From Street Markets to Shopping Malls: The Modern Service Multiplier*

Matthew Schwartzman†

February, 2026

Click here for the latest version

Abstract: Over the development process, modern service firms like supermarkets and restaurants replace traditional micro-entrepreneurs like street vendors and food hawkers. I argue that this transformation drives service-led growth. Empirically, modern services are more productive, and consumer demand for them rises with income. In a novel theory of the service sector, these two features imply a modern service multiplier: developments that favor service modernization set off a self-reinforcing growth cycle. As workers transition to the modern sector, they increase aggregate income and redirect demand to modern services, which pulls in more workers and generates further growth. I estimate the model with local-level data from Brazil and show that demand effects amplified the technology-driven rise of modern services by 25% between 2000 and 2010, generating service-led growth by shifting workers into more productive jobs. Demand effects also amplify misallocation, doubling the output loss to distortions in the poorest regions of Brazil.

*First and foremost, I thank Michael Peters, Fabrizio Zilibotti, and Kevin Donovan for their guidance on this project. I also thank Giuseppe Moscarini, Pascual Restrepo, Mushfiq Mobarak, Rohini Pande, Lucas Conwell, Alvaro Cox, Tim Hersey, Bernardo Ribeiro, Lindsey Uniat, Emily Breza, Paco Buera, Doug Gollin, Josh Hausman, Joe Kaboski, Renata Narita, Daniel Ramos-Menchelli, Federico Rossi, Todd Schoellman, Matthew Shapiro, Yongseok Shin, Gabriel Ulyssea, and numerous seminar participants for their insightful comments. I am grateful for the financial support of the Sylff Fellowship and the STEG PhD Research Grant.

†University of Michigan. Email: schwam@umich.edu.

1 Introduction

Most of the world's workers are service workers. Data from the International Labour Organization shows that in low- and middle-income countries, services surpassed agriculture as the largest sector of employment in 2012. The rise of services in the developing world has led to intense interest in how to foster a productive service sector and create good service jobs. Recent evidence is encouraging: services in low-income countries show signs of strong productivity growth (Fan et al., 2023; Nayyar et al., 2021). But theories of what drives service productivity growth are scarce, especially for the developing world, where the sector primarily consists of basic consumer services like retail and food service. Without a theory of service-led growth, economists do not have a framework to extrapolate beyond the limited empirical evidence and guide policymakers on how to foster service-sector growth.

In this paper, I argue that the key to understanding service-led growth in developing countries is to recognize services as a sector in transition. I focus on the transformation that occurs within labor-absorbing consumer service industries like retail, hospitality, and transportation. Over the development process, traditional micro-entrepreneurs in these industries – the street vendors, food hawkers, and rickshaw drivers that dominate services in low-income countries – are almost totally replaced by modern service firms like supermarkets, restaurants, and rideshare apps. The shift from individual service providers to larger firms is not a superficial change. It is a wholesale reorientation of services to more sophisticated forms of production.

To make this argument, I pair micro-data with a novel theory of the service sector. In the data and the theory, when workers take up modern service jobs, they generate first-order productivity gains in the style of Lewis (1954) and Song et al. (2011). Due to income effects on demand, the resulting growth causes consumers to redirect expenditure towards modern services and pull even more workers into the productive modern sector, setting off a virtuous cycle of productivity growth. To quantify the strength of this virtuous cycle, I define and characterize the *modern service multiplier*, which summarizes the self-reinforcing nature of service modernization. I then estimate the model's key parameters from micro-data and use the model to obtain a measure of the modern service multiplier in all 557 microregions of Brazil. I find that demand effects significantly accelerated service modernization in the Brazilian economy, particularly in the poorest regions. In turn, service modernization generated Lewisian service-led growth by reallocating workers to more productive jobs.

The paper begins by documenting key empirical features of services, both worldwide and in detailed micro-data on Brazilian workers and consumers. I first document service modernization as a development indicator: the composition of consumer service employment shifts from traditional self-employment without hired workers – own-account work – to wage employment at modern firms. The size of the shift is large enough that essentially all of the decline in urban own-account work with development happens within consumer service industries. Panel data on Brazilian workers shows that traditional services are less productive than modern services. Workers who transition from traditional to modern services earn a wage premium, and this premium is higher in poorer places.

On the consumer side, household expenditure data shows that rich households consume modern services more intensely than poor households. In fact, service consumption is highly segregated – a large share of households do not shop at modern service establishments at all. These patterns suggest that modern services are a luxury. Since service firms sell to local consumers, the modern service sector is sensitive to the level of local income. Modern services can only thrive when local household income is high enough to support demand. They also are sensitive to the size of the local market: even conditional on local income, modern services are more dominant in populous, highly urbanized markets.

Guided by the empirical facts, I develop a novel theory that micro-founds the modernization of services in decisions by heterogeneous consumers. The model breaks from most of the literature on modernization by introducing income effects on consumer demand. Many existing models of problems like modern technology adoption assume a single homogeneous final good. This assumption is appropriate to study a sector like manufacturing in a small open economy, but it is a poor fit for services. Service providers sell a differentiated final product to local consumers. The micro-data shows that those consumers' demand for modern services is sensitive to their level of income, and that some of them do not even use modern services. I therefore draw on Lagakos (2016) to model the demand side, and assume that consumers must pay a fixed access cost to enjoy modern services, capturing the idea that these services are complemented by consumer-side investments in products like cars. The fixed cost generates the empirically relevant extensive margin of modern service consumption and implies non-homothetic demand: it is a more substantial barrier to poor consumers and therefore has more bite on service consumption when aggregate income is low.

On the labor market side, modern services use wage labor, which is subject to frictions. Traditional services have the advantage of using own-account labor, which acts as a frictionless outside option to wage work. Frictions therefore create a wedge between the marginal revenue product of labor in modern and traditional services. The wedge summarizes the wide range of extra costs that modern firms face in hiring workers, ranging from search frictions to taxes, minimum wages, and other regulations. When this wedge is present, aggregate productivity is determined in part by the size of the modern sector, as in the canonical Lewis (1954) model.

The central insight of the model is that when the economy departs from homothetic demand and frictionless labor markets, service modernization becomes self-reinforcing. Labor market frictions and income effects interact. The same workers who are locked out of wage work by frictions turn to less productive traditional service work and become poor consumers with weak demand for modern services. When workers and firms produce modern services together, they generate Lewisian productivity gains, increasing aggregate income, reducing the bite of consumer fixed costs, and raising demand for more modern service production.

These demand effects amplify the model's comparative statics: they imply that service modernization begets more service modernization. To summarize the degree to which modern services pull themselves into existence, I introduce and define the modern service multiplier. Similar to the amplification rate in Buera et al. (2023), the modern service multiplier quantifies the strength of the model's amplifying forces. I show that the same conditions that imply Lewisian gains also increase the economy's multiplier when demand effects are present. When the model's amplifying forces are strong enough, the economy can support multiple equilibria.

The key elements of the model – income effects on demand and labor market frictions – imply that service modernization is a powerful and potentially self-sustaining engine of growth. In response to a growth catalyst such as improvements in technology or declining frictions, service modernization can propel the economy to even higher levels of income. In the quantitative part of the paper, I document that modern service technology was the main driver of service modernization in Brazil from 2000-2010, that a substantial portion of the gains it delivered were Lewisian reallocation gains, and that its impact was amplified by demand effects.

I estimate the model in order to measure Lewisian gains and the modern service multiplier in every microregion in Brazil. Because services respond to local demand, these objects must be measured at the local level. The key data sources are Brazilian Census data on the local allocation of employment and earnings among service workers, panel data on worker employment and earnings dynamics, and consumer expenditure data. I use the expenditure data to estimate the size of consumer income effects, as parameterized by fixed access costs. Then I pair the model with Census data to quantify labor market frictions, the technology level of non-service production, and the technology level of modern services for every microregion in Brazil in 2000 and 2010. Together, these local economic primitives allow me to compute the local multiplier.

I validate the estimated model using several untargeted moments, including the empirical response of the service sector to a natural experiment that shocked local aggregate income. For the natural experiment, I follow the work of Dix-Carneiro and Kovak (2017) and many others to estimate the local effects of exposure to import competition in the wake of Brazil's unilateral trade liberalization. Regions that were more exposed to import competition experienced a drop in the value of their industrial output, a negative shock to local income that originated outside the service sector. Consistent with the model, the income shock was a negative demand shock for modern services: it caused decreases in the modern employment share and modern wage premium among service workers.

In the final part of the paper, I use counterfactual simulations of the estimated model to quantify the drivers and amplifiers of service modernization and growth within Brazil. I find that technological progress in modern services was the most important primitive driver of service modernization. One-third of the service productivity gains it delivered were Lewisian gains from the reallocation of economic resources into modern services. Demand effects amplify the impact of modern service technology: without them, the modern service technology would have needed to grow 25% faster to generate the same degree of service modernization.

These aggregate effects mask significant local-level heterogeneity across Brazil. Lewisian gains and amplification are especially strong in Brazil's poorest microregions, where labor market frictions and fixed costs to access modern services bite the hardest. Because these forces widen the productivity gap between modern and traditional services, they create the conditions that most favor self-reinforcing productivity growth from service modernization. Richer microregions have mostly escaped these conditions, meaning there is less potential for their service economies to benefit from Lewisian gains and multiplier effects.

Labor market frictions did not drive significant service modernization from 2000 to 2010, but demand effects amplify their role in explaining productivity differences across space. In the fully estimated model, frictions generate output losses in the poorest microregions that are twice as large as in a counterfactual model with demand effects shut down. Models of modernization that abstract away from demand-side considerations are therefore likely to underestimate the aggregate cost of distortions.

Related literature I study the productivity of the service sector in developing countries, a subject of intense interest in recent economic research (Nayyar et al., 2021; Rodrik and Sandhu, 2024). Fan et al. (2023) measure significant increases of service TFP in Indian districts but are agnostic about its sources. In higher-income settings, Buera and Kaboski (2012) present a theory where the growth of market services stems from rising demand for more specialized and skill-intensive output. Eckert et al. (2022) argue that a decline in the price of information and communications technology (ICT) led to productivity growth in business service firms. Hsieh and Rossi-Hansberg (2023) propose a theory of an "industrial revolution in services," also precipitated by ICT, and reflected in the spatial expansion of large, modern consumer service firms. These recent theories break with the older Baumol (1967) view that services have low capacity for productivity growth.

My paper provides a bridge between the empirical literature on service productivity growth in low-income countries and theories of service-led growth in high-income countries. In my model, technological progress that favors modern firms, like ICT, also sets off service growth, but the growth mechanism is different. Rather than a mechanical effect from raising productivity in the entire service sector, gains come largely from reallocating labor into more productive modern jobs, generating productivity gains on the margin in the style of Lewis (1954). This mechanism is similar to the reallocation of capital to better manufacturing firms in Song et al. (2011), or to the expansion of productive firms due to management capital (Akcigit et al., 2021; Hjort et al., 2022; Cox, 2023; Engbom et al., 2024). I focus my analysis on labor-absorbing consumer services, since skilled service professions comprise only a small fraction of employment in low-income settings. While I focus on modernization in these industries, these are also the same services that are subject to marketization as female labor force participation rises (Ngai and Petrongolo, 2017).

As in studies of the "big push" like Rosenstein-Rodan (1943), Murphy et al. (1989), Matsuyama (1991), and Buera et al. (2023), the model generates a demand externality to participating in the modern sector, which can support multiple equilibria in extreme cases. In most of these papers, the demand externality comes from fixed costs of modern production that can only be recouped with high enough aggregate demand. In my model, it comes from workers getting richer, spending more on modern services as consumers, and drawing other workers' labor into the modern sector, so that those workers also get richer. It is a particular instance of the point made by Cooper and John (1988): strategic complementarities can generate multiple equilibria and multipliers. In recent work, Walker et al. (2024) develop a model where input indivisibilities on the supply side amplify comparative statics in low-income settings. My paper focuses instead on the role of income effects on the demand side.

In the micro-foundations of the theory, I model own-account work as a result of frictions and subsistence concerns, in the vein of research on informality and subsistence self-employment (Ulyssea, 2010; Meghir et al., 2015; Narita, 2020; Poschke, 2023; Donovan et al., 2023; Schoar, 2010; Poschke, 2013; Breza et al., 2021; Herreño and Ocampo, 2023), and as a result of worker sorting (Gollin, 2008; Feng et al., 2024). I therefore capture two roles of own-account work: some workers turn to it out of necessity, while for others it is the best use of their labor. My micro-foundation of consumer demand includes a fixed cost of modern shopping as in Lagakos (2016) and Ramos-Menchelli and Sverdlin-Lisker (2022); Bronnenberg and Ellickson (2015) also discuss how modern retailers benefit from consumer-side investments in goods like cars and refrigerators. In my empirics, I use Bachas et al. (2023)'s classification of establishments as modern or traditional. As in that paper and Atkin et al. (2018), I find that richer consumers disproportionately shop from modern firms. The model also implies that exogenous increases in local income have a formalizing effect on the local economy, as in Gerard et al. (2021).

The rest of the paper is as follows. In Section 2, I document the shift to modern services in aggregate, and use micro-data to show that modern service firms have higher marginal productivity and sell to richer consumers. Section 3 lays out the model and defines the economy's equilibrium, while Section 4 explores the central implications of the theory. Section 5 estimates the model at the level of local Brazilian economies and validates the estimation against several untargeted moments. In Section 6, I run counterfactual simulations of the estimated model to determine the magnitude of Lewisian gains and amplification. Section 7 concludes.

2 Motivating Facts

Before formalizing my argument in a model, I present five facts about modern services that motivate the analysis. A coherent picture emerges from the facts: as the aggregate economy transitions to a modern service sector, individual service workers make a parallel transition out of subsistence work as micro-entrepreneurs into new careers as wage workers. When they make this transition, workers' labor becomes significantly more productive and their income rises. In turn, as richer consumers they make a consumption transition to more intensive use of modern services, so that the shift to modern services is partially self-sustaining. Market size effects – which are stronger in modern than traditional services – further amplify the transition to modern services.

2.1 Service Transformation Along the Development Path

The modern service share is a robust indicator of economic development, both worldwide and in Brazil. Figure 1 shows how the composition of the service sector varies with aggregate income. I focus on those consumer service industries that employ a large share of workers even in low-income countries: wholesale and retail trade, hospitality and food service, transportation, and other personal services.[1] Figure 1a plots the share of these basic service workers in modern employment – defined as formal wage employment[2] – against the log of aggregate real GDP per capita for 199 country-years as observed in data from IPUMS and the Penn World Table. Figure 1b plots the same relationship among Brazilian microregions in 2000 and 2010, using local employment and income data from the decennial Census. Microregions in Brazil are a collection of municipalities that constitute a single labor market; they are analogous to commuting zones in the United States.

Across both country-years and Brazilian microregions, the modern service share rises steeply with aggregate income. The modern sector accounts for approximately 100% of service employment in the world's richest countries, but a drastically lower share in poorer places. The same patterns hold within Brazil: the modern service share is tightly linked with aggregate income in the cross-section, and it increased by an average of 10 percentage points from 2000-2010. Notably, the rise of modern services was balanced across Brazil: poor microregions and rich microregions alike saw the same magnitude of service transformation. In the quantitative section, I document that this balanced transformation reflects across-the-board improvements in the technology used by modern service firms.

[1] In Brazil, these industries account for 35% of urban employment.

[2] A worker is counted as formal if they are part of the social security system.

Figure 1: Service Modernization and Development, Worldwide and in Brazil

(a) Service Modernization Worldwide **(b)** Service Modernization in Brazil by Year

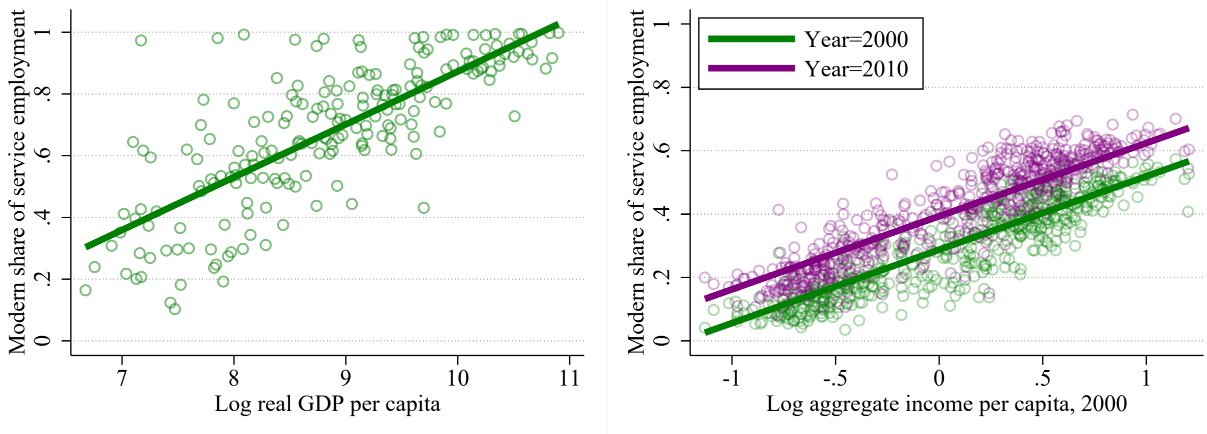

Note: Left panel: Employment data is from IPUMS data on 199 country-years spanning 76 unique countries; real GDP per capita data is from the Penn World Table for the same country-years. Right panel: data on both employment and income from the long form of Brazil's decennial Census. "Consumer services" refers to wholesale and retail trade, hospitality and food service, transportation, and other personal services. Modern employment is work for wages with pension contributions. All aggregations within a region-year use person weights provided by either IPUMS or the Brazilian Census.

2.2 Service Transformation Drives the Decline of Own-Account Work

The modern service share is a clear development indicator. Another more widely-known development indicator is the share of self-employed or own-account workers in the economy: richer economies have a much smaller share of the labor force operating their own production unit without outside employees (Gollin, 2008; Feng et al., 2024). Figure 2 shows that these two indicators reflect the same underlying trend: the exit of service micro-entrepreneurs along the development path. Each panel of Figure 2 presents two relationships. First, the red solid line displays own-account employment as a share of total non-agricultural employment. Second, the blue dashed line displays own-account work specifically within consumer services, again as a share of total employment outside of agriculture. Panel 2a plots these relationships across country-years, while Panel 2b does so across microregion-years within Brazil.

In both cases, most of the decline in own-account work with development is accounted for by the decline in consumer service own-account work. Within Brazil, the decreasing share of consumer service micro-entrepreneurs accounts for the entire gradient of own-account work with respect to aggregate income; own-account work in other sectors does not decline at all with development. If high rates of own-account work in low-income settings reflect subsistence self-employment, then this fact shows that urban subsistence labor pools specifically in the service sector.[3] As modern service businesses replace traditional micro-enterprises, the workers who

[3] For more thorough discussions and evidence of subsistence self-employment, see Schoar (2010), Poschke

Figure 2: Own-account work, services, and development

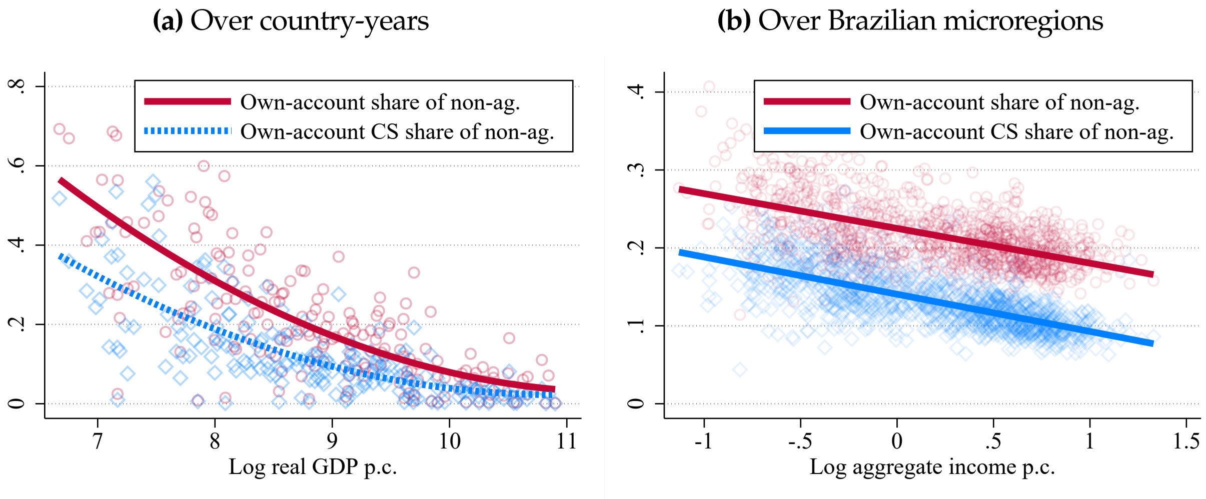

Note: Same data sources as in Figure 1: IPUMS + Penn World Table for country-years, Brazilian Census for microregion-years. Trend lines in the left panel use a quadratic fit, while in the right panel they use a linear fit.

operated those micro-enterprises exit subsistence labor and transition to the pool of wage labor.

2.3 Modern Services Are More Productive

The facts in Figures 1 and 2 do not necessarily imply that the shift to modern service employment causes development in and of itself. The high share of traditional and own-account service work in poor economies could simply indicate that modern services there are unproductive. If this is true, the composition of the service sector might be efficient: workers on the margin are no more productive in modern services than traditional services, so moving them into the modern sector does nothing to improve either individual income or aggregate output.

In Figure 3, I use panel data on Brazilian workers to test and reject this explanation for the high traditional share of services in poor regions. Specifically, I regress log earnings on an indicator for modern employment in a monthly panel of Brazil's urban workers, the Pesquisa Mensal de Emprego (PME).[4] I restrict the sample to service workers to test the hypothesis that service employment is efficiently allocated between traditional and modern services. If it is, then there should be no significant difference between modern and traditional earnings for workers on the margin between these sectors.

I test this hypothesis in two sub-samples: the red bars of Figure 3 present estimated sectoral earnings gaps for service workers in Brazil's underdeveloped Northeastern region,

(2013), Breza et al. (2021), Herreño and Ocampo (2023), and Donovan et al. (2023).

[4]The PME was collected through monthly interviews with workers in six of Brazil's largest metropolitan areas: Recife, Salvador, Belo Horizonte, Rio de Janeiro, Sao Paulo, and Porto Alegre. Workers are interviewed in a maximum of eight months over a sixteen-month interval.

while blue bars present estimates for the more developed South and Southeast.[5] For each sub-sample, I use three different specifications of controls. First, I control only for city-month fixed effects to compare all modern and traditional service workers within that city-month. Next, I add controls for observable worker characteristics – demographics, education, and so on – to account for differences in the composition of modern and traditional workers. Finally, I use a specification controlling for individual fixed effects, so that the earnings gap is identified from workers who transition between modern and traditional services. I provide more detail about how I estimate the wage premium in Appendix A.2.

Figure 3: Worker earning premium from modern employment, by region

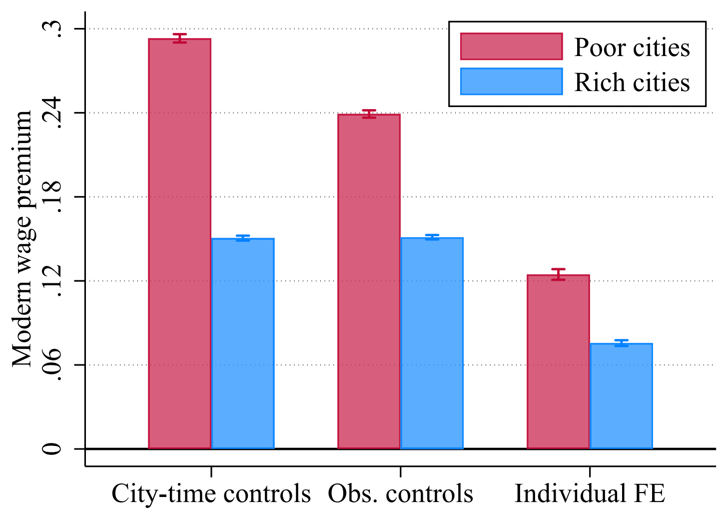

Note: Data on earnings and employment status from Brazil's PME. Bars represent the coefficient from regressing the log of worker earnings on a dummy for modern employment, which refers to wage employment with pension contributions. Capped spikes represent 95% confidence intervals. "City-time controls" include fixed effects for metropolitan region interacted with month and year fixed effects; "observable controls" include the worker's age, age squared, gender, race, and education level, as well as the industry of employment; "individual FE" includes individual worker fixed effects. "Poor cities" refers to cities in Brazil's less developed Northeast, while "rich cities" refers to cities in the more developed Southeast and South.

