.

We find that CS are luxuries: high-income households spend a higher share on CS. Quantitatively, we find the elasticity between the spending share on CS and individual income to be between 0.25 and 0.3 for the OLS specification and around 0.55 in the IV case. We can then use the CS expenditure share data to validate our estimates of regional CS productivity A_{rCS}. Our theory implies that, conditional on total expenditure, CS shares should be large in regions where prices are low, that is, where A_{rCS} is large relative to the local wage. This suggests that the regional fixed effects δ_r should be positively correlated with A_{rCSt}.

In the left panel of Figure C-2 we depict this correlation as scatter plot between $\hat{\delta}_r$ and our estimates $\ln A_{rCSt}$. There is a robust positive relationship; that is, in regions that we estimate to be productive in the CS sector, consumers spend a large fraction of their income on CS holding income constant.

As alluded to in the main text, we can also 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

$$\ln p_{nr} = \delta_r + \delta_n + u_{nr}, \tag{C-2}$$

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 the right panel of Figure C-2 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-5 Outliers in Quantitative Analysis

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

This is seen Table C-3, where we report the upper and lower quantiles of the regional distribution of welfare changes for the different counterfactuals. Consider for example the agricultural sector. If agricultural productivity had not grown since 1987, the most adversely affected region would have seen its welfare decline by 55.9% in terms of an equivalent variation. Conversely, some regions would have seen their welfare increase. The last row of Table C-3 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 also report the estimated distribution of the welfare effects in our baseline analysis, where we truncate the productivity growth distribution at the top and bottom 3%. This has large effects on the distribution of welfare effects in the right tail of the distribution.

	Regional Welfare Changes (%)									
	Min	1%	2%	3%	5%	95%	97%	98%	99%	Max
Agriculture	-55.9	-44.9	-43.0	-42.8	-39.5	2.3	6.3	15.3	16.6	45.7
Industry	-34.5	-28.0	-27.1	-25.4	-23.8	-6.3	-4.1	-3.1	1.2	25.8
Cons. Serv.	-98.9	-96.3	-89.8	-86.6	-77.4	18.4	44.4	135.4	320.5	1498.7
Cons. Serv. (Baseline)	-93.8	-93.0	-88.0	-85.9	-77.0	18.3	35.6	40.0	71.7	91.1

Table C-3: 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-4 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 -21.7%, due to the extreme outliers reported in Table C-4. Once such outliers are truncated, we recover our baseline results of a welfare loss of about -24.7%. In the last row of Table C-4 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	-21.7%	-23.4%	-24.2%	-24.7%	-25.0%	-25.1%
Employment Share	0	0.5%	1.9%	3.2%	5.4%	8.0%

Table C-4: 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-6 Details of Robustness Analysis (Section 7)

In Figure C-3 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-3, we assume our baseline estimate of $\varepsilon = 0.395$ in Bangalore and $\varepsilon = 0.29$ in rural Bankura as suggested by column 7 of Table 4. 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, again motivated by the results reported in Table 4, we assume that individuals above (below) the median have an elasticity of 0.418 and 0.265 respectively. The left panel of Figure C-3 highlights that this *amplifies* the differential welfare impact of service-led growth between rich and poor households.

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-4, 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.

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-4 shows that the implied welfare impact of sectoral productivity growth is essentially

