

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<https://www.econometricsociety.org/>

Econometrica, Vol. 91, No. 4 (July, 2023), 1457–1494

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

TIANYU FAN

Department of Economics, Yale University

MICHAEL PETERS

Department of Economics, Yale University and NBER

FABRIZIO ZILIBOTTI

Department of Economics, Yale University and NBER

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

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

TIANYU FAN

Department of Economics, Yale University

MICHAEL PETERS

Department of Economics, Yale University and NBER

FABRIZIO ZILIBOTTI

Department of Economics, Yale University and NBER

Structural transformation in most currently developing countries takes the form of a rapid rise in services but limited industrialization. In this paper, we propose a new methodology to structurally estimate productivity growth in service industries that circumvents the notorious difficulties in measuring quality improvements. In our theory, the expansion of the service sector is both a consequence—due to income effects—and a cause—due to productivity growth—of the development process. We estimate the model using Indian household data. We find that productivity growth in nontradable consumer services such as retail, restaurants, or residential real estate was an important driver of structural transformation and rising living standards between 1987 and 2011. However, the welfare gains were heavily skewed toward high-income urban dwellers.

KEYWORDS: Consumer services, economic growth, India, inequality, non-homothetic preferences, productivity, spatial equilibrium, structural change, welfare.

1. INTRODUCTION

URBANIZATION AND STRUCTURAL CHANGE ARE TRANSFORMING THE LIVES of hundreds of millions of people across the globe. Consider India, the most populous country in the world: Thirty years ago, only a quarter of the population resided in urban areas, and almost two-thirds of the labor force was employed in agriculture. Today, the share of

Tianyu Fan: tianyu.fan@yale.edu

Michael Peters: m.peters@yale.edu

Fabrizio Zilibotti: fabrizio.zilibotti@yale.edu

This project stems from early discussions with Philippe Aghion and Robin Burgess. A previous version was circulated under the title: *Service-Led or Service-Biased Growth? Equilibrium Development Accounting Across Indian Districts*. We are grateful to three referees and to Klaus Desmet and Sebastian Sotelo for their thoughtful discussions at the NBER SI and the ASSA. We thank Daron Acemoglu, Simon Alder, Manuel Amador, Marios Angeletos, Treb Allen, David Atkin, Johannes Boehm, Timo Boppart, Fabian Eckert, Reto Foellmi, Penny Goldberg, Doug Gollin, Cormac O’Dea, Gordon Hanson, Peter Klenow, Samuel Kortum, Ramon Marimon, Marc Melitz, Giuseppe Moscarini, Andreas Müller, Juan Pablo Nicolini, Rachel Ngai, Ezra Oberfield, Richard Rogerson, María Sáez Martí, Ludwig Straub, Todd Schoellman, Kjetil Storesletten, Aleh Tsyvinski, Danyang Xie, and seminar participants at the ASSA Meeting 2021, the Cowles Macro Conference, the Minneapolis FED, CREi, Dartmouth, Harvard University, HEC Lausanne, IIES, Johns Hopkins SAIS, the NBER Summer Institute, New York University, MIT, Penn State, Peking University, RIDGE, STEG, SED, the Università di Bologna, the Università 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, Dalton Zhang, and Huihuang Zhu for excellent research assistance. Michael Peters and Fabrizio Zilibotti thank the “Minnesota Opportunity and Inclusive Growth Institute” for its generous hospitality. Zilibotti acknowledges the support of the National Science Foundation, grant 1949228.

people living in cities has increased by 10 percentage points (p.p.), and the employment share of agriculture is down to 42%.

In this paper, we argue that productivity growth in the service sector has played a key role in this transformation and in the accompanying rise in living standards. We focus on nontraded services that serve final consumers, such as retail, restaurants, local transportation, or residential real estate. We refer to such services as consumer services (CS). Employment in CS has risen dramatically in recent decades and now accounts for one-third of aggregate employment in India; this share increases to almost two-thirds in urban districts such as Delhi or Mumbai.

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

We use the estimated model to infer the heterogeneous welfare effects associated with the process of structural change across both localities and the income distribution, building a bridge between economic growth and economic development.

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

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

¹Failure to account for quality changes can bias price indices upwards and lead one to underestimate productivity growth in services. Suppose, for instance, that improvements in logistics reduce the cost of home delivery, which makes these services accessible to more consumers. But, online shopping is more expensive than traditional retail. In that case, the average price paid by consumers for the service would grow. The increase, however, reflects a convenience value for which consumers are willing to pay.

To discipline this elasticity, we estimate households' Engel curves using micro data on consumption expenditures. We parametrize preferences by an indirect utility function in the Price Independent Generalized Linear (PIGL) class. [Muellbauer \(1976\)](#) first introduced this preference class, and [Boppart \(2014\)](#) recently revamped it in the growth literature. PIGL has two important properties. First, it features aggregation: the choice of a set of agents endowed with PIGL preferences facing a common price vector can be rationalized as the choice of a representative agent whose preferences also fall into the PIGL class. Second, we prove that, under some conditions, PIGL preferences enable one to seamlessly go back and forth between preferences defined over final expenditure and over sectoral value-added. This novel theoretical result is important because, as shown by [Herrendorf, Rogerson, and Valentinyi \(2013\)](#), the mapping between the parameters of the value-added demand system and the ones derived from preferences over final goods generally depends on the entire input–output matrix. We formally establish that under PIGL preferences, the key parameter governing the income elasticity is common to both demand systems at the individual and aggregate level. This makes it legitimate to use micro data on household expenditure to estimate the income elasticity of the aggregate value-added demand system, which is our elasticity of interest.

We apply our methodology to India, which is a rapidly growing economy with an annual GDP per capita growth rate of 4.2% during 1987–2011. Our estimation exploits individual geolocalized consumption and employment data, and we estimate sectoral productivity growth for 360 Indian districts. Our measurement of CS employment is consistent with the assumptions that such activities are nontradable and 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, to a large extent, producer services (PS) such as legal services, ICT, or consulting serve as inputs to the industrial sector and as such, their value-added can be shipped across locations.² Leveraging this distinction, our estimation exploits novel micro data on service-sector firms that report whether firms sell to consumers or to other firms.

Our analysis delivers four main findings. First, at the spatial level, there are large sectoral productivity differences, and CS shows the largest productivity gap between urban and rural districts. Thus, cities in India have a higher service employment share not only because their residents are richer, but also because CS are provided more efficiently. Second, service-led growth played an important role in economic development. At the aggregate level, the rising productivity of CS accounts for almost one-third of the increase in welfare between 1987 and 2011. Third, and most importantly, service-led growth has yielded strikingly unequal welfare effects. Productivity growth in CS was the main source of welfare gains for richer households living in urbanized districts. By contrast, living standards improved for poorer households from rural districts mostly due to productivity growth in agriculture. For instance, the average resident of districts in the top quintile of urbanization would have been better off taking a 37% income cut in 2011 than moving back to the 1987 productivity level of the CS sector. For the residents of districts in the three bottom quintiles of urbanization, the corresponding figure is a mere 13%. Finally, productivity growth in CS was a key driver of structural change. The agricultural employment share would have declined by just 9 p.p. (instead of 18 p.p. as it actually did) if productivity in CS had remained at its 1987 level.

²The stark assumption that CS are consumed locally is in line with the findings of [Gervais and Jensen \(2019\)](#), who estimate sector-specific trade costs and conclude that PS are as tradable as tangible goods, whereas trade costs in CS activities are substantially higher.

We carry out the main analysis under a set of stark assumptions aimed at retaining 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 in which India is an open economy with international trade flows calibrated to the data. In particular, we zoom in on the growing role of ICT services exports. Second, we relax the assumption that skills are perfect substitutes and assume, instead, that labor inputs provided by people with different educational attainment are imperfect substitutes. Moreover, we allow skill intensities to vary across sectors (e.g., agriculture is less skill-intensive), districts, and time (skill-biased technical change.) In this extension, changes in educational attainment are an engine of structural change and local comparative advantage. Finally, we allow for labor mobility across districts. While the quantitative results change to some extent in each extension, the broad picture is consistent and robust: service-led growth is a prominent feature of the Indian economy with major implications for both aggregate growth and the distribution of welfare gains.

Related Literature. Our paper contributes to the literature on structural transformation including, among others, Kongsamut, Rebelo, and Xie (2001), Ngai and Pissarides (2007), Herrendorf, Rogerson, and Valentinyi (2013), Gollin, Lagakos, and Waugh (2014), and García-Santana, Pijoan-Mas, and Villacorta (2021).

A recent strand of this literature focuses on the service sector. Buera and Kaboski (2012) emphasize the importance of skill-intensive services in the US since 1950. Hsieh and Rossi-Hansberg (2023) argue that in more recent years, ICT has been a major source of productivity growth. Their view is echoed by Eckert, Ganapati, and Walsh (2022). Chatterjee, Giannone, and Kuno (2023) associate rising productivity in services with regional divergence. A few studies focus on services in the developing world. Among them, Duarte and Restuccia (2010) document large cross-country productivity differences, Gollin, Jedwab, and Vollrath (2016) emphasize the relationship between urbanization and consumption of nontradable services, and, most recently, Nayyar, Hallward-Driemeier, and Davies (2021), use cross-country data to highlight the promise of service-led growth in today's developing world. Desmet, Ghani, O'Connell, and Rossi-Hansberg (2015), Dehejia and Panagariya (2016), and Lamba and Subramanian (2020) study aspects of the recent economic development of India, and document an important role for the service sector and cities. Jedwab, Ianchovichina, and Haslop (2022) analyze the link between premature deindustrialization and the growth of consumption cities characterized by high employment shares of urban nontradables. Their work is part of a broader debate on consumer cities, a notion stretching back to Max Weber, that was revived by Glaeser, Kolko, and Saiz (2001). While they emphasize the amenity value of cities, these authors also point at the efficiency gains of locating local services such as restaurants and theaters close to affluent consumers in urban areas. Atkin, Faber, and Gonzalez-Navarro (2018) study the welfare gains associated with the entry of global retail chains in Mexico that stem from pro-competitive effects on the prices charged by domestic stores. Finally, Chen, Pei, Song, and Zilibotti (2023) adopt the methodology of our paper to document the growing importance of productivity growth in services for China during the last 10 years.

On the methodological side, we build on the large literature on development accounting; see, for example, Caselli (2005), Hall and Jones (1999), and Klenow and Rodríguez-Clare (1997). This literature postulates aggregate production functions and uses information on the accumulation of productive factors to fit the data. Our paper is closer to the structural approach of Gancia, Müller, and Zilibotti (2013), who exploit the restrictions imposed by an equilibrium model to identify sectoral productivity. We perform our

accounting exercise in the context of a model with interregional trade linkages, which is commonly used in the economic geography literature; see, for example, Redding and Rossi-Hansberg (2017) or Allen and Arkolakis (2014). Budí-Ors and Pijoan-Mas (2022) link, as we do, spatial inequality with the process of structural change.

Nonhomothetic preferences play a central role in our analysis. Our paper is especially close to Boppart (2014) and Alder, Boppart, and Müller (2022), who propose PIGL preferences to study the process of structural transformation. Eckert and Peters (2022) incorporate these preferences into a spatial model of structural change. Comin, Lashkari, and Mestieri (2021) and Matsuyama (2019) build, instead, on the class of generalized CES preferences postulated by Sato (1975). In our paper, we use PIGL preferences because of their tractable aggregation properties. Our results on the unequal gains from service growth are reminiscent of Fajgelbaum and Khandelwal (2016), who measure the unequal gains from trade in a setting with nonhomothetic 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 Supplemental Appendix (Fan, Peters, and Zilibotti (2023b)) contains details of the theoretical and empirical analysis. A Web Appendix, which is available from the working paper version of this article (Fan, Peters, and Zilibotti (2023a)), contains additional results.

2. STRUCTURAL CHANGE IN INDIA AND THE DEVELOPING WORLD

Between 1987 and 2011, India experienced fast economic development: income per capita grew by a factor of three and the employment structure changed markedly. The upper left panel of Figure 1 highlights the pattern of structural change with low industrialization: most of the transformation took the form of an outflow out of agriculture and an inflow into services and construction whose employment shares increased by 9 and 7 p.p., respectively. By contrast, manufacturing employment was stagnant. Today, the service sector accounts for about one-third of aggregate employment.

A large part of this expansion originated in services that facilitate consumers' access to final consumption goods. The upper-right panel of Figure 1 decomposes the service sector into four subsectors.³ The first group serves mostly consumers. These service industries grew significantly after 1987 and employed almost 55% of all Indian service workers in 2011. The second group, which sells a large part of their services to industrial firms, also grew substantially but only accounted for a tenth of service employment. For instance, ICT, a fast-growing industry, accounts for less than 1% of total employment in 2011. Transport services, which serves both consumers and industries, also expanded. Finally, the employment share of mostly government-run activities such as public administration and education remained constant over time. Figure 1 also shows that all service activities are much more prevalent in urban areas.