In all specifications and sub-samples, there is a positive and significant premium on modern work. Even after controlling for individual fixed effects, marginal workers earn 7-12% more in the modern sector – rejecting the hypothesis that their labor is equally productive across sectors. Since wage workers do not necessarily capture the full marginal product of their labor, this number represents a lower bound on the gap in marginal productivity between modern and traditional services. The gap is significantly larger in poorer regions. This fact indicates that workers in poor places face more severe barriers to joining the modern sector, and that the sizable pool of own-account service work in those places largely reflects subsistence labor.

[5] In 2008, the Northeast accounted for 29% of Brazil's population but only 12% of its GDP. The Southeast accounted for 38% of population and 49% of GDP, while the South accounted for 12.5% of population and 15% of GDP.

However, it also implies that shifting marginal workers into more productive modern firms represents an opportunity for significant growth.

2.4 Consumer Demand: Modern Services As a Luxury

Figures 2 and 3 show that barriers to wage labor supply contribute to the low modern service share in poor economies. I now use data from Brazil's consumer expenditure survey (POF) to show that consumer demand also plays a significant role. For most purchases, the POF asks about the type of establishment where the purchase took place. I exploit this feature to classify each consumption outlay of a household as modern or traditional consumption, using the classification of consumer-facing establishments developed by Bachas et al. (2023).[6] For instance, a food hawker would count as traditional consumption, while a sit-down restaurant would count as modern consumption. I successfully classify consumption as modern or traditional for 95% of relevant household expenditures.[7]

Figure 4 highlights two notable features of modern consumption at the household level. Panel 4a presents a binned scatter plot that traces an Engel curve of the modern expenditure share against the log of monthly household income per person. Modern services are clearly a "luxury," with their consumption following a steep income gradient. Even conditional on a battery of controls, including geography,[8] households in the richest income bin spend about 12 percentage points more of their expenditure on modern consumption than households in the poorest income bin. Appendix A.3 provides more details about the Engel curve estimation.

Panel 4b provides context to interpret the income gradient in Panel 4a. It presents a histogram of the modern share of consumption expenditure by household, revealing a striking pattern: there are large mass points near 0% and 100%. Most households consume almost exclusively from either traditional or modern establishments. I show in Appendix A.4 that this stark pattern is not explained by the presence or lack of modern establishments in the household's local economy. If it was, the location-level histogram of modern service expenditure shares would resemble the household-level histogram; Appendix A.4 shows that it does not.

As households move along the Engel curve in Panel 4a, their consumption behavior shifts dramatically. Rather than gradually increasing their share of modern services, households are more likely to make a discrete jump in their modern service use as income rises. This observation motivates my modeling choice to introduce a fixed cost of consuming modern

[6] Bachas et al. (2023) construct this classification separately for a large number of countries' consumer expenditure surveys. I use the classification they developed specifically for Brazil.

[7] This excludes certain expenditure categories, such as rent and utilities, that do not take place at an establishment; the modules of the POF that deal with these expenditures do not ask about the location of purchase.

[8] Specifically, I control for the household's primary sampling unit (PSU). PSUs are a small geographic unit that serve the same role in Brazilian national surveys as a Census tract in the US. There are 12,800 PSUs in Brazil, of which 4,696 were sampled for the POF, giving each PSU an average population of 15,000 in 2008. By population, a PSU is equivalent to about 1.5 US ZIP codes.

Figure 4: Modern consumption expenditure among households

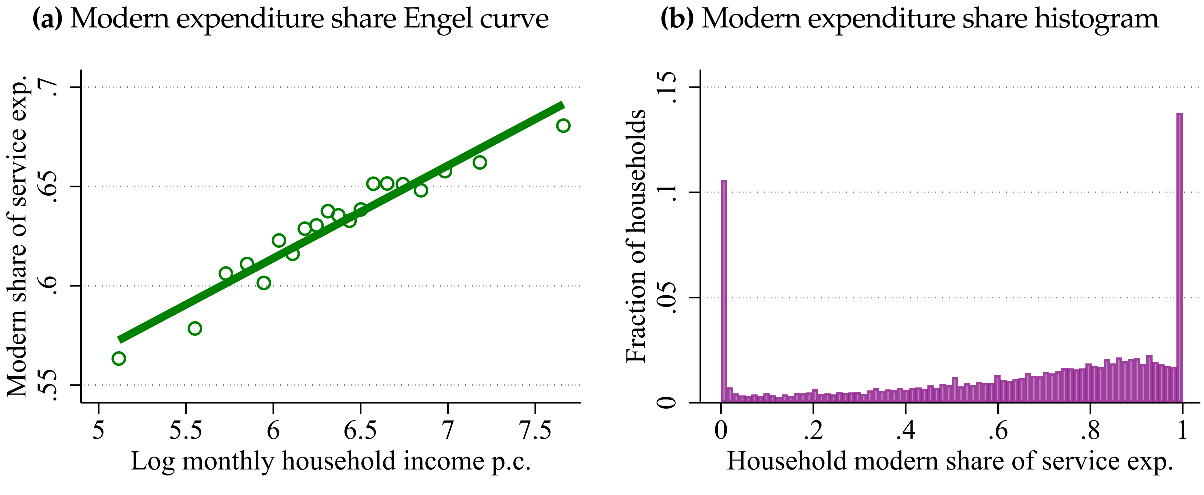

Note: Data on household income and consumption establishment from Brazil's POF. I use the Bachas et al. (2023) classification of establishments to classify consumption as modern or traditional. To compute the modern expenditure share, I use household expenditures aggregated over the previous year. The binned scatter in the Engel curve includes household-level controls for number of household members, number of earners, household head's race, gender, and age by 10-year dummies. It also includes primary sampling unit (PSU) fixed effects to control for geographic location.

services, which disproportionately constrains poor households.

The facts about consumption of modern services are consistent with those found by Lagakos (2016), who documents that the county- and country-level share of modern retail – the largest consumer service industry – increases with aggregate income. Combined with the productivity gap between modern and traditional services, they imply that service transformation can be self-reinforcing. As workers move to the modern sector, their labor becomes discretely more productive and their income rises. In the product market, they become richer consumers with higher demand for modern services. This shift of demand increases the marginal revenue product of modern service labor and pulls even more workers into the modern sector, setting off a virtuous cycle of reallocation-driven growth and rising demand.

2.5 Modern Services Respond to Market Size

Figure 4 establishes a key source of amplification in the process of service transformation: income growth re-directs demand towards the more productive modern service sector, generating further income growth. The next and final fact establishes an additional source of amplification: modern services have greater returns to scale at the sectoral level than traditional services. In particular, *market size effects* favor the modern service sector. To demonstrate the importance of market size for modern services, Figure 5 plots the modern service share against the log of microregion-level population density in the 2000 and 2010 Brazilian Census. While the relationship with population density is noisier than with aggregate income, modern services clearly dominate in more densely populated regions.

Figure 5: Service Modernization and Population Density

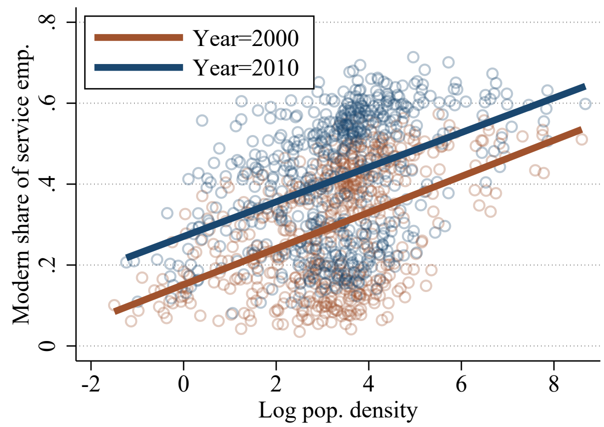

Note: Data from the long form of Brazil's decennial Census. All aggregations within microregion-year use person weights provided by the Census.

Unsurprisingly, part of this relationship is driven by underlying variation in income: the microregions with the highest population density also tend to have higher average income. In Table A.1 of the Appendix, I regress the modern service share against population density while controlling for local income per capita. While the relationship between modern services and population density is somewhat attenuated after controlling for income, it remains positive and statistically significant.

The descriptive evidence therefore supports market size effects that favor the modern service sector. It aligns with the tendency for modern services to agglomerate in shopping malls and commercial districts. Evidence from natural experiments also lends support to the hypothesis that modern services benefit from a larger market. In particular, Imbert and Ulyssea (2023) use a shift-share instrument to show that migration-induced population growth in Brazil increases the formal share of employment in the long run. Their work does not focus on the composition of the service sector, but since a large share of informal employment is in providing traditional services, their results suggest that exogenous increases in market size re-orient the service sector towards greater provision of modern services.

3 Theory

The previous section showed that consumer services shift to modern service firms with development, that services also host a traditional "subsistence" sector of own-account workers, that modern services pay a wage premium relative to own-account work, they are a luxury preferred by richer consumers, and that they benefit from market size effects. I now provide a novel theory of the service sector that accounts for these facts. In the model, modern service consumption carries a fixed cost that makes it a luxury, while wage employment is subject to frictions that generate a wage premium and a pool of subsistence own-account workers.

In order to highlight how key features of the modern service sector – the wage premium it pays over traditional services, its status as a luxury sector, and its gains from agglomeration – give rise to a process of amplified modernization and growth in services, I first present the model with homogeneous labor. Before taking the model to the data, I show how to enrich it with worker heterogeneity, which acts as a dampening force, and with trade in goods.

3.1 Households

The economy consists of a measure L of households that are identical except for the cost they face to access modern services. Each household inelastically supplies a measure 1 of labor to the economy and owns a share $\frac{1}{L}$ of a mutual fund that runs all firms in the economy and distributes firm profits to the household. The household spends all labor income and profits as nominal expenditure e. Household utility is linear in net consumption C_i of the composite final consumption commodity.

Households i source consumption from three sectors: the goods sector, denoted $C_{i,X}$, and the modern and traditional service sectors, denoted $C_{i,M}, C_{i,T}$, and combine them according to the following production function to create the final consumption commodity Y_i:

$$Y_i = \frac{C_{i,X}^{1-\phi}\left(\omega_M^{\frac{1}{\xi}} C_{i,M}^{\frac{\xi-1}{\xi}} + \omega_T^{\frac{1}{\xi}} C_{i,T}^{\frac{\xi-1}{\xi}}\right)^{\frac{\xi}{\xi-1}\phi}}{(1-\phi)^{1-\phi}\phi^{\phi}}, \tag{1}$$

where $\omega_M + \omega_T = 1$.

Gross final consumption Y_i therefore combines goods and services in a Cobb-Douglas aggregator that assigns weight $1-\phi$ to goods and weight ϕ to services. Services are themselves a CES composite of modern and traditional services, with elasticity of substitution ξ. I assume going forward that $\xi > 1$, so that modern and traditional services are gross substitutes. This production structure implies the final price index P, which is a function of the goods price P_X and the composite service price P_S, which in turn depends on the modern and traditional service prices P_M, P_T:

$$P = P_X^{1-\phi} P_S^{\phi}; \quad P_S = \left(\omega_M P_M^{1-\xi} + \omega_T P_T^{1-\xi}\right)^{\frac{1}{1-\xi}}. \tag{2}$$

Key assumption: Fixed cost of modern consumption The previous section documented two key features of modern service demand. First, modern services are a luxury: as incomes rise, so does the share of modern services in household consumption. Second, the extensive margin of modern service consumption is empirically relevant: a large share of households do not consume modern services at all. To capture both of these features, I assume that households must pay a fixed cost denominated in the final good to consume a positive quantity of modern services. In particular, the household's utility-relevant *net* consumption is

$$C_i = Y_i - \kappa_i \mathbf{1}\{C_{i,M} > 0\}, \tag{3}$$

where κ_i varies across households according to the CDF $F_\kappa(\kappa_i)$.

The fixed cost captures the idea that the modern service sector requires a complementary input from consumers. This input includes transportation and storage technologies, like the crucial role of cars emphasized by Lagakos (2016) or fuel in Ramos-Menchelli and Sverdlin-Lisker (2022). Fixed costs generate income effects on modern service demand, implying an upward-sloping modern service Engel curve.

To see this more formally, consider the problem facing household i, who chooses whether or not to source any consumption from the modern service sector:

$$U_i = C_i = \max\left\{ \frac{\omega_T^{\frac{\phi}{\xi-1}} e}{P_X^{1-\phi} P_T^{\phi}}, \frac{e}{P} - \kappa_i \right\}, \tag{4}$$

where the first expression in the braces represents real consumption if the household chooses to forgo modern services, and the second expression represents consumption if the household pays the fixed cost to access a full consumption bundle, including modern services. Households choose to consume modern services if and only if their fixed cost κ_i is below some threshold κ^*, the value of κ at which households are indifferent about including modern services in their consumption bundle:

$$C_{i,M} > 0 \Longleftrightarrow \kappa_i \leq \kappa^*\left(e, \vec{P}\right) = \frac{e}{P}\left(1 - \left[\frac{\omega_M}{\omega_T}\left(\frac{P_M}{P_T}\right)^{1-\xi} + 1\right]^{\frac{\phi}{1-\xi}}\right). \tag{5}$$

The household's modern service consumption increases in real purchasing power $y = \frac{e}{P}$, which allows them to overcome the fixed cost and fully participate in the market for services. The fixed cost also implies consumption segregation: those households whose purchasing power is not high enough to overcome the fixed cost κ_i are locked out of participation in the modern service economy, limiting potential demand for modern services. Note that the threshold κ^* can be written as $\kappa^*\left(y, \frac{P_M}{P_T}\right)$, where $\frac{\partial \kappa^*}{\partial y} > 0$ and $\frac{\partial \kappa^*}{\partial (P_M/P_T)} < 0$.

Aggregate expenditure shares Because final consumption is Cobb-Douglas over goods and services, a share $1-\phi$ of expenditure is devoted to goods, while a share ϕ is devoted to services. Among service expenditures, a share ϑ_M are devoted to modern services, where ϑ_M satisfies

$$\vartheta_M = \underbrace{F_\kappa\left(\kappa^*\left(y,\frac{P_M}{P_T}\right)\right)}_{\text{Potential market}} \times \underbrace{\omega_M\left(\frac{P_M}{P_S}\right)^{1-\xi}}_{\text{Share of potential market}}. \tag{6}$$

The variable ϑ_M, the modern share of service expenditures, is the key endogenous object capturing the degree of service modernization. The market size of modern services is limited by the fixed cost κ: only a share $F_\kappa(\kappa^*)$ of households are in the market for modern services. In this sense, the fixed cost formulation of consumer demand shares some features with Stone-Geary preferences, with modern services as a luxury: holding prices fixed, households only consume modern services once their income passes a certain threshold. But unlike Stone-Geary, those households who enter the market for modern services make a discrete jump to consuming a positive quantity of modern services $C_{i,M} >> 0$. In particular, those households who participate in the market for modern services give it the standard CES share of service expenditures, $\omega_M\left(\frac{P_M}{P_S}\right)^{1-\xi}$. This formulation is consistent with the bifurcated distribution of modern service expenditures in Figure 4b.

With this demand system, aggregate expenditure E_k on each sector $k \in \{X,M,T\}$ satisfies

$$E_X = P_X Y_X = (1-\phi)eL; \quad E_M = P_M Y_M = \phi\vartheta_M eL; \quad E_T = P_T Y_T = \phi(1-\vartheta_M)eL. \tag{7}$$

3.2 Production

Labor is the only factor of production. In the goods and modern service sectors, labor is hired by firms for a wage w to produce differentiated varieties n of goods and services. In the traditional service sector, workers engage in informal production to produce an undifferentiated service, and earn the revenues from selling their services to consumers at price P_T. In particular, aggregate sectoral production is

$$\begin{aligned} Y_X &= \left(\int_0^{N_X} y_X(n_X)^{\frac{\sigma-1}{\sigma}} dn_X\right)^{\frac{\sigma}{\sigma-1}}, \quad y_X(n_X) = z_X h_X^P(n_X) \\ Y_M &= \left(\int_0^{N_M} y_M(n_M)^{\frac{\sigma-1}{\sigma}} dn_M\right)^{\frac{\sigma}{\sigma-1}}, \quad y_M(n_M) = z_M h_M^P(n_M) \\ Y_T &= z_T H_T, \end{aligned} \tag{8}$$

where $h_k^P(n_k)$ denotes production labor used for variety n_k in sector $k \in \{X,M\}$, and H_T denotes aggregate labor in the traditional service sector. Each sector $k \in \{X,M,T\}$ has a technology level

z_k that transforms one unit of production labor into z_k units of output.

Operating a firm in the goods and modern service sectors requires a fixed cost of $f_k, k \in \{X, M\}$ units of overhead labor. Aggregate labor in these sectors satisfies:

$$\begin{aligned} H_X &= \int_0^{N_X} \left(f_X + h_X^P(n_X)\right) dn_X; \\ H_M &= \int_0^{N_M} \left(f_M + h_M^P(n_M)\right) dn_M, \end{aligned} \tag{9}$$

and by the resource constraint, $H_X + H_M + H_T = L$.

Key assumption: Love-of-variety The assumption of love-of-variety ($\sigma < \infty$) in goods and modern services is a theoretical micro-foundation of the empirical finding that employment in traditional services declines in aggregate market size, as supported by both descriptive empirics (see Figure 5) and natural experiments such as Imbert and Ulyssea (2023). Love-of-variety implies that modern services have increasing returns to scale at the sectoral level, which acts as an additional amplifying force that accelerates service modernization.

Key assumption: Labor market friction An additional robust empirical finding is that workers' marginal product of labor is greater in modern services than in traditional services, supporting the notion of traditional service employment as "subsistence entrepreneurship." This form of employment emerges in response to frictions that limit workers' ability to take up jobs in modern services. I capture these frictions with a generic wedge τ that renders the marginal revenue product of labor unequal between traditional service self-employment and wage work:

$$w = (1+\tau) P_T z_T = (1+\tau) \text{MRPL}_T. \tag{10}$$

Rather than taking a stance on the specific friction that curtails formal wage employment, τ encapsulates a wide range of phenomena that limit workers' and firms' ability to produce together in the formal sector, such as taxes, search frictions, minimum wages, other regulations on the formal sector, efficiency wages stemming from agency problems, or worker disamenities from wage work.

Firm optimization and sectoral supply Firms in the goods and modern service sectors charge the standard markup over their constant marginal cost: $p_k(n_k) = \frac{\sigma}{\sigma-1}\frac{w}{z_k}$ for $k \in \{X, M\}$. Firm revenues are then $r_k(n_k) = \frac{\sigma}{\sigma-1} w h_k^P(n_k)$, and operating profits gross of fixed costs are $\frac{1}{\sigma-1} w h_k^P(n_k)$. Since firms enter until profits net of fixed costs are zero, $h_k^P(n_k) = (\sigma-1) f_k \forall n_k \in [0, N_k]$. Combined with labor aggregation over all firms from (9), this implies that aggregate sectoral labor is proportional to the sector's measure of active firms:

$$H_X = \sigma f_X N_X; \quad H_M = \sigma f_M N_M. \tag{11}$$

Each of these sectors has the standard CEX price index:

$$P_X = \left(\int_0^{N_X} \left(\frac{\sigma}{\sigma-1} \frac{w}{z_X} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} = \frac{\sigma}{\sigma-1} \frac{w}{z_X} N_X^{\frac{1}{1-\sigma}} = \frac{\sigma^{\frac{\sigma}{\sigma-1}}}{\sigma-1} \frac{w}{z_X} \left(\frac{f_X}{H_X} \right)^{\frac{1}{\sigma-1}};$$
$$P_M = \left(\int_0^{N_M} \left(\frac{\sigma}{\sigma-1} \frac{w}{z_M} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} = \frac{\sigma}{\sigma-1} \frac{w}{z_M} N_M^{\frac{1}{1-\sigma}} = \frac{\sigma^{\frac{\sigma}{\sigma-1}}}{\sigma-1} \frac{w}{z_M} \left(\frac{f_M}{H_M} \right)^{\frac{1}{\sigma-1}}, \tag{12}$$

while the price index in traditional services is straightforward and reflects the labor wedge τ:

$$P_T = \frac{w}{(1+\tau) z_T}. \tag{13}$$

3.3 Equilibrium

In equilibrium, households maximize utility, implying the sectoral demand system outlined in (7). Firms maximize profits and workers maximize income – subject to the labor market wedge – implying the sectoral prices in (12) and (13). In addition, markets must clear so that consumer expenditures equal revenues in every sector. As I show in Appendix B.1, market clearing and the lack of equilibrium profits imply that the economy's labor allocation $\vec{H}$ depends only on service modernization ϑ_M and exogenous parameters:

$$H_X(\vartheta_M) = \frac{1-\phi}{1+\tau\phi(1-\vartheta_M)} L;$$
$$H_M(\vartheta_M) = \frac{\phi\vartheta_M}{1+\tau\phi(1-\vartheta_M)} L; \tag{14}$$
$$H_T(\vartheta_M) = \frac{(1+\tau)\phi(1-\vartheta_M)}{1+\tau\phi(1-\vartheta_M)} L.$$

By the system of prices implied by the production side of the economy in (12) and (13), real income $y = \frac{e}{P}$ and relative prices are a function of parameters and the labor allocation $\vec{H}$. That is, they can be written as $y\left(\vec{H}(\vartheta_M)\right), \frac{P_M}{P_T}\left(\vec{H}(\vartheta_M)\right), \frac{P_M}{P_S}\left(\vec{H}(\vartheta_M)\right)$. In Appendix B.1, I provide the particular expressions for real income and relative prices as functions of $\vec{H}(\vartheta_M)$. Using the allocation in (14), and real income and relative prices as a function of that allocation, the equilibrium is characterized by a fixed point in the degree of service modernization ϑ_M:

$$\vartheta_M = \underbrace{F_\kappa\left(\kappa^*\left(y\left(\vec{H}(\vartheta_M)\right), \frac{P_M}{P_T}\left(\vec{H}(\vartheta_M)\right) \right) \right) \times \omega_M \left(\frac{P_M}{P_S}\left(\vec{H}(\vartheta_M)\right) \right)^{1-\xi}}_{T(\vartheta_M)}, \tag{15}$$

where $T(\vartheta_M)$ summarizes the transformation of ϑ_M on the right-hand side of the equation. Note that $T(\vartheta_M)$ is continuous and differentiable at all ϑ_M such that $F_\kappa(\kappa)$ is continuous and differentiable at $\kappa^*(\vartheta_M)$.

4 Theoretical Analysis and Extensions

Having established the economic environment and defined the equilibrium, I now analyze the model's key properties. First, because frictions imply an unequal marginal product of labor across sectors, modernization generates Lewisian reallocation gains. Second, developments that shift workers out of subsistence and into the modern sector are amplified by the response of consumer demand, leading to further modernization and Lewisian income growth. Modernization therefore adds to service-led growth and is self-reinforcing.

4.1 Modernization Gains and Amplification

To see the gains from modernization, note that in equilibrium, y can be written in terms of ϑ_M and economic primitives: $y^* = y(\vartheta_M|\vec{a})$, where $\vec{a}$ represents the vector of primitives. Consider the response of y to a change in any individual primitive a, in particular the comparative static $\frac{d\log y}{d\log a}$. The effects of a on real income y can be decomposed into a direct effect and a modernization effect:

$$\frac{d\log y}{d\log a} = \overbrace{\frac{\partial \log y}{\partial \log a}}^{\text{Direct effect}} + \overbrace{\frac{\partial \log y}{\partial \log \vartheta_M}\frac{d\log \vartheta_M}{d\log a}}^{\text{Modernization}} \tag{16}$$

The modernization effect is the product of two terms. The first is the modernization gain $\frac{\partial \log y}{\partial \log \vartheta_M}$: the relative growth of real income that follows directly from a relative increase in modernization. The second is the modernization elasticity $\frac{d\log \vartheta_M}{d\log a}$: the relative equilibrium response of modernization to a relative change in the primitive a. The following proposition establishes that in the presence of either labor market frictions or love-of-variety, the modernization gain is strictly positive.

Proposition 1. *Assume that $\tau \geq 0$ and $\sigma > 1$, and the equilibrium features a positive modern share ($\vartheta_M > 0$). Then the modernization gain $\frac{\partial \log y}{\partial \log \vartheta_M}$ is strictly positive if and only if $\tau > 0$ or $\sigma < \infty$.*

Proof. The partial derivative of $\log y$ with respect to $\log \vartheta_M$ can be written as:

$$\frac{\partial \log y}{\partial \log \vartheta_M} = h_M(\vartheta_M)\left(\tau + \frac{\tau(1-\phi)}{\sigma - 1} + \frac{1+\tau\phi}{F_\kappa(\kappa^*)(\sigma-1)}\right), \tag{17}$$

which is strictly positive when either $\tau > 0$ or $\sigma < \infty$, and zero when $\tau = 0$ and $\sigma \to \infty$. Here $h_M(\vartheta_M) = \frac{H_M(\vartheta_M)}{L}$. See Appendix B.2 for details on how to derive Equation (17). ■

Proposition 1 establishes that two of the most empirically relevant features of modern services – labor market frictions and agglomeration economies – imply that the process of service modernization itself contributes to growth of real income. Fundamentally, these are Lewisian gains from service modernization: the marginal product of labor is greater outside of traditional services due to both distortions ($\tau > 0$) and differences in return to scale ($\sigma < \infty$), so that reallocation of "surplus labor" out of traditional services contributes to income growth.