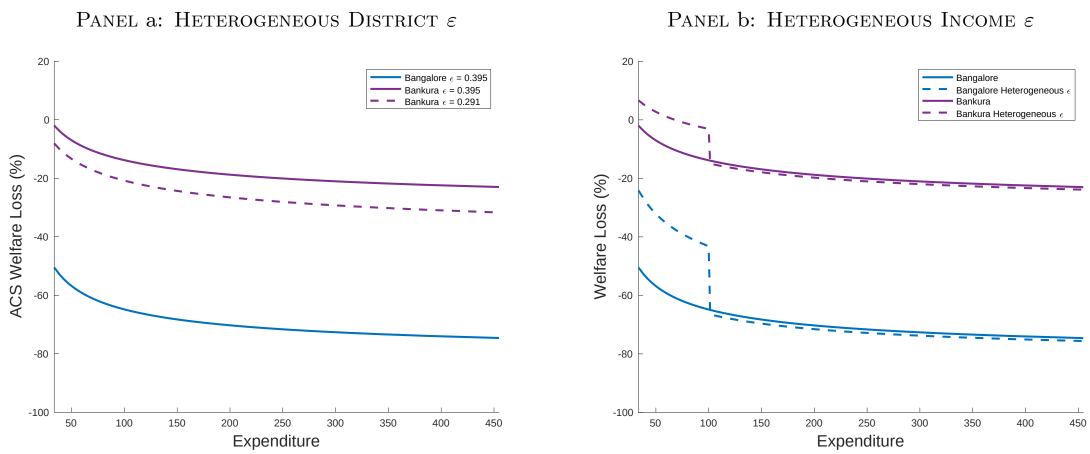

Figure C-3: Heterogeneous Engel Elasticities. In the left panel we allow for heterogeneous ε across locations. We assume that ε of individuals in Bangalore (Bankura) is 0.395 (0.291), which is in line with the results reported in Table 4. In the right panel we allow for different ε across individuals. In line with Table 4, we assume that individuals above (below) the median income have ε of 0.415 (0.218).

independent of ω_F.[14]

Finally, in panel (c) of Figure C-4 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-4 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-4 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-5 we report the analogue to Table 10, that is the welfare effects of agricultural and industrial productivity growth. Table C-5 shows that our baseline results are not significantly affected by either the alternative modelling assumptions or the alternative measurement choices.

[14] Because the regression coefficient between log foodshares and log expenditure only identifies the Engel elasticity ε for $\omega_F \approx 0$, for this exercise we re-calibrate ε so that our model is consistent with an estimated coefficient of food-shares with respect to spending of 0.395, when we run the regression in the model.

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

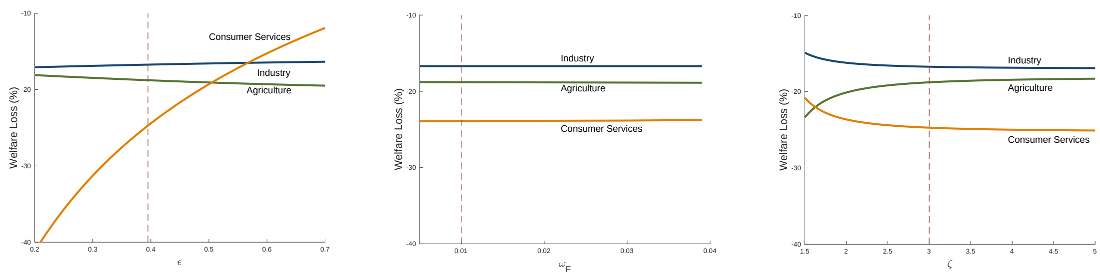

Figure C-4: Robustness Analysis. Panels (a), (b), and (c) show the aggregate welfare effects as a function of the preference parameters ε, ω_F, and the tail parameter of the skill distribution ζ. The vertical dashed line corresponds to the parameter value in our benchmark analysis.