³ Using the official NIC classification, the four subsectors contain the following industries: (i) wholesale and retail trade; repair of motor vehicles and motorcycles; accommodation and food services; health and social work; arts and entertainment; other service activities, (ii) finance and insurance; ICT; real estate; professional, scientific, and technical activities; administrative and support services; publishing, (iii) transport and storage, and (iv) education and public administration.

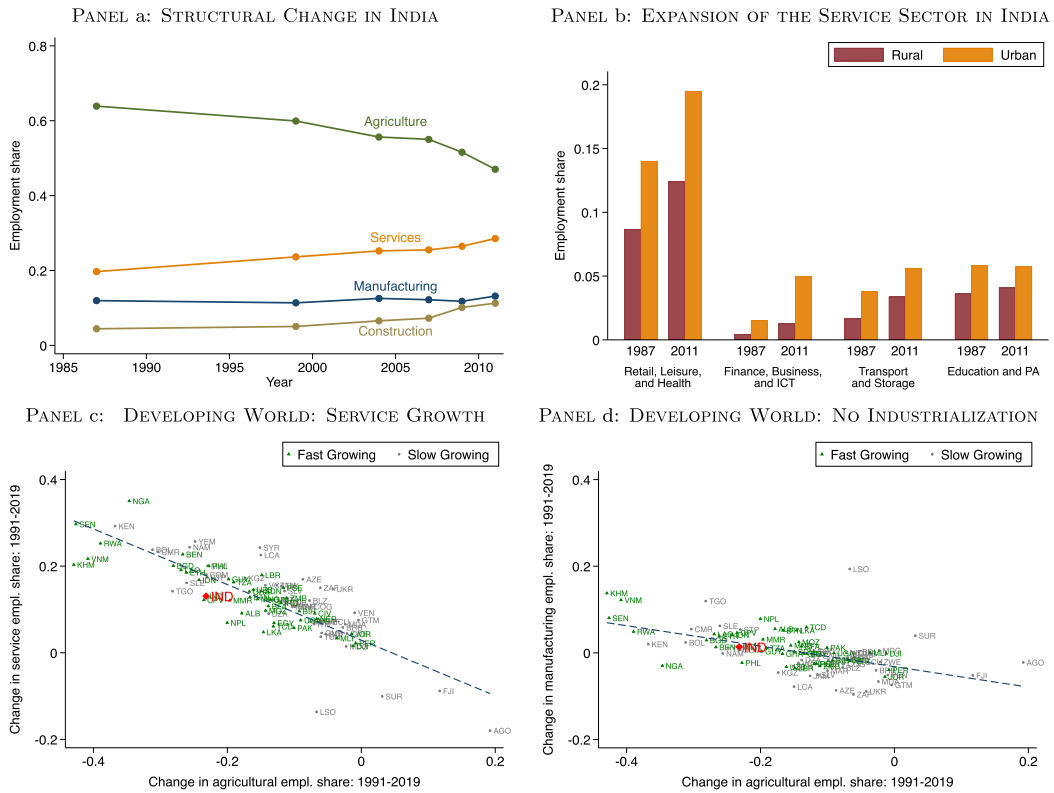


FIGURE 1.—Structural Change toward Services in India and in the Developing World. The upper-left panel shows the evolution of sectoral employment shares in India. The upper-right panel shows employment shares for different service industries (see footnote 3 for details), separately for rural and urban districts. We split India into rural and urban districts so that half of the population belongs to each type of district. The lower-left (lower right) panel shows the cross-country correlation between changes in agricultural employment shares and changes in service employment shares (manufacturing employment shares) between 1991–2019 for all non-OECD countries that are poorer than China in 2019. Countries with an average growth rate exceeding 4% are labeled in green. Panels (a) and (b) are constructed from the micro data from India NSS (see Section 4); Panels (c) and (d) use data from the International Labor Organization and the Penn World Tables.

India’s pattern of a decline in agriculture with low industrialization is by no means exceptional in today’s developing world. In the lower panels of Figure 1, we display the cross-country relationship between the change in the employment share of agriculture and those of services (left panel) and manufacturing (right panel) during 1991–2019.

To home in on the developing world, we include all non-OECD countries whose income per capita was below that of China in 2019. The left panel shows a strong negative relationship: a 10 p.p. reduction in the agriculture share is matched on average by a 6.4 p.p. increase in the service share. The right panel shows that the relationship, albeit negative, is substantially weaker for the industrial sector: a 10 p.p. reduction in the agriculture share is associated with a 2.4 p.p. increase in the manufacturing share.

Crucially, the low speed of industrialization is *not* a mark of lackluster development. In Figure 1, we indicate *fast-growing* countries (which we define as countries with an annual growth rate of at least 4%) with green labels. While these countries experienced faster declines in the agricultural employment share, they still saw a substantial expansion of

the service sector: on average, the agricultural employment share declined by 18 p.p. and the employment share of services grew by 13 p.p. Moreover, Figure 1 shows that the typical developing country indeed grew *like India*: the observation for India, highlighted in red, is not far from the regression line.⁴ Nor is the predominance of CS relative to PS a special feature of India: in Appendix Figure B-1, we show that the pattern of panel (b) of Figure 1 is perfectly in line with the international evidence.

3. THEORY

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

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

3.1. Technology

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

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

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

where x_{rFt} and x_{rGt} denote the inputs of food and goods, respectively; H_{rCS} 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 CS content (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 ($\lambda_{nCS} \approx 1$).

⁴Services also play an increasingly dominant role in advanced economies. The main difference is that in richer nations the service sector mostly grows at the expense of manufacturing rather than agriculture. Even in a country like China, whose stellar growth has been led for decades by the manufacturing sector, services have gained significant ground in the last 10 years while the employment share of manufacturing has been shrinking (Chen et al. (2023)).

⁵As we describe in more detail below, we assume that the industrial sector employs both manufacturing and PS workers. Because the value-added of, say, corporate lawyers and consultants is embodied in industrial goods, PS are ultimately tradable.

⁶The representation of technology in (1) is akin to the Cobb–Douglas input–output structure commonly assumed in the production network literature; see Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi (2012). The scalar $\tilde{\lambda}_n \equiv \lambda_{nF}^{-\lambda_{nF}} \lambda_{nG}^{-\lambda_{nG}} \lambda_{nCS}^{-\lambda_{nCS}}$ is an inconsequential normalization to simplify expressions.

The tradable food and industrial good are CES aggregates of regional varieties:

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

which are produced according to the linear technologies

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

where sectoral productivities A_{rst} can differ across regions. We refer to A_{rmt} in (1) as CS productivity even though it applies to all inputs. The assumption that CS must be supplied locally allows us to separately identify A_{rmt} from A_{rFt} and A_{rGt} .

Nontradable CS versus Tradable PS. In our theory, tradability is the key difference between CS and PS. While CS value-added can only be consumed locally, the PS value-added is embodied in goods and is ultimately tradable.

When mapping the model to the data, we include the value-added of PS in the industrial sector; namely, we let $H_{rGt} = H_{rMt} + H_{rPS_t}$.⁷ This specification does *not* restrict manufacturing and PS workers to being perfect substitutes. To see why, suppose industrial firms combine the inputs of manufacturing workers and PS to produce industrial goods using the technology $y_{rGt} = g_{rt}(H_{rMt}, H_{rPS_t})$, where g_{rt} is a linearly homogeneous function. As long as firms maximize profits, the marginal products of H_{rMt} and H_{rPS_t} 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 such as Delhi or Bangalore with a comparative advantage in tradable PS like finance or ICT to export the value-added of PS to the rest of India (and, in Section 7, even internationally).

3.2. Preferences and Demand System

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

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

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

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

⁸Linear homogeneity allows us to write $y_{rGt} = g_{rt}(1 - s_{rPS_t}, s_{rPS_t})H_{rGt}$, where $s_{rPS_t} = H_{rPS_t}/H_{rGt}$. 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 $g = [(A_{rMt}H_{rMt})^{(\sigma-1)/\sigma} + (A_{rPS_t}H_{rPS_t})^{(\sigma-1)/\sigma}]^{\sigma/(\sigma-1)}$. Then, $A_{rGt} = (A_{rMt}^{\sigma-1} + A_{rPS_t}^{\sigma-1})^{1/(\sigma-1)}$.




FIGURE 2.—Engel curves. The figure shows the good-specific expenditure share as a function of income e (see (3)).

where e denotes total spending and \mathbf{p}_r is the vector of prices in region r . The mnemonic FE is a reminder that the indirect utility function in (2) 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 $B(\mathbf{p}_r) = \exp(\int_{n=0}^1 \beta_n \ln p_{rn} dn)$ and $D(\mathbf{p}_r) = (\int_{n=0}^1 \kappa_n \ln p_{rn} dn)$, where $\int_0^1 \beta_n dn = 1$ and $\int_0^1 \kappa_n dn = 0$.⁹

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

$$\vartheta_n^{FE}(e, \mathbf{p}_r) = \beta_n + \kappa_n \left(\frac{e}{\exp\left(\int_n \beta_n \ln p_{rn} dn\right)} \right)^{-\varepsilon}. \tag{3}$$

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

3.2.1. Final Expenditure and Value-Added

Equation (3) defines the expenditure shares over final products. For our purposes, it is essential to derive a demand system for the value-added produced by the three grand sectors F , G , and CS , because we estimate our model using data on sectoral employment. To derive this value-added demand system, note first that the prices of tradable goods are

⁹Our functional form for $D(\mathbf{p}_r)$ is more restrictive than the one in Boppart (2014). In Section 7.3, we generalize the preference structure along the lines of his original contribution.

given by the usual CES price indices

$$P_{rst}^{1-\sigma} = \sum_{j=1}^R \tau_{rj}^{1-\sigma} A_{jst}^{\sigma-1} w_{jt}^{1-\sigma}, \quad \text{for } s \in \{F, G\}, \tag{4}$$

where $\tau_{rj} \geq 1$ is the iceberg cost of shipping variety j to region r . The price of final good n in region r is then given by $p_{rnt} = P_{rFt}^{\lambda_{nF}} P_{rGt}^{\lambda_{nG}} (\mathcal{A}_{rnt}^{-1} w_{rt})^{\lambda_{nCS}}$, where w_{rt} denotes the wage in region r . Plugging this expression into the indirect utility function, (2) yields a representation of consumers’ preferences over sectoral value-added aggregates.

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

$$\mathcal{V}(e, \mathbf{P}_r) = \frac{1}{\varepsilon} \left(\frac{e}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} P_{rCS}^{\omega_{CS}}} \right)^\varepsilon - \sum_{s \in \{F, G, CS\}} \nu_s \ln P_{rst}, \tag{5}$$

where $\mathbf{P}_r = (P_{rFt}, P_{rGt}, P_{rCS})$, $P_{rCS} \equiv A_{rCS}^{-1} w_{rt}$, P_{rFt} and P_{rGt} are given by (4), and

$$\omega_s \equiv \int_n \lambda_{ns} \beta_n dn, \quad \nu_s \equiv \int_n \lambda_{ns} \kappa_n dn, \quad \text{and} \quad \ln A_{rCS} \equiv \int_n \frac{\beta_n \lambda_{nCS}}{\omega_{CS}} \ln \mathcal{A}_{rnt} dn. \tag{6}$$

The associated value-added expenditure shares are given by

$$\vartheta_{rst}(e, \mathbf{P}_r) = \omega_s + \nu_s \left(\frac{e}{P_{rFt}^{\omega_F} P_{rGt}^{\omega_G} P_{rCS}^{\omega_{CS}}} \right)^{-\varepsilon}. \tag{7}$$

PROOF: See Appendix A-1.

Q.E.D.

Proposition 1 states three important properties of our theory. First, the indirect utility function defined over value-added also falls into the PIGL class and has the same functional form as the corresponding expressions over final products (2). In particular, the expenditure share over sectoral value-added, ϑ_{rst} in (7), features the same Engel elasticity ε as in (3). This result enables us to estimate ε from micro data for household expenditure shares on final products and then use it in the value-added demand system.

Second, the regional CS productivity index A_{rCS} , which is akin to the average CS productivity of all final products, \mathcal{A}_{rnt} , weighted by their CS content λ_{nCS} and their asymptotic spending share β_n , is a sufficient statistic for the local CS sector. Because preferences are nonhomothetic and CS are provided locally, productivity growth yields heterogeneous welfare effects. If goods with a high CS content are luxuries, productivity growth in CS is skewed toward rich consumers.¹⁰ Moreover, given its nontradable nature, CS productivity growth predominantly benefits local residents. Thus, if urban districts experience faster productivity growth, city dwellers are going to be the main beneficiaries of service-led growth. In contrast, the benefits from productivity growth in tradable sectors diffuse spatially through trade.