The modernization process also improves labor productivity in the service sector. To see this, denote revenue per worker in services deflated by the service price index as $A_S = \frac{E_S/H_S}{P_S}$, where $H_S = H_M + H_T$. Like y, the equilibrium value of A_S can be written in terms of ϑ_M (see Appendix B.2 for the full expression of $A_S(\vartheta_M)$). Also like y, A_S features Lewisian modernization gains when $\tau > 0$ or $\sigma < \infty$:

$$\frac{\partial \log A_S}{\partial \log \vartheta_M} = \vartheta_M \left(\frac{\tau}{1+\tau(1-\vartheta_M)} + \frac{1+\tau}{F_\kappa(\kappa^*)(\sigma-1)(1+\tau\phi(1-\vartheta_M))} \right). \tag{18}$$

Frictions and love-of-variety imply that any change in economic primitives that fosters a shift of expenditures towards the modern service sector will also foster productivity growth in the service sector, because service modernization involves a reallocation of resources to more productive uses.

Proposition 1 establishes conditions under which modernization gains are positive. In fact, the same conditions that imply modernization gains, i.e. $\frac{\partial \log y}{\partial \log \vartheta_M} > 0$, also increase the modernization elasticity $\frac{d \log \vartheta_M}{d \log a}$. That is, these conditions *amplify* the transition to modern services and the associated Lewisian gains. Consider the fixed point of ϑ_M characterized by Equation (15), and take the total derivative with respect to a (in logs) on both sides to find the modernization elasticity.

$$\begin{aligned}
\frac{d \log \vartheta_M}{d \log a} &= \frac{\partial \log T(\vartheta_M|a)}{\partial \log a} + \frac{\partial \log T(\vartheta_M|a)}{\partial \log \vartheta_M} \frac{d \log \vartheta_M}{d \log a} \\
&= \frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)} \left(\frac{\partial \log y}{\partial \log a} + \frac{\partial \log \kappa^*}{\partial \log P_M/P_T} \frac{\partial \log P_M/P_T}{\partial \log a} \right) + (1-\xi) \frac{\partial \log P_M/P_S}{\partial \log a} \\
&\quad + \left(\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)} \left(\frac{\partial \log y}{\partial \log \vartheta_M} + \frac{\partial \log \kappa^*}{\partial \log P_M/P_T} \frac{\partial \log P_M/P_T}{\partial \log \vartheta_M} \right) + (1-\xi) \frac{\partial \log P_M/P_S}{\partial \log \vartheta_M} \right) \frac{d \log \vartheta_M}{d \log a}.
\end{aligned} \tag{19}$$

Equation (19) is the key equation of the model. It demonstrates both how the introduction of income effects through fixed costs augments the comparative statics of the model, and how the rise of modern services is *self-reinforcing*.

To see the importance of income effects (which operate through fixed costs), note that most of the terms in (19) depend on $\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)}$. In a standard model without fixed costs to consume modern services, $\kappa_i \equiv 0$, so that $f_\kappa(\kappa) \equiv 0 \, \forall \kappa > 0$, and these terms would disappear. In the model with income effects, the modern service sector grows partially through demand generation: the expansion of the potential market for modern services. A portion of demand generation happens through direct impacts on κ^*. For example, increases in goods productivity z_X will grow the market for modern services through the mechanical effect of z_X on real incomes and hence on the potential market for modern services, because $\frac{\partial \log y}{\partial \log z_X} > 0$.

But under the right conditions, such as when modernization gains $\frac{\partial \log y}{\partial \log \vartheta_M}$ are positive, service modernization is self-reinforcing through the demand channel. A crucial term in (19) is $\frac{f_\kappa(\kappa^*)}{F_\kappa(\kappa^*)}\frac{\partial \log y}{\partial \log \vartheta_M}\frac{d\log \vartheta_M}{d\log a}$. This term captures the extent to which the modern service sector *generates its own demand* by delivering income gains to consumers, who in turn expand the potential market for modern services. At the same time, supply-side forces that limit the size of the modern service sector – such as the distortions parameterized by τ – echo into the demand side of the economy by reducing purchasing power and shifting demand towards the traditional sector.

More generally, all of the terms multiplying $\frac{d\log \vartheta_M}{d\log a}$ on the right-hand side of (19) are non-negative, and strictly positive under certain conditions on parameters. This feature of the model shows that modernization begets more modernization – to a significant extent, modern services pull themselves into existence. To understand the extent to which modernization is self-reinforcing, it is useful to define the modern service multiplier:

Definition 1. *The **modern service multiplier** μ_M is defined as:*

$$\mu_M = \frac{1}{1-T'(\vartheta_M)} = \frac{1}{1-\left(\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)}\left(\frac{\partial \log y}{\partial \log \vartheta_M}+\frac{\partial \log \kappa^*}{\partial \log P_M/P_T}\frac{\partial \log P_M/P_T}{\partial \log \vartheta_M}\right)+(1-\xi)\frac{\partial \log P_M/P_S}{\partial \log \vartheta_M}\right)}, \tag{20}$$

so that the modernization elasticity can be written as:

$$\frac{d\log \vartheta_M}{d\log a} = \mu_M \times \left(\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)}\left(\frac{\partial \log y}{\partial \log a}+\frac{\partial \log \kappa^*}{\partial \log P_M/P_T}\frac{\partial \log P_M/P_T}{\partial \log a}\right)+(1-\xi)\frac{\partial \log P_M/P_S}{\partial \log a}\right) \tag{21}$$

The next proposition writes out the expression for the modern service multiplier in terms of primitives, which clarifies the conditions that amplify the economy's comparative statics.

Proposition 2. *Given parameters, for any equilibrium $\vartheta_M > 0$ and corresponding $y(\vartheta_M)$, $h_M(\vartheta_M)$, $\kappa^*(\vartheta_M)$ s.t. $F(\kappa^*) > 0$, the modern service multiplier μ_M can be written as follows:*

$$\begin{aligned}\mu_M^{-1} = 1 &- \frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)}h_M(\vartheta_M)\left(\tau+\frac{\tau(1-\phi)}{\sigma-1}+\frac{1+\tau\phi}{F_\kappa(\kappa^*)(\sigma-1)}\right)\\ &- \frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)}\phi\frac{y-\kappa^*}{\kappa^*}\frac{\vartheta_M}{F_\kappa(\kappa^*)}\frac{1}{\sigma-1}\frac{1+\tau\phi}{1+\tau\phi(1-\vartheta_M)}\\ &-(\xi-1)\left(1-\frac{\vartheta_M}{F_\kappa(\kappa^*)}\right)\frac{1}{\sigma-1}\frac{1+\tau\phi}{1+\tau\phi(1-\vartheta_M)},\end{aligned} \tag{22}$$

so that μ_M^{-1} is less than 1 if $\tau\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)} > 0$ or $\sigma < \infty$, and equal to 1 if $\tau\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)} = 0$ and $\sigma \to \infty$.

Proof. See Appendix B.3 for details on how to derive (22). ■

Once again, the same conditions that deliver modernization gains – distortions τ and agglomeration economies $\sigma < \infty$ – also amplify the modernization process. In the case where $\sigma \rightarrow \infty$, non-homothetic demand $\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)} > 0$ is also necessary for distortions τ to act as an amplifier: the modernization gains implied by τ feed into further modernization only when rising incomes shift demand towards modern services. Under these conditions, a small change in primitives can trigger dramatic reallocation of expenditure and resources into the modern service sector, with every bit of reallocation driving Lewisian gains in a partially self-sustaining feedback loop.

This definition of the modern service multiplier relates closely to the amplification rate examined by Buera et al. (2023). The amplification rate in this context is $1 - \mu_M^{-1}$: when this expression is positive, the economy features modernization complementarities and amplification.

Visualizing amplification Figure 6 illustrates how distortions and agglomeration in the modern service sector amplify the economy's comparative statics. Both Figure 6a and Figure 6b plot the economy's equilibrium before and after a 10% increase in z_M, the technology level in the modern service sector. The black line in each is simply the 45-degree line of identity for ϑ_M. The green lines represent the transformation $T(\vartheta_M)$ from Equation (15); equilibrium is characterized by the fixed point where $T(\vartheta_M)$ intersects with the 45-degree line. The solid green line shows $T(\vartheta_M)$ before the 10% increase in z_M, while the dotted green line shows $T(\vartheta_M)$ after the increase in z_M. Both figures have the same baseline equilibrium before the increase in z_M, with an equilibrium value of $\vartheta_M = 0.57$.

In the economy with distortions ($\tau > 0$), an upward-sloping Engel curve ($\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)} > 0$), and love-of-variety ($\sigma < \infty$) plotted in Figure 6a, the transformation $T(\vartheta_M)$ is upward-sloping. Because, in general, the modern service multiplier $\mu_M = 1/(1 - T'(\vartheta_M))$, the positive slope of $T(\vartheta_M)$ amplifies the response to an increase in modern service technology z_M. This amplification reflects the Lewisian gains that shift demand further towards modern services as the sector grows. In the figure, amplification is illustrated by movement along the upward-sloping $T(\vartheta_M)$ curve towards the new equilibrium $\vartheta_M = 0.79$. In contrast, when these modernization gains are shut down in Figure 6b, $T(\vartheta_M)$ is flat and the economy does not feature any amplification: movement along $T(\vartheta_M)$ does not increase the equilibrium modernization share, and the modern service multiplier μ_M is equal to 1, causing only a modest increase to $\vartheta_M = 0.63$.

Potential multiplicity Proposition 2 lays out conditions under which the equilibrium features amplification, but does not guarantee that the equilibrium is unique. The model can in fact support multiple equilibria, in particular when it has an equilibrium in which $\mu_M^{-1} < 0$. As in Buera et al. (2023), this is the case when the amplification rate $1 - \mu_M^{-1}$ exceeds 1 for some values of ϑ_M. I establish this result more formally and discuss multiple equilibria further in Appendix B.4.

Figure 6: Equilibrium Comparative Statics With and Without Amplification

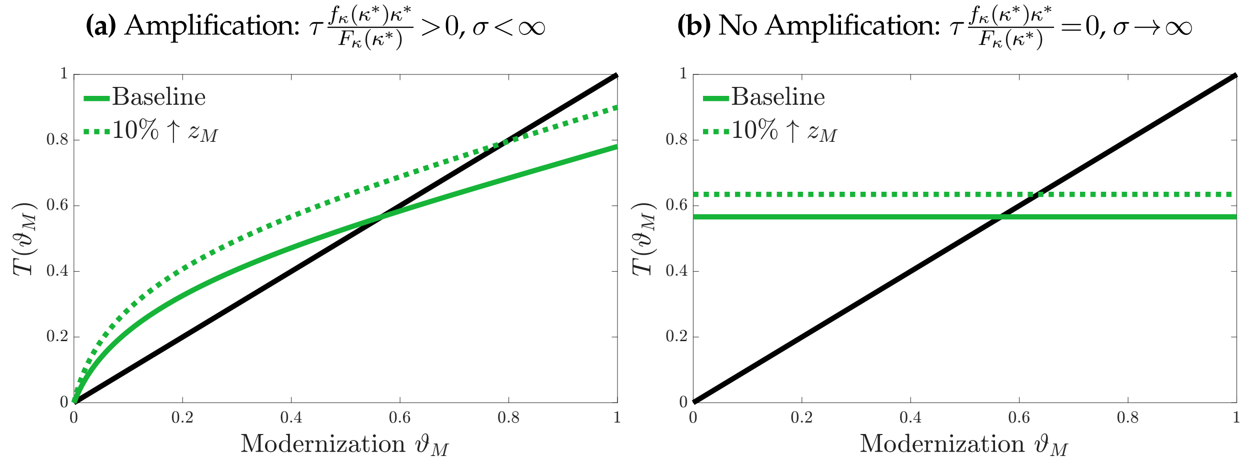

Note: Both figures plot $T(\vartheta_M)$ from Equation (15) under different parameter specifications. In Figure 6a, τ is set to 0.5, σ is set to 6, and $F_\kappa(\kappa)$ is non-degenerate. In Figure 6b, $\tau\frac{f_\kappa(\kappa)\kappa}{F_\kappa(\kappa)}=0$ and $\sigma\to\infty$. In both figures, the baseline equilibrium is at $\vartheta_M=0.57$.

4.2 Drivers of Service Modernization and Growth

By Equation (19), any primitive a such that $\frac{\partial \log T(\vartheta_M|a)}{\partial \log a}\neq 0$ contributes to modernization. And Proposition 1 establishes that under generic conditions, the modernization generated by these primitives creates Lewisian gains to real income and service productivity.[9] Guided by the model, I focus my attention on four primitives that can act as catalysts of the modernization process. The quantitative section of the paper directly measures each primitive's contribution to modernization within Brazil.

Goods technology z_X: The level of technology in goods z_X – which acts as a stand-in for all real income shifters outside the service sector – drives service modernization whenever a positive density of households are on the extensive margin of consuming modern services. In particular, $\frac{\partial \log T(\vartheta_M|z_X)}{\partial \log z_X}=(1-\phi)\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)}$. Propositions 1 and 2 highlight that under the right conditions, the initial shock to real income implied by an increase in z_X sets off a virtuous cycle of service modernization and further income growth. In other words, income effects on demand – captured here by $f_\kappa(\kappa^*)>0$ – mean that developments *outside* the service sector can lead to modernization and productivity growth *within* the service sector.

Total market size L: Because modern services and goods both have increasing returns to scale at the sector level whenever $\sigma<\infty$, total market size L effectively raises productivity in these sectors, thereby encouraging service modernization. L leads to service modernization both by raising real income and by reducing the relative price of modern services. In other words, $\frac{\partial \log T(\vartheta_M|L)}{\partial \log L}>0$, so that population growth and urbanization can act as another catalyst

[9] Many primitives also have a direct effect on real income and service productivity that does not operate through the modernization channel.

of service-led growth.

Labor market frictions τ: Because the labor market friction τ distorts the allocation of labor towards traditional services, reductions in τ naturally lead labor to reallocate towards modern services. This prediction would exist even in more standard models that omit income effects and scale effects in modern services. However, because these additional model ingredients contribute to the modern service multiplier of Definition 1, they imply that labor market frictions are substantially more costly – and their alleviation leads to substantially larger reallocation and gains – than in a standard model. Labor market frictions do not merely act as a supply shifter away from modern services; with non-homothetic demand they have the added effect of suppressing demand and shrinking the market for the more productive modern service sector.

Modern service technology z_M: Similarly, because modern and traditional services are gross substitutes ($\xi > 1$), technological improvements within modern services such as ICT also lead to service modernization and real income growth. Again, this is true whether or not the model has distortions, income effects, or scale effects. But the role that these key ingredients play in *amplifying* the comparative statics of the model means that even modest changes in technology can lead to dramatic differences in the composition and productivity of the service sector. In other words, even if z_M is a residual of the model, the model's novel ingredients shrink the variation in this residual that is necessary to explain variation in modern service shares.

4.3 Enriching the Model

The model presented in Section 3 made two simplifying assumptions: a homogeneous labor force and a closed economy. However, the descriptive facts from worker panel data show that the unconditional modern wage premium partially reflects worker selection. And as model validation, I use a trade shock to show that trade-induced income shocks act as demand shocks for modern services. I therefore enrich the model to accommodate worker heterogeneity and selection, and trade in goods. This extended model is the one that I take to the data in Section 5.

4.3.1 Worker Heterogeneity

To allow for worker heterogeneity and sorting, I nest a Roy model with selection into wage work within the model laid out above. I assume that each of the L households in the economy contains a measure 1 of individual workers. Each worker supplies all of their labor either to wage work in modern services and tradable goods, or to own-account work in traditional services. Workers are heterogeneous along two dimensions: their ability as a wage worker, and their ability as an own-account worker. In particular, a worker j can supply h_j efficiency units of wage labor or $s_j h_j$ efficiency units of own-account labor, so that h_j captures the worker's absolute ability, while s_j captures the worker's relative ability as an own-account worker. Within the household, h and s are distributed jointly according to the distribution function $G(h,s)$. This distribution is identical across households, each of which pools the labor income of all the workers it contains, so that all households earn the representative level of nominal income e.

A type-(h,s) worker earns wh if she supplies her labor to wage work, and $P_T z_T sh$ if she supplies her labor to own-account work.[10] Workers therefore sort between wage work and own-account work on s: those with s below a threshold s_w specialize in wage work, while those with $s > s_w$ become own-account workers. In a frictionless world, $s_w = \frac{w}{P_T z_T}$ so that the marginal worker with $s = s_w$ has identical earnings in wage and own-account work. The model with heterogeneous labor preserves the wedge τ as a distortion relative to efficient sorting. In particular, $s_w = \frac{w}{(1+\tau)P_T z_T}$, so that workers at the margin s_w earn strictly more from wage work.[11]

Let $\tilde{G}(s) = \lim_{h\to\infty} G(s,h)$ be the distribution function of s over workers. Then the measures of effective wage and own-account labor, respectively H_w, H_T, are

$$H_w = \tilde{G}(s_w)\mathbb{E}[h|s \leq s_w]L; \quad H_T = \Big(1 - \tilde{G}(s_w)\Big)\mathbb{E}[sh|s > s_w]L. \tag{23}$$

Because effective labor allocations and relative prices now depend on the sorting threshold s_w, adding worker heterogeneity requires some modification of the market-clearing conditions of the model. In Appendix B.5, I provide these market-clearing conditions and characterize the model's equilibrium under worker heterogeneity. Once again, equilibrium is characterized by a fixed point $\vartheta_M = T^G(\vartheta_M)$, where $T^G(\cdot)$ represents the transformation of ϑ_M under worker heterogeneity, when the distribution function $G(\cdot)$ is non-degenerate.

The next proposition establishes how worker heterogeneity modifies the model's key properties.

Proposition 3. *Assume that $\tau \geq 0$, $\sigma > 1$, and the equilibrium features a positive modern share ($\vartheta_M > 0$, $\tilde{G}(s_w) > 0$). Then:*

1. *The modernization gain $\frac{\partial \log y}{\partial \log \vartheta_M}$ remains strictly positive if $\tau > 0$ or $\sigma < \infty$, and the local density $\tilde{g}(s_w)$ is sufficiently large.*
2. *Given $(\vartheta_M, s_w(\vartheta_M))$, the slope $T^{G'}(\vartheta_M)$ increases in the density of relative ability s around s_w, $\tilde{g}(s_w)$. More disperse worker ability, i.e. lower $\tilde{g}(s_w)$, therefore acts as a dampening force in the model. When $\tilde{g}(s_w)$ is small, $T^{G'}(\vartheta_M)$ may be negative, so that $\mu_M < 1$.*

Proof. See Appendix B.6. ∎

Proposition 3 establishes that while Lewisian gains from modernization are preserved for moderate levels of worker heterogeneity, the degree of heterogeneity – i.e. the disper-

[10] An own-account worker produces $z_T sh$ units of traditional service output, which she sells at price P_T.

[11] The wedge τ can be micro-founded, for example as an efficiency wage, where firms must pay workers a minimum premium over own-account work to prevent shirking.

sion of worker ability – dampens the comparative statics of the model. Graphically, worker heterogeneity reduces the slope of $T^G(\vartheta_M)$, implying that the multiplier $\frac{1}{1-T^{G'}(\vartheta_M)}$ declines in the dispersion of ability. In economic terms, worker heterogeneity makes sectoral labor supply less elastic: as wage employment expands, the marginal wage worker must be more generously compensated to draw her out of traditional services. The modern sector then faces an upward-sloping labor supply curve rather than a flat one. The result that heterogeneous efficiency of productive units acts as a dampening force also exists in the model of Buera et al. (2023), who instead consider firm heterogeneity.

4.3.2 Trade in Goods

The quantitative section of the paper makes comparisons across regions and uses a trade-induced demand shock for the local region's tradable goods as a validation exercise. Conducting these exercises requires a theory of (1) how price levels across regions relate to one another, and (2) how demand shocks for the local tradable good translate into local purchasing power.

I therefore assume that the model laid out above applies to every microregion r of Brazil, each of which is endowed with its own technology level to produce goods and modern services $(z_{X,r}, z_{M,r})$ and its own degree of labor market frictions τ_r. Each microregion is a small open economy that produces a local tradable good $Y_{X,r}$ and trades it at price $P_{X,r}$ for a world consumption good Y_w at price P_w. Households value the world good in their consumption bundle: $Y_i = \frac{C_{i,w}^{1-\phi} C_{i,S}^{\phi}}{(1-\phi)^{1-\phi}\phi^{\phi}}$, where $C_{i,S}$ is the same service composite as above. I normalize the world price P_w to 1. By balanced trade, the value of local consumption of the world good, $Y_{w,r}$, must equal the value of local tradable production. $Y_{w,r}$ is therefore:

$$Y_{w,r} = P_{X,r} Y_{X,r} = P_{X,r} z_{X,r} \left(\int_0^{N_X} \left(h_X^P(n_X) \right)^{\frac{\sigma-1}{\sigma}} dn_X \right)^{\frac{\sigma}{\sigma-1}}, \tag{24}$$

so that the effective goods productivity term $P_{X,r} z_{X,r} = \tilde{z}_{X,r}$ plays the same role that z_X plays in the closed economy model. A demand shock for the local tradable good – that is, a shock to $P_{X,r}$ – is therefore isomorphic to a shock to the tradable technology $z_{X,r}$.

5 Estimation and Validation

The previous section explores the theoretical Lewisian gains from service modernization. This section estimates the real-world drivers of service modernization. As Section 4.2 shows, modern services could have expanded due to general income growth from outside the service sector ($\tilde{z}_X$), population growth (L), an abatement of labor market frictions (τ), or technological progress in modern services (z_M). I pair data on the Brazilian service transformation with the fully parameterized model to disentangle these explanations for every microregion in Brazil from 2000-2010. In the following section, I use the results to show that the modern service technology z_M was the main driver of service modernization within Brazil from 2000-2010. Because each microregion's economy is characterized by these four local fundamentals $(\tilde{z}_X, L, \tau, z_M)$, the results also allow me to compute model objects such as Lewisian modernization gains and the modern service multiplier μ_M for every microregion in Brazil. Section 6 shows that these model objects vary substantially within Brazil, establishing the importance of considering each microregion as its own local economy.

5.1 Data

I use three main data sources: the Brazilian Census, to obtain local measures of modern service employment, modern and traditional service earnings, and aggregate income; consumer expenditure data, to measure the strength of non-homotheticity; and worker panel data, to measure employment and earnings dynamics. The data can be downloaded from the website of Brazil's national statistical office, IBGE. To process all data in the paper, I use the Data Zoom package developed by the Economics Department at PUC-Rio.

Brazilian Census I use the long form of the Census for the years 2000 and 2010. For each individual, the long-form Census includes information on employment status, industry of employment, and take-home income, as well as geographic information on the municipality and microregion where the individual lives.

Using person weights provided by Brazil's national statistical office (IBGE), I aggregate individual information up to the microregion level to extract four statistics for each microregion-year: the share of service workers in modern employment, the cross-sectional wage premium of modern service work, aggregate income per capita, and population density.[12] These four statistics are crucial to separately identify modern service technology from labor market frictions, market size effects, and productivity growth outside services. I also use the Census for aggregate national statistics that I use in the calibration, such as the average employment share of services and the average within-microregion dispersion of modern and traditional earnings.

Consumer expenditure survey (POF) I use the 2008 vintage of Brazil's consumer expenditure survey to estimate the distribution $F_\kappa(\cdot)$ of the modern shopping fixed cost. The POF records the type of physical establishment where each household purchase took place, as the "location of purchase."[13] I group these locations into "modern" and "traditional" final consumption establishments using the classification developed by Bachas et al. (2023), and aggregate modern

[12] As in Section 2, "services" refers specifically to consumer services: retail and wholesale trade, hospitality and food service, transportation, and other personal services. Modern employment refers to wage employment with pension contributions.

[13] Location of purchase is not recorded for certain recurring payments such as rent, utilities, and loan repayments.

and traditional expenditure over the past year for each household.[14] A household's modern share of expenditure is total modern expenditure divided by the sum of modern and traditional expenditure, as seen in the histogram in Figure 4b. Besides categorized expenditure, the POF also collects information on household income, demographics, and geography.

Worker panel survey (PME) I use the PME for the years 2002-2016 as a source of panel data on workers. It is a rotating monthly panel survey of urban workers in six large Brazilian metropolitan areas that records employment status and income.[15] I use the PME to estimate the worker ability distribution $G(\cdot)$. The panel dimension of the PME is especially useful for this part of the estimation. Because the PME allows me to observe workers in different employment states, I can identify the correlation between own-account and wage employment ability.