	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	**-18.8**	**-21.5**	**-14.5**	**-22.2**	**-15.2**	**-16.7**	**-11.9**	**-21.9**	**-12.4**	**-20.6**
					Alternative calibrations of ε (Section 7.2)					
$\varepsilon = 0.321$ **(OLS estimator)**	**-18.5**	**-21.3**	**-14.4**	**-21.6**	**-15.3**	**-16.8**	**-12.0**	**-21.9**	**-12.7**	**-20.6**
$\varepsilon = 0.426$ **(Match CS Prod. Growth in ETD)**	**-18.9**	**-21.7**	**-14.6**	**-22.4**	**-15.2**	**-16.7**	**-11.8**	**-21.9**	**-12.3**	**-20.7**
					Alternative measurement choices (Section 7.2)					
Allocate PS share based on WIOD	**-18.9**	**-22.0**	**-15.1**	**-21.9**	**-15.7**	**-18.7**	**-12.8**	**-24.9**	**-13.6**	**-23.5**
Allocate ICT & Business to PS	**-19.1**	**-22.3**	**-15.5**	**-22.2**	**-16.0**	**-18.0**	**-12.0**	**-24.3**	**-12.5**	**-22.9**
Allocate Construction to Industry	**-18.3**	**-21.9**	**-11.9**	**-23.1**	**-13.9**	**-21.9**	**-12.8**	**-32.3**	**-13.5**	**-29.3**
					Alternative modeling assumptions (Section 7.3)					
Open economy	**-18.9**	**-21.9**	**-14.8**	**-22.2**	**-15.6**	**-19.0**	**-14.7**	**-23.6**	**-15.1**	**-22.4**
Open economy (large ICT)	**-18.3**	**-22.1**	**-14.5**	**-21.2**	**-15.5**	**-19.2**	**-14.8**	**-23.7**	**-15.6**	**-22.5**
Imperfect skill substitution	**-23.5**	**-28.1**	**-17.7**	**-25.6**	**-20.4**	**-16.4**	**-10.6**	**-22.5**	**-9.8**	**-22.5**
Spatial labor mobility	**-18.7**	**-21.2**	**-14.7**	**-22.1**	**-15.3**	**-16.7**	**-12.0**	**-21.6**	**-12.6**	**-20.5**

Table C-5: The importance of service-led growth—Robustness.

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

by Tianyu Fan, Michael Peters, and Fabrizio Zilibotti

July 7, 2022

- Not for Publication Unless Requested -

OA-1 Additional theoretical results

OA-1.1 CES Preferences

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

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

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

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

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

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

This implies that

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

and

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

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

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

where

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

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

$$\vartheta_{rst} = \int_n \beta_n \varsigma_{rnt}^{s}(\mathbf{P}_{rt})\,dn + \left(\int_n \kappa_n \varsigma_{rnt}^{s}(\mathbf{P}_{rt})\,dn\right)\left(\frac{e}{B(\mathbf{P}_{rt})}\right)^{-\varepsilon}, \tag{OA-2}$$

where $\varsigma_{rnt}^{s}(\mathbf{P}_{rt})$ are the sectoral cost shares for good n given in (OA-1). The notation $\varsigma_{rnt}^{s}(\mathbf{P}_{rt})$ stresses that these shares depend on the regional prices of tradable goods and CS. Equation (OA-2) is a direct generalization of the Cobb-Douglas structure considered in the main text. There, the spending shares $\varsigma_{rnt}^{s}(\mathbf{P}_{rt})$ are constant and given by $\varsigma_{rnt}^{s}(\mathbf{P}_{rt}) = \lambda_{ns}$. 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, (OA-2) can be written as

$$\vartheta_{rst} = \omega_{rst} + \nu_{rst}\left(\frac{e}{B(\mathbf{P}_{rt})}\right)^{-\varepsilon}, \tag{OA-3}$$

where $\omega_{rst} \equiv \int_n \beta_n \varsigma_{rnt}^{s}(\mathbf{P}_{rt})\,dn$ and $\nu_{rst} \equiv \int_n \kappa_n \varsigma_{rnt}^{s}(\mathbf{P}_{rt})\,dn$. 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.

OA-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

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

Then,

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

and

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

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

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

OA-2 Generalizations of Theory: Formal Details

In this section we provide additional formal details for the extension of our theory discussed in Sections 7.3 in the main text and A-5 in the Appendix.