¹⁰In fact, the expenditure share $\vartheta_{rst}(e, \mathbf{P}_r)$ exactly measures the welfare exposure of a change in prices at the individual level. Formally, letting $e(\mathbf{P}_r, \mathcal{V})$ denote the expenditure function associated with the utility level \mathcal{V} given the price vector \mathbf{P}_r , $\partial \ln e(\mathbf{P}_r, \mathcal{V}) / \partial \ln P_{rst} = \vartheta_{rst}(e, \mathbf{P}_r)$.

Third, whether sectoral value added is a luxury or a necessity depends on the *correlation* of the good-specific demand parameters κ_n with their factor intensities λ_{ns} . The expenditure share for sectoral value-added is rising in income if and only if $\nu_s < 0$, that is, if income-elastic *products* have a large *sectoral* input requirement. By contrast, if all goods were produced with equal factor proportions, or more generally if λ_{ns} were orthogonal to κ_n , the demand for sectoral value-added would be homothetic even though the underlying demand for final products is nonhomothetic. However, the demand system is fully determined by the parameters ν_s and ω_s and the aggregate CS productivity index $A_{rCS,t}$, and does not separately depend on the preference parameters defined over final goods $[\beta_n, \kappa_n]_{n=0}^1$, nor on the product-specific CS productivities $[\mathcal{A}_{rnt}]_{n=0}^1$.

The closed-form expression of the mapping from the final expenditure to the value-added demand system in Proposition 1 hinges on the assumption that the final good production function is Cobb–Douglas (cf. equation (1)).¹¹ In Section WA-1.1 of the Web Appendix, we extend our analysis to a setting where (1) takes a CES form. In this case, we can still obtain an analytical characterization where the final expenditure and value-added representations share the same Engel elasticity ε , that is, we can derive the analogue of equation (7). However, estimating the CES model would require additional data about the expenditure on individual final goods.

3.2.2. Heterogeneity and Aggregate Demand

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

Suppose individuals differ in their human capital that determines the number of efficiency units of labor supplied to the market. Individual h 's income is then given by $e_{rt}^h = q^h w_{rt}$, 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 empirically relate to the regional data on educational attainment.

Because our analysis abstracts from savings and capital accumulation, income equals expenditure. Defining with slight abuse of notation the expectation operator $\mathbb{E}_{rt}[x] \equiv \mathbb{E}[x; F_{rt}(x)]$, equation (7) implies that the *aggregate* spending share on value-added produced in sector s by consumers residing in region r is given by

$$\bar{\vartheta}_{rst} \equiv \frac{L_{rt} \int \vartheta_{rst}(q w_{rt}, \mathbf{P}_{rt}) q w_{rt} dF_{rt}(q)}{L_{rt} \int q w_{rt} dF_{rt}(q)} = \omega_s + \bar{\nu}_{rst} \left(\frac{\mathbb{E}_{rt}[q] w_{rt}}{P_{rF,t}^{\omega_F} P_{rG,t}^{\omega_G} P_{rCS,t}^{\omega_{CS}}} \right)^{-\varepsilon}, \tag{8}$$

where

$$\bar{\nu}_{rst} \equiv \frac{\mathbb{E}_{rt}[q^{1-\varepsilon}]}{\mathbb{E}_{rt}[q]^{1-\varepsilon}} \nu_s. \tag{9}$$

¹¹More formally, using the expression for p_{rnt} , we can express $B(\mathbf{p}_r)$ as

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

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

Comparing (8) with (7) clarifies the sense in which PIGL allows for a representative household: the *aggregate* demand system in (8) is isomorphic to that of a consumer in region r who earns the average income $\mathbb{E}_{r_t}[q]w_{r_t}$ and has the inequality-adjusted preference parameter \bar{v}_{rst} in (9). Crucially, the Engel elasticity of the aggregate demand system, ε , is the same as at the individual level.

The inequality adjustment term $\mathbb{E}_{r_t}[q^{1-\varepsilon}]/\mathbb{E}_{r_t}[q]^{1-\varepsilon}$, depends, in general, on the distribution of efficiency units F_{r_t} .¹² The analysis further simplifies if we assume q follows a Pareto distribution with c.d.f. $F_{r_t}(q) = 1 - (\underline{q}_{r_t}/q)^\zeta$. In this case, equation (9) boils down to

$$\bar{v}_{rst} = \bar{v}_s = \frac{\zeta^\varepsilon (\zeta - 1)^{1-\varepsilon}}{\zeta + \varepsilon - 1} \nu_s.$$

Thus, if income is Pareto distributed with a common tail parameter ζ , \bar{v}_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 Engel elasticity (ε). Given \bar{v}_s , the distribution F_{r_t} only enters through the average income term $\mathbb{E}_{r_t}[q]w_{r_t} = \frac{\zeta}{\zeta-1} \underline{q}_{r_t} w_{r_t}$.

3.2.3. Welfare and Inequality

The aggregation properties of PIGL come in especially handy for welfare analysis. To this aim, define the utilitarian welfare function at the regional level as $\mathcal{U}_{r_t}(w_{r_t}, \mathbf{P}_{r_t}) \equiv \int \mathcal{V}(qw_{r_t}, \mathbf{P}_{r_t}) dF_{r_t}(q)$. Plugging in the indirect utility function in (5) yields

$$\mathcal{U}_{r_t}(w_{r_t}, \mathbf{P}_{r_t}) = \frac{\zeta^{1-\varepsilon} (\zeta - 1)^\varepsilon}{\zeta - \varepsilon} \times \left(\frac{1}{\varepsilon} \left(\frac{\mathbb{E}_{r_t}[q]w_{r_t}}{P_{rF_t}^{\omega_F} P_{rG_t}^{\omega_G} P_{rCS_t}^{\omega_{CS}}} \right)^\varepsilon - \sum_{s \in \{F, G, CS\}} \nu_s^\mu \ln P_{rst} \right), \quad (10)$$

where $\nu_s^\mu \equiv \bar{v}_s \times ((\zeta - \varepsilon)(\zeta - (1 - \varepsilon)))/(\zeta(\zeta - 1))$. 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 $\mathbb{E}_{r_t}[q]w_{r_t}$ and the inequality-adjusted taste parameter ν_s^μ .

3.3. Equilibrium

We can now characterize the competitive equilibrium.

PROPOSITION 2: *The sectoral labor allocations $\{H_{rF_t}, H_{rG_t}, H_{rCS_t}\}_r$ and local wages $\{w_{r_t}\}$ are determined by the following equilibrium conditions:*

1. *Market clearing for local CS:*

$$w_{r_t} H_{rCS_t} = \left(\omega_{CS} + \bar{v}_{CS} \left(\frac{A_{rCS_t}^{\omega_{CS}} \mathbb{E}_{r_t}[q] w_{r_t}^{1-\omega_{CS}}}{P_{rF_t}^{\omega_F} P_{rG_t}^{\omega_G}} \right)^{-\varepsilon} \right) w_{r_t} H_{r_t}, \quad (11)$$

where P_{rF_t} and P_{rG_t} are given by (4).

¹²Note that $\mathbb{E}_{r_t}[q^{1-\varepsilon}]/\mathbb{E}_{r_t}[q]^{1-\varepsilon} \equiv 1 - (\text{ATK}_\varepsilon(q^h))^{1-\varepsilon}$, where ATK_ε denotes the Atkinson index. Thus, if $\varepsilon \in (0, 1)$, $\mathbb{E}_{r_t}[q^{1-\varepsilon}]/\mathbb{E}_{r_t}[q]^{1-\varepsilon} \in [0, 1]$ is an inverse measure of income inequality. Moreover, $\bar{v}_{rst} \leq \nu_s$, that is, the aggregate expenditure share varies less than the underlying individual share with total expenditure. The gap between \bar{v}_{rst} and ν_s increases with inequality. Thus, a mean-preserving spread in district-level income reduces \bar{v}_{rst} and the extent to which the district-level expenditure changes with income. Intuitively, more inequality increases the weight on the expenditure of richer households whose preferences are closer (under our PIGL representation) to homothetic. Note that inequality does not affect the Engel elasticity ε .

2. *Market clearing for tradable goods:*

$$w_{rt}H_{rst} = \sum_{j=1}^R \pi_{rsjt} \left(\omega_s + \bar{v}_s \left(\frac{A_{jCS}^{\omega_{CS}} \mathbb{E}_{jt}[q] w_{jt}^{1-\omega_{CS}}}{P_{jFt}^{\omega_F} P_{jGt}^{\omega_G}} \right)^{-\varepsilon} \right) w_{jt}H_{jt}, \tag{12}$$

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

Proposition 2 characterizes the sectoral employment allocations and equilibrium wages across space. The contrast between equations (11) and (12) reflects the tradable nature of food and goods versus the nontradable nature of CS. The demand for CS value-added hinges on both local income and local CS productivity. For instance, the retail sector could be large in urban districts either because local consumers are more educated and richer or because more-efficient department store chains open branches in large cities. Instead, the demand for tradable goods originates from all localities.

4. EMPIRICAL ANALYSIS: DATA AND MEASUREMENT

Our analysis relies on five data sets: (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). We defer a more detailed description of these data sets to Appendix B-2.

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

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

The NSS Consumer-Expenditure Schedule contains information on households’ expenditure on different categories of final goods that we use to estimate the Engel elasticity ε .

¹³In Section WA-5.3 in the Web Appendix, we document that average expenditure is strongly correlated with average wages and average income per capita at the district level.

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 Service Survey was conducted in 2006 and is representative of India's service sector. It covers almost 200,000 private enterprises subdivided into seven service industries.¹⁴ Finally, we rely on ETD for measuring the average relative price of agricultural goods (while we do not use any published price index for services).

Geography. To compare spatial units over time, we create a time-invariant definition of Indian districts. Appendix B-3 describes in detail how we construct this crosswalk. Because the boundaries of several districts changed over time, we harmonized them using GIS software, relying on maps for the years 1991, 2001, and 2011. We exclude two small districts that existed in 2011 but did not exist in 1987. We also exclude districts with less than 50 observations because they do not allow us to precisely estimate sectoral employment shares. In the end, we obtain 360 regions that cover the vast majority of the Indian territory.

Consumer versus Producer Services. A key step in our measurement is to distinguish between CS, that is, nontradable services catering to consumers, and PS, that is, services which are used as intermediate inputs. To perform this split, we combine information from the EC and the Service Survey.

We aim to assign firms to the CS sector if they sell to consumers and to the PS sector if they sell to other firms. Ideally, we would use firm-level input-output matrices. To the best of our knowledge, this information is not available in India for the time period of our study. We therefore leverage micro data on firms' downstream trading partners contained in the Service Survey, which reports whether a firm sells mostly to consumers or to other firms. The Service Survey contains too few observations to precisely estimate the employment shares of firms selling to consumers in 360 districts within narrowly defined industries. We therefore rely on the fact that the propensity to sell to other firms is highly correlated with firm size. As Table I shows, only 6% of firms with three employees sell to other firms, while the share increases to 43% for firms with more than 50 employees.

TABLE I
SHARE OF PRODUCER SERVICES BY FIRM SIZE.

	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	97,337	46,571	13,227	5156	2777	4841	2830	601	403

Note: The table reports the share of firms selling to firms (rather than private individuals) in different size categories.

¹⁴These 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, and (vii) other personal service activities. In Appendix B-2.3, we compare the Service Survey with the EC and document that it is indeed representative of the distribution of firm size in India.

TABLE II
SHARE OF CONSUMER SERVICES WITHIN SERVICE EMPLOYMENT.

	Overall	In Selected Categories				Across Space	
		Retail, Leisure, and Health	Finance and Business	ICT	Transport and Storage	Urban	Rural
Share of CS	89	97	82	47	70	88	91

Note: The table reports the share of service employment allocated to the CS sector. To aid readability, we aggregate the service industries into four categories.

We use the pattern in Table I in the following way. First, we estimate the CS employment share by firm size for different service industries.¹⁵ Then we use the *district*-specific size distribution from the EC to infer the aggregate CS employment share in district r . More formally, we compute the CS employment share in service industry k in region r as $s_{rk}^{CS} = \sum_b \omega_{kb}^{CS} \ell_{kbr}$, where ω_{kb}^{CS} is the share of employment in firms selling to consumers in service industry k in size class b , and ℓ_{kbr} is the employment share of firms of size b in service industry k in region r . The spatial variation in CS employment thus stems from differences in: (i) total service employment, (ii) the relative importance 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 II, we report the resulting allocation of employment to CS. At the aggregate level, our procedure allocates 89% of service employment to CS and 11% to PS. This allocation differs across service industries. For instance, within the retail and restaurant industry, 97% of workers are employed by establishments catering to consumers. Instead, in the ICT sector, less than half of employment caters to consumers.¹⁶ Moreover, the share of CS within services is smaller in urban areas, reflecting the more prominent role of ICT and business activities in cities.