5.2 Quantitative Approach

Because consumer services are consumed locally and therefore are non-tradable, the relevant notion of an economy is a local labor market, i.e. a microregion in the Brazilian context. I index each of the 557 microregions by $r = 1,...,557$ and consider them in two years, $t = 2000$ and $t = 2010$.

The model is parameterized by a set of structural parameters that I assume to be constant over time and space. Consumer preferences depend on three parameters: the weight ω_M on modern services, the modern/traditional elasticity of substitution ξ, and the service Cobb-Douglas weight ϕ. I take four parameters from the production side of the economy – traditional service technology z_T, fixed costs f_M, f_X, and the elasticity of substitution across varieties σ – to be time- and space-invariant.[16] Besides these consumption and production parameters, the structure of the economy depends on two distribution functions: the distribution of fixed shopping costs $F(\kappa)$, and the joint distribution of worker ability $G(h,s)$. In the next sub-section, I specify functional forms for these distributions and show how to estimate their parameters from micro-data on consumers and workers.

While these underlying structural parameters are constant across time and space, each microregion-year is characterized by four local fundamentals: market size L_{rt}, labor market frictions τ_{rt}, effective goods productivity $\tilde{z}_{X,rt}$, and modern service technology $z_{M,rt}$. I draw on the development accounting literature to measure each of these fundamentals at the local level (Hall and Jones, 1999; Caselli, 2005; Fan et al., 2023). L_{rt} can be taken directly from the data.[17]

[14] Following Bachas et al. (2023), I classify specialized shops, large stores, service institutions, and entertainment institutions as sources of modern consumption. Traditional consumption comes from non-market sources, vendors without a storefront, convenience/corner shops, individual service providers, and informal entertainers.

[15] Workers are interviewed up to eight times over a period of sixteen months. In the absence of attrition, workers are interviewed for four consecutive months, then not interviewed for eight consecutive months, then interviewed again for four more consecutive months.

[16] Note that whenever $\sigma < \infty$, fixed costs are isomorphic to the technology level z_M, z_X from the perspective of consumers, and when $\sigma \to \infty$, fixed costs are irrelevant.

[17] Because L_{rt} essentially captures how much economic activity is concentrated within a given space, I measure

I infer the remaining local fundamentals to match the three microregion-level moments I take from the Census: the share of service workers employed in the modern sector, denoted $\hat{\ell}_{M,rt}$, the cross-sectional modern wage premium $\hat{w}_{M,rt}$, and the average local wage $\overline{y}^{w}_{rt}$. I illustrate how these parameters vary with regional income in Appendix C.1.

5.3 Location-Invariant Parameters

I use two approaches to discipline the location-invariant parameters that determine the underlying structure of the economy. First, I specify parametric forms for the two distribution functions and estimate key parameters from consumer and worker micro-data. Second, I calibrate preference and production parameters to match aggregate moments from Brazil.

5.3.1 Distribution of Fixed Shopping Costs

I specify the distribution of fixed costs to shop for modern services as following a logistic function:

$$F(\kappa)=\frac{1}{1+\exp(-\varepsilon_1\log\kappa-\varepsilon_0)}. \tag{25}$$

Therefore the probability of modern shopping is a logistic function of $\log\kappa^*(y_i,\vec{P})$. Since the effect of household income y_i on κ^* is multiplicatively separable from the effects of prices, the probability of modern shopping is itself a logistic function of $\log y_i$ and an additional term capturing the effect of prices. In particular:

$$\mathbf{Pr}\Big[\text{Modern shopping}_i|y_i,\vec{P}\Big]=\frac{1}{1+\exp\Big(-\varepsilon_1\log y_i-\tilde{\varepsilon}_0(\vec{P})\Big)}. \tag{26}$$

I directly estimate ε_1 from the POF, using a logistic regression with an indicator for being a modern shopping household on the left hand side, and with per capita household income as the explanatory variable of interest.[18] I control for differences in the prices faced by households by adding fixed effects for the local geographic area in which the household was sampled. ε_1 is therefore identified by comparing richer and poorer households in the same local economy, i.e. facing the same prices. As in Fan et al. (2023), I use income as predicted by the household head's occupation, to reduce potential attenuation bias from measurement error in the explanatory variable.

The logistic regression yields an estimate of $\varepsilon_1=0.68$. This estimate implies that for households with a 50% probability of modern shopping, a 1% increase in income increases the probability of modern shopping by 0.17 percentage points.

it as population density.

[18] Because the distribution of modern expenditure shares is so polarized, most shares between 0 and 1 have very low density, so I classify a household as a modern shopping household if their modern share exceeds 50%.

5.3.2 Distribution of Worker Ability

Next, I specify two-dimensional worker ability (h,s) as following a bivariate log-normal distribution:

$$\begin{bmatrix}\log h\\ \log s\end{bmatrix}\sim\mathcal{N}\left(\begin{bmatrix}\mu_h\\ \mu_s\end{bmatrix},\begin{bmatrix}\sigma_h^2 & \rho\sigma_h\sigma_s\\ \rho\sigma_h\sigma_s & \sigma_s^2\end{bmatrix}\right). \tag{27}$$

I normalize μ_h and μ_s so that the population-level average ability is 1 in both wage employment and own-account work.[19] I then estimate σ_h, σ_s, and ρ from worker micro-data.

In particular, I assume that for a worker i in region r at time t, I observe either wage earnings $y^w_{i,rt}$ or own-account earnings $y^T_{i,rt}$, where observed earnings satisfy:

$$\begin{aligned}\log y^w_{i,rt} &= \log h_i + \log w_{rt};\\ \log y^T_{i,rt} &= \log h_i + \log s_i + \log(P_{T,rt} z_{T,rt}) + \log u^T_{i,rt}.\end{aligned} \tag{28}$$

w_{rt} and $P_{T,rt} z_{T,rt}$ are skill prices common to a region, time, and employment type. $u^T_{i,rt}$ is potential measurement error in traditional earnings, capturing the fact that business income in the informal sector is likely not known with precision. I assume that $\log u^T_{i,rt}$ is uncorrelated with other variables.

Only one of $y^w_{i,rt}$ or $y^T_{i,rt}$ is observed in a given period, with selection on s_i, which makes estimation more complicated than a multivariate log-normal distribution with all variables observed. I present the formal details of the estimation procedure in Section C.2 of the appendix. However, it is clear which moments are most informative about which parameters. Dispersion of wage earnings helps to identify σ_h; dispersion of traditional earnings, after adjusting for measurement error, helps to identify σ_s. The association between wage workers' earnings and their transition rate into own-account employment is the main moment used to identify ρ.

5.3.3 Preference and Production Parameters

I set ϕ, the Cobb-Douglas weight of services in consumer preferences, to directly match the service share of the aggregate Brazilian wage bill in the Census, 0.63. To set ω_M, the weight of modern services in the service composite, I normalize $P_M = P_T$ for aggregate Brazil in the POF and match the modern expenditure share among modern shopping households, yielding $\omega_M = 0.9$.[20] The location parameter of $F(\kappa)$, ε_0, is calibrated to 1.49 to match the aggregate share of modern shoppers in the POF (70%).

[19] This normalization is without loss, since average ability is not separately identified from technology levels. It implies that $\mu_h = -\frac{\sigma_h^2}{2}$, and $\mu_s = -\frac{(\sigma_s + 2\rho\sigma_h)\sigma_s}{2}$.

[20] Since ω_M is a demand shifter, its level is not separately identified from productivity without separate price and quantity data.

Pre-set parameters I externally impose values on three parameters of the model. First, I assume that the traditional service technology z_T is constant across time and space and normalize its value to 1.0.[21] Second, I take values from the literature for two of the model's elasticities. Following Feng et al. (2024)'s calibration of demand for modern and traditional consumption across countries, I set the elasticity of substitution between modern and traditional services to $\xi = 3.97$. For the elasticity of substitution across varieties within sector – which also governs love-of-variety – I take a value of $\sigma = 5.02$ from Peters (2022), who assesses short- and long-run effects of an exogenous change in market size in a model with love-of-variety. Both elasticities contribute to the model's comparative statics: higher values of ξ and lower values of σ will tend to magnify the response of service modernization to shocks. The elasticities are difficult to identify without clean variation in aggregate shocks to the local economy. However, I show below that with these benchmark values, the economy's response to aggregate shocks is in line with what has been found from natural experiments in Brazil.

Table 1: Parameter Values

Parameter	Description	Value	Target Moment	Moment Value
ε_1	Inverse scale param. of $F(\kappa)$	0.68	Logit coef: $\frac{\partial \log \frac{\mathbf{Pr}[\text{Modern shopping}_i]}{1-\mathbf{Pr}[\text{Modern shopping}_i]}}{\partial \log y_i}$	0.68
σ_h	Dispersion of $\log h$	0.67	SD($\log y_{it}^w$)	0.67
σ_s	Dispersion of $\log s$	0.68	SD($\log y_{it}^T$)	0.89
σ_u	Dispersion of $\log u$	0.43	SD($\Delta \log y_{it}^T$)	0.40
ρ	Correlation of $\log h, \log s$	0.0376	$\mathbf{Pr}$[WT trans.$\vert \log y_{it}^w$] probit coef	0.056
ϕ	Cobb-Douglas weight of services	0.63	Service share of wage bill	0.63
ω_M	Preference weight for M	0.90	Modern share ∣ Modern shopping	0.90
ε_0	Location parameter of $F(\kappa)$	1.49	Unconditional $\mathbf{Pr}$[Modern shopping]	0.70
τ_{rt}	Labor market friction	rt-specific	Modern service wage premium	
$\tilde{z}_{X,rt}$	Effective prod. in X	rt-specific	Average wage	
$z_{M,rt}$	Technology in M	rt-specific	Modern share of service emp.	
Parameter	**Description**	**Value**	**Source**	
z_T	Technology in T	1.0	Normalization	
ξ	M-T elasticity of substitution	3.97	Feng et al. (2024)	
σ	Cross-variety elasticity of sub.	5.02	Peters (2022)	

5.4 Development Accounting Procedure

With the economy's structure fully specified, there is a direct mapping between local fundamentals $(\tau_{rt}, \tilde{z}_{X,rt}, z_{M,rt})$ and local moments from the Census: the modern share of service employment $\hat{\ell}_{M,rt}$, the modern wage premium $\hat{w}_{M,rt}$, and the average local wage $\overline{y}_{rt}^w$. The development accounting method that I outline below uses a sequential procedure to infer fundamentals.

[21] The actual revenue productivity of labor in traditional services can still vary significantly, since it depends on both labor market sorting and consumer demand.

First, given the ability distribution $G(\cdot,\cdot)$, the allocation of labor and the wage premium $\hat{w}_{M,rt}$ identify the labor market friction τ_{rt} and two endogenous objects: the sorting threshold $s_{w,rt}$ and modern service share $\vartheta_{M,rt}$. In particular, $\vartheta_{M,rt}$ can first be inferred by comparing total earnings of modern and traditional service workers, after adjusting for measurement error in the traditional sector. Simultaneously, the marginal distribution $\tilde{G}(s)$ can be inverted to infer $s_{w,rt}$ from the reallocation of labor. $1+\tau_{rt}$ is the wedge that rationalizes the modern wage premium after accounting for measurement error, relative skill prices, and selection implied by $s_{w,rt}$:

$$\begin{aligned} \frac{\vartheta_{M,rt}}{1-\vartheta_{M,rt}} &= \hat{w}_{M,rt} \frac{\hat{\ell}_{M,rt}}{1-\hat{\ell}_{M,rt}} \mathbb{E}\left[u^T\right]; \\ s_{w,rt} &= \tilde{G}^{-1}(\ell_{w,rt}); \\ \underbrace{\frac{w_{rt}}{P_{T,rt} z_T}}_{\text{Rel. skill price}} = (1+\tau_{rt}) s_{w,rt} &= \underbrace{\hat{w}_{M,rt}}_{\text{Skill premium}} \underbrace{\frac{\mathbb{E}[sh|s > s_{w,rt}]}{\mathbb{E}[h|s \le s_{w,rt}]} \mathbb{E}\left[u^T\right]}_{\text{Selection / Meas. err.}}. \end{aligned} \tag{29}$$

Second, conditional on $s_{w,rt}$, the nominal wage in *reais* identifies effective goods productivity $\tilde{z}_{X,rt}$. Inference of $\tilde{z}_{X,rt}$ must take into account selection and total efficiency units $H_{X,rt}$ used for goods production, which contribute to market size effects when $\sigma < \infty$:

$$\frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}} \frac{\tilde{z}_{X,rt}}{f_X^{\frac{1}{\sigma-1}}} = \underbrace{\frac{\overline{y}^w_{rt}}{\mathbb{E}[h|s \le s_{w,rt}]}}_{\text{Selection}} \underbrace{H_{X,rt}^{\frac{-1}{\sigma-1}}}_{\text{Market size}} = \frac{\overline{y}^w_{rt}}{\mathbb{E}[h|s \le s_{w,rt}]} \left[\frac{1-\phi}{1-\phi(1-\vartheta_{M,rt})} \ell_{w,rt} \mathbb{E}[h|s \le s_{w,rt}] L_{rt} \right]^{\frac{-1}{\sigma-1}}. \tag{30}$$

To aid with interpretability, I express nominal earnings in terms of the average wage for all of Brazil in 2000, i.e. $\overline{y}^w_{BR,2000} = 1$. Since nominal earnings in my framework represent purchasing power over goods, I deflate year-2010 earnings by the change in the price of Brazilian goods from 2000 to 2010 according to the Economic Transformation Database (Kruse et al., 2022).[22] I also normalize $f_X = 1$ since it is isomorphic to $\tilde{z}_X$ when $\sigma < \infty$.

Finally, given the distribution of fixed shopping costs $F(\cdot)$ and consumer demand parameters, the modern service share $\vartheta_{M,rt}$ identifies the modern service technology $z_{M,rt}$. In particular, conditional on nominal earnings and $P_{T,rt}$ – both a function of sorting $s_{w,rt}$ and the labor market

[22] I compute this as the 2010/2000 ratio of total nominal value-added divided by total real value-added in agriculture, mining, and manufacturing. Per the Economic Transformation Database, the price of these tradable goods increased by a factor of 2.43 from 2000 to 2010.

friction τ_{rt} – there is a monotonic negative mapping from $P_{M,rt}$ to $\vartheta_{M,rt}$ that can be inverted.

$$\vartheta_{M,rt} = \vartheta_M(P_{M,rt}|e_{rt}, P_{T,rt}) = F_\kappa\left(e_{rt}\frac{1-\left(\frac{\omega_M}{1-\omega_M}\left(\frac{P_{M,rt}}{P_{T,rt}}\right)^{1-\xi}+1\right)^{\frac{\phi}{1-\xi}}}{\left(\omega_M P_{M,rt}^{1-\xi}+(1-\omega_M)P_{T,rt}^{1-\xi}\right)^{\frac{\phi}{1-\xi}}}\right)\frac{\omega_M P_{M,rt}^{1-\xi}}{\omega_M P_{M,rt}^{1-\xi}+(1-\omega_M)P_{T,rt}^{1-\xi}}$$
$$\implies P_{M,rt} = \vartheta_M^{-1}(\vartheta_{M,rt}|e_{rt}, P_{T,rt}). \tag{31}$$

The modern service technology is then immediately implied by $P_{M,rt}$:

$$\frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\frac{z_{M,rt}}{f_M^{\frac{1}{\sigma-1}}} = \frac{w_{rt}}{P_{M,rt}}H_{M,rt}^{\frac{-1}{\sigma-1}} = \frac{w_{rt}}{P_{M,rt}}\left[\frac{\phi\vartheta_{M,rt}}{1-\phi(1-\vartheta_{M,rt})}\ell_{w,rt}\mathbb{E}[h|s\leq s_{w,rt}]L_{rt}\right]^{\frac{-1}{\sigma-1}}, \tag{32}$$

where, similar to f_X, I normalize $f_M = 1$. In the results section, I plot effective modern service productivity $\tilde{z}_{M,rt} = \frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\left(\frac{L_{rt}}{f_M}\right)^{\frac{1}{\sigma-1}} z_{M,rt}$ to capture all variation in modern service productivity due to exogenous forces.

Since $\tilde{z}_{M,rt}$ is the last fundamental to be identified, and it principally explains variation in service modernization $\vartheta_{M,rt}$, it functions as the structural residual of the accounting procedure. I show below that one virtue of my theory – which incorporates realistic features of demand for services – is that it reduces the variation in this structural residual that is required to explain service modernization across space and time.

5.5 Results: Local Productivity and Labor Market Frictions

In Figure 7, I plot binned scatters of two key estimated fundamentals – modern service productivity $\tilde{z}_M$ and labor market frictions τ – against the log of microregion-level income per capita in 2000. Panel 7a plots the log of effective labor market productivity $\tilde{z}_M$ on the same scale as log aggregate income. The covariance of $\log\tilde{z}_M$ with aggregate income is remarkably weak: the gap in modern productivity between the richest and poorest regions of Brazil is a small fraction of the gap in income. Despite the enormous differences in modern service employment between rich cities like Sao Paulo and impoverished rural areas, modern firms are not massively more productive in rich regions. On the other hand, $\log\tilde{z}_M$ does have a clear time-series pattern, increasing by about 0.24 log points from 2000 to 2010 in a balanced way across Brazil. These results suggest that there is a strong pan-Brazil component to the modern service technology z_M that may stem from economy-wide changes like advances in communications.

Panel 7b shows the estimated labor market frictions τ across microregions in 2000 and 2010. Frictions do follow a significant gradient with aggregate income: they are substantially more

Figure 7: Estimated Primitives by Microregion Income

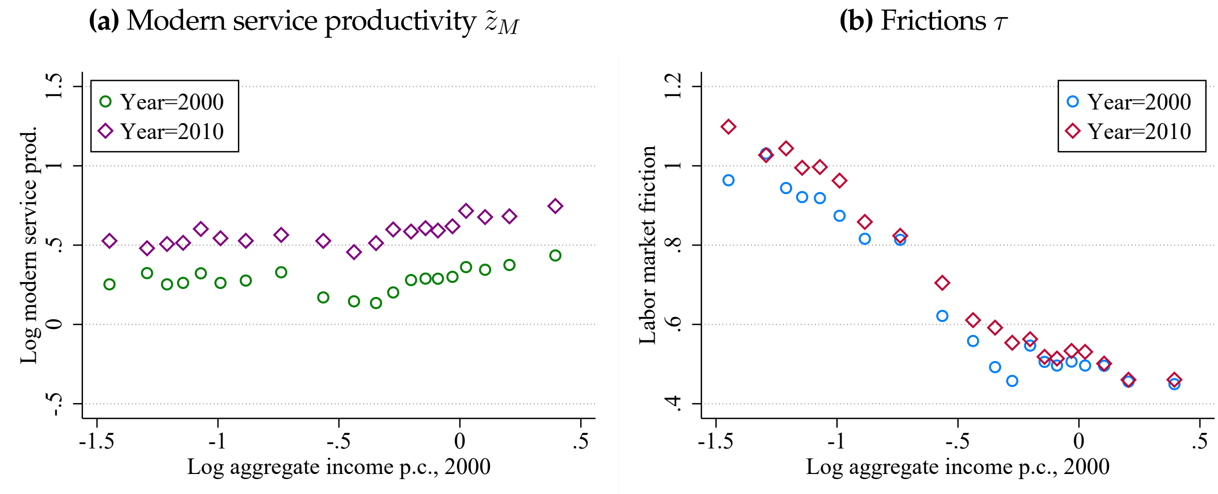

Note: Estimation results for $\log\tilde{z}_M$ and τ, using the method described in 5.4.

severe in poor microregions, which helps explain why poor parts of Brazil have low modern employment despite a high modern wage premium. Unlike $\tilde{z}_M$, labor market frictions show no signs of substantial improvement from 2000 to 2010 – in some parts of Brazil, they actually got more severe. These estimates follow naturally from the fact that modern service employment expanded from 2000 to 2010 while the modern wage premium rose rather than fell.

5.6 Untargeted Moments

By design, the estimation strategy exactly matches the target moments. Now I test it against untargeted moments. At the micro-level, the model hinges on non-homothetic consumer behavior and worker sorting between traditional and modern work. I document that the model performs very well at capturing moments relevant to these key consumer and worker decisions.

Modern shopping behavior The model uses labor market data to infer the modern expenditure share $\vartheta_{M,rt}$ and the share of modern-shopping households $F(\kappa^*_{rt})$ in every microregion in each year. The consumer expenditure data does not report microregions, but I can use it to aggregate consumer behavior to the state level, then aggregate model-generated data to the state level for comparison.

In Figure 8, I compare model and data estimates of the state-level Engel curve for the modern expenditure share and the extensive margin of modern shopping. Panel 8a plots these Engel curves for overall modern expenditure. There is a moderate level difference where the model systematically underestimates the modern expenditure share as measured in the data. This would happen if some establishments that are classified as modern employ "traditional" workers. For example, a bakery run by a self-employed entrepreneur might show up as modern expenditure because it is a specialized store, but in labor market data I would classify the baker as a traditional service worker because she is self-employed. Despite the level difference, the model-implied slope of modern expenditure with respect to income is almost parallel with the state-level Engel curve in the data.

In Panel 8b, I conduct the same exercise for the extensive margin of modern shopping, represented in the model by $F(\kappa^*)$. In this instance, the levels align well, and the slope with respect to income is slightly flatter in the model than in the data. Overall, Figure 8 shows that the model does well at capturing the systematic co-variation of consumer behavior and aggregate income.

Figure 8: State-Level Engel Curves: Model vs. Data (POF)

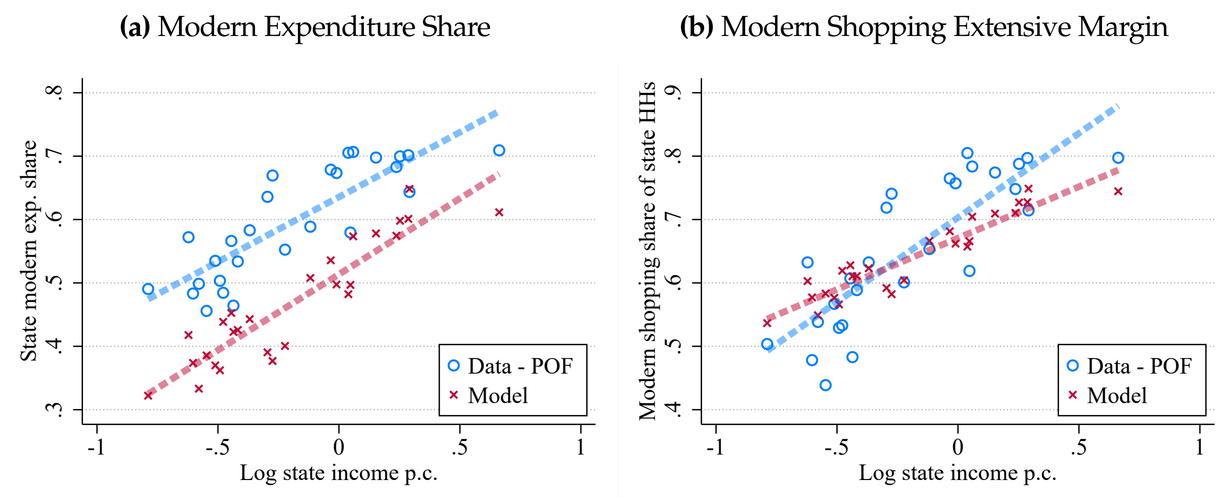

Note: ϑ_M and $F(\kappa^*)$ in the model are estimated at the microregion level and then aggregated up to the state level. In the data, I use the 2008 POF with household weights to compute the modern share of all household expenditures within a state. The plots present raw state-level scatters and an unweighted line of best fit.

Own-account work and income In my model, individuals could be traditional own-account workers for two reasons. Either they are gifted entrepreneurs (they have high s), or labor market frictions τ push them into own-account work for subsistence income. These two types of own-account workers are visible in the data on individual earnings and own-account employment status. In particular, in Figure 9, I examine own-account employment over the individual income distribution of service workers. I place these workers into bins based on their income, after residualizing on microregion fixed effects. Then within each income bin, I plot the probability that a worker in that bin is employed in own-account work, also after residualizing on microregion fixed effects.

Consistent with the model, Figure 9 shows that the relationship between own-account employment and income – plotted by blue dots – has a pronounced U-shape. Among the lowest-earning service workers, over 80% are in own-account work – these are subsistence workers. In the middle of the income distribution this number drops below 40%. But in the right tail among the richest workers, almost 60% are own-account workers – the gifted entrepreneurs.