OA-2.1 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

$$P_G = \left(p_{G,D}^{1-\eta} + \varphi^{\eta} p_{G,ROW}^{1-\eta} \right)^{\frac{1}{1-\eta}} . \tag{OA-4}$$

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 (OA-4) 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.[15] We assume that the ROW buys a bundle of regional varieties of ICT services

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

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

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

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:

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

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

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:

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

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 $A_{ICTt} = 1$.[16]

[16] 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.

2. To identify Υ_{ICT} we use that

$$\begin{aligned}\sum_r w_r H_{rICTt} &= \sum_r \left(\frac{w_{rt}^{1-\sigma} A_{rICTt}^{\sigma-1}}{\sum_{j=1}^R w_{jt}^{1-\sigma} A_{jICTt}^{\sigma-1}}\right) \left(\sum_j w_{jt}^{1-\sigma} A_{jICTt}^{\sigma-1}\right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICTt} \\ &= \left(\sum_j w_j^{1-\sigma} a_{jICTt}^{\sigma-1}\right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{ICTt}. \quad \text{(OA-5)}\end{aligned}$$

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

$$M_{1t} = \frac{\left(\sum_j w_{jt}^{1-\sigma} A_{jGt}^{\sigma-1}\right)^{\frac{1-\eta}{1-\sigma}} \Upsilon_{Gt}}{\sum_r w_{rt} H_{rGt}} = \frac{P_{G,IND}^{1-\eta} \Upsilon_{Gt}}{\sum_r w_{rt} H_{rGt}} = \frac{\Upsilon_{Gt}}{\sum_r w_{rt} H_{rGt}}. \quad \text{(OA-6)}$$

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}.

OA-2.2 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,

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

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

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

where $\vartheta_s^h\left(qw_{rt}^-, P_{rt}\right)$ 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

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

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

$$w_{rt}^{-}H_{rCSt}^{-}+w_{rt}^{+}H_{rCSt}^{+}=\vartheta_{CS}\left(\underline{q}_{rt}^{-}w_{rt}^{-},\underline{q}_{rt}^{+}w_{rt}^{+},s_{rt}^{Y,-},\mathbf{p_{rt}}\right)\overline{w}_{rt}L_{rt}. \tag{OA-7}$$

Finally, labor market clearing implies

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

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

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 7.3.

OA-2.3 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.[17] 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

$$L_{rt} = \frac{\left(\hat{e}_{rt} \mathcal{B}_{rt}\right)^{\eta}}{\sum_j \left(\hat{e}_{jt} \mathcal{B}_{jt}\right)^{\eta}} L. \tag{OA-8}$$

where $\hat{e}_{rt} = E_{rt}[q] w_{rt} \left(1 + \overline{\varpi}_{rt}\left(\hat{w}_{rt}, \hat{\mathbf{P}}_{\mathbf{rst}} | w_{rt}, \mathbf{P}_{\mathbf{rst}}\right)\right)$. Holding $\sum_j \left(\hat{e}_{jt} \mathcal{B}_{jt}\right)^{\eta}$ constant, the partial elasticity with respect to the money-metric utility is given by η.[18] 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 (OA-8), we can identify $\mathcal{B}_{rt}$ given the observed allocation of labor and wages. This also implies that we cannot separately identify η without

[17] 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.

[18] 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.

additional information. Because we are specifically 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 a counterfactual is the same as the one observed in the data between 1987 and 2011.

If we counterfactually keep all productivity vectors at their 1987 value, we require $\eta = 0.61$ to match the observed standard deviation of population growth between 1987 and 2011. To generate twice the standard deviation, we would require $\eta = 2.05$.

OA-3 Additional empirical results

OA-3.1 Growth Without Industrialization: Country-Specific Results

In Table OA-1, which has the same structure as Table 1 in the main text, we report the change in sectoral employment shares and income per capita for 27 developing countries. Where there are, of course, idiosyncratic differences across countries, the broad pattern of "growth without industrialization" is observed in most of the developing world.