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 Nonagricultural Enterprises Survey 1999–2000” data set, which covers the construction sector and also reports whether a firm sells to consumers or other firms. These data imply that 13% of private sector construction employment is associated with producer services; see Appendix B-4.3.

In Section 7, we show that our results are robust to alternative measurement strategies, such as (i) allocating ICT and business services entirely to PS, (ii) splitting PS and CS according to aggregate Input–Output tables, and (iii) allocating construction to the industrial sector, rather than services.

Human Capital. Consistent with our theory, we measure each district’s endowment of human capital, $F_{rt}(q)$, and its distribution across sectors in terms of efficiency units of labor. We classify people into four educational groups: (i) less than primary school, (ii) primary and upper primary/middle school, (iii) secondary school, and (iv) more than secondary school. We associate each step in the education ladder with 3 extra years of

¹⁵We split the service sector into seven categories: “Retail and wholesale,” “Hotels and restaurants,” “Transport,” “Finance,” “Business services and ICT,” “Health,” and “Community services.”

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

education, consistent with the organization of schools in India, and measure the effect of each additional year by an estimated Mincerian return to schooling ρ (see Section 5.1 below).

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

5. ESTIMATION: IDENTIFICATION AND RESULTS

We now turn to the estimation of the model. Our approach is in the tradition of development accounting; see, for example, Caselli (2005), Hall and Jones (1999), and Gancia, Müller, and Zilibotti (2013)). Whereas those studies infer productivity from an aggregate production function, we rely on the equilibrium structure of our model and estimate the entire distribution of productivity $\{A_{rst}\}$ across sectors, space, and time.

The model has eight preference parameters and two parameters for the skill distribution: $\Omega = \{(\varepsilon, \nu_{CS}, \nu_F, \nu_G, \omega_{CS}, \omega_F, \omega_G, \sigma), (\rho, \zeta)\}$. In addition, each region is characterized by a 3-tuple of regional productivity levels in agriculture, industry, and CS: $\mathbf{A}_{rt} = \{A_{rFt}, A_{rGt}, A_{rCS t}\}$. Given the parameter vector Ω , there exists a one-to-one mapping from equilibrium skill prices $\{w_{rt}\}$ and sectoral employment allocations $\{H_{rst}\}$ to the underlying productivity fundamentals in \mathbf{A}_{rt} . In Section 5.1, we describe how we estimate the vector of structural parameters Ω . 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 is the crucial parameter in our analysis. It determines how fast the expenditure on food shrinks and, conversely, how fast it expands for CS as income rises. To estimate ε , we use the cross-sectional relationship between household income and expenditure shares on food.

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

$$\vartheta_{\mathcal{F}}^{\text{FE}}(e, \mathbf{p}_r) = \beta_{\mathcal{F}} + \kappa_{\mathcal{F}} \left(\frac{e}{\exp\left(\int_n \beta_n \ln p_{rn} dn\right)} \right)^{-\varepsilon}, \tag{13}$$

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 (13) yields a log-linear

relationship between household income and expenditure shares:¹⁷

$$\ln \vartheta_{\mathcal{F}}^{\text{FE}}(e, \mathbf{p}_r) \approx \varepsilon \left(\int_n \beta_n \ln p_{rn} dn \right) - \varepsilon \times \ln e + \ln \kappa_{\mathcal{F}}. \quad (14)$$

We can then estimate ε from the linear regression

$$\ln \vartheta_{\mathcal{F}}^h = \delta_r + \varepsilon \times \ln e_h + x'_h \psi + u_{rh}, \quad (15)$$

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 effect, and x_h is a set of household characteristics that could induce a correlation between total spending $\ln e_h$ and food shares. Comparing (15) with (14), it is apparent that the terms $(\int_n \beta_n \ln p_{rn} dn)$ and $\ln(\kappa_{\mathcal{F}})$ are absorbed in the region fixed effects δ_r .

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

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

In column 6, we present the results from an IV regression addressing concerns about measurement error and unobserved income shocks that could bias the estimate. We instrument total expenditure with a full set of three-digit occupation fixed effects.¹⁹ The exclusion restriction is that occupations only affect spending shares on food through income. The instruments have a strong predictive power in the first-stage regression (F -Stat = 62). The IV estimate of 0.395 is larger than the OLS estimate.

In Figure 3, we show a binscatter plot of the data for log food expenditure shares versus log expenditure after absorbing district-food item fixed effects, that is, corresponding to specification (5). Consistent with our PIGL specification, the relationship is indeed

¹⁷The assumption that $\beta_{\mathcal{F}}$ is small is convenient but inconsequential. In Appendix C-1.1, we estimate ε from (13) without imposing this restriction. We find that $\beta_{\mathcal{F}} = 0$ is, in fact, the best estimate. Moreover, we also estimate ε for a range of value of $\beta_{\mathcal{F}}$ and find that they are very similar to the ones reported in Table III. In column 3 of Table III, we report the estimate for $\beta_{\mathcal{F}} = 0.05$.

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

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

TABLE III
ESTIMATES OF THE ENGEL ELASTICITY ε .

	Food Exp. Share								Pooled Data	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\ln \varepsilon$	-0.332 (0.008)	-0.321 (0.007)	-0.369 (0.009)	-0.313 (0.008)	-0.334 (0.007)	-0.395 (0.013)			-0.230 (0.007)	-0.392 (0.023)
$\ln \varepsilon \times$ below median							-0.218 (0.010)			
$\ln \varepsilon \times$ above median							-0.415 (0.011)			
$\ln \varepsilon \times$ low urbanization								-0.291 (0.007)		
$\ln \varepsilon \times$ high urbanization								-0.358 (0.012)		
Trim (top and bottom 5%)		✓	✓	✓	✓	✓	✓	✓	✓	✓
$\beta_{\mathcal{F}} = 0.05$			✓						✓	✓
Serv. Categories									✓	✓
Addtl. Controls				✓	✓	✓	✓	✓	✓	✓
IV						✓				✓
N	101,650	91,492	91,488	91,443	1,129,730	85,919	91,443	91,443	182,068	171,190
R^2	0.476	0.425	0.417	0.437	0.635	0.197	0.446	0.439	0.822	0.032

Note: The table shows the estimated coefficient ε of the regression (15). In columns 1–4 and 6–8, the dependent variable is the income share spent by each household on a set of 17 items classified as “food.” these are: Beverages; cereals; cereal substitutes; dry fruit, edible oil; egg, fish, and meat; fresh fruit; intoxicants; milk and milk products; pan; packaged processed food products; pulses and products; salt and sugar; served processed food; spices; tobacco; vegetables. In column 5, the dependent variable is the income share spent on each of these 17 items. In all specifications, we control for a (within-district) urban/rural dummy, a set of fixed effects for household size, and the number of workers within the household. All regressions include region fixed effects; region-food item fixed are included in the fifth column. In columns 6 and 10, we instrument expenditure with a set of occupation fixed effects. In columns 9 and 10, we consider a pooled regression, where the dependent variables are $\ln(\vartheta_{\mathcal{F}}^h - \beta_{\mathcal{F}})$ for food items and $\ln(\beta_{\mathcal{S}} - \vartheta_{\mathcal{S}}^h)$ for service items. Standard errors, clustered at the district level, are in parentheses.




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-item 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. Compare the regression in column 5 of Table III.

approximately log-linear. However, careful scrutiny reveals some mild concavity suggesting a higher elasticity for high-income consumers. In column 7 of Table III, we therefore allow for different elasticities for households above and below the median income. The estimated elasticity is somewhat larger for high-income households.

In column 8, we allow the elasticity to differ between rural and urban districts. We define all districts in the top quartile of the distribution of urbanization as urban. While urban locations have higher elasticities, the differences are moderate.

Even though estimating ε from the expenditure system for food is consistent with our theory, we can also use the information for other expenditure categories. The expenditure survey contains information on spending on some consumer service categories, such as domestic servants, barber shops, or tailor services; see Appendix C-1.2. In columns 9 and 10, we pool the expenditure shares on these services with those on food items and estimate ε using both sources of variation. More formally, we estimate (15) using as dependent variable $\ln(\vartheta_{\mathcal{F}}^h - \beta_{\mathcal{F}})$ for food items and $\ln(\beta_S - \vartheta_S^h)$ for services. Note that $\beta_S > \vartheta_S^h$ if services are luxuries.²⁰

We set β_S to match the expenditure share of the 99% quantile of the observed distribution in India. This yields $\beta_S = 0.2$.²¹ For food items, we set $\beta_{\mathcal{F}} = 0.05$ as in column 3. In these regressions, we control for a full set of interactions of district-item fixed effects to account for price differences across both locations and types of final goods or services. While the OLS elasticity is smaller in column 9 than in column 4, the estimated coefficient in the IV regression of column 10 is almost identical to its analogue in column 5. We conclude that the results are robust to the inclusion of expenditure on some services.²²

For our baseline analysis, we set the Engel elasticity ε equal to the IV estimate of 0.395. As we show in Section 7, this turns out to be a conservative choice because the welfare gains attributed to CS productivity growth are decreasing in ε , implying that the effects we emphasize would be larger if we relied on the OLS rather than the IV estimates. Moreover, this estimate is closer to the estimates for rich households and urban locations where concerns about nonmeasured subsistence food consumption are less salient.

Other Preference Parameters. We estimate the six remaining parameters of the demand system, ω_s and \bar{v}_s , directly from the equilibrium conditions.²³ In Appendix A-2, we show that the market clearing conditions imply

$$\sum_{r=1}^R w_{rt} H_{rFt} = \omega_F \sum_{r=1}^R w_{rt} H_{rt} + \bar{v}_F \sum_{r=1}^R \left(\omega_{CS} - \frac{H_{rCS_t}}{H_{rt}} \right) w_{rt} H_{rt}. \tag{16}$$

²⁰Equation (13) implies that $\ln(s_{nrt}^h) = v_n + \varepsilon \exp(\int_n \beta_n \ln p_{rn} dn) - \varepsilon \ln e^h$, where for a necessity, $s_{nrt}^h = \vartheta_{nrt}^h - \beta_n$ and $v_n = \ln(\kappa_n)$ and, for a luxury, $s_{nrt}^h = \beta_n - \vartheta_{nrt}^h$ and $v_n = \ln(-\kappa_n)$.

²¹In principle, one could estimate β_S and ε jointly. However, β_S would solely be identified from the shape of Engel curves of consumers with expenditure shares below β_S . In addition, it needs to satisfy the theoretical restriction of describing the asymptotic spending share on categories, we have measures for (i.e., domestic servants, barber shops, tailor services, etc.). We therefore prefer to directly rely on the 99% quantile of the observed expenditure shares in our data.

²²For our main specification, we rely exclusively on food expenditure data for two reasons. First, we believe they are more precisely measured. Second, we are guided by a precise prior on the asymptotic expenditure share. In Appendix C-1.2, we estimate ε using service expenditure alone. Reassuringly, the IV estimate of the Engel elasticity is not significantly different from that of column 6.

²³The market-level demand system depends on the aggregate preference parameters \bar{v}_s , which are related to the primitive micro-level preference parameters v_s via (9). Identifying v_s is only required to quantify the welfare consequences of service-led growth, not to estimate the model.

TABLE IV
STRUCTURAL PARAMETERS

Parameter	Target	Value	
Preference parameters	ε	Engel elasticity (Table III, column 6)	0.395
	ω_F	Agricultural spending share US	0.01
	ω_{CS}	Equation (16), $t \in \{1987, 2011\}$	0.692
	ω_G	Implied by $\sum_s \omega_s = 1$	0.298
	\bar{v}_F	Equation (16), $t \in \{1987, 2011\}$	1.276
	\bar{v}_{CS}	Normalization	-1
	\bar{v}_G	Implied by $\sum_s \bar{v}_s = 0$	-0.276
Skill parameters	σ	Set exogenously	5
	ρ	Mincerian schooling returns	0.056
	ζ	Earnings distribution within regions	3

Note: The table summarizes the estimated structural parameters. The details of the estimation are discussed in the text.

Since these equations must hold for both $t = 1987$ and $t = 2011$, they represent two moment conditions for the three parameters ω_F , ω_{CS} , and \bar{v}_F . Note that these equations are independent of ε , trade costs, the elasticity of substitution σ , and the skill distribution. To attain identification, we exploit that ω_F pins down the asymptotic value-added share of the agricultural sector. In the US, the agricultural employment share (as well as its value-added share) is about 1%. Hence, we set $\omega_F = 0.01$ and use (16) for $t = 1987$ and $t = 2011$ to identify \bar{v}_F and ω_{CS} .