In red X-marks, I show the same relationship within my model, using the model-generated income and employment distribution for a representative microregion that has the same modern employment share and wage premium as the aggregate Brazilian economy. The model's calibration of the worker ability distribution did not use any information from the Census, instead relying on the PME (monthly worker panel data), nor did it target the probability of own-account work conditional on income. The model distribution has the same pronounced U-shape as is present in the data. It puts more mass on the left side of the U, suggesting that a higher share of own-account workers are subsistence workers than are gifted entrepreneurs.

Figure 9: Own-account employment vs. income among service workers

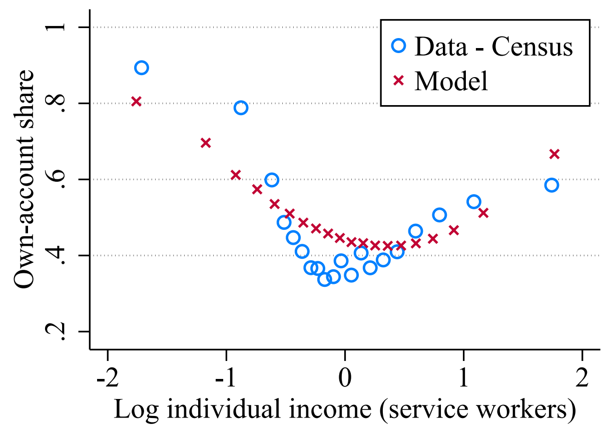

Note: Both model and data distribution have microregion fixed effects stripped out. Specifically, in the data I regress both log individual income and own-account employment on fixed effects for each microregion, then I plot the individual-level residuals against one another.

Natural experiment: trade liberalization Finally, I offer additional validation from a natural experiment. The central driver of amplification in the model is that rising incomes shift aggregate demand to modern services even when the modern service technology z_M is held constant. An ideal experiment to test this feature would be a shock to $\tilde{z}_X$: this shock holds modern technology fixed but should still affect employment and expenditure shares due to income effects.

The Brazilian context provides a natural experiment of just this type. I use Brazil's unilateral trade liberalization in the early 1990s, when the government cut tariffs on a large set of tradable goods. Microregions were differentially exposed to import competition based on their industry mix, generating quasi-random variation. A large literature has studied the local impacts of this exposure, including Kovak (2013), Dix-Carneiro and Kovak (2017), Dix-Carneiro and Kovak (2019), Ponczek and Ulyssea (2022), Dix-Carneiro et al. (2021), and Felix (2021).

Trade liberalization generated exogenous microregion-level variation in the price of locally produced tradable goods. As shown in Section 4.3.2, this variation represents a shock to $\tilde{z}_{X,rt}$. In the model, when $f_\kappa(\kappa^*) > 0$ so that demand is non-homothetic, a shock to $\tilde{z}_{X,rt}$ acts as a demand shock for modern services. Exogenous increases in $\tilde{z}_{X,rt}$ should therefore lead to a higher modern share of service employment. They should also increase the modern wage premium within services, to draw more workers into wage employment in the modern sector.

To test this prediction, I measure local exposure to import competition using the "regional tariff reductions" computed by Dix-Carneiro and Kovak (2017).[23] I use tariff reductions as an instrument for an exogenous income shock originating outside the service sector. Using two-stage least-squares, I estimate the impact of this shock on microregion-level modern employment and the modern wage premium among service workers in 2000, the first Census year following the shock. In all empirical regressions, I control for local aggregate income, the modern service share, and the modern wage premium in the Census year prior to the shock (1991).

I present results from these regressions in Table 2. The results confirm one of the model's central predictions: even though the income shock came from outside the service sector, it was a demand shock for modern services. A 10% trade-induced increase in income induces 0.5% of service workers to switch to modern services, and increases the modern wage premium in services by 2.1%. The fact that the modern wage premium moves in the same direction as modern service employment supports the interpretation of the trade shock as a *demand shock* to modern services rather than a *supply shock* to informal labor. If instead the results were driven by the tradable sector shedding labor into traditional services in badly-shocked regions, then those regions would see the modern service share and modern wage premium move in opposite directions. The flood of new traditional entrepreneurs would decrease the modern share of services, but because of these entrepreneurs' subsistence-level earnings, it would push up the average modern wage premium.

Table 2: Impact of Trade-Driven Income Shock on Local Service Sector in 2000

	Modern share of service emp.	Modern wage premium
Log income shock	0.051^{***}	0.211^{**}
	[0.013,0.090]	[0.028,0.393]
Observations	409	409
State FE	Yes	Yes
Lag Controls	Yes	Yes
Log Income IV	Tariff shock	Tariff shock
1st Stage F-Stat	85.1	85.1

Note: 95% CI from robust standard errors in brackets. Outcomes are in 2000, the first Census year post-trade liberalization. All regressions control for 1991 modern employment share, log income, and modern wage premium. Microregions are weighted by 1991 population. I present the first stage regression in Appendix C.3.

The results in Table 2 are not only qualitatively in line with the theory; the model generates quantitatively similar labor reallocation. To show this, I introduce a small shock to $\tilde{z}_X$ for every

[23] Dix-Carneiro and Kovak (2017) use a differently-harmonized microregion definition that leaves them with 409 microregions, rather than the 557 I use for my main results.

microregion in the model and compare the resulting change in the modern service share against the change in nominal income, producing the model equivalent of the IV coefficient above. In Figure 10, I compare the empirical point estimate and 95% CI (in blue) against the model's average effect and 95% range of effects (in red).

Figure 10: Impact of Aggregate Income on Modern Service Share: Model vs. Data

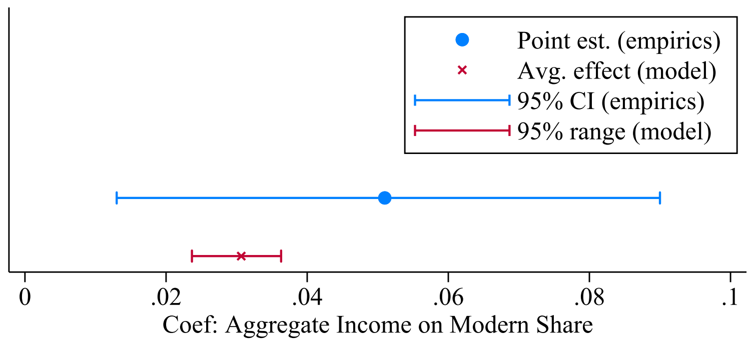

Note: "Empirics" displays point estimate and 95% CI from Table 2. The equivalent effect in the model is $\frac{d\log \ell_M}{d\log \tilde{z}_X} \Big/ \frac{d\log e_G}{d\log \tilde{z}_X}$. For the model results, the figure presents the average and the range between the 2.5% and 97.5% percentile of the distribution of microregion-level effects.

On average, the model predicts that a 10% trade-induced income shock leads 0.31% of service labor to reallocate from traditional to modern services. The model is slightly conservative relative to the empirical results, but well within the 95% confidence intervals of the empirical estimates. It leaves open the possibility for other mechanisms, such as production linkages, to partially explain the impact of trade on the service sector (Dix-Carneiro et al., 2021). In Appendix C.4, I conduct a similar exercise where I compare the model's response to a population shock against the results in Imbert and Ulyssea (2023). Again, the model's predictions are qualitatively consistent and quantitatively conservative relative to empirical results. I conclude that the model has predictive power, and that parameters governing its response to aggregate shocks fall within a sensible range.

6 Quantitative Exploration

Finally, I use the estimated and validated model to quantify the drivers of service modernization and service-led growth within Brazil from 2000 to 2010. I show that modern service technology was the main catalyst that set off service modernization in Brazil, and that a significant portion of the service-led growth it generated was Lewisian gains from reallocation to the more productive modern sector. Then I show that the realistic demand effects introduced in the model – income effects and love-of-variety – significantly amplify comparative statics in many regions of Brazil. Regions with the greatest amplification, defined by the modern service multiplier μ_M, are the places with the largest Lewisian gains and the greatest output losses to labor market frictions.

6.1 Service-Led Growth: Catalyst and Mechanisms

To measure the contribution of the modern service technology z_M to the transformation of the Brazilian economy, I simulate a counterfactual Brazil in which $z_{M,r}$ is held fixed at its 2000 level for every microregion, while the other three time-varying fundamentals $(\tau_r, z_{X,r}, L_r)$ are allowed to evolve according to their measured values in 2000 and 2010. I then compare the transformation of the service sector in this counterfactual Brazil against the actual changes in the Brazilian economy from 2000-2010. I focus on the contribution of z_M to two service-sector outcomes: first, the degree of service modernization $\Delta\vartheta_{M,r}$, and second, growth of service-sector productivity $\Delta \log A_{S,r}$. Taking into account worker heterogeneity, I define service productivity $A_{S,r}$ as real service output per efficiency unit of labor in the service sector.[24]

I compute these changes across all Brazil (weighting each microregion by population) and present the results in Figure 11a. The full change in outcomes from changing all primitives is represented by the blue bar on the left, while the contribution of z_M is represented by the red bar on the right. Modern service technology z_M explains the vast majority of development in the Brazilian service sector: 10.2 out of 11.6 percentage points of service modernization, and 21.6% out of the 23.8% 10-year growth rate of service productivity.

While it is natural that improvements in the modern service technology would increase service-sector productivity, Section 4 showed that two distinct mechanisms underpin service-led growth. These are direct effects – growth that would occur while holding the modernization margin ϑ_M fixed – and Lewisian modernization effects – growth that is driven by the increase in ϑ_M, i.e. the reallocation of economic activity to the more productive modern sector. As suggested by Equation (16), I compute Lewisian growth for each microregion by calculating the modernization gain $\frac{\partial \log A_{S,r}}{\partial \log \vartheta_{M,r}}$ and multiplying it by actual service modernization $\Delta \log \vartheta^{2010}_{M,r,2000}$. I compare Lewisian growth against the population-weighted average change in $\log A_{S,r}$ in Figure 11b. The Lewisian modernization channel was responsible for just over one-third of 2000-2010 service-sector productivity growth. Service modernization was not merely a response to rising incomes: the redirection of resources toward modern services substantially contributed to economic growth.

6.2 Amplification and its Implications

As discussed in Section 4, besides implying Lewisian gains from service modernization, income effects and love-of-variety on the demand side of the economy generate amplified comparative

[24]This implies that $A_{S,r} = \frac{E_{S,r}/P_{S,r}}{H_{M,r}+H_{T,r}}$, where $E_{S,r}$ is local expenditure on services, $P_{S,r}$ is the local service price index, and $H_{M,r}, H_{T,r}$ are the local supply of efficiency units to modern and traditional services, respectively.

Figure 11: Drivers of Service Modernization and Growth

(a) The Contribution of z_M **(b)** The Lewisian Component of Growth

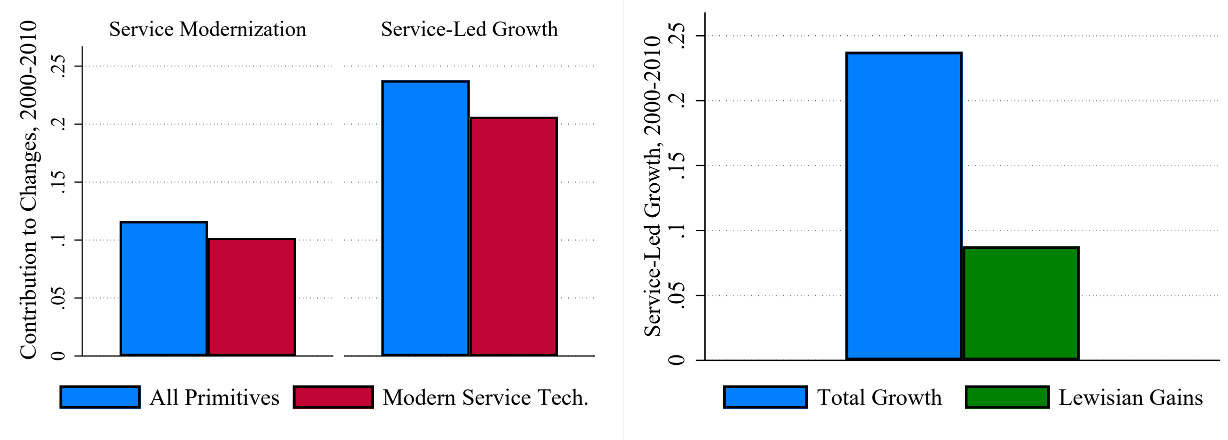

statics. Not only does reallocation of resources to the modern sector generate growth, it has the potential to crowd in more reallocation by redirecting demand towards modern services. To demonstrate the role these demand-side effects play in amplifying the model's comparative statics, I compare results from the full model against a counterfactual model that shuts down these channels of redirected demand. In the counterfactual model I set $\varepsilon_0 \to \infty$ so that $F_\kappa(\kappa) \equiv 1 \forall \kappa > 0$, shutting down income effects, and $\sigma \to \infty$, shutting down love-of-variety.

Table 3 presents summary statistics on model objects inferred from the accounting procedure under the full model and the "no demand effects" counterfactual model. The full model requires less variation in the structural residual – exogenous differences in effective modern service productivity $\tilde{z}_M$ – to explain the data. In the cross-section, demand effects reduce the inferred dispersion of $\log \tilde{z}_M$ by approximately one-third, from 0.311 to 0.208 in 2000 and from 0.236 to 0.166 in 2010. In the time series, demand effects reduce the inferred average change in $\log \tilde{z}_M$ from 2000 to 2010 by one-fifth, implying $\tilde{z}_M$ needed to grow by 24% rather than 30% to achieve the same change in service modernization. By including realistic demand effects, the model shifts responsibility for service modernization from exogenous forces to an endogenous market response.

Table 3: Variation in $\tilde{z}_M$ With and Without Demand Effects

	$\mathrm{SD}(\log \tilde{z}_{M,2000})$	$\mathrm{SD}(\log \tilde{z}_{M,2010})$	$\overline{\Delta \log \tilde{z}^{2010}_{M,2000}}$	Share $\mu_{M,2000} > 1$	Share $\mu_{M,2010} > 1$
Full Model	0.208	0.166	0.240	46%	21%
No Demand Effects	0.311	0.236	0.300	0%	0%

In other words, forces internal to the model amplify the economy's response to changes in primitives. A straightforward way to see this amplification is by examining μ_M, the modern service multiplier. As shown above, μ_M represents the degree to which service modernization ϑ_M –

the key endogenous variable that characterizes equilibrium – crowds in further service modernization. When the model features worker heterogeneity, μ_M can be either less than or greater than 1: worker heterogeneity acts as a dampening force, but income effects and love-of-variety act as amplifying forces. With these amplifying demand effects shut down, μ_M is strictly less than 1 regardless of the aggregate state of the economy. But when demand effects are present, a significant portion of Brazilian microregions are in the amplification region of the parameter space (Buera et al., 2023), with a multiplier that exceeds 1. Indeed, the full model implies that 46% of microregions in 2000 and 21% in 2010 have a modern service multiplier $\mu_M > 1$.

Heterogeneous Amplification by Initial Income The result that the multiplier μ_M exceeds 1 in some, but not all, of Brazil's microregions hints at the dramatic heterogeneity in μ_M across Brazil. Figure A.8a presents a histogram of the values of μ_M in 2000 and 2010. While the modal value of μ_M is below 1 in both years, there is a long tail of microregions with μ_M well above 1, particularly in 2000. This variation is closely linked with the significant heterogeneity in economic development across regions within Brazil. As Figure A.8b makes clear, μ_M tends to be much higher in lower-income microregions. These are the same microregions with severe labor market frictions and many households on the margin of entering the market for modern services – conditions that favor amplification. But as incomes grow and more households make the transition into modern shopping, the scope for further demand-driven modernization diminishes. For this reason, richer microregions tend to have a lower multiplier, and almost all microregions have weaker amplification in 2010 than 2000, after 10 years of income growth and modernization.

Figure 12: Heterogeneity in Modern Service Multiplier μ_M

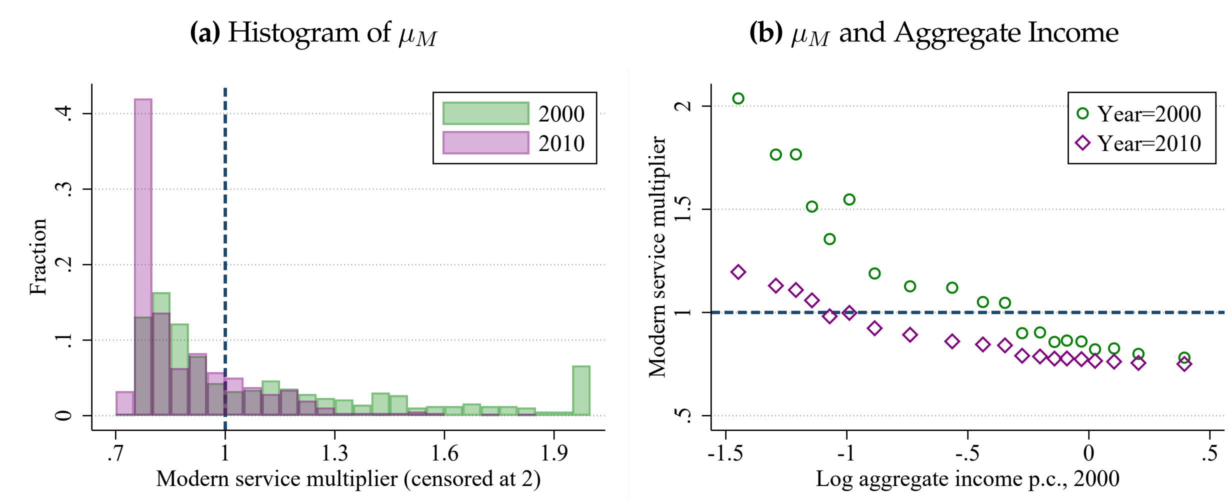

Heterogeneous Lewisian Gains Just as the scope for demand-driven amplification of service modernization tends to decline with aggregate income, so does the scope for Lewisian gains to contribute to growth. In Figure 13a, I plot Lewisian service gains from 2000-2010 – that is, $\frac{\partial \log A_{S,r}}{\partial \log \vartheta_{M,r}} \times \Delta \log \vartheta^{2010}_{M,r,2000}$ – against the log of year-2000 aggregate income using a binned scatter plot. Low-income microregions, whose economies have the greatest modern service multipliers, also saw the greatest service-sector gains through the Lewisian reallocation channel. Together with Figure A.8b, the results suggest that self-reinforcing Lewisian gains are a viable channel for service-led growth at early stages of economic development, but that these potential gains become exhausted at higher levels of income and modernization. They also help to explain income convergence among Brazil's microregions from 2000 to 2010.

The Cost of Distortions While demand effects amplify the scope for self-reinforcing Lewisian gains, they also amplify the aggregate costs of labor market frictions that limit the size of the modern service sector. To illustrate how severe the impact of frictions is, in both the full model and the "no demand effects" counterfactual, I compare real output in each microregion against its simulated level in a benchmark economy with low labor market frictions. For a realistic benchmark economy, I set labor market frictions τ to their 5th-percentile level across Brazilian microregions for all microregions where τ exceeds this level.[25] I then compute the loss of log real output to frictions as $\log(y_r/y_r^{\text{benchmark}})$ under both models, and plot this loss against microregion level income in the binned scatter of Figure 13b.

Losses to frictions roughly double for the poorest microregions under demand effects, from a 15% to a 30% loss of output. In the model with demand effects, income losses to misallocation beget further income losses. Labor market frictions force workers into subsistence traditional work, turning them into poorer consumers. As poorer consumers, their capacity to shop for modern services decreases, and demand redirects towards the traditional sector. This shrinks the effective market size of the modern sector, generating further output losses due to love-of-variety. This process reduces demand for modern wage labor, forcing more workers into the less productive traditional sector and reinforcing the feedback loop of losses. The same micro-level frictions have much greater macro-level consequences when demand for modern services is sensitive to local income.

6.3 Discussion: Additional Mechanisms

This sub-section discusses possible forces behind service modernization and amplification that are not explicitly included in the model. I discuss concrete drivers of modernization and how they would be interpreted through the model framework.

Human Capital I model the worker ability distribution as fixed over time and space. This means that increases in human capital will be interpreted through the technology and wedges that I use the model to measure. If human capital makes large firms more productive due to better management and worker efficiency (Engbom et al., 2024), then the model would interpret

[25]Note that measurement of τ relies only on inversion of the worker ability distribution and not on any specification of the demand-side of the economy, so that τ_r is measured at the same level under the full model and the "no demand effects" counterfactual.

Figure 13: The Consequences of Amplification

(a) The Magnitude of Lewisian Gains **(b)** The Amplified Real Output Loss to Frictions τ

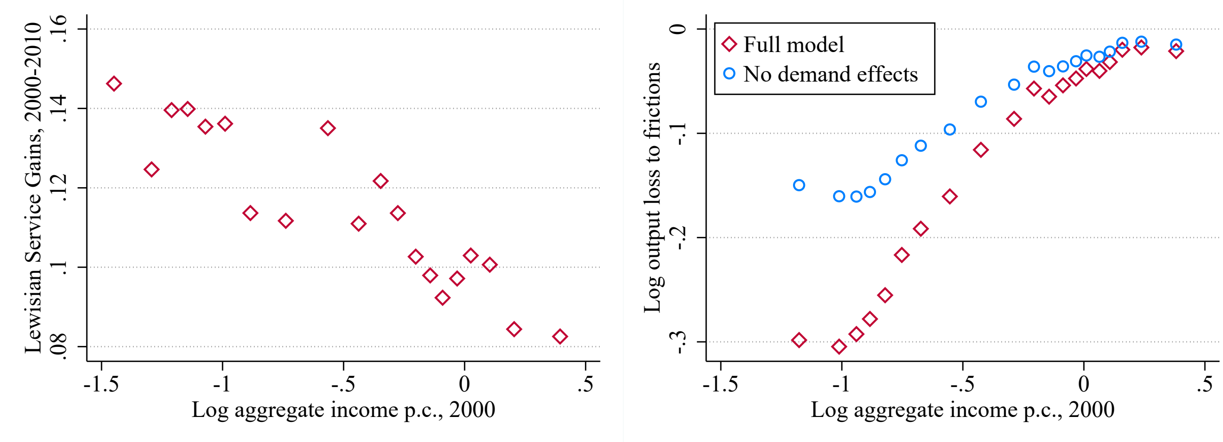

this productivity improvement through shifts of technology $\tilde{z}_{X,rt}$, $\tilde{z}_{M,rt}$. The measured rise in $\tilde{z}_M$ is consistent with improved access to higher education across Brazil, as documented by Cox (2023). If better access to human capital also brings better access to formal jobs, the model would measure this change as a reduction in barriers to wage employment, i.e. a decrease in τ_{rt}.

Migration Migration has two economically relevant effects for service modernization: it changes the local market size and may generate sorting across space. I take local market size L_{rt} as given and directly measure it. Sorting across space would imply local differences in the worker ability distribution. Imbert and Ulyssea (2023) argue that intra-Brazilian migrants do not appear to be selected on ability, although they are selected on age and gender. If migrants were selected on ability, then as in the human capital discussion, ability differences across space would mainly be measured as differences in effective technology $\tilde{z}_{X,rt}$, $\tilde{z}_{M,rt}$. However, one effect missing from the model is that endogenous migration may amplify the response to local shocks, potentially explaining why shocks have quantitatively smaller effects in my model than in natural experiments, as shown in Section 5.

Fiscal Transfers In the model, all income is earned. Transfers would allow local spending to differ from local earnings. Since the measurement procedure looks at earnings rather than spending, transfer income is an unobserved demand shifter for modern services in the model with income effects. Regions that are net recipients of transfers will shift demand towards the modern sector, pushing up the modern service share and the modern wage premium. These changes would be reflected in the model's structural residual $\tilde{z}_{M,rt}$. Gerard et al. (2021) find that fiscal transfers through the Bolsa Familia program had significant formalizing effects on recipient regions, so it is plausible that explicitly considering transfers would further reduce inferred variation in $\tilde{z}_{M,rt}$.

Infrastructure Fixed costs to access modern services limit the modern sector's market size. Infrastructure that relieves these fixed costs – such as transportation improvements, or Internet access that facilitates e-commerce – increases demand for modern services. In the logistic specification of the fixed cost distribution, this infrastructure is akin to an increase in the intercept parameter ε_0: it shifts the distribution of fixed costs to the left. While I treat ε_0 as time- and space-invariant, an increase in ε_0 is isomorphic to an increase in $\tilde{z}_M$ for the accounting procedure. The final step of the procedure infers $\tilde{z}_M$ from ϑ_M and is the only step that uses the fixed cost distribution. For every ε_0, there exists a $\tilde{z}_M$ that is consistent with ϑ_M, and $\frac{d\vartheta_M}{d\varepsilon_0}$ and $\frac{d\vartheta_M}{d\tilde{z}_M}$ are both strictly positive, so an unobserved increase in ε_0 will be inferred as an increase in $\tilde{z}_M$.