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 OA-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.

OA-3.2 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 OA-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 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 OA-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 OA-3 and OA-4 we report the more detailed classification of the consumer service (category 24) and entertainment spending (category 20) categories.

In Table OA-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 OA-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 OA-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 OA-1. Through the lens of our theory, this dispersion is generated through heterogeneity in income and relative prices.

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

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

In Table OA-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 OA-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

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

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

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 OA-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 OA-7.

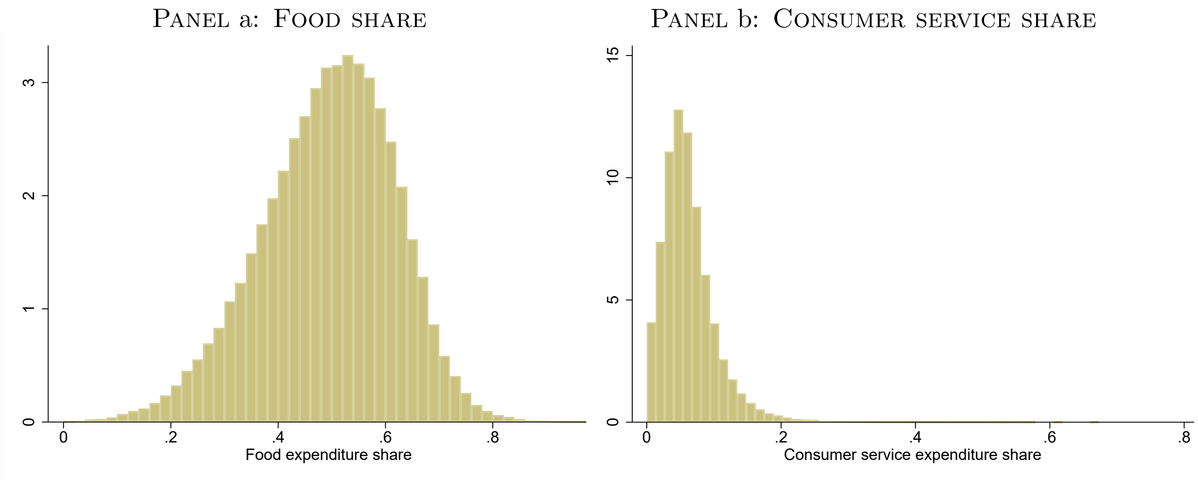

Figure OA-1: Distributon of food and consumer service expenditure shares. The figure shows the unconditional distribution of the expenditure shares for food (left panel) and consumer services (right panel).

firms are more likely to sell to firms as opposed to consumers. In Figure OA-2 we 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 OA-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 OA-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.

OA-3.3 Urbanization and Aggregate Growth

In Figure OA-3 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 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.

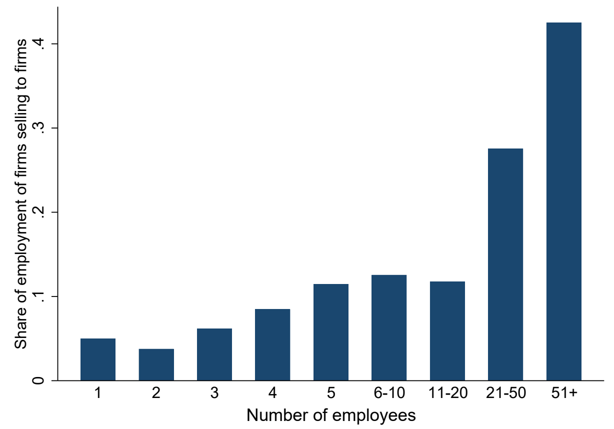

Figure OA-2: Producer service share by firm size. The figure shows the share of service firms whose main customers are other firms (as opposed to private individuals) with a breakdown by firm size.