As we show in Appendix A-2, \bar{v}_{CS} is not separately identified from A_{rCS_t} . The average level of A_{rCS_t} plays no role in our analysis. Under the assumption of stable preferences, we can still calculate the growth over time of A_{rCS_t} , which is our main object of interest. Therefore, without loss of generality, we set $\bar{v}_{CS} = -1$. The remaining parameters ω_G and \bar{v}_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, which is a consensus estimate in the literature.

In the first panel of Table IV, we report the resulting estimates. The implied 70% asymptotic value-added share of CS, ω_{CS} , is reasonable.²⁴ For instance, the value-added share of the service sector in the US (that is not a targeted moment and includes PS and CS) has averaged 77% throughout the last decade. The asymptotic value-added share of the good-producing sector (that includes both manufacturing and PS) is 30%. Moreover, $\bar{v}_G = -0.276$, which implies that industrial goods are also luxuries, albeit with a smaller income elasticity than CS.

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

²⁴Our model implies that the regional CS income share cannot exceed ω_{CS} . For $\omega_{CS} = 0.692$, four small districts violate the constraint. In these cases, we topcode the share of CS and split the excess proportionally between the other two sectors. In practice, this issue is inconsequential because these districts account for a mere 0.15% and 0.23% of Indian value-added in 1987 and 2011, respectively.

an average annual rate of return of 5.6%, which is on the lower end of standard Mincerian regressions, although broadly in line with the findings of recent studies for India using the NSS; see [Singhari and Madheswaran \(2016\)](#). In Appendix C-7, we show that our results are robust to assuming a higher return to education. Given the estimate of ρ , we then calculate the average amount of human capital per region as $\mathbb{E}_{rt}[q] = \sum_s \exp(\rho \times s) \ell_r(s)$, where $\ell_r(s)$ denotes the share of people in region r with s years of education.

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

Trade Costs τ . We calibrate the matrix of trade costs based on two recent studies. First, we leverage [Alder’s \(2023\)](#) estimates of travel times along the most efficient route between the centroids of each pair of Indian districts. Then we transform these travel times into trade costs so as to match the average trade costs across Indian states estimated by [Van Leemput \(2021\)](#). We describe the details of this procedure in Appendix Section B-5.²⁵

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

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

Consider first the identification of A_{rCS_t} . The CS market clearing condition (11) implies that, for each region r , the local CS employment share is given by

$$\frac{H_{rCS_t}}{H_{rt}} = \omega_{CS} + \bar{\nu}_{CS} \times \left(\underbrace{P_{rFt}^{-\omega_F} P_{rGt}^{-\omega_G}}_{\text{Prices}} \times \underbrace{\mathbb{E}_{rt}[q]}_{\text{Skills}} \times \underbrace{w_{rt}^{1-\omega_{CS}}}_{\text{Wages}} \times \underbrace{A_{rCS_t}^{\omega_{CS}}}_{\text{Productivity}} \right)^{-\varepsilon}, \tag{17}$$

where $\bar{\nu}_{CS} < 0$ and $H_{rCS_t}/H_{rt} < \omega_{CS}$, since CS are luxuries. Equation (17) highlights the role of demand (through wages, tradable prices, and the local supply of skills) and productivity in determining the employment share of CS. Inverting the relationship yields a unique solution for A_{rCS_t} as a function of observables and parameters. Given local wages, skills and tradable prices, A_{rCS_t} is increasing in the observed employment share H_{rCS_t}/H_{rt} . Conversely, given the employment share H_{rCS_t}/H_{rt} , A_{rCS_t} is decreasing in the determinants of local demand. This structural decomposition of the observed variation in CS employment shares into income effects and service-led growth is a key step of our methodology. Note that the estimates of A_{rCS_t} do not rely on any published CS price index—an important advantage given the notorious measurement difficulties.

²⁵We thank Simon Alder for sharing his data with us. The results are not sensitive to changes in the target average trade costs. In a previous version of this paper, we used a set of gravity equations in which we assumed trade costs to be a power function of distance, with the distance elasticity calibrated to trade flows within the US. The two approaches yield very similar quantitative results; see Section WA-4 in the Web Appendix.

The procedure to estimate productivity in the tradable sectors is different. Equation (12) implies relative productivity across two locations in sector s is given by (see Appendix A-2 for the derivation)

$$\frac{A_{rs}}{A_{js}} = \left(\frac{H_{rs}}{H_{js}}\right)^{\frac{1}{\sigma-1}} \times \left(\frac{w_r}{w_j}\right)^{\frac{\sigma}{\sigma-1}} \times \left(\frac{\sum_{d=1}^R \tau_{rd}^{1-\sigma} P_{dst}^{\sigma-1} \bar{\vartheta}_{dst} w_{dt} H_{dt}}{\sum_{d=1}^R \tau_{jd}^{1-\sigma} P_{dst}^{\sigma-1} \bar{\vartheta}_{dst} w_{dt} H_{dt}}\right)^{\frac{1}{1-\sigma}}. \quad (18)$$

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

Equations (17)–(18) 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. To this aim, we target two moments; see Appendix A-2. First, we target a 4.2% annual growth rate for real income per capita, which matches real GDP per capita growth in the World Bank data (WDI) using the industrial good as the numeraire.²⁶ Second, we target the change in the price of agricultural goods relative to industrial goods as reported in the ETD. Empirically, agricultural prices rose by a factor of 1.52 relative to prices in the industrial sector.²⁷ Given these moments, our model identifies the full set of sector-region productivities A_{rst} in both 1987 and 2011.²⁸

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

Interestingly, both the productivity dispersion and its correlation with urbanization are highest in the CS sector. Hence, the large employment share of CS in urban locations is not a mere consequence of high wages or an abundance of human capital; it also reflects high CS productivity. Among the tradable goods, productivity is significantly more dispersed in industry than in agriculture. To understand why, note that a district's relative

²⁶We take GDP in terms of industrial goods as our measure of real GDP because industrial goods are tradable. When we compute real GDP using a chained Fisher index, we obtain a growth rate of 4.6%.

²⁷The ETD data covers the time period between 1990 and 2011. We combine ETD's precursor (the 10-sector database by the GGDC) to get a relative price change of 1.52.

²⁸We keep trade costs, τ , constant over time. Allen and Atkin (2022) document a 20% decline of transport time between 1987 and 2011 owing to improvements in Indian infrastructure. As we show in Section WA-4 in the Web Appendix, assuming a reduction in trade costs consistent with their estimate has negligible effects on the estimates of productivity growth in CS, while it slightly reduces those in the tradable sectors. Therefore, one should interpret our estimates of productivity growth in the tradable sectors as inclusive of reductions in trade costs.




FIGURE 4.—Estimated Sectoral Productivities. The figure shows a binscatter plot of the estimated (logarithms of) productivity in agriculture, CS, and industry, $\ln A_{r,s,t}$, across urbanization-rate bins for 2011. In each panel, the sectoral mean is normalized to unity.

productivity reflects its sectoral employment share relative to its skill price (see equation (18)). 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 V. Two patterns are salient. First, in most districts CS productivity grew. In the median region, it grew by 2.6% annually between 1987 and 2011—less than productivity growth in the industrial sector and more than in agriculture. Second, productivity growth in CS was highly unequal across space, with the top 10% of locations experiencing growth above 11%. When we aggregate across regions, we find an average productivity growth in CS about 4%, larger than in the two tradable sectors.²⁹

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

TABLE V
REGIONAL DISTRIBUTION OF SECTORAL PRODUCTIVITY GROWTH.

	Sectoral Productivity Growth					Aggregate
	10th	25th	50th	75th	90th	
Consumer Services (g_{rCS})	-1.3	0.3	2.6	6.4	11.1	4.0
Agriculture (g_{rF})	0.3	1.1	1.8	2.6	3.3	2.0
Industry (g_{rG})	1.8	2.6	3.5	4.4	5.1	3.6

Note: The table reports moments of the distribution of sectoral productivity growth. These growth rates are annualized and calculated as $g_{r,s} = \frac{1}{2011-1987} (\ln A_{r,s,2011} - \ln A_{r,s,1987})$. Columns 1-5 report different quantiles. Column 6 reports the population-weighted average in 2011.

²⁹To account for measurement error, we winsorize the top and bottom 3% of the estimated distribution of productivity growth in CS. Appendix C-6 discusses the details and reports robustness results for these choices.

TABLE VI
ANNUAL PRODUCTIVITY GROWTH (ETD).

Annual Growth of Real Value-Added per Worker in the Published Data (1990–2010)				
Agriculture	Manufacturing	Mining	Finance and Business	Trade, Restaurants, and Hotels
2.6%	5.3%	4.2%	4.1%	4.2%

Note: The table reports the annual growth of real value-added per worker in India for the period 1990–2010, broken down by sectors and service industries. The data are from the ETD (published by the GGDC).

5.3. Nontargeted Moments

In this section, we compare the predictions of our model to some nontargeted moments. We summarize the main findings here and defer the details to Appendix C-5.

Nationwide Sectoral Productivity Growth. Our methodology allows us to recover sectoral productivity estimates for all Indian districts. We are not aware of alternative estimates at the sector-region level. However, the ETD provides estimates of nationwide growth in real value added per worker for 12 sectors. The service industry “trade, restaurants, and hotels” is the best match to our notion of CS.

In Table VI, we report annual sectoral productivity growth according to the ETD. The ETD data confirms the important role of the service sector for Indian growth (an annual growth of 4.2% for the Indian retail sector). The ETD data also confirms that productivity in manufacturing grew faster than in agriculture. Overall, the ETD figures are broadly in line with our estimates reported in Table V, although our methodology assigns a more salient role to service-led growth.

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

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

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

In Table VII, we report the elasticities of substitution and the sectoral spending elasticities in rural and urban districts. Our estimates imply that CS and industrial goods are complements, with an elasticity of substitution between 0.4 and 0.9, that agricultural and CS value-added are substitutes with an elasticity between 1.2 and 1.7, and that agricultural and industrial output are also substitutes, but with a smaller elasticity.

We find these results economically plausible. As the (quality-adjusted) price of CS-intensive restaurants declines, individuals substitute away from home-cooked meals, making agricultural and CS value-added substitutes. Similarly, falling prices of industrial value-added increase the spending share on CS value-added if consumers reallocate their spending to products that heavily rely on CS. The results in Table VII are also broadly in line with the existing literature. A number of papers document evidence of complementarity between goods and services either in two-sector models or in three-sector

TABLE VII
ELASTICITIES OF SUBSTITUTION AND INCOME ELASTICITIES.

Urbanization quantile	Elasticities of Substitution			Spending Elasticities		
	Agr. and CS	Ind. and CS	Agr. and Ind.	Agr.	CS	Ind.
1 (Rural)	1.7	0.4	1.2	0.6	1.7	1.3
5 (Urban)	1.2	0.9	1.1	0.6	1.2	1.1

Note: 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.

models where all elasticities are forced to be identical; see Herrendorf, Rogerson, and Valentinyi (2014), Comin, Lashkari, and Mestieri (2021), and Duernecker, Herrendorf, and Valentinyi (2017). Given the small size of the agricultural sector in the US, this is consistent with our finding that industrial goods and services are complements. In terms of spending elasticities, we estimate CS and industrial goods to be luxuries and agricultural output to be a necessity. This is consistent with Comin, Lashkari, and Mestieri (2021), who report spending elasticities of 0.57, 1.15, and 1.29 for Tanzania.

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

6. THE UNEQUAL EFFECTS OF SERVICE-LED GROWTH

We now turn to our two main questions of interest: How important was productivity growth in the service sector for the rise of living standards in India? How skewed were these benefits across space and the income distribution?

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

As in Baqaee and Burstein (2023), we measure welfare changes in terms of equivalent variations relative to the status quo in 2011. In other words, we calculate what share of its 2011 income a household residing in region r endowed with human capital q would be willing to forego to avoid the change of prices and wages associated with a counterfactual return of productivity in sector s to the 1987 level in all Indian districts. More formally, let $x_r = (w_r, \mathbf{P}_r)$ and $\hat{x}_r = (\hat{w}_r, \hat{\mathbf{P}}_r)$ 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 . For instance, if $\varpi^q = -20\%$, the consumer would be indifferent between giving up 20% of her 2011 income and the counterfactual allocation. Using the indirect utility function \mathcal{V} given in (2), $\varpi^q(\hat{x}_r|x_r)$ is implicitly defined by

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

In Appendix A-4, we derive an analytical expression for $\varpi^q(\hat{x}_r|x_r)$. Following a similar procedure, and exploiting the aggregation properties of PIGL preferences, we also calculate equivalent variations for the utilitarian welfare functions at the regional level.

TABLE VIII
THREE INDIAN DISTRICTS.