6.3.1 Alternate Sources of Amplification

Non-Homothetic Preferences The present version of the model micro-founds non-homothetic demand for modern services in fixed costs to consume them. An alternate micro-foundation is to directly assume non-homothetic preferences over the two service sectors. While this alternate specification does not capture the extensive margin of modern service consumption, it preserves most other features of the model, including the potential for amplification. In Appendix C.5, I follow Comin et al. (2021) and implement a model with non-homothetic CES (NHCES) preferences within the service sector. The NHCES specification implies that $\frac{\vartheta_M}{1-\vartheta_M} = \frac{\omega_M}{1-\omega_M}\left(y^{-\varepsilon}\frac{P_M}{P_T}\right)^{1-\xi}$. I estimate the key non-homotheticity parameter ε from consumer demand for modern services and repeat the quantitative procedure from Section 5. The quantitative results are broadly similar to those in Section 6. Like the baseline model, non-homothetic CES shrinks the residual variation in $\log \tilde{z}_M$ required to match the data, and implies a multiplier greater than 1 in a significant share of microregions.

Other Values of σ Since I use an external estimate of the within-sector elasticity of substitution σ, in Appendix C.6 I present quantitative results under other values, ranging from $\sigma = 3$ to $\sigma = 10$. As expected, lower values of σ imply stronger amplification in the model. Appendix C.6 also shows that the model requires a value of σ between 3 and 4 to match the empirical response of employment shares to trade and migration shocks. Relative to these values, the benchmark value of $\sigma = 5.02$ delivers conservative results on the degree of amplification in the economy.

7 Conclusion

One of the most salient differences between developing and mature economies is the service sector. As a country develops, modern service firms replace traditional service micro-entrepreneurs. This paper uses theory and data to argue that this is more than a superficial change. It is a central and partially self-sustaining driver of service-led growth. The theory is rooted in two basic facts about services: modern service firms are more productive on the margin and in higher demand among rich consumers. As development raises aggregate income, it also raises demand for modern services, pulling additional workers into the more productive modern sector and generating further income gains.

In Brazil, technological progress in modern service firms like supermarkets and chain restaurants catalyzed a process of service-led growth through modernization. Demand effects are strong enough that this process features amplification, as measured by the modern service multiplier, in a substantial share of microregions. Developments that improve productivity in service firms, such as expansions of communications infrastructure, are likely to have larger impacts in low-income settings due to feedback effects on demand. A significant share of service-led growth comes from the Lewisian process of reallocation to more productive jobs. Active labor market policies that help workers transition into these jobs have the potential to generate significant growth in general equilibrium. In contrast, policy stances that implicitly protect traditional service jobs from competition, such as lax enforcement of taxes and labor regulations, have a large aggregate cost.

This paper leaves open several questions that future research could address. For tractability, I use a representative household framework and analyze aggregate consumption gains. If the impact of service transformation on consumers is heterogeneous, so that poor consumers benefit less than rich consumers, then this would explain why policymakers take a permissive stance towards the traditional service sector. I abstract from firm heterogeneity and firm dynamics to focus on workers and consumers, but firm behavior is another important component of service transformation. I also omit human capital accumulation from my analysis, which may contribute to the shift to modern services. Despite these limitations, I believe that my theory of service transformation highlights powerful amplification and reallocation forces for researchers and policymakers to consider.

References

[1] AKCIGIT, U., H. ALP, AND M. PETERS (2021): "Lack of selection and limits to delegation: firm dynamics in developing countries," *American Economic Review*, 111, 231–75.

[2] ATKIN, D., B. FABER, AND M. GONZALEZ-NAVARRO (2018): "Retail globalization and household welfare: Evidence from mexico," *Journal of Political Economy*, 126, 1–73.

[3] BACHAS, P., L. GADENNE, AND A. JENSEN (2023): "Informality, consumption taxes, and redistribution," *Review of Economic Studies*.

[4] BAUMOL, W. J. (1967): "Macroeconomics of unbalanced growth: the anatomy of urban crisis," *The American economic review*, 57, 415–426.

[5] BREZA, E., S. KAUR, AND Y. SHAMDASANI (2021): "Labor rationing," *American Economic Review*, 111, 3184–3224.

[6] BRONNENBERG, B. J. AND P. B. ELLICKSON (2015): "Adolescence and the path to maturity in global retail," *Journal of Economic Perspectives*, 29, 113–134.

[7] BUERA, F. J., H. HOPENHAYN, Y. SHIN, AND N. TRACHTER (2023): "Big push in distorted economies," Tech. rep., National Bureau of Economic Research.

[8] BUERA, F. J. AND J. P. KABOSKI (2012): "The rise of the service economy," *American Economic Review*, 102, 2540–2569.

[9] CASELLI, F. (2005): "Accounting for cross-country income differences," *Handbook of economic growth*, 1, 679–741.

[10] COMIN, D., D. LASHKARI, AND M. MESTIERI (2021): "Structural change with long-run income and price effects," *Econometrica*, 89, 311–374.

[11] COOPER, R. AND A. JOHN (1988): "Coordinating coordination failures in Keynesian models," *The Quarterly Journal of Economics*, 103, 441–463.

[12] COX, A. (2023): "From Classroom to Prosperity: Fostering Development Through Higher Education," Tech. rep., Working Paper. Yale University.

[13] DIX-CARNEIRO, R., P. K. GOLDBERG, C. MEGHIR, AND G. ULYSSEA (2021): "Trade and informality in the presence of labor market frictions and regulations," Tech. rep., National Bureau of Economic Research.

[14] DIX-CARNEIRO, R. AND B. K. KOVAK (2017): "Trade liberalization and regional dynamics," *American Economic Review*, 107, 2908–2946.

[15] ——— (2019): "Margins of labor market adjustment to trade," *Journal of International Economics*, 117, 125–142.

[16] DONOVAN, K., W. J. LU, T. SCHOELLMAN, ET AL. (2023): "Labor market dynamics and development," *The Quarterly Journal of Economics*, Forthcoming.

[17] ECKERT, F., S. GANAPATI, AND C. WALSH (2022): "Urban-biased growth: a macroeconomic analysis," Tech. rep., National Bureau of Economic Research.

[18] ENGBOM, N., H. MALMBERG, T. PORZIO, F. ROSSI, AND T. SCHOELLMAN (2024): "Economic Development According to Chandler," Tech. rep., mimeo, New York University.

[19] FAN, T., M. PETERS, AND F. ZILIBOTTI (2023): "Growing Like India—the Unequal Effects of Service-Led Growth," *Econometrica*, 91, 1457–1494.

[20] FELIX, M. (2021): "Trade, labor market concentration, and wages," *Job Market Paper*.

[21] FENG, Y., D. LAGAKOS, AND J. E. RAUCH (2024): "Unemployment and development," *The Economic Journal*, 134, 614–647.

[22] GERARD, F., J. NARITOMI, AND J. SILVA (2021): "Cash transfers and formal labor markets: Evidence from Brazil," .

[23] GOLLIN, D. (2008): "Nobody's business but my own: Self-employment and small enterprise in economic development," *Journal of Monetary Economics*, 55, 219–233.

[24] HALL, R. E. AND C. I. JONES (1999): "Why do some countries produce so much more output per worker than others?" *The quarterly journal of economics*, 114, 83–116.

[25] HERREÑO, J. AND S. OCAMPO (2023): "The macroeconomic consequences of subsistence self-employment," *Journal of Monetary Economics*, 136, 91–106.

[26] HJORT, J., H. MALMBERG, AND T. SCHOELLMAN (2022): "The missing middle managers: Labor costs, firm structure, and development," Tech. rep., National Bureau of Economic Research.

[27] HSIEH, C.-T. AND E. ROSSI-HANSBERG (2023): "The industrial revolution in services," *Journal of Political Economy Macroeconomics*, 1, 3–42.

[28] IMBERT, C. AND G. ULYSSEA (2023): "Rural migrants and urban informality: Evidence from brazil," Tech. rep.

[29] KAPLAN, G. AND G. MENZIO (2016): "Shopping externalities and self-fulfilling unemployment fluctuations," *Journal of Political Economy*, 124, 771–825.

[30] KOVAK, B. K. (2013): "Regional effects of trade reform: What is the correct measure of liberalization?" *American Economic Review*, 103, 1960–1976.

[31] KRUSE, H., E. MENSAH, K. SEN, AND G. DE VRIES (2022): "A manufacturing (re) naissance? Industrialization in the developing world," *IMF Economic Review*, 71, 439.

[32] LAGAKOS, D. (2016): "Explaining cross-country productivity differences in retail trade," *Journal of Political Economy*, 124, 579–620.

[33] LEWIS, W. A. (1954): "Economic development with unlimited supplies of labour," .

[34] MATSUYAMA, K. (1991): "Increasing returns, industrialization, and indeterminacy of equilibrium," *The Quarterly Journal of Economics*, 106, 617–650.

[35] MEGHIR, C., R. NARITA, AND J.-M. ROBIN (2015): "Wages and informality in developing countries," *American Economic Review*, 105, 1509–46.

[36] MURPHY, K. M., A. SHLEIFER, AND R. W. VISHNY (1989): "Industrialization and the big push," *Journal of political economy*, 97, 1003–1026.

[37] NARITA, R. (2020): "Self-employment in developing countries: A search-equilibrium approach," *Review of Economic Dynamics*, 35, 1–34.

[38] NAYYAR, G., M. HALLWARD-DRIEMEIER, AND E. DAVIES (2021): *At your service?: the promise of services-led development*, World Bank Publications.

[39] NGAI, L. R. AND B. PETRONGOLO (2017): "Gender gaps and the rise of the service economy," *American Economic Journal: Macroeconomics*, 9, 1–44.

[40] PETERS, M. (2022): "Market size and spatial growth—evidence from Germany's post-war population expulsions," *Econometrica*, 90, 2357–2396.

[41] PONCZEK, V. AND G. ULYSSEA (2022): "Enforcement of labour regulation and the labour market effects of trade: Evidence from Brazil," *The Economic Journal*, 132, 361–390.

[42] POSCHKE, M. (2013): "'Entrepreneurs out of necessity': A snapshot," *Applied Economics Letters*, 20, 658–663.

[43] ——— (2023): "Wage employment, unemployment and self-employment across countries," Tech. rep., Working Paper. McGill University.

[44] RAMOS-MENCHELLI, D. AND D. SVERDLIN-LISKER (2022): "Fragmented Markets and the Proliferation of Small Firms: Evidence from Mom-and-Pop Shops in Mexico," Tech. rep., MIT.

[45] RODRIK, D. AND R. SANDHU (2024): "Servicing Development: Productive Upgrading of Labor-Absorbing Services in Developing Economies," Tech. rep., National Bureau of Economic Research.

[46] ROSENSTEIN-RODAN, P. N. (1943): "Problems of industrialisation of eastern and south-eastern Europe," *The economic journal*, 53, 202–211.

[47] SCHOAR, A. (2010): "The divide between subsistence and transformational entrepreneurship," *Innovation policy and the economy*, 10, 57–81.

[48] SONG, Z., K. STORESLETTEN, AND F. ZILIBOTTI (2011): "Growing like china," *American economic review*, 101, 196–233.

[49] ULYSSEA, G. (2010): "Regulation of entry, labor market institutions and the informal sector," *Journal of Development Economics*, 91, 87–99.

[50] WALKER, M. W., N. SHAH, E. MIGUEL, D. EGGER, F. S. SOLIMAN, AND T. GRAFF (2024): "Slack and Economic Development," Tech. rep., National Bureau of Economic Research.

A Empirical Appendix

A.1 Panel variation: services and development

Figure A.1 is the first-difference version of Figure 1a, comparing first differences in log real GDP per capita and the modern share of service employment across country-years. First differences of both variables are residualized against initial log real GDP per capita, initial modern share of service employment, and the time gap between observations in years, since not all countries run their surveys with the same regularity.

Figure A.1: Service modernization: first differences

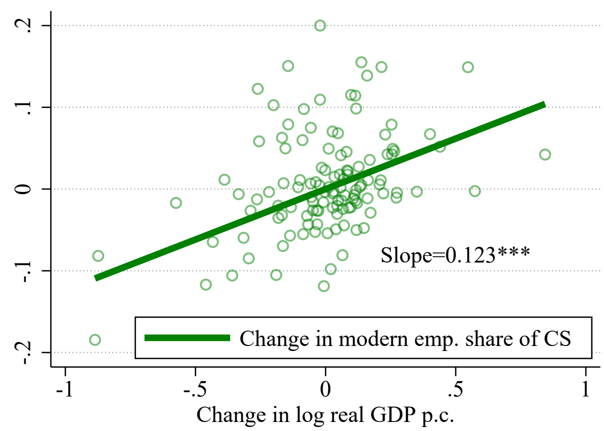

Figure A.2 shows that the overall share of consumer services as a percentage of non-agricultural work is relatively flat with respect to development, and in the poorest countries it is just above 40%. The fact that it represents two-thirds of own-account work in poor countries therefore shows that it is over-represented among own-account work.

A.2 Estimating the wage premium

To produce Figure 3, I run regressions of the following form among service workers:

$$\log(\text{Earnings}_{ict}) = \beta_0 + \beta_1 \text{Modern Employment}_{ict} + \alpha_{ct} + \vec{\gamma} X_{ict} + u_{ict}, \tag{33}$$

Figure A.2: Consumer services as a share of non-agricultural work and development

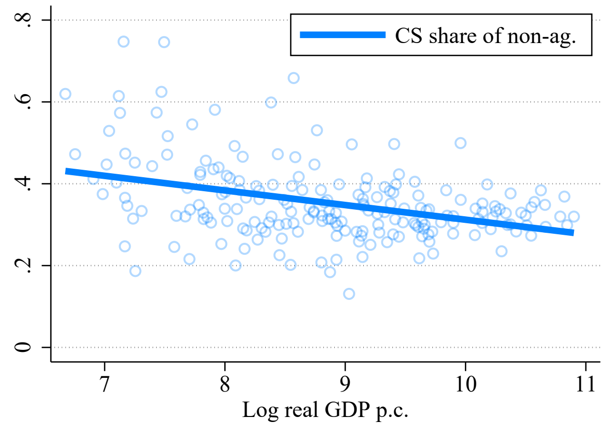

where modern employment is an indicator for wage employment with pension contributions, and α_{ct} is a city-time fixed effect. In all regressions, I use person weights provided by IBGE.

Figure 3 presents different estimates of β_1 for different sub-samples and specifications of controls X_{ict}. The red bars correspond to estimates in the Northeastern cities of Salvador and Recife; the blue bars correspond to estimates in the Southeastern and Southern cities of Sao Paulo, Rio de Janeiro, Belo Horizonte, and Porto Alegre.

The first pair of bars is the "raw" version with only city-time fixed effects; the second pair adds controls for observable characteristics: the worker's sex, race, level of education, age (quadratic), and the industry of the worker's job in that period. The third pair of bars controls for worker individual fixed effects as well as city-time fixed effects and the industry of work. It therefore identifies the effect of modern employment on earnings within worker, among traditional-modern transitioners, after stripping away industry effects. The capped spikes on the bars correspond to robust standard errors of the estimates.

A.3 Estimating the Engel curve

To estimate the Engel curve in Figure 4a, I first construct the modern expenditure share of the household. I use the classification of Bachas et al. (3) and aggregate all consumption expenditures over the previous year. In this classification, the traditional sector consists of non-market consumption sources, vendors withouat a storefront, convenience/corner shops, individual service providers, and informal entertainers. The modern sector consists of specialized shops, large stores, service institutions, and entertainment institutions. The POF has two components: an "expenditure booklet" where households record all expenditures for a week, and a retrospective interview over larger purchases over the previous year. I combine both of these in computing the modern expenditure share, making sure to use the "annualization factor" provided by IBGE: multiplying weekly purchases by 52, monthly purchases by 12, and so on.

The binned scatter of the Engel curve presents the relationship between the modern expenditure share and log income after residualizing both on a battery of controls. These controls are: the number of household residents, the number of earners, the household's age in ten-year dummies, the head's race, the head's gender, the head's education level, and the primary sampling unit (PSU) where the household lives. I provide more detail on PSUs in the next sub-section of this Appendix. In constructing the binned scatter, I use the household weights provided by IBGE.

To compute the income gradient that I match in the estimated model, I use a similar procedure but use an instrumental variable for household income. Specifically, I use the set of 3-digit codes for household head's occupation to predict income. I do this to eliminate potential attenuation bias from measurement error on income. The resulting income gradient is a statistically significant 0.14 with a robust standard error of 0.01.

A.4 Location-level modern expenditure shares

One hypothesis for why some households have zero modern expenditures is that there are simply no modern establishments in the local economy. If this were the main reason, it would imply that the modern expenditure share averaged over all local households should be close to zero for some places in Brazil. I test this implication by examining the histogram of modern expenditure shares at the level of primary sampling units (PSUs) of the POF. A PSU is analogous to a census tract (albeit larger in size), and Brazil's national statistical office (IBGE) uses them in two-stage stratified sampling to select the final sample of households in POF. There are 12,800 PSUs in Brazil, and 4,696 of them were randomly sampled for the POF. Within each PSU, 12 to 16 households were selected to be surveyed. In 2008, the average population of a PSU was 15,000, which is equivalent to about 1.5 US ZIP codes. On average, there are two to three PSUs per municipality in Brazil.

Figure A.3 presents the histogram of PSU-level average modern expenditure shares. To compute the average PSU-level share, I simply aggregate modern and total expenditure over all households within a given PSU and take their ratio. The PSU-level distribution has a predictable unimodal shape, and notably has very few PSUs with an average modern expenditure share near zero. I therefore conclude that it is unlikely that households consume zero modern services due to having no modern service establishments in the local economy.

A.5 Regressions of the Modern Service Share on Population Density

Figure A.3: Histogram: PSU-level average modern expenditure share

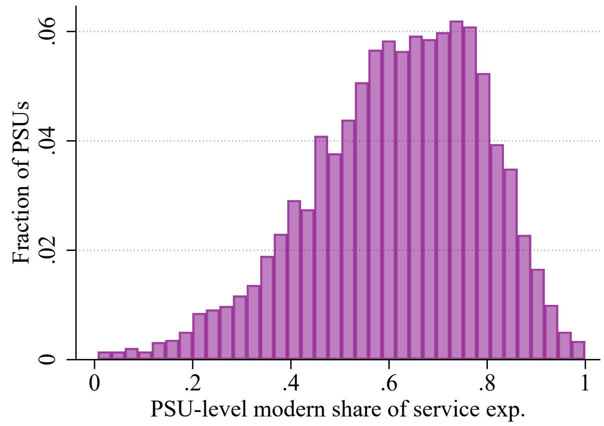

Note: POF sample is 48,706 households from 4,694 primary sampling units (PSUs). PSU-level modern expenditure share is calculated by averaging over all households in the PSU.

Table A.1: Modern Services and Log Population Density

	Modern serv. shr.	Modern serv. shr.
Year=2000 × Log pop. density	0.045***	0.017***
	[0.039,0.050]	[0.014,0.021]
Year=2010 × Log pop. density	0.043***	0.008***
	[0.037,0.049]	[0.004,0.012]
Year=2000 × Log income p.c.		0.225***
		[0.216,0.233]
Year=2010 × Log income p.c.		0.312***
		[0.300,0.323]
Year=2010	0.119***	0.091***
	[0.090,0.148]	[0.073,0.110]
Observations	1114	1114

Note: Data from 2000 and 2010 Brazilian Censuses. 95% confidence intervals from robust standard errors in brackets.

B Theoretical Appendix

B.1 Additional Equations and Market-Clearing Conditions

Markets must clear so that expenditures equal revenues in every sector. And since there are zero profits in equilibrium, revenues of each sector equal labor compensation. Therefore, from (7),

$$\begin{aligned} wH_X &= (1-\phi)eL \\ wH_M &= \phi\vartheta_M eL \\ \frac{w}{1+\tau}H_T &= \phi(1-\vartheta_M)eL \end{aligned} \tag{34}$$

Denoting each sector's employment share by $h_k = \frac{H_k}{L}$, the equations in (34) then imply:

$$\begin{aligned} e &= w(1-h_T) + \frac{w}{1+\tau}h_T \\ h_T &= \frac{(1+\tau)\phi(1-\vartheta_M)}{1+\tau\phi(1-\vartheta_M)} \\ \implies e &= \frac{w}{1+\tau\phi(1-\vartheta_M)} \end{aligned} \tag{35}$$

By (34) and (35), the economy's labor allocation $\vec{H}$ depends only on service modernization ϑ_M and exogenous parameters, yielding Equation (14).

In turn, Equation (14) and the expressions for goods prices, service prices, and nominal expenditures in (12), (13), and (35) yield the following expressions for real income and relative prices as functions of $\vec{H}(\vartheta_M)$:

$$
\begin{aligned}
y\left(\vec{H}(\vartheta_M)\right) =& \overbrace{\frac{1}{1+\tau\phi(1-\vartheta_M)}}^{e/w} \times \overbrace{\left(z_X \frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\left(\frac{H_X}{f_X}\right)^{\frac{1}{\sigma-1}}\right)^{1-\phi}}^{(P_X/w)^{\phi-1}} \\
& \underbrace{\times\left(\omega_M\left(z_M \frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\left(\frac{H_M}{f_M}\right)^{\frac{1}{\sigma-1}}\right)^{\xi-1} + \omega_T((1+\tau)z_T)^{\xi-1}\right)^{\frac{\phi}{\xi-1}}}_{(P_S/w)^{-\phi}} \\
=& \frac{1}{1+\tau\phi(1-\vartheta_M)}\left(z_X \frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\left(\frac{(1-\phi)L}{f_X(1+\tau\phi(1-\vartheta_M))}\right)^{\frac{1}{\sigma-1}}\right)^{1-\phi} \\
& \times\left(\omega_M\left(z_M \frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\left(\frac{\phi\vartheta_M L}{f_M(1+\tau\phi(1-\vartheta_M))}\right)^{\frac{1}{\sigma-1}}\right)^{\xi-1} + \omega_T((1+\tau)z_T)^{\xi-1}\right)^{\frac{\phi}{\xi-1}} \\
\frac{P_M}{P_T}\left(\vec{H}(\vartheta_M)\right) =& (1+\tau)\frac{z_T}{z_M}\frac{\sigma^{\frac{\sigma}{\sigma-1}}}{\sigma-1}\left(\frac{f_M}{H_M}\right)^{\frac{1}{\sigma-1}} = (1+\tau)\frac{z_T}{z_M}\frac{\sigma^{\frac{\sigma}{\sigma-1}}}{\sigma-1}\left(\frac{f_M(1+\tau\phi(1-\vartheta_M))}{\phi\vartheta_M L}\right)^{\frac{1}{\sigma-1}} \\
\frac{P_M}{P_S}\left(\vec{H}(\vartheta_M)\right) =& \frac{\frac{P_M}{P_T}\left(\vec{H}(\vartheta_M)\right)}{\left(\omega_M\left(\frac{P_M}{P_T}\left(\vec{H}(\vartheta_M)\right)\right)^{1-\xi} + \omega_T\right)^{\frac{1}{1-\xi}}}
\end{aligned}
\tag{36}
$$

B.2 Extended Proof of Proposition 1

Using (36), we have:

$$
\begin{aligned}
y\left(\vec{H}(\vartheta_M)\right) =& \frac{1}{1+\tau\phi(1-\vartheta_M)}\left(z_X \frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\left(\frac{(1-\phi)L}{f_X(1+\tau\phi(1-\vartheta_M))}\right)^{\frac{1}{\sigma-1}}\right)^{1-\phi} \\
& \times\left(\omega_M\left(z_M \frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\left(\frac{\phi\vartheta_M L}{f_M(1+\tau\phi(1-\vartheta_M))}\right)^{\frac{1}{\sigma-1}}\right)^{\xi-1} + \omega_T((1+\tau)z_T)^{\xi-1}\right)^{\frac{\phi}{\xi-1}}
\end{aligned}
\tag{37}
$$

The partial derivative of $\log y$ with respect to $\log\vartheta_M$ is then:

$$\begin{aligned}\frac{\partial \log y}{\partial \log\vartheta_M} =& \frac{\tau\phi\vartheta_M}{1+\tau\phi(1-\vartheta_M)}+\frac{1-\phi}{\sigma-1}\frac{\tau\phi\vartheta_M}{1+\tau\phi(1-\vartheta_M)} \\ &+\frac{\phi}{\sigma-1}\frac{1+\tau\phi}{1+\tau\phi(1-\vartheta_M)}\frac{\omega_M\left(z_M\frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\left(\frac{\phi\vartheta_M L}{f_M(1+\tau\phi(1-\vartheta_M))}\right)^{\frac{1}{\sigma-1}}\right)^{\xi-1}}{\omega_M\left(z_M\frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\left(\frac{\phi\vartheta_M L}{f_M(1+\tau\phi(1-\vartheta_M))}\right)^{\frac{1}{\sigma-1}}\right)^{\xi-1}+\omega_T((1+\tau)z_T)^{\xi-1}} \\ =& h_M(\vartheta_M)\left(\tau+\frac{\tau(1-\phi)}{\sigma-1}+\frac{1+\tau\phi}{F_\kappa(\kappa^*)(\sigma-1)}\right),\end{aligned} \tag{38}$$

which is exactly the expression that appears in the main text of Proposition 1. Note that the final line used $h_M(\vartheta_M)=\frac{H_M}{L}=\frac{\phi\vartheta_M}{1+\tau\phi(1-\vartheta_M)}$ and the fact that $\omega_M\left(\frac{P_M}{P_S}\right)^{1-\xi}=\frac{\vartheta_M}{F_\kappa(\kappa^*)}$.

