Distribution Costs

Alessandra Peter* Cian Ruane†

May 11, 2026

Abstract

We provide the first direct estimates of distribution expenses incurred by manufacturing plants and quantify their importance for aggregate consumption and measured misallocation. Using a novel measure from the Indian Annual Survey of Industries, we document three facts: distribution expenses amount to over half of labor costs, are over three times larger as a share of sales for plants in the largest decile relative to the smallest, and declined by one third from 2000 to 2010. We develop a model of heterogeneous manufacturing firms that rely on distribution services to sell across space. The improvements in distribution over that time period raised manufacturing consumption by 24.5%. The gains materialize quickly, but unevenly: large firms expand while many small local firms shrink or exit. Distribution costs also matter for measured misallocation: standard TFPR measures generate spurious dispersion and a positive relationship with size. In the ASI, accounting for distribution costs lowers measured TFPR dispersion by 5.1% and the elasticity of TFPR with respect to plant size by 7.0%.

*Department of Economics, New York University; alessandra.peter@nyu.edu.

†Central Bank of Ireland; cian.ruane@centralbank.ie.

We would like to thank Pete Klenow, Nick Bloom, Chad Jones, Melanie Morten and Monika Piazzesi for their continued support and guidance on this project. Thank you to Manuel Garcia-Santana for discussing the paper. Thanks also for the insightful comments provided by all the participants of the Virtual International Trade and Macro Seminar, STEG, the Stanford Macroeconomics Lunch, the Stanford Trade Workshop, the Stanford Development Tea, and EAGLS. Arnav Deshpande provided excellent research assistance. We gratefully acknowledge financial support from the Stanford Institute for Innovation in Developing Economies (SEED). Cian additionally acknowledges financial support from the Stanford Institute for Economic Policy Research. The work in this paper is partly supported by the Macroeconomic Research in Low-Income Countries program of the IMF and the UKs Foreign, Commonwealth and