District	Urban Share	Population	Avg. Income	Emp. Share (%)			Prod. Growth (%)		
				Agr.	Ind.	CS	Agr.	Ind.	CS
Bangalore	0.77	10.6	3781	8	36	56	3.4	5.9	10.7
Chengalpattu	0.67	8.1	2807	12	37	51	2.8	4.9	8.7
Bankura	0.07	3.0	1597	64	7	28	1.5	2.1	2.4

Note: The table reports descriptive economic and demographic statistics in 2011 for the selected districts discussed in the text. The figures for productivity growth are from our estimates.

6.1. Sources of Welfare Growth in India

To highlight the unequal effects of service-led growth, we first zoom in on three districts. Then we consider different 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.³⁰ Bankura is a rural district in West Bengal, which mostly relies on agriculture. Table VIII provides some descriptive statistics for these districts. Household income is significantly higher in Bangalore and Chengalpattu. Both the patterns of sectoral specialization and the estimated productivity growth are markedly diverse. In 2011, the employment share of CS was about 56% in Bangalore, 51% in Chengalpattu, and 28% in Bankura. There were large differences in CS productivity growth ranging from 2.4% in Bankura to 11% in Bangalore. Industrial productivity growth was 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 was lower in all sectors in Bankura.

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

For comparison, in the right panel, we depict the equivalent variations of agricultural productivity growth. For very poor households in Bankura, the equivalent variation is 25%. The bulk of the welfare gains is due to agricultural productivity growth in the entirety of India (rather than from changes in local productivity), which reduces food prices overall. The diffusion of the effects of productivity growth in agriculture via trade explains why the spatial differences are small.

³⁰We use the border of Chengalpattu in 1987. This district was split into Kancheepuram and Thiruvallur between 1991 and 2001 and later 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 median income in each district.

Average Effects. To draw more general lessons, we compute average welfare effects at different levels of aggregation. In Figure 6, we depict the population-weighted average equivalent variation in different urbanization quintiles (left panel) and in different percentiles of the income distribution (right panel).³¹ Because the welfare results are based on an estimated model, they entail sampling uncertainty. To quantify this uncertainty, we




FIGURE 6.—The Unequal Effects of Service-Led Growth. The figure displays the average percentage welfare losses (using district population as weights) associated with counterfactually setting productivity in agriculture, CS, and industry, to the respective 1987 level, broken down by urbanization quintile in 2011 (panel (a)) and by the 10th, 25th, 50th, 75th, 90th, 95th, and 99th percentiles of the income distribution in 2011 (panel (b)). We compute the distribution of such welfare losses using a nonparametric bootstrap. The respective boxes cover the 25%–75% quantile of the bootstrap distribution. The horizontal lines on the top and bottom refer to the 5% and 95% quantiles of the bootstrap distribution.

³¹The interpretation of the average welfare effects is subject to the usual caveat (see Appendix Section A-4). In particular, the formal aggregation properties of the model only apply to people living in the same district who face the same price vector. Nevertheless, they are informative statistics.

estimate the distribution of the welfare effects using a nonparametric bootstrap procedure (Horowitz (2019)); see Section WA-6 in the Web Appendix. In Figure 6, we report these distributions as a boxplot. Each box shows the 25%–75% quantiles of the distribution of welfare gains. The line within the box indicates the median, and the two vertical lines on the top and the bottom indicate the 5% and 95% quantiles.

The left panel of Figure 6 shows that the benefits of agricultural productivity growth are larger in rural districts like Bankura than in urbanized districts. For households in the four lowest quintiles of urbanization, the average equivalent variation is about 20%. For the top quintile of urbanization, it drops to 15%. By contrast, the gains from productivity growth in CS and industry are skewed toward urban locations. The average equivalent variation for CS is a staggering 37% for the most urbanized quintile.

In the right panel, we focus on the income distribution, showing the welfare effects at the 10th, 25th, 50th, 75th, 90th, 95th, and 99th percentiles. The benefits of productivity growth in CS and industry are increasing in income, whereas the pattern is the opposite for agriculture. In the case of CS, the equivalent variation for the top decile of the income distribution is very large and comparable to that for the top quintile of urbanization. Interestingly, for households below the median income, the welfare effects of productivity growth in CS and in the industrial sector are roughly equal, both being smaller than those from agriculture.

In the left panel of Figure 7, we report the population-weighted average equivalent variation across all Indian districts. On average, Indians would have been willing to sacrifice 20% of their income in lieu of giving up the observed productivity growth originating in the CS sector. To put this number into perspective, the equivalent variation from all sources of productivity growth in India since 1987 is 64%. Hence, productivity growth in the CS sector accounts for roughly one-third of the increase in economic well-being. Productivity growth in agriculture and industry were also important sources of welfare improvement, albeit smaller than CS.

In summary, productivity growth in CS played an important role for economic development in India. In urban areas and for rich households, growth in CS was the dominant

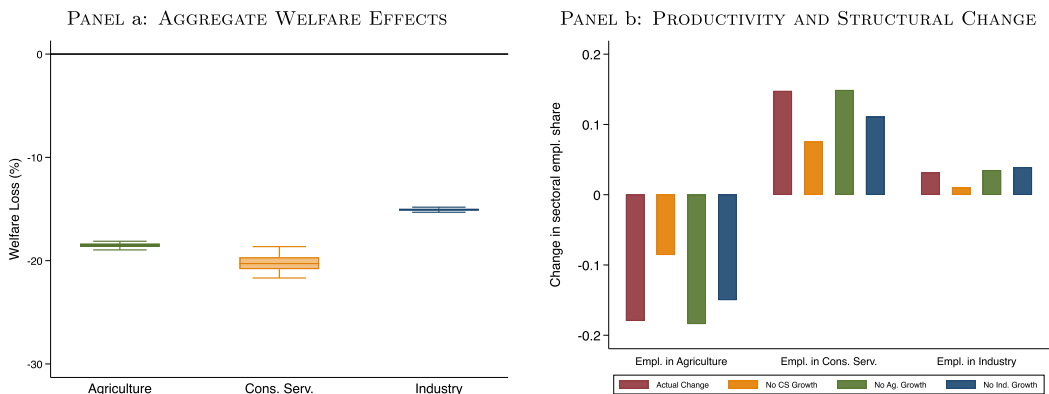


FIGURE 7.—Aggregate Welfare Effects and Structural Change. In the left panel, we show the analogue of Figure 6 with welfare effects aggregated up to the nationwide level. In the right panel, we show changes in sectoral employment (in efficiency units). We depict the actual change in India (red bars)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

source of rising living standards. By contrast, technical progress in agriculture was the most important source of welfare gains for below-median households.

6.2. Structural Change

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

Structural Change in the Theory. We first consider how prices and wages affect sectoral spending shares. Differentiating equation (8) for any two sectors s and k yields

$$\frac{\partial \bar{\vartheta}_{rst}}{\partial \ln P_{rkt}} = \varepsilon \omega_k (\bar{\vartheta}_{rst} - \omega_s) \quad \text{and} \quad \frac{\partial \bar{\vartheta}_{rst}}{\partial \ln w_{rt}} = -\varepsilon (\bar{\vartheta}_{rst} - \omega_s). \tag{19}$$

Because food is a necessity, whereas industrial goods and CS are luxuries, $\bar{\vartheta}_{rFt} > \omega_F$, whereas $\bar{\vartheta}_{rGt} < \omega_G$, and $\bar{\vartheta}_{rCS_t} < \omega_{CS}$. Thus, falling prices in *any* sector increase the expenditure share on goods and CS and decrease the expenditure share on food. Similarly, higher wages increase spending on goods and CS and reduce spending on food. In the case of nontradable CS, $\bar{\vartheta}_{rCS_t} = H_{rCS_t}/H_{rt}$. Thus, productivity growth in *any* sector increases the employment share of CS both due to falling prices and higher wages. However, the price impact in (19) depends on the sectoral origin of productivity growth. In particular, $\frac{\partial \bar{\vartheta}_{rCS_t} / \partial \ln P_{rCS_t}}{\partial \bar{\vartheta}_{rCS_t} / \partial \ln P_{rFt}} = \frac{\omega_{CS}}{\omega_F}$, which, according to our calibration, is a large number. Hence, productivity growth in CS causes significantly faster structural change than productivity growth in agriculture.

To gauge the magnitude of the difference, consider the Indian economy in 1987. A hypothetical 10% increase in A_{rCS} in *all* districts changes the employment shares of F, G, and CS by -1.5, 0.3, and 1.2 p.p., respectively. Note that this split, whereby 80% of the decline in agriculture gets absorbed in the service sector, is quantitatively in line with the experience of most developing countries, documented in Figure 1. By contrast, a 10% increase in A_{rF} in all districts yields much smaller changes of -0.023, 0.005, and 0.018 p.p. While uniform productivity growth in agriculture drives some structural change, its quantitative importance is small in our calibration.

It is useful to contrast these results with the case in which productivity increases in a *single* region. Suppose, for instance, productivity grows only in Bankura. A 10% increase in CS productivity reduces employment in agriculture and industry by 1.2 and 0.2 p.p., and increases the CS sector by 1.4 p.p. Hence, the employment effects of a local CS shock are relatively similar to the effects of an aggregate increase in CS productivity. By contrast, a 10% increase in agricultural productivity in Bakura alone *increases* employment in agriculture by 3.2 p.p. and in CS by 0.4 p.p., while decreasing industrial employment by 3.6 p.p. Two observations are in order. First, our model predicts sectoral specialization based on comparative advantage: rising productivity shifts employment toward agriculture. Second, agricultural employment is dissociated from agricultural spending. While productivity growth induces specialization in agriculture, it also shifts the spending of Bankura’s residents away from food.

These different implications of local versus aggregate productivity shocks in agriculture are related to a debate in the empirical literature. The prediction that a positive *local* productivity shock in agriculture slows structural change out of agriculture and causes deindustrialization is in line with the findings of recent papers of Asher, Campion, Gollin,

and Novosad (2022), who study the long-run impacts of irrigation canals on structural change in India, and Kelly, Mokyr, and Ó Gráda (2023), who document a negative effect of local agricultural productivity on the onset of the British Industrial Revolution. The effects of *aggregate* shocks are less clear. On the one hand, Gollin, Hansen, and Wingender (2021) find that the adoption of high-yielding crop varieties (the Green Revolution) sped up the decline of agriculture. On the other hand, Moscona (2019), relying on an identification strategy that exploits exogenous variation in ecological characteristics, finds that productivity growth in agriculture slows urbanization and industrial development.

Structural Change in the Estimated Model. We now consider the impact of productivity growth we inferred from the calibrated model. The right panel of Figure 7 shows the sectoral reallocation between 1987 and 2011. All figures are in effective units of labor. In contrast to the welfare analysis, sampling variation plays a minor role for these results and we do not include the standard errors to improve readability.

For each sector, we depict four bars. The leftmost bar shows the actual data for India: agricultural employment declined by 18 p.p. and CS increased by 15 p.p. The industrial sector, which contains PS, only increased 3 p.p. The remaining three bars depict the counterfactual change in the sectoral employment shares when we shut down (one at a time) productivity growth in CS, agriculture, and industry, respectively. Productivity growth in CS was responsible for the lion's share of the structural transformation. In its absence, the agricultural employment share would have only declined by about 9 p.p. (instead of 18 p.p.) and the rise in CS employment would have only been 8 p.p. (instead of 15 p.p.). Hence, our theory does recognize the importance of income effects that originate from productivity growth in other sectors in shifting labor from agriculture to services. However, quantitatively, these effects explain only half of the observed structural transformation in India.

In line with the theoretical results discussed above, the effects of agricultural productivity growth (green bars) are modest. If anything, productivity growth in agriculture appears to have marginally *increased* employment in agriculture and slowed structural change. This reflects both the small effect of average productivity growth and the significant heterogeneity in the estimated productivity changes across districts.

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

7. ROBUSTNESS

In this section, we discuss the robustness of our results. In Section 7.1, we study the sensitivity of our results to changes in structural parameters, most notably, the Engel elasticity ε . In Section 7.2, we revisit some measurement choices concerning the split between CS and PS. In Section 7.3, we generalize our preference structure. In Section 7.4, we study various generalizations of the model (open economy, skill heterogeneity, spatial mobility). For each experiment, Table IX reports the welfare effects associated with productivity growth in CS at the aggregate level (Figure 7) and by percentile of urbanization and income (Figure 6). We defer the corresponding results for agricultural and industrial productivity growth to Appendix C-7.

TABLE IX
THE IMPORTANCE OF SERVICE-LED GROWTH—ROBUSTNESS.