Similar results for A_S The same Lewisian gains apply to modern service productivity $A_S=\frac{E_S/P_S}{H_S}$. In particular:

$$\begin{aligned}A_S=&\frac{\phi e/w}{h_M+h_T}\frac{w}{P_S} \\ =&\frac{1}{1+\tau(1-\vartheta_M)}\left(\omega_M\left(z_M\frac{\sigma-1}{\sigma^{\frac{\sigma}{\sigma-1}}}\left(\frac{\phi\vartheta_M L}{f_M(1+\tau\phi(1-\vartheta_M))}\right)^{\frac{1}{\sigma-1}}\right)^{\xi-1}+\omega_T((1+\tau)z_T)^{\xi-1}\right)^{\frac{1}{\xi-1}} \\ \implies \frac{\partial\log A_S}{\partial\log\vartheta_M}=&\vartheta_M\left(\frac{\tau}{1+\tau(1-\vartheta_M)}+\frac{1+\tau}{F_\kappa(\kappa^*)(\sigma-1)(1+\tau\phi(1-\vartheta_M))}\right)\end{aligned} \tag{39}$$

B.3 Details of Proposition 2

Proposition 2 uses the following expression for the modern service multiplier:

$$\begin{aligned}\mu_M^{-1}=&1-\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)}h_M(\vartheta_M)\left(\tau+\frac{\tau(1-\phi)}{\sigma-1}+\frac{1+\tau\phi}{F_\kappa(\kappa^*)(\sigma-1)}\right) \\ &-\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)}\phi\frac{y-\kappa^*}{\kappa^*}\frac{\vartheta_M}{F_\kappa(\kappa^*)}\frac{1}{\sigma-1}\frac{1+\tau\phi}{1+\tau\phi(1-\vartheta_M)} \\ &-(\xi-1)\left(1-\frac{\vartheta_M}{F_\kappa(\kappa^*)}\right)\frac{1}{\sigma-1}\frac{1+\tau\phi}{1+\tau\phi(1-\vartheta_M)},\end{aligned} \tag{40}$$

This follows from the fact that $T'(\vartheta_M)=\frac{\partial\log T(\vartheta_M)}{\partial\log\vartheta_M}$, and

$$\frac{\partial\log T(\vartheta_M)}{\partial\log\vartheta_M}=\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)}\left(\frac{\partial\log y}{\partial\log\vartheta_M}+\frac{\partial\log\kappa^*}{\partial\log P_M/P_T}\frac{\partial\log P_M/P_T}{\partial\log\vartheta_M}\right)+(1-\xi)\frac{\partial\log P_M/P_S}{\partial\log\vartheta_M} \tag{41}$$

The proof of Proposition 1 established the expression for $\frac{\partial\log y}{\partial\log\vartheta_M}$, which appears in the first line of (40). For the second line, it is possible to show that $\frac{\partial \log \kappa^*}{\partial \log P_M/P_T} = -\phi \frac{y-\kappa^*}{\kappa^*} \frac{\vartheta_M}{F_\kappa(\kappa^*)}$ and $\frac{\partial \log P_M/P_T}{\partial \log \vartheta_M} = -\frac{1}{\sigma-1} \frac{1+\tau\phi}{1+\tau\phi(1-\vartheta_M)}$. For the third line, it is similarly possible to show that $\frac{\partial \log P_M/P_S}{\partial \log \vartheta_M} = -\left(1 - \frac{\vartheta_M}{F_\kappa(\kappa^*)}\right) \frac{1}{\sigma-1} \frac{1+\tau\phi}{1+\tau\phi(1-\vartheta_M)}$.

B.4 Potential Multiplicity

While Proposition 2 lays out conditions under which $\mu_M^{-1} < 1$, it does not guarantee that $\mu_M > 1$. Indeed, when τ is very large or σ is close to 1, it is possible that $\mu_M^{-1} < 0$ so that the sign of μ_M flips. Since $\mu_M = \frac{1}{1-\partial \log T(\vartheta_M)/\partial \log \vartheta_M} = \frac{1}{1-T'(\vartheta_M)}$ at any equilibrium, $\mu_M < 0$ corresponds to an equilibrium in which $T'(\vartheta_M) > 1$, so that $T(\vartheta_M)$ crosses the 45-degree line from below. The next proposition establishes that this condition implies multiple equilibria with $\vartheta_M > 0$.

Proposition 4. *Suppose that $F_\kappa(\cdot)$ is non-degenerate and there exists an equilibrium $\vartheta_M^* > 0$ such that F_κ is differentiable at $\kappa^*(\vartheta_M)$ and $\mu_M(\vartheta_M^*) < 0$. Then there exists $\vartheta_M^{**} > \vartheta_M^*$ such that $\vartheta_M^{**} = T(\vartheta_M^{**})$, i.e. ϑ_M^{**} is also an equilibrium of the economy.*

Proof. Let $\tilde{T}(\vartheta_M) = T(\vartheta_M) - \vartheta_M$, so that equilibrium is defined by $\tilde{T}(\vartheta_M) = 0$. If $\mu_M(\vartheta_M^*) < 0$, then $T'(\vartheta_M^*) > 1$ and $\tilde{T}'(\vartheta_M^*) > 0$, so that there must exist a $\tilde{\vartheta}_M > \vartheta_M^*$ in the neighborhood of ϑ_M^* such that $\tilde{T}(\tilde{\vartheta}_M) > 0$. Next, note that $T(1) < 1$ and hence $\tilde{T}(1) < 0$, since $F_\kappa(\cdot) \leq 1$ and $\omega_M \left(\frac{P_M}{P_S}\right)^{1-\xi} < 1$ for finite values of P_T. Therefore, by the intermediate value theorem and continuity of $\tilde{T}(\vartheta_M)$, there exists a $\vartheta_M^{**} \in (\tilde{\vartheta}_M, 1)$ such that $\tilde{T}(\vartheta_M^{**}) = 0$, i.e. ϑ_M^{**} is an equilibrium. ■

The same conditions that generate amplification – those that increase the slope $T'(\vartheta_M)$ – have the potential to make the multiplier μ_M negative if they are strong enough. Proposition 4 establishes that under these conditions, the economy has multiple equilibria. The model therefore shares the key feature of Buera et al. (7): the conditions on primitives that contribute to amplification can also generate multiplicity.

Figure A.4 depicts an example of an economy with potential multiplicity. In the baseline parametrization, it features an equilibrium ϑ_M^* in which $T'(\vartheta_M^*) > 1$. As laid out by Proposition 4, this feature implies an additional equilibrium $\vartheta_M^{**} > \vartheta_M^*$ in which $T'(\vartheta_M^{**}) < 1$. (There is also a lower equilibrium ϑ_M in this case, although that is not guaranteed by the proposition.) However, technological progress in the form of increasing z_M allows the economy to escape multiplicity. With a large enough rise in z_M, the lower equilibria are eliminated and only the highest-modernization equilibrium remains. In this case, where technological progress can eliminate low-modernization equilibria, discrete jumps from one equilibrium to another are possible in response to changes in primitives, rather than the local comparative statics established in the previous subsection. Multiplicity therefore emerges as an extreme case of amplification.

Figure A.4: An Example of Potential Multiplicity

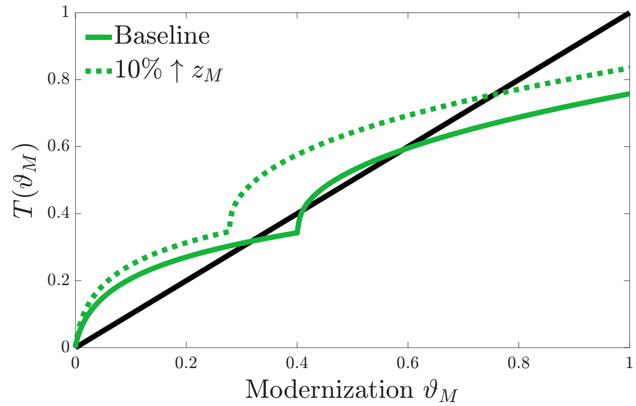

Note: As in Figure 6a, τ is set to 0.5 and σ is set to 6. In order to generate the S-shape of $T(\vartheta_M)$, F_κ is specified as a bimodal mixture of two unimodal distributions.

B.5 Market-clearing under worker heterogeneity

The market-clearing equations from (14) also require some modification when workers are heterogeneous. Most notably, the market-clearing equation for traditional labor H_T is now:

$$P_T z_T H_T = \frac{w}{(1+\tau)s_w} H_T = \phi(1-\vartheta_M)eL \tag{42}$$

The endogenous threshold s_w can therefore be expressed as a function of ϑ_M by combining the expressions for $\frac{H_T}{H}$ implied by (23) and (42). In particular, $s_w(\vartheta_M)$ is the s_w that solves:

$$\begin{aligned} \frac{\Big(1-\tilde{G}(s_w)\Big)\mathbb{E}[sh|s>s_w]}{\Big(1-\tilde{G}(s_w)\Big)\mathbb{E}[sh|s>s_w]+\tilde{G}(s_w)\mathbb{E}[h|s\leq s_w]} &= \frac{(1+\tau)s_w\phi(1-\vartheta_M)}{(1+\tau)s_w\phi(1-\vartheta_M)+1-\phi(1-\vartheta_M)} \\ \iff s_w\frac{\tilde{G}(s_w)\mathbb{E}[h|s\leq s_w]}{\Big(1-\tilde{G}(s_w)\Big)\mathbb{E}[sh|s>s_w]} &= \frac{1-\phi(1-\vartheta_M)}{(1+\tau)\phi(1-\vartheta_M)} \end{aligned} \tag{43}$$

The equations that pin down equilibrium are now similar to those in the basic model, but key endogenous variables such as labor allocations, the traditional price index, and total expenditures depend on $s_w(\vartheta_M)$. I denote the functions for these endogenous variables under worker heterogeneity with the G superscript, since they represent the model when the ability distribution function $G(h,s)$ is non-degenerate.

$$\begin{aligned}
H_X^G(\vartheta_M) &= \frac{1-\phi}{(1+\tau)s_w(\vartheta_M)\phi(1-\vartheta_M)+1-\phi(1-\vartheta_M)} H(s_w(\vartheta_M)) \\
H_M^G(\vartheta_M) &= \frac{\phi\vartheta_M}{(1+\tau)s_w(\vartheta_M)\phi(1-\vartheta_M)+1-\phi(1-\vartheta_M)} H(s_w(\vartheta_M)) \\
H_T^G(\vartheta_M) &= \frac{(1+\tau)s_w(\vartheta_M)\phi(1-\vartheta_M)}{(1+\tau)s_w(\vartheta_M)\phi(1-\vartheta_M)+1-\phi(1-\vartheta_M)} H(s_w(\vartheta_M)) \\
P_T^G(\vartheta_M) &= \frac{w}{(1+\tau)z_T s_w(\vartheta_M)} \\
e^G(\vartheta_M)L &= w\left(H_w^G(\vartheta_M) + \frac{1}{(1+\tau)s_w(\vartheta_M)} H_T^G(\vartheta_M)\right)
\end{aligned} \tag{44}$$

Other endogenous variables have the same form with respect to $\vec{H}(\vartheta_M), P_T(\vartheta_M), eL$:

$$\begin{aligned}
P_M^G(\vartheta_M) &= \frac{\sigma^{\frac{\sigma}{\sigma-1}}}{\sigma-1} \frac{w}{z_M} \left(\frac{f_M}{H_M^G(\vartheta_M)}\right)^{\frac{1}{\sigma-1}} \\
P_X^G(\vartheta_M) &= \frac{\sigma^{\frac{\sigma}{\sigma-1}}}{\sigma-1} \frac{w}{z_X} \left(\frac{f_X}{H_X^G(\vartheta_M)}\right)^{\frac{1}{\sigma-1}} \\
P_S^G(\vartheta_M) &= \left(\omega_M P_M^G(\vartheta_M)^{1-\xi} + \omega_T P_T^G(\vartheta_M)^{1-\xi}\right)^{\frac{1}{1-\xi}} \\
P^G(\vartheta_M) &= P_X^G(\vartheta_M)^{1-\phi} P_M^G(\vartheta_M)^{\phi} \\
y^G(\vartheta_M) &= \frac{e^G(\vartheta_M)}{P^G(\vartheta_M)}
\end{aligned} \tag{45}$$

Therefore the model with worker heterogeneity can be summarized by the modified transformation function $T^G(\vartheta_M)$:

$$T^G(\vartheta_M) = F_\kappa\left(\kappa^*\left(y^G(\vartheta_M), \frac{P_M^G}{P_T^G}(\vartheta_M)\right)\right) \times \omega_M \left(\frac{P_M^G}{P_S^G}(\vartheta_M)\right)^{1-\xi}, \tag{46}$$

with $\vartheta_M = T^G(\vartheta_M)$ characterizing equilibrium.

B.6 Proof of Proposition 3

Part 1: Lewisian gains Take the partial derivative of $\log y^G$ with respect to $\log\vartheta_M$:

$$\frac{\partial \log y^G}{\partial \log \vartheta_M} = \overbrace{\frac{H_w^G \frac{\partial \log H_w^G}{\partial \log s_w} + \frac{H_T^G}{(1+\tau)s_w}\left(\frac{\partial \log H_T^G}{\partial \log s_w} - 1\right)}{H_w^G + \frac{H_T^G}{(1+\tau)s_w}} \frac{\partial \log s_w}{\partial \log \vartheta_M}}^{\frac{\partial \log e^G/w}{\partial \log \vartheta_M}}$$
$$+ \underbrace{\frac{1}{\sigma-1}\left((1-\phi)\frac{d\log H_X^G}{d\log\vartheta_M} + \phi\frac{\vartheta_M}{F_\kappa(\kappa^*)}\frac{d\log H_M}{d\log\vartheta_M}\right) + \phi\left(1-\frac{\vartheta_M}{F_\kappa(\kappa^*)}\right)\frac{\partial \log s_w}{\partial \log \vartheta_M}}_{-\frac{\partial \log P^G/w}{\partial \log \vartheta_M}} \tag{47}$$

Because $H_X^G(\vartheta_M) = (1-\phi)e^G(\vartheta_M)/w$ and $H_M^G(\vartheta_M) = \phi\vartheta_M e^G(\vartheta_M)/w$, the derivative representing the gains from modernization can be re-written as:

$$\frac{\partial \log y^G}{\partial \log \vartheta_M} = \left(1 + \frac{1}{\sigma-1}\left(1-\phi+\phi\frac{\vartheta_M}{F_\kappa(\kappa^*)}\right)\right)\frac{H_w^G \frac{\partial \log H_w^G}{\partial \log s_w} + \frac{H_T^G}{(1+\tau)s_w}\left(\frac{\partial \log H_T^G}{\partial \log s_w} - 1\right)}{H_w^G + \frac{H_T^G}{(1+\tau)s_w}}\frac{\partial \log s_w}{\partial \log \vartheta_M}$$
$$+ \frac{1}{\sigma-1}\frac{\phi\vartheta_M}{F_\kappa(\kappa^*)} + \phi\left(1-\frac{\vartheta_M}{F_\kappa(\kappa^*)}\right)\frac{\partial \log s_w}{\partial \log \vartheta_M} \tag{48}$$

From Equation (43), $\frac{\partial \log s_w}{\partial \log \vartheta_M}$ is generically positive. The expression $1+\frac{1}{\sigma-1}\left(1-\phi+\phi\frac{\vartheta_M}{F_\kappa(\kappa^*)}\right)$ is also positive, as are the terms on the second line above. Therefore a sufficient condition for Lewisian gains $\left(\frac{\partial \log y^G}{\partial \log \vartheta_M} > 0\right)$ is that $H_w^G \frac{\partial \log H_w^G}{\partial \log s_w} + \frac{H_T^G}{(1+\tau)s_w}\left(\frac{\partial \log H_T^G}{\partial \log s_w} - 1\right) \geq 0$. Substituting in expressions for $\frac{\partial \log H_w^G}{\partial \log s_w}$ and $\frac{\partial \log H_T^G}{\partial \log s_w}$, this condition is equivalent to:

$$\frac{\tau}{1+\tau}s_w\tilde{g}(s_w)\mathbb{E}[h|s=s_w] - \frac{H_T^G}{(1+\tau)s_w} \geq 0, \tag{49}$$

which is clearly satisfied when the density $\tilde{g}(s_w)$ is sufficiently large.

Note also that the expression for $\frac{\partial \log s_w}{\partial \log \vartheta_M}$ is as follows:

$$\frac{\partial \log s_w}{\partial \log \vartheta_M} = \left[1 + \frac{s_w\tilde{g}(s_w)\mathbb{E}[h|s=s_w]}{H_w^G} + s_w\frac{\tilde{g}(s_w)s_w\mathbb{E}[h|s=s_w]}{H_T^G}\right]^{-1}\frac{\vartheta_M}{1-\vartheta_M}\frac{1}{1-\phi(1-\vartheta_M)}, \tag{50}$$

so that the Lewisian gains in Proposition 1 from the homogeneous labor case emerge in the limit where $s_w \equiv 1$, $\mathbb{E}[h|s=s_w] \equiv 1$, and $\tilde{g}(s_w) \to \infty$.

Part 2: $T^{G'}(\vartheta_M)$ and the multiplier μ_M Note that at any equilibrium, where $\vartheta_M = T^G(\vartheta_M)$, $T^{G'}(\vartheta_M) = \frac{\partial \log T^G(\vartheta_M)}{\partial \log \vartheta_M}$. Considering the expression for $T^G(\vartheta_M)$, this elasticity satisfies:

$$\frac{\partial \log T^G(\vartheta_M)}{\partial \log \vartheta_M} = \frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)} \frac{\partial \log y^G}{\partial \log \vartheta_M} + \left(\frac{f_\kappa(\kappa^*)\kappa^*}{F_\kappa(\kappa^*)} \frac{\partial \log \kappa^*}{\partial \log P_M/P_T} + (1-\xi)\left(1 - \frac{\vartheta_M}{F_\kappa(\kappa^*)}\right)\right) \frac{\partial \log P_M^G/P_T^G}{\partial \log \vartheta_M} \tag{51}$$

The proof of Part 1 already established that the Lewisian gains $\frac{\partial \log y^G}{\partial \log \vartheta_M}$ increase in the density $\tilde{g}(s_w)$. Furthermore, the partial derivative $\frac{\partial \log \kappa^*}{\partial \log P_M/P_T}$ is independent of $\tilde{g}(s_w)$, and the coefficient on $\frac{\partial \log P_M^G/P_T^G}{\partial \log \vartheta_M}$ is negative. Therefore, to establish the result that $\tilde{g}(s_w)$ increases the slope $T^{G'}(\vartheta_M)$, it is sufficient to show that $\frac{\partial \log P_M^G/P_T^G}{\partial \log \vartheta_M}$ decreases in $\tilde{g}(s_w)$. To show this, note from (44) and (45) that:

$$\frac{P_M^G}{P_T^G}(\vartheta_M) = \frac{\sigma^{\frac{\sigma}{\sigma-1}}}{\sigma-1} \frac{(1+\tau) z_T s_w(\vartheta_M)}{z_M} \left(\frac{f_M}{H_M^G(\vartheta_M)}\right)^{\frac{1}{\sigma-1}} \tag{52}$$

Substituting in the expression that $H_M^G(\vartheta_M) = \phi \vartheta_M e^G(\vartheta_M)/w$ and taking partial derivatives in logs yields:

$$\frac{\partial \log P_M^G/P_T^G}{\partial \log \vartheta_M} = \frac{\partial \log s_w}{\partial \log \vartheta_M} - \frac{1}{\sigma-1}\left(1 + \frac{\partial \log e^G/w}{\partial \log \vartheta_M}\right) \tag{53}$$

The proof of part 1 of the proposition already establishes that $\frac{\partial \log e^G/w}{\partial \log \vartheta_M}$ increases in $\tilde{g}(s_w)$. And from equation (50), $\frac{\partial \log s_w}{\partial \log \vartheta_M}$ also clearly decreases in $\tilde{g}(s_w)$. Therefore $\frac{\partial \log P_M^G/P_T^G}{\partial \log \vartheta_M}$ decreases in $\tilde{g}(s_w)$ and hence $T^{G'}(\vartheta_M)$ increases in $\tilde{g}(s_w)$, i.e. it is large when the distribution of s is denser around s_w. The economic interpretation of the result is that more disperse worker ability acts as a dampening force in the model. ∎

Figure A.5: Amplification in the Case of Worker Heterogeneity

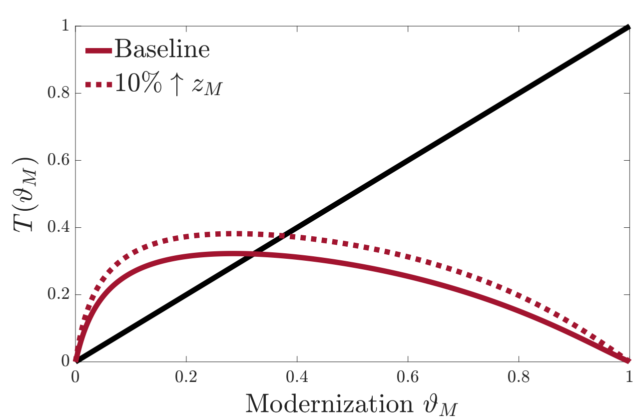

B.7 Equations for s_w under Log-Normal Distribution

The key equation implicitly defining s_w as a function of ϑ_M is

$$s_w \frac{\tilde{G}(s_w)\mathbb{E}[h|s\le s_w]}{\left(1-\tilde{G}(s_w)\right)\mathbb{E}[sh|s>s_w]} = \frac{1-\phi(1-\vartheta_M)}{(1+\tau)\phi(1-\vartheta_M)} \tag{54}$$

Under the assumption that (h,s) are log-normally distributed, the expressions for the different pieces of the LHS are:

$$\begin{aligned}
\tilde{G}(s_w) &= \Phi\left(\frac{\log s_w - \mu_s}{\sigma_s}\right) \\
1-\tilde{G}(s_w) &= \Phi\left(\frac{\mu_s - \log s_w}{\sigma_s}\right) \\
\mathbb{E}[h|s] &= \exp\left(\mu_h + \frac{(1-\rho^2)\sigma_h^2}{2} - \frac{\rho\sigma_h}{\sigma_s}\mu_s\right) s^{\frac{\rho\sigma_h}{\sigma_s}} \\
\implies \mathbb{E}[h|s\le s_w] = \mathbb{E}[\mathbb{E}[h|s]|s\le s_w] &= \exp\left(\mu_h + \frac{(1-\rho^2)\sigma_h^2}{2} - \frac{\rho\sigma_h}{\sigma_s}\mu_s\right)\mathbb{E}\left[s^{\frac{\rho\sigma_h}{\sigma_s}}|s\le s_w\right] \\
&= \exp\left(\mu_h + \frac{\sigma_h^2}{2}\right)\frac{\Phi\left[\frac{\log s_w-\mu_s}{\sigma_s} - \rho\sigma_h\right]}{\Phi\left[\frac{\log s_w-\mu_s}{\sigma_s}\right]}, \qquad (55) \\
\mathbb{E}[sh|s>s_w] = \mathbb{E}[s\mathbb{E}[h|s]|s>s_w] &= \exp\left(\mu_h + \frac{(1-\rho^2)\sigma_h^2}{2} - \frac{\rho\sigma_h}{\sigma_s}\mu_s\right)\mathbb{E}\left[s^{1+\frac{\rho\sigma_h}{\sigma_s}}|s> s_w\right] \\
&= \exp\left(\mu_h + \mu_s + \frac{\sigma_h^2+\sigma_s^2+2\rho\sigma_h\sigma_s}{2}\right)\frac{\Phi\left[\frac{\mu_s-\log s_w}{\sigma_s} + \sigma_s + \rho\sigma_h\right]}{\Phi\left[\frac{\mu_s-\log s_w}{\sigma_s}\right]} \\
\implies s_w \frac{\tilde{G}(s_w)\mathbb{E}[h|s\le s_w]}{\left(1-\tilde{G}(s_w)\right)\mathbb{E}[sh|s>s_w]} &= s_w \frac{\exp\left(\mu_h + \frac{\sigma_h^2}{2}\right)\Phi\left[\frac{\log s_w-\mu_s}{\sigma_s} - \rho\sigma_h\right]}{\exp\left(\mu_h + \mu_s + \frac{\sigma_h^2+\sigma_s^2+2\rho\sigma_h\sigma_s}{2}\right)\Phi\left[\frac{\mu_s-\log s_w}{\sigma_s} + \sigma_s + \rho\sigma_h\right]}
\end{aligned}$$

This implies that the equation is solved by a fixed point:

$$s_w = \frac{\exp\left(\mu_h + \mu_s + \frac{\sigma_h^2+\sigma_s^2+2\rho\sigma_h\sigma_s}{2}\right)\Phi\left[\frac{\mu_s-\log s_w}{\sigma_s} + \sigma_s + \rho\sigma_h\right]}{\exp\left(\mu_h + \frac{\sigma_h^2}{2}\right)\Phi\left[\frac{\log s_w-\mu_s}{\sigma_s} - \rho\sigma_h\right]} \frac{1-\phi(1-\vartheta_M)}{(1+\tau)\phi(1-\vartheta_M)} \tag{56}$$

C Structural Estimation Appendix

C.1 Variation in Local Moments

In Figure A.6, I present binned scatters showing how the three local target moments vary across microregions of different income levels and over time. Panel A.6a is a binned scatter version of Figure 1b, showing service modernization over time and space. Panel A.6b presents the cross-sectional modern wage premium in 2000 and 2010 against the log of per capita income in 2000. Richer microregions tend to have a lower wage premium, although the relationship is noisier than for employment shares. And for a large share of microregions, the modern wage premium increased from 2000 to 2010. Panel A.6c plots aggregate income in 2000 and 2010 against income in 2000. The 2000 points therefore follow the 45-degree line, while the 2010 points show that Brazil experienced regional convergence: the poorest microregions had the highest 2000-2010 income growth.