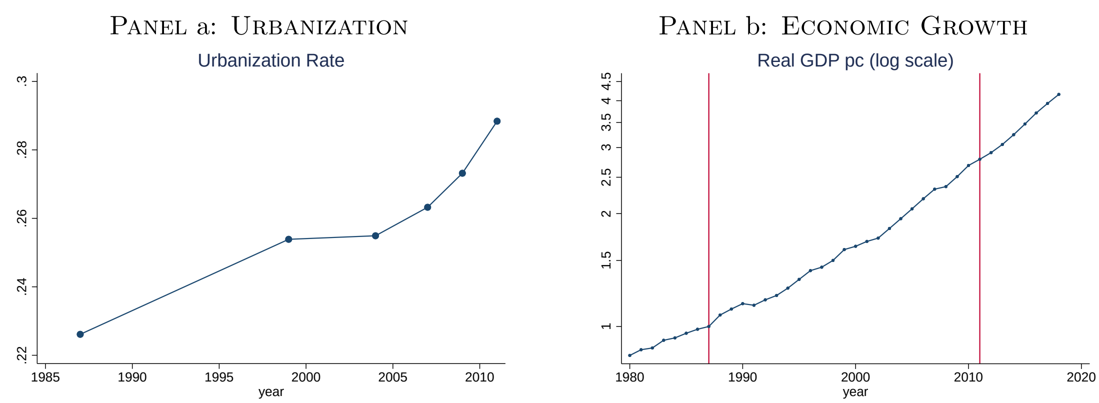

Figure OA-3: Economic Growth in India: 1987 - 2011. This figure shows the evolution of the urbanization rate (Panel a) and income per capita (Panel b). The urbanization rate is the share of population living in urban areas according to the definition of the NSS. Income per capita stems from World Bank.

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 OA-4 shows that there is a strong positive correlation between urbanization and expenditure per capita in the NSS data in 2011.

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 OA-5 we report sectoral income shares by urbanization quintiles in 1987 (Panel a) and in 2011 (Panel b). If anything, the patters we describe

	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 OA-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

in Figure B-3 are more pronounced because earnings are higher in service industries and in cities.

OA-4 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:

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

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

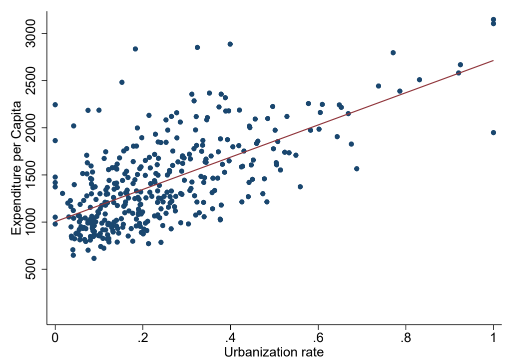

Figure OA-4: Expenditure per capita vs. Urbanization. The figure shows a scatter plot of the average expenditure per capita in the NSS data across district-level urbanization rates in 2011.

as the current dataset.[19]

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 expenditure of food reported in Table 4) and the preference parameters ν_F and ω_{CS} (as explained in Section 5),

[19] 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.

Panel a: Sectoral Income by Urbanization (1987) Panel b: Sectoral Income by Urbanization (2011)