	Aggregate Effects	Urbanization Quintiles		Income Quintiles		
		1st	5th	10th	50th	90th
Baseline	-20.5	-13.1	-36.8	-13.7	-14.6	-37.7
<i>Alternative calibrations of ε (Section 7.1)</i>						
$\varepsilon = 0.415$ (High Income Households)	-19.5	-12.3	-35.4	-12.7	-13.5	-36.4
$\varepsilon = 0.321$ (OLS estimate)	-25.2	-17.1	-42.7	-17.9	-19.1	-43.4
<i>Alternative measurement choices (Section 7.2)</i>						
Allocate PS share based on WIOD	-18.8	-13.5	-31.3	-14.0	-14.0	-31.9
Allocate ICT and Business to PS	-17.0	-15.3	-23.6	-14.2	-12.2	-24.0
Allocate Construction to Industry	-12.5	-2.5	-31.7	-4.6	-8.7	-23.3
<i>Alternative modeling assumptions (Section 7.4)</i>						
Open economy	-17.7	-11.7	-31.5	-12.5	-12.1	-31.6
Imperfect skill substitution	-19.8	-9.8	-37.5	-9.8	-11.4	-40.1
Spatial labor mobility	-18.4	-13.4	-29.9	-	-	-

Note: 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.

7.1. Sensitivity to Structural Parameters

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

For our analysis, we rely on the IV estimate of $\varepsilon = 0.395$ (column 6 in Table III). In the second row of Table IX, we present an alternative calibration based on the elasticity estimated for the sample of high-income households, $\varepsilon = 0.415$, which is the largest elasticity in Table III. The effects are marginally smaller but very similar to the baseline results. In the third row, we set $\varepsilon = 0.321$, the OLS estimate of the Engel elasticity. This change reduces the income effects and magnifies the importance of service-led growth, especially in cities.³² Finally, in Appendix C-7, we allow ε to be larger in urban districts, according to the estimates of column 8 in Table III. This only leads to a marginal reduction in the inequality of welfare effects across districts. In summary, our main results are robust to the entire range of ε estimated in Table III.³³

In Appendix C-7, we also discuss the sensitivity of our results to changes in other parameters: the asymptotic food share ω_F , the tail of the skill distribution ζ , the educational return ρ , and the elasticity of substitution across local varieties σ (all other parameters

³²Boppart (2014) estimates Engel elasticities for the US from CEX and PSID. His estimates range between 0.22 and 0.29. Because his model has only two sectors, the estimates are not directly comparable. Nevertheless, the results in Table IX indicate that lower income elasticities would magnify the welfare effects associated with CS growth.

³³We also consider a calibration where we do not estimate ε but calibrate it by targeting the aggregate productivity growth of the Indian retail sector (4.2%) according to ETD (see Table VI). This yields $\varepsilon = 0.385$, which is smaller than our baseline estimate. The resulting welfare gains are slightly larger.

are either point-identified in our theory or pinned down by normalization.) The effects of these changes are quantitatively small and do not affect our conclusions.

7.2. Measurement: The PS–CS Split

We split employment in the service sector into PS and CS according to whether firms in different service industries sell more to firms or to consumers; see Table II. Our data-driven approach could underestimate the PS sector if some firms reported sales to small firms as sales to individuals. To address this concern, we consider two alternative classifications.

First, we use aggregate Input–Output tables from the WIOD to measure the share of service output that is used as an intermediate input in the industrial and agricultural sectors. In India, this number is about 20%. Thus, we increase the relative size of the PS sector so that it accounts for 20% of value-added in the service sector. 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. This as a generous upper bound as in reality many law and financial firms sell their services to consumers (e.g., savings banks or divorce lawyers). Under the alternative classification, PS account for 22% of service employment. Because the employment share of business services and ICT is especially large in cities, assigning them to PS reduces the share of CS mostly in urban areas.

Rows 4 and 5 of Table IX report the results. As expected, both reclassifications reduce the estimated productivity growth of CS. The associated welfare effects decline by 1.7 and 3.5 p.p., respectively. Nonetheless, they remain large. At the spatial level, the welfare effects of service-led growth become less unequal, but overall CS productivity growth continues to mostly benefit urban dwellers. Overall, neither reclassification of PS alters the broad picture.

Finally, we turn our attention to the construction sector. In our main analysis, we merge residential construction with the CS sector because it produces nontradable goods. However, the conventional classification regards construction as part of the industrial sector. For this reason, we analyzed how our results would change if (inconsistently with our theory) we merged the whole construction activity with the manufacturing sector. We report the results in row 6 of Table IX. The reclassification of construction activities increases the average welfare effect of productivity growth in the industrial sector. While CS continue to contribute significantly to aggregate welfare growth, the magnitude is appreciably smaller. Interestingly, the welfare effects of service-led growth become even more skewed in favor of urban districts than in our baseline analysis, because the construction sector is relatively more salient in rural areas. The smaller aggregate welfare effect is therefore mostly driven by rural districts, where construction accounts for the bulk of nontradable activities. By contrast, service-led growth in urban locations is not primarily driven by construction activities.

7.3. Generalized PIGL Preferences

In our analysis, we parametrized the indirect utility function by setting $D(\mathbf{P}_r) = \sum_s \nu_s \ln P_{rs}$ in the value-added representation. In this section, we generalize the approach

of Boppart (2014) to a three-sector environment. We assume a CES function:³⁴

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

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

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

This new specification flexibly adjusts the weights of the pseudo-price index by a term that depends on the new parameter γ that is set to zero in our baseline specification. In other words, the parameter γ affects the strength of relative price effects. In particular, equation (19) now reads

$$\frac{\partial \bar{\vartheta}_{rst}}{\partial \ln P_{rk}} = (\gamma \nu_k + \varepsilon \omega_k) \times (\bar{\vartheta}_{rst} - \omega_s) \quad \text{and} \quad \frac{\partial \bar{\vartheta}_{rst}}{\partial \ln w_{rt}} = -\varepsilon (\bar{\vartheta}_{rst} - \omega_s). \quad (20)$$

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

To see how γ affects the welfare effects of service-led growth, consider again the three districts of Bangalore, Chengalpattu, and Bankura. The left panel of Figure 8 shows how the welfare effects associated with the estimated productivity growth in CS over the period 1987–2011 vary as functions of γ .³⁵ The special case of $\gamma = 0$ corresponds to our baseline analysis. The welfare effects are increasing in γ . As $\gamma \rightarrow \gamma^*$, the model requires larger and larger variations in CS prices (hence, productivities) to rationalize the observed variation in employment shares. Over time, it requires a larger productivity growth in CS, which magnifies the welfare effects. Note that, while the welfare effects grow unboundedly large as $\gamma \rightarrow \gamma^*$, they decline only slowly in the range of negative γ . Figure WA-2 in the Web Appendix shows how changes in γ affect the distribution of productivity growth

³⁴This specification preserves the isomorphism between the expenditure and the value-added approach; see Section WA-2.1 in the Web Appendix. For simplicity, we directly write the value-added functions.

³⁵For a given value of γ , we can identify our model with preferences based on (7.3) from exactly the same moments as our baseline model. We always recalibrate all other parameters when varying γ .




FIGURE 8.—Welfare Effects of CS Productivity Growth as a function of γ . In the left panel, we depict the welfare effect of CS growth as a function of γ for the representative consumer in three Indian districts. We depict the range of γ where $0 \leq \text{EOS}_{\text{CS},G} \leq 1$ for 90% of districts as the shaded areas. In the right panel, we show the aggregate welfare effect and the welfare effect for the 1st and 5th quantiles of the urbanization rate.

in CS in the region where $\gamma < \gamma^*$. Increasing γ raises both the average and the spread of productivity growth.

We can further discipline the range of plausible γ 's by considering the implied Allen–Uzawa elasticities of substitution between G and CS (see Section WA-2.3 in the Web Appendix for details). The estimates in the literature suggest that goods and CS are closer complements than under Cobb–Douglas preferences. Thus, we focus on the range of γ such that $\text{EOS}_{\text{CS},G} \in (0, 1)$ in at least 90% of the Indian districts, which yields $\gamma \in [-0.02, 0.05]$.³⁶

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

7.4. Other Generalizations of the Theory

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

Open Economy. Our main analysis treats India as a closed economy. However, international trade, in particular exports of ICT services, has become increasingly important. To incorporate these dimensions, we extend our model to allow for international trade. We assume households, both in India and in the rest of the world, consume differentiated industrial goods sourced from many countries. To capture India's comparative advantage in ICT, we assume India is an ICT exporter and exports the entirety of its ICT value-added. We 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 and 1.56% of total earnings in

³⁶This range is also consistent with the aggregate rate of CS productivity growth. If we calibrate γ to match the rate of 4.2% as reported in Table VI, we find $\gamma = 0.02$.

2011 (in 1987, it was 0.11%). Given the small size of the ICT sector in 1987, we assume it was zero in 1987 and target the earnings share in 2011. We calibrate the parameters so as to generate trade flows like in the data. As seen in row 7 in Table IX, international trade, especially recognizing the tradable nature of ICT services, mildly reduces the welfare effect of productivity growth in CS, especially in cities, which (as shown in Figure 1) saw the fastest increase in ICT employment. Nevertheless, CS continue to play an important role for aggregate growth and for urban areas in particular.

Imperfect Substitution and Skill Bias in Technology. Our analysis assumes that all workers' efficiency units are perfect substitutes. We now generalize our model assuming workers with different educational attainments are imperfect substitutes. Because agricultural workers have, on average, lower educational attainment, an increase in the skill endowment could be responsible for the reallocation of workers from agriculture to CS (see, e.g., Porzio, Rossi, and Santangelo (2022) or Hendricks and Schoellman (2023)).

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

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

where H^+ and H^- denote high- and low-skilled workers, respectively. Note that the technology admits differences in both Hicks-neutral TFP (A_{rst}) and skill bias (Z_{rst}) across sectors, districts, and time. We calibrate the elasticity of substitution between high- and low-skilled workers, ρ , to 1.8, a standard estimate in the literature. The results in row 8 in Table IX show that the quantitative role for the CS sector is very similar to the one of our baseline calibration. If anything, the unequal effects across the income ladder are more pronounced because skilled individuals are more likely to work in the CS sector.

This extension yields two additional findings. First, across districts, Z_{rs} increases in the level of urbanization for all sectors. This increase reflects the empirical observation that the skill premium is higher in urban than in rural districts. Second, we find evidence for skill-biased technical change: over time, Z_{rs} increases in all sectors. Although our accounting approach cannot uncover causal links, these patterns are consistent with models of directed technical change and directed technology adoption such as in Acemoglu and Zilibotti (2001) and Gancia, Müller, and Zilibotti (2013).

Spatial Mobility. In our baseline model, we assumed people to be spatially immobile. However, a counterfactual decline in CS productivity could prompt people to move out of cities. Labor mobility could then work as a form of insurance, thereby reducing the associated welfare losses. To gauge the quantitative importance of labor mobility, we reestimate our model in the presence of an endogenous location choice, which we model as a discrete choice, where individuals receive idiosyncratic preference shocks and locations differ in amenities. We set local amenities so that the spatial equilibrium matches the spatial distribution of the Indian population in both year.

Allowing for an endogenous location choice does not affect the estimation of the parameters nor the productivities. However, labor mobility affects the counterfactuals. We calibrate the elasticity of labor mobility so that, holding amenities fixed, resetting the productivities 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. To calculate the welfare effects, we sample one million fictitious households and associate each of them with a vector of realizations of the geographic preference shock (one per district). Then we counterfactually reset the CS productivity distribution to the 1987 level,

allowing people to relocate optimally to their preferred district. Finally, we calculate the equivalent variation for each household.

In the last row of Table IX, we report the results of this experiment. We do not report the results by income because individuals draw their human capital after moving. As expected, labor mobility lowers the equivalent variation of productivity growth in CS, but the difference is moderate—from an average of 20.5% to 18.4%. The effect is somewhat more conspicuous for households that chose to reside in urban areas in the baseline economy of 2011. Intuitively, resetting CS productivity to the 1987 level reduces the economic appeal of urban areas. The option to migrate allows some households to partially offset these economic losses by moving to districts that better suit their geographic preferences. Altogether, empirically plausible migration responses to changes in the economic environment do not alter the broad picture.

8. CONCLUSION

Service-led growth is a widespread feature of the contemporary world. The classic argument of Baumol (1967) suggests that this trend could lead to economic stagnation. This view has been recently echoed by Rodrik (2016) who expresses concern for the premature deindustrialization of many developing countries. In this paper, we develop a novel methodology to structurally estimate productivity growth in services and assess its role as an engine of growth. The methodology lends itself to a quantitative analysis of the welfare effects of service-led growth across space and the income ladder.

Our application to India delivers two main results. First, productivity growth in consumer services such as retail, restaurants, or residential real estate, was both fast and important for welfare, accounting for one-third of the improvement in living standards between 1987 and 2011. Second, service-led growth had unequal welfare consequences: it disproportionately benefited the urban middle-class while being far less important for poor people living in rural India. This happened for two reasons: (i) consumer services are locally provided and their productivity grew particularly fast in urban areas, and (ii) richer households spend more on service-intensive goods owing to nonhomothetic preferences. While our analysis suggests that low employment growth in the manufacturing sector could be less of a threat to the sustainability of future growth than economists previously thought, it also raises novel concerns about inequality that remain invisible in aggregate statistics.