Figure A.6: Identifying Variation: Employment Shares, Wage Premium, and Income in 2000 and 2010

(a) Modern share of services $\hat{L}_{M,rt}$ **(b)** Modern wage premium $\hat{w}_{M,rt}$ **(c)** Log aggregate income p.c. $\bar{y}_{rt}$

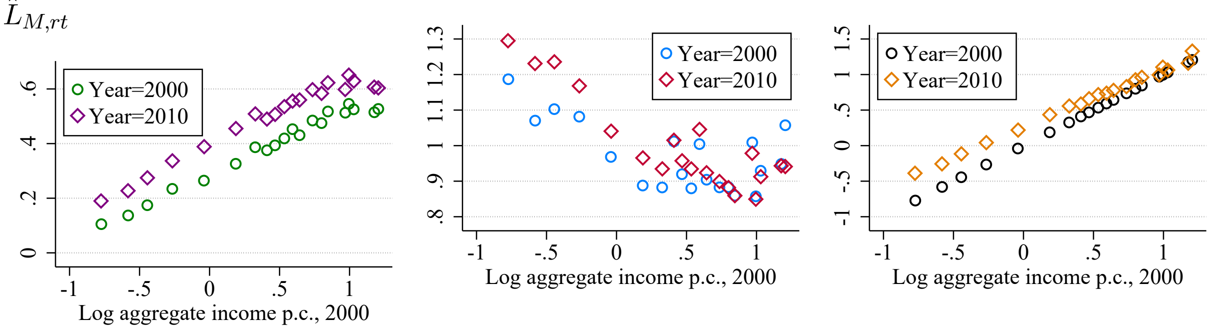

Note: Data from the 2000 and 2010 Brazilian Census. Aggregate income per capita is denominated in units of the real 2000 Brazilian minimum wage. All aggregation within microregions uses individual person weights as provided by the Census, and microregions are weighted by population in the binned scatters.

C.2 Identifying the Worker Ability Distribution

Because there is selection into the employment state under which I observe workers, all moments from the earnings distribution must involve a selection correction. First consider the variance of wage earnings $y_{i,rt}^{w}$ within a given local economy rt. This is equal to **Var**$[\log h_i | s_i \leq s_{w,rt}]$. By properties of the bivariate normal distribution and the truncated normal distribution, the moment has the following expression:

$$\mathbf{Var}[\log h_i | s_i \leq s_{w,rt}] = \sigma_h^2 \left[1 - \rho^2 \frac{\varphi\left(\frac{\log s_{w,rt} - \mu_s}{\sigma_s}\right)}{\Phi\left(\frac{\log s_{w,rt} - \mu_s}{\sigma_s}\right)} \left(\frac{\log s_{w,rt} - \mu_s}{\sigma_s} + \frac{\varphi\left(\frac{\log s_{w,rt} - \mu_s}{\sigma_s}\right)}{\Phi\left(\frac{\log s_{w,rt} - \mu_s}{\sigma_s}\right)} \right) \right] \tag{57}$$

The variance of log wage earnings is therefore informative about σ_h^2 but must correct for selection around $s_{w,rt}$. In the data, the standard deviation of wage earnings within location-time is 0.671, implying a variance of 0.450.

Next consider the variance of traditional earnings $\log y_{i,rt}^T$ within rt. Because traditional earnings are subject to measurement error, their variance reflects both worker heterogeneity and idiosyncratic error:

$$\mathbf{Var}\left[\log y_{i,rt}^T\right] = \mathbf{Var}[\log s_i + \log h_i | s_i > s_w] + \mathbf{Var}\left[\log u_{it}^T\right] \tag{58}$$

It is therefore necessary to account for $\mathbf{Var}\left[\log u_i^T\right] = \sigma_u^2$ before inferring dispersion of s, σ_s^2, from the variance of log traditional earnings. To make progress, I assume that $\log u_{it}^T$ follows an AR(1) process with normally distributed innovations ζ_{it}:

$$\log u_{it}^T = \varrho \log u_{it-1}^T + \zeta_{it} \tag{59}$$

This implies that for the ergodic distribution of u_{it}^T, $\sigma_u^2 = \frac{\sigma_\zeta^2}{(1-\varrho^2)}$, where σ_ζ^2 is the variance of the innovations ζ.[26] To estimate σ_ζ^2, note that the variance of one-period earnings growth in traditional earnings (after stripping away region-time effects) is $\frac{2}{1+\varrho}\sigma_\zeta^2$, and that of two-period earnings growth is $2\sigma_\zeta^2$. These moments yield an estimate of $\varrho = 0.074$, indicating weak persistence of measurement error, $\sigma_\zeta^2 = 0.184$, and hence $\sigma_u^2 = 0.185$.

Again using properties of the bivariate normal distribution and the truncated normal distribution, the variance of log traditional earnings satisfies:

$$\begin{aligned}\mathbf{Var}\left[\log y_{i,rt}^T\right] - \sigma_u^2 = {} & \sigma_h^2 + \sigma_s^2 + 2\rho\sigma_h\sigma_s \\ & + \left[\sigma_s^2 + 2\rho\sigma_h\sigma_s + \rho^2\sigma_h^2\right] \frac{\varphi\left(\frac{\log s_{w,rt} - \mu_s}{\sigma_s}\right)}{1 - \Phi\left(\frac{\log s_{w,rt} - \mu_s}{\sigma_s}\right)} \left(\frac{\log s_{w,rt} - \mu_s}{\sigma_s} - \frac{\varphi\left(\frac{\log s_{w,rt} - \mu_s}{\sigma_s}\right)}{1 - \Phi\left(\frac{\log s_{w,rt} - \mu_s}{\sigma_s}\right)} \right)\end{aligned} \tag{60}$$

Therefore the dispersion of log traditional earnings is informative about σ_s, but must account for measurement error, correlation ρ between h and s, and selection around $s_{w,rt}$. In

[26] I also assume that in their first period of traditional employment, workers draw an initial u_{it}^T from the ergodic distribution.

the data, the variance of log traditional earnings within location-time is 0.789, so that after subtracting the variance of u, the left-hand-side variance is 0.604.

Finally, the main moment I use to identify the correlation coefficient ρ is the probability of transition from wage employment to traditional employment, conditional on initial wage earnings. I assume that while in wage employment, workers make new draws of s from the conditional distribution $G(s|h_i)$ that may induce them to transition into traditional employment.[27] Note that the conditional distribution of s_i satisfies:

$$\log s_i | \log h_i \sim \mathcal{N}\left(\mu_s + \frac{\rho\sigma_s}{\sigma_h}(\log h_i - \mu_h), (1-\rho^2)\sigma_s^2\right) \tag{61}$$

Recall also that $\log y^w_{i,rt} = \log h_i + \log w_{rt}$. Therefore the probability of a wage to traditional employment transition WT, conditional on wage earnings, is:

$$\begin{aligned} \mathbf{Pr}\left(WT|\log y^w_{i,rt}\right) &= 1 - \Phi\left[\frac{\log s_{w,rt} - \left(\mu_s + \frac{\rho\sigma_s}{\sigma_h}\left(\log y^w_{i,rt} - \log w_{rt} - \mu_h\right)\right)}{\sigma_s\sqrt{1-\rho^2}}\right] \\ &= \Phi\left[\frac{\rho}{\sigma_h\sqrt{1-\rho^2}}\log y^w_{i,rt} + \mu_{0,rt}\right], \end{aligned} \tag{62}$$

where $\mu_{0,rt}$ is a location-specific fixed effect. The coefficient $\frac{\rho}{\sigma_h\sqrt{1-\rho^2}}$ can therefore be directly estimated from a probit model. The resulting coefficient is 0.056, suggesting weakly positive correlation between h_i and s_i. If one calls this equation β_1, then note that $\rho = \sqrt{\frac{\beta_1^2\sigma_h^2}{1+\beta_1^2\sigma_h^2}}$.

The final moment I need to identify the worker ability distribution is the share of workers in modern wage employment. In the PME, this share is $\ell_w = 0.61$. This information is in fact sufficient to solve for $\sigma_h^2\sigma_s^2$, because $\frac{\log s_w - \mu_s}{\sigma_s} = \Phi^{-1}[\ell_w]$. Plugging this in and solving the system of equations for σ_h, σ_s, and ρ yields $\sigma_h = 0.67, \sigma_s = 0.68, \rho = 0.038$.

C.3 Trade shock

First stage regression I present the first stage regression of the trade shock IV in Table A.2.

C.4 Population Shocks in the Model vs. Empirical Evidence

Figure A.7 conducts a similar exercise to Figure 10. In this case, I reproduce a population shock in the style of Imbert and Ulyssea (28) in the model, and compare the model's effects against the effects that they estimate. Imbert and Ulyssea (28) use a shift-share IV to generate exogenous variation in cities' migration rate, which is defined as the ratio between migrants

[27] Note that because of the wedge τ, these transitions can be – and often are – accompanied by a loss of earnings.

Table A.2: First Stage of 2SLS IV Regression: Trade Shock

	Log income shock
Tariff shock (DK17)	-3.702***
	[-4.491,-2.913]
Observations	409
State FE	Yes
Lag Controls	Yes

Note: 95% CI from robust standard errors in brackets. Log income is in 2000, the first Census year post-trade liberalization. Regression controls for 1991 modern employment share, log income, and modern wage premium. Microregions are weighted by 1991 population. For details on construction of the tariff shock, see Dix-Carneiro and Kovak (14).

and initial population. To reproduce this shock in the model, I instead increase the local market size L_{rt} by 1% and compute the semi-elasticity $\frac{d}{d\log L}$. I focus on results for the growth of formal wage employment, which corresponds roughly to ℓ_w in my model. Imbert and Ulyssea (28) find that a 1 p.p. increase in the migration rate increases formal wage employment by 0.271 percentage points. In my model, a 1% increase in L generates, on average, a 0.08 percentage point increase in formal wage employment. Again, the aggregate response to shocks is more conservative in my model than in empirical evidence, suggesting the elasticities governing the model's aggregate behavior fall well within the plausible range.

Figure A.7: Population Shocks and Labor Reallocation: Model vs. Data

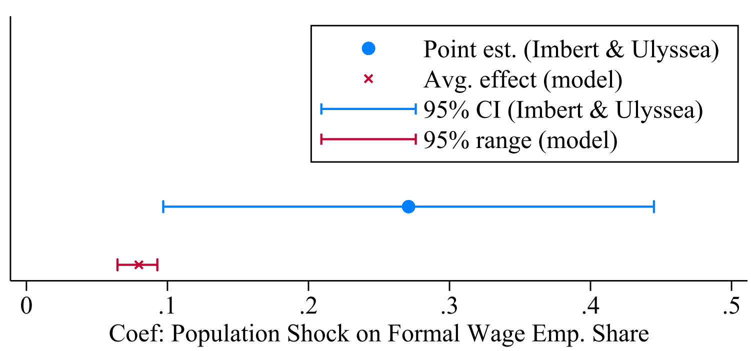

Note: "Empirics" displays point estimate and 95% CI from Imbert and Ulyssea (28). "Model" displays the average and the range between the 2.5% and 97.5% percentile of the distribution of microregion-level effects $\frac{d\ell_w}{d\log L}$.

C.5 Non-Homothetic CES Extension

Under the non-homothetic CES specification, the production side of the economy is identical, i.e. frictions, worker heterogeneity, and trade in goods are left unchanged. On the household side I now eliminate fixed costs to access modern services and specify household utility/real income y_i implicitly as follows:

$$1=\frac{(C_{i,X}/y_i)^{1-\phi}}{\phi^{\phi}(1-\phi)^{1-\phi}}\left(\omega_M^{\frac{1}{\xi}}\left(\frac{C_{i,M}}{y_i^{\varepsilon_M}}\right)^{\frac{\xi-1}{\xi}}+(1-\omega_M)^{\frac{1}{\xi}}\left(\frac{C_{i,T}}{y_i^{\varepsilon_T}}\right)^{\frac{\xi-1}{\xi}}\right)^{\phi\frac{\xi}{\xi-1}} \tag{63}$$

Faced with prices P_X,P_M,P_T and expenditure e_i, this demand system implies:

$$\begin{aligned}P_X C_{i,X}&=(1-\phi)e_i\\ \frac{P_M C_{i,M}}{P_T C_{i,T}}=\frac{\vartheta_{i,M}}{1-\vartheta_{i,M}}&=\frac{\omega_M}{1-\omega_M}\left(y_i^{\varepsilon_M-\varepsilon_T}\frac{P_M}{P_T}\right)^{1-\xi}\end{aligned} \tag{64}$$

I specify $\varepsilon_M=1-\frac{\varepsilon}{2}$,$\varepsilon_T=1+\frac{\varepsilon}{2}$, so that there is one non-homotheticity parameter ε to estimate. Demand is homothetic in the case where $\varepsilon=0$, and modern services are a luxury when $\varepsilon>0$.

Real income y_i can also be written implicitly as a function of expenditure e_i and sectoral prices:

$$y_i=\frac{e_i}{P_X^{1-\phi}}\left(\omega_M\left(\frac{P_M}{y_i^{\frac{\varepsilon}{2}}}\right)^{1-\xi}+(1-\omega_M)\left(\frac{P_T}{y_i^{\frac{-\varepsilon}{2}}}\right)^{1-\xi}\right)^{\frac{-\phi}{1-\xi}} \tag{65}$$

Therefore the transformation $T(\vartheta_M)$ now uses the same relative expenditure and prices as a function of ϑ_M as implied in the baseline model – $\frac{e}{w}(\vartheta_M),\frac{P_X}{w}(\vartheta_M),\frac{P_M}{w}(\vartheta_M),\frac{P_T}{w}(\vartheta_M)$ – but then computes real income $y(\vartheta_M)$ based on Equation (65) and implies ϑ_M based on Equation (64):

$$\vartheta_M=T(\vartheta_M)=\frac{\frac{\omega_M}{1-\omega_M}\left(y(\vartheta_M)^{\varepsilon_M-\varepsilon_T}\frac{P_M}{P_T}\right)^{1-\xi}}{1+\frac{\omega_M}{1-\omega_M}\left(y(\vartheta_M)^{\varepsilon_M-\varepsilon_T}\frac{P_M}{P_T}\right)^{1-\xi}} \tag{66}$$

Given primitives, the new formulation of $T(\vartheta_M)$ is sufficient to solve for equilibrium under the NHCES specification. I now discuss how to implement this specification in the quantitative part of the paper.

C.5.1 Quantitative Procedure under Non-Homothetic CES

First, since the demand system differs from the one in the baseline model with fixed costs, it is necessary to estimate the new non-homotheticity parameter ε. To do this, I fit Equation (64) to data on expenditure shares and income:

$$\log\left(\frac{\vartheta_{i,M}}{1-\vartheta_{i,M}}\right)=\varepsilon(\xi-1)\log y_i+\beta' K_i, \tag{67}$$

where K is a vector of household-level controls, including geography to control for local prices. Since in the individual-level microdata, $\vartheta_{i,M}$ clusters at zero and one, I use binscatter to compute the average ϑ_M within income bins, then find the ε that best fits Equation (67). This procedure yields an estimate of $\varepsilon(\xi-1)=0.404$, implying $\varepsilon=0.136$.

The first two steps of the sequential accounting procedure are identical to the baseline model. These steps yield values for $\vartheta_{M,rt}$, $s_{w,rt}$, τ_{rt}, and $\tilde{z}_{X,rt}$, which also imply values of $P_{T,rt}$,e_{rt},w_{rt}.

Using Equation (64), $P_{M,rt}$ can be written as a function of y_{rt} and prices:

$$P_{M,rt}=y_{rt}^{\varepsilon}P_{T,rt}\left(\frac{\omega_M}{1-\omega_M}\frac{1-\vartheta_{M,rt}}{\vartheta_{M,rt}}\right)^{\frac{1}{\xi-1}} \tag{68}$$

Plugging this expression into the equation for y in (65), it is possible to directly infer y from e_{rt}, $\vartheta_{M,rt}$, and $P_{T,rt}$:

$$y_{rt}=\left[e_{rt}P_{T,rt}^{-\phi}\left(\frac{1-\omega_M}{1-\vartheta_{M,rt}}\right)^{\frac{\phi}{\xi-1}}\right]^{\frac{1}{1+\phi\frac{\varepsilon}{2}}} \tag{69}$$

$P_{M,rt}$ is then immediate from (68). Finally, $z_{M,rt}$ can be inferred from $P_{M,rt}$ in the same way as the baseline model.

C.5.2 Quantitative Results under Non-Homothetic CES

After performing the quantitative procedure above under the non-homothetic CES specification, I compare the results against those in the baseline specification. I examine the inferred variation in $\log\tilde{z}_M$, compute the multiplier μ_M, and simulate trade and migration shocks.

The results are remarkably similar to those in the baseline model with fixed costs to consume modern services. Table A.3 reproduces Table 3 from the main text, adding a line for results from the non-homothetic CES model. It shows that like the baseline model, the non-homothetic CES model shrinks the variation in effective modern service productivity $\log\tilde{z}_M$ required to explain the data, in fact by slightly more than the baseline model. The baseline and non-homothetic CES models also imply almost exactly the same share of microregions with conditions for amplification, i.e. $\mu_M>1$.

Table A.3: Variation in $\tilde{z}_M$ and μ_M under Baseline and Non-Homothetic CES

	SD($\log\tilde{z}_{M,2000}$)	SD($\log\tilde{z}_{M,2010}$)	$\overline{\Delta\log\tilde{z}_{M,2000}^{2010}}$	Share $\mu_{M,2000}>1$	Share $\mu_{M,2010}>1$
Baseline	0.208	0.166	0.240	46%	21%
No Demand Effects	0.311	0.236	0.297	0%	0%
Non-Homothetic CES	0.206	0.145	0.205	46%	20%

Figure A.8 displays the heterogeneity in μ_M under the non-homothetic CES model. Again the heterogeneity is similar, but relative to the baseline model, the non-homothetic CES specification implies fewer extremely high values of μ_M and lower average μ_M among the poorest microregions. This is because the baseline model implies an income elasticity of modern service demand that is more variable than under non-homothetic CES.

Figure A.8: Heterogeneity in μ_M under Non-Homothetic CES

(a) Histogram of μ_M (NHCES) **(b)** μ_M and Aggregate Income (NHCES)

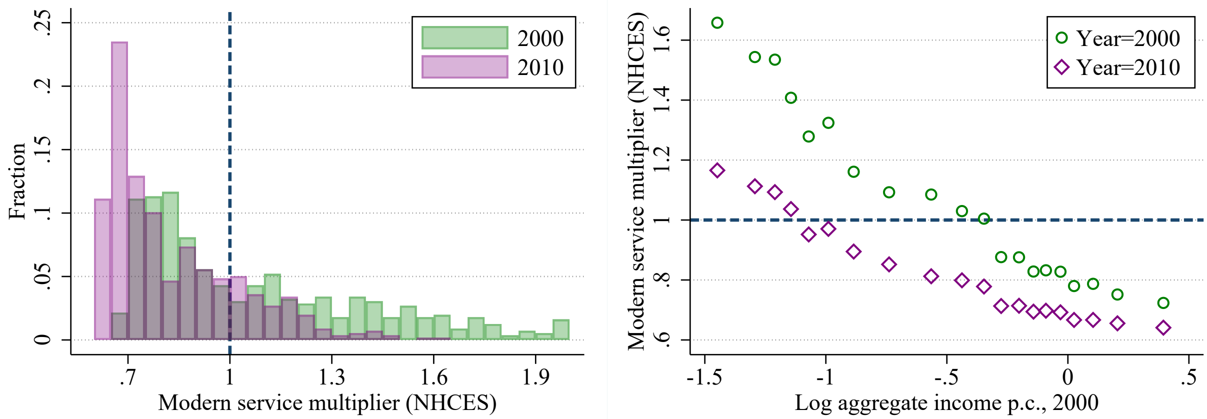

In Table A.4, I compare the response to simulated trade and migration shocks in the baseline and non-homothetic CES models against the data. The two models imply almost exactly the same response to migration shocks, both of which are significantly smaller than the empirical effects found in Imbert and Ulyssea (28). For the trade shock, the non-homothetic CES model implies a smaller response of modern service employment than the baseline model, which is in turn smaller than the response in the data. The baseline model therefore performs slightly better at replicating the empirical response to the shock, even if both models are conservative relative to the data.

Table A.4: Effect of aggregate shocks by different specifications of non-homothetic demand

	Data	Baseline (fixed costs)	Non-homothetic CES
Trade shock: $\frac{d\ell_M}{d\log\tilde{z}_X} \Big/ \frac{d\log e_G}{d\log\tilde{z}_X}$	0.051	0.0307	0.0226
Migration shock: $\frac{d\ell_w}{d\log L}$	0.271	0.076	0.076

C.6 Alternate values of σ

In Table A.5, I reproduce the results from Table 3 under three alternative values of σ: $\sigma=3$, $\sigma=4$, and $\sigma=10$. The results broadly confirm that lower values of σ – that is, greater love-of-variety in modern services – act as an amplifying force in the model. Lower values of σ tend to imply less inferred variation in $\log\tilde{z}_M$, and a greater share of microregions with multipliers $\mu_M>1$. As

σ increases, the results get closer to the "no demand effects" counterfactual with more disperse $\log\tilde{z}_M$ and no amplification.

Table A.5: Variation in $\tilde{z}_M$ and μ_M with Different Levels of σ

	SD($\log\tilde{z}_{M,2000}$)	SD($\log\tilde{z}_{M,2010}$)	$\overline{\Delta\log\tilde{z}^{2010}_{M,2000}}$	Share $\mu_{M,2000} > 1$	Share $\mu_{M,2010} > 1$
Baseline ($\sigma = 5.02$)	0.208	0.166	0.240	46%	21%
No Demand Effects	0.311	0.236	0.297	0%	0%
$\sigma = 3$	0.217	0.141	0.165	75%	59%
$\sigma = 4$	0.200	0.150	0.215	61%	38%
$\sigma = 10$	0.244	0.204	0.281	12%	0.4%

Since the quantitative results are somewhat sensitive to the value of σ, in Table A.6 I verify that the model's response to aggregate shocks is closer to the data under lower rather than higher values of σ. The baseline value of $\sigma = 5.02$ produces smaller model responses than observed in the data for both the trade shock (as shown in Section 5.6) and the migration shock (28). Reducing the value of σ increases the magnitude of the model's response; for both shocks, a σ between 3 and 4 is required to match the empirical response. In contrast, a greater value of $\sigma = 10$ diminishes the model's response and brings it further from the data. The benchmark value of $\sigma = 5.02$ is therefore conservative relative to a value of σ that explicitly targeted the empirical response to trade and migration shocks.

Table A.6: Effect of aggregate shocks by different values of σ

	Data	Baseline ($\sigma = 5.02$)	$\sigma = 3$	$\sigma = 4$	$\sigma = 10$
Trade shock: $\frac{d\ell_M}{d\log\tilde{z}_X} \Big/ \frac{d\log e_G}{d\log\tilde{z}_X}$	0.051	0.0307	0.092	0.043	0.024
Migration shock: $\frac{d\ell_w}{d\log L}$	0.271	0.076	0.343	0.130	0.026