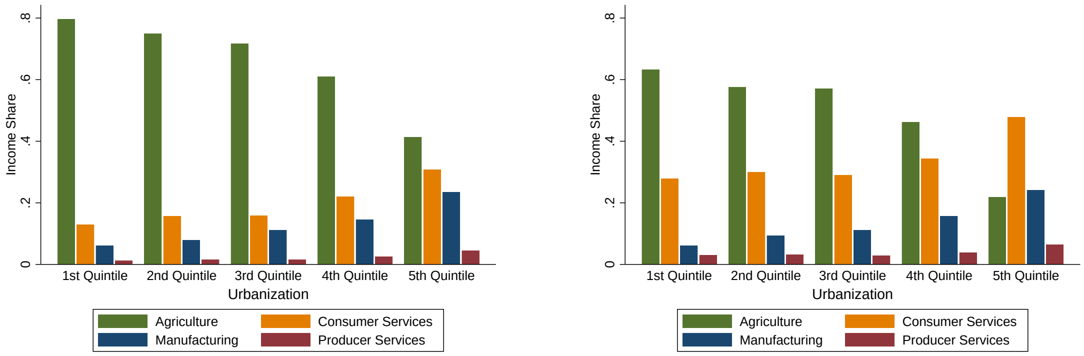

Figure OA-5: Spatial Structural Change in India. The figure plots plots the sectoral income shares by urbanization quintile in 1987 and 2011.

(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-5).

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

$$\left\{\Delta\varpi_r^{q(b)}, \Delta\overline{\varpi}_r^{(b)}, \Delta\overline{\varpi}^{(b)}, L_{s2011}^{CF_F,(b)}, L_{s2011}^{CF_CS,(b)}, L_{s2011}^{CF_I,(b)}\right\}_{b=1}^{B}. \tag{OA-9}$$

In practice we take $B = 200$.

6. From OA-9 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 OA-6 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.

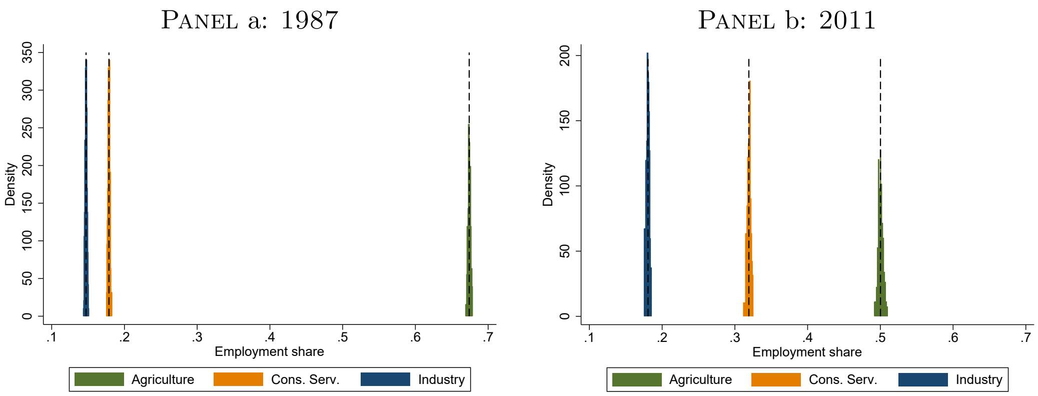

Figure OA-6: BOOTSTRAP DISTRIBUTION OF AGGREGATE EMPLOYMENT SHARES. The figure shows the bootstrap distribution of the aggregate sectoral employment share in 1987 (left panel) and 2011 (right panel). The vertical dashed line corresponds to the empirically observed value.

In Figure OA-7 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 (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.