Our framework has some limitations that future research should address. First, understanding the determinants of productivity growth in services is of first-order importance, especially for policy guidance. Second, service-led growth has implications on other dimensions of inequality such as gender disparity. Third, our approach ignores frictions in mobility across sectors that may be important in reality. In spite of these and other limitations, we believe our portable methodology will be useful to study structural change and the role of service-led growth in the development experience of other countries.

REFERENCES

- ACEMOGLU, DARON, AND FABRIZIO ZILIBOTTI (2001): “Productivity Differences,” *Quarterly Journal of Economics*, 116 (2), 563–606. [1491]
- ACEMOGLU, DARON, VASCO M. CARVALHO, ASUMAN OZDAGLAR, AND ALIREZA TAHBAZ-SALEHI (2012): “The Network Origins of Aggregate Fluctuations,” *Econometrica*, 80 (5), 1977–2016. [1463]
- ALDER, SIMON (2023): “Chinese Roads in India: The Effect of Transport Infrastructure on Economic Development,” SSRN Working Paper 2856050. [1477]

- ALDER, SIMON, TIMO BOPPART, AND ANDREAS MÜLLER (2022): “A Theory of Structural Change That Can Fit the Data,” *American Economic Journal: Macroeconomics*, 14 (2), 160–206. [1461]
- ALLEN, TREB, AND COSTAS ARKOLAKIS (2014): “Trade and the Topography of the Spatial Economy,” *Quarterly Journal of Economics*, 129 (3), 1085–1140. [1461]
- ALLEN, TREB, AND DAVID ATKIN (2022): “Volatility and the Gains From Trade,” *Econometrica*, 90 (5), 2053–2092. [1478]
- ASHER, SAM, ALISON CAMPION, DOUGLAS GOLLIN, AND PAUL NOVOSAD (2022): “The Long-Run Development Impacts of Agricultural Productivity Gains: Evidence From Irrigation Canals in India,” Report, Johns Hopkins. [1485,1486]
- ATKIN, DAVID, BENJAMIN FABER, AND MARCO GONZALEZ-NAVARRO (2018): “Retail Globalization and Household Welfare: Evidence From Mexico,” *Journal of Political Economy*, 126 (1), 1–73. [1460]
- BAQAEE, DAVID R., AND ARIEL BURSTEIN (2023): “Welfare and Output With Income Effects and Taste Shocks,” *The Quarterly Journal of Economics*, 138 (2), 769–834. [1481]
- BAUMOL, WILLIAM (1967): “Macroeconomics of Unbalanced Growth: The Anatomy of the Urban Crisis,” *American Economic Review*, 57 (3), 415–426. [1492]
- BOPPART, TIMO (2014): “Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences,” *Econometrica*, 82 (6), 2167–2196. [1459,1461,1464,1465,1487,1489]
- BUDÍ-ORS, TOMÁS, AND JOSEP PIJOAN-MAS (2022): “Macroeconomic Development, Rural Exodus, and Uneven Industrialization,” SSRN Working Paper 4069894. [1461]
- BUERA, FRANCISCO J., AND JOSEPH P. KABOSKI (2012): “The Rise of the Service Economy,” *American Economic Review*, 102 (6), 2540–2569. [1460]
- CASELLI, FRANCESCO (2005): “Accounting for Cross-Country Income Differences,” in *Handbook of Economic Growth*, Vol. 1A, ed. by Philippe Aghion and Steven Durlauf. Elsevier B.V., 679–741. [1460,1472]
- CHATTERJEE, SHOUMITRO, ELISA GIANNONE, AND KAN KUNO (2023): “Unequal Global Convergence,” Report, Pennsylvania State University. [1460]
- CHEN, XILU, GUANGYU PEI, ZHENG MICHAEL SONG, AND FABRIZIO ZILIBOTTI (2023): “Tertiarization Like China,” *Annual Review of Economics*, 15. <https://doi.org/10.1146/annurev-economics-071122-030026>. [1460, 1463]
- COMIN, DIEGO, DANIAL LASHKARI, AND MARTÍ MESTIERI (2021): “Structural Change With Long-Run Income and Price Effects,” *Econometrica*, 89 (1), 311–374. [1461,1481]
- DE VRIES, GAAITZEN, LINDA ARFELT, DOROTHEA DREES, MAREIKE GODEMANN, CALUM HAMILTON, BENTE JESSEN-THIESEN, AHMET IHSAN KAYA, HAGEN KRUSE, EMMANUEL MENSAH, AND PIETER WOLTER (2021): “The Economic Transformation Database (ETD): Content, Sources, and Methods,” WIDER Technical Note 2/2021. [1469]
- DEHEJIA, RAJEEV, AND ARVIND PANAGARIYA (2016): “The Link Between Manufacturing Growth and Accelerated Services Growth in India,” *Economic Dev. and Cultural Change*, 64 (2), 221–264. [1460]
- DESMET, KLAUS, EJAZ GHANI, STEPHEN O’CONNELL, AND ESTEBAN ROSSI-HANSBERG (2015): “The Spatial Development of India,” *Journal of Regional Science*, 55 (1), 10–30. [1460]
- DUARTE, MARGARIDA, AND DIEGO RESTUCCIA (2010): “The Role of the Structural Transformation in Aggregate Productivity,” *Quarterly Journal of Economics*, 125 (1), 129–173. [1460]
- DUERNECKER, GEORG, BERTHOLD HERRENDORF, AND AKOS VALENTINYI (2017): “Structural Change Within the Service Sector and the Future of Baumol’s Disease,” SSRN Working Paper 3082293. [1481]
- ECKERT, FABIAN, AND MICHAEL PETERS (2022): “Spatial Structural Change,” Working Paper 30489, National Bureau of Economic Research. [1461]
- ECKERT, FABIAN, SHARAT GANAPATI, AND CONOR WALSH (2022): “Urban-Biased Growth: A Macroeconomic Analysis,” Working Paper 30515, National Bureau of Economic Research. [1460]
- FAJGELBAUM, PABLO D., AND AMIT K. KHANDELWAL (2016): “Measuring the Unequal Gains From Trade,” *The Quarterly Journal of Economics*, 131 (3), 1113–1180. [1461]
- FAN, TIANYU, MICHAEL PETERS, AND FABRIZIO ZILIBOTTI (2023a): “Growing Like India: The Unequal Effects of Service-Led Growth,” Working Paper 28551, National Bureau of Economic Research. [1461]
- (2023b): “Supplement to ‘Growing Like India—the Unequal Effects of Service-Led Growth’,” *Econometrica Supplemental Material*, 91, <https://doi.org/10.3982/ECTA20964>. [1461]
- GANCIA, GINO, ANDREAS MÜLLER, AND FABRIZIO ZILIBOTTI (2013): “Structural Development Accounting,” in *Advances in Economics and Econometrics: Theory and Applications (10th World Congress of the Econometric Society)*, ed. by Daron Acemoglu, Manuel Arellano, and Eddie Dekel. Cambridge Univ. Press, 373–418. [1460,1472,1491]
- GARCÍA-SANTANA, MANUEL, JOSEP PIJOAN-MAS, AND LUCCIANO VILLACORTA (2021): “Investment Demand and Structural Change,” *Econometrica*, 89 (6), 2751–2785. [1460]

- GERVAIS, ANTOINE, AND J. BRADFORD JENSEN (2019): "The Tradability of Services: Geographic Concentration and Trade Costs," *Journal of International Economics*, 118, 331–350. [1459]
- GLAESER, EDWARD L., JED KOLKO, AND ALBERT SAIZ (2001): "Consumer City," *Journal of Economic Geography*, 1 (1), 27–50. [1460]
- GOLLIN, DOUGLAS, CASPER WORM HANSEN, AND ASGER MOSE WINGENDER (2021): "Two Blades of Grass: The Impact of the Green Revolution," *Journal of Political Economy*, 129 (8), 2344–2384. [1486]
- GOLLIN, DOUGLAS, REMI JEDWAB, AND DIETRICH VOLLRATH (2016): "Urbanization With and Without Industrialization," *Journal of Economic Growth*, 21 (1), 35–70. [1460]
- GOLLIN, DOUGLAS, DAVID LAGAKOS, AND MICHAEL WAUGH (2014): "The Agricultural Productivity Gap," *Quarterly Journal of Economics*, 129 (2), 939–993. [1460]
- HALL, ROBERT, AND CHARLES JONES (1999): "Why Do Some Countries Produce so Much More Output per Worker Than Others?" *Quarterly Journal of Economics*, 114 (1), 83–116. [1460,1472]
- HENDRICKS, LUTZ, AND TODD SCHOELLMAN (2023): "Skilled Labor Productivity and Cross-Country Income Differences," *American Economic Journal: Macroeconomics*, 15, (1), 240–268. [1491]
- HERRENDORF, BERTHOLD, RICHARD ROGERSON, AND AKOS VALENTINYI (2013): "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, 103 (7), 2752–2789. [1459,1460]
- (2014): "Growth and Structural Transformation," in *Handbook of Economic Growth*, Vol. 2. Elsevier, 855–941. [1481]
- HOROWITZ, JOEL L. (2019): "Bootstrap Methods in Econometrics," *Annual Review of Econ.*, 11, 193–224. [1484]
- HSIEH, CHANG-TAI, AND ESTEBAN ROSSI-HANSBERG (2023): "The Industrial Revolution in Services," *Journal of Political Economy Macroeconomics* 1 (1), 3–42. [1460]
- JEDWAB, REMI, ELENA IANCHOVICHINA, AND FEDERICO HASLOP (2022): "Consumption Cities vs. Production Cities: New Considerations and Evidence," SSRN Working Paper 4166684. [1460]
- KELLY, MORGAN, JOEL MOKYR, AND CORMAC Ó GRÁDA (2023): "The Mechanics of the Industrial Revolution," *Journal of Political Economy*, 131 (1), 59–94. [1486]
- KLENOW, PETER, AND ANDRÉS RODRÍGUEZ-CLARE (1997): "The Neoclassical Revival in Growth Economics: Has It Gone Too Far?" in *NBER Macroeconomics Annual 1997*, Vol. 12. National Bureau of Economic Research, Inc., 73–103. [1460]
- KONGSAMUT, PIYABHA, SERGIO REBELO, AND DANYANG XIE (2001): "Beyond Balanced Growth," *Review of Economic Studies*, 68 (4), 869–882. [1460]
- LAMBA, ROHIT, AND ARVIND SUBRAMANIAN (2020): "Dynamism With Incommensurate Development: The Distinctive Indian Model," *Journal of Economic Perspectives*, 34 (1), 3–30. [1460]
- MATSUYAMA, KIMINORI (2019): "Engel's Law in the Global Economy: Demand-Induced Patterns of Structural Change, Innovation, and Trade," *Econometrica*, 87 (2), 497–528. [1461]
- MOSCONA, JACOB (2019): "Agricultural Development and Structural Change, Within and Across Countries," Report, MIT. [1486]
- MUELLBAUER, JOHN (1976): "Community Preferences and the Representative Consumer," *Econometrica*, 44 (5), 979–999. [1459]
- NAYYAR, GAURAV, MARY HALLWARD-DRIEMEIER, AND ELWYN DAVIES (2021): *At Your Service?: The Promise of Services-Led Development*. World Bank Publications. [1460]
- NGAI, L. RACHEL, AND CHRISTOPHER A. PISSARIDES (2007): "Structural Change in a Multisector Model of Growth," *American Economic Review*, 97 (1), 429–443. [1460]
- PORZIO, TOMMASO, FEDERICO ROSSI, AND GABRIELLA SANTANGELO (2022): "The Human Side of Structural Transformation," *American Economic Review*, 112 (8), 2774–2814. [1491]
- REDDING, STEPHEN J., AND ESTEBAN ROSSI-HANSBERG (2017): "Quantitative Spatial Economics," *Annual Review of Econ.*, 9 (1), 21–58. [1461]
- RODRIG, DANI (2016): "Premature Deindustrialization," *Journal of Economic Growth*, 21 (1), 1–33. [1492]
- SATO, RYUZO (1975): "The Most General Class of CES Functions," *Econometrica*, 43, 999–1003. [1461]
- SINGHARI, SMRUTIREKHA, AND S MADHESWARAN (2016): "The Changing Rates of Return to Education in India: Evidence From NSS Data," Institute for Social and Economic Change Bangalore Working Paper 358. [1477]
- VAN LEEMPUT, EVA (2021): "A Passage to India: Quantifying Internal and External Barriers to Trade," *Journal of International Economics*, 131, 103473. [1477]

Co-editor Charles I. Jones handled this manuscript.

Manuscript received 8 July, 2022; final version accepted 28 February, 2023; available online 16 March, 2023.