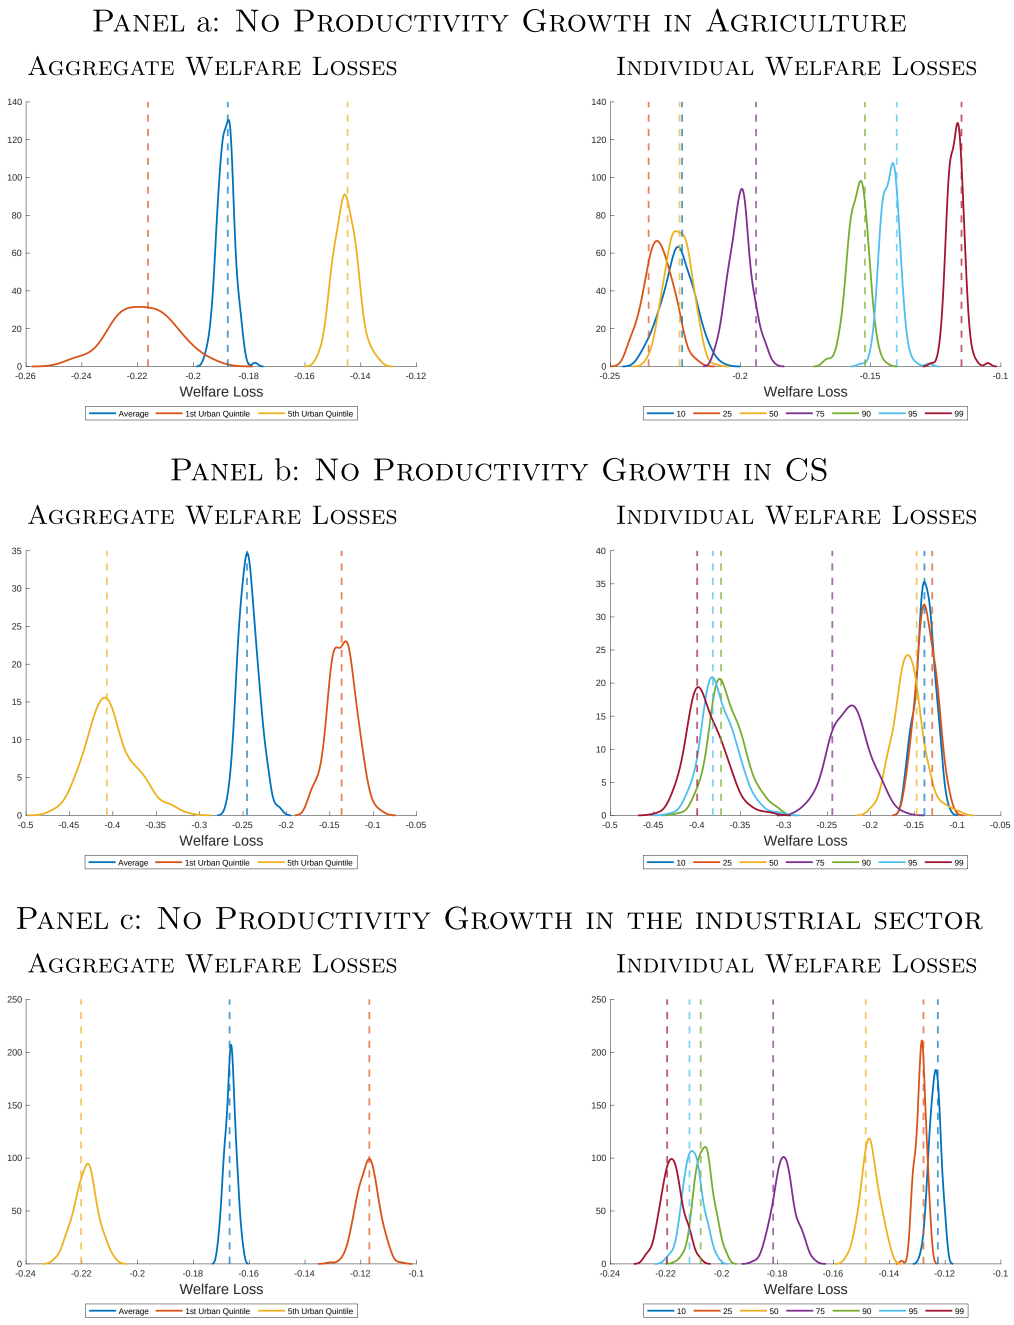

Figure OA-7: Bootstrap Distribution of Welfare Losses. The figure shows the bootstrap distribution of the welfare losses when we counterfactually set sectoral productivity in 2011 to its level in 1987. In panel (a) we shut down productivity growth in agriculture, in panel (b) we shut down productivity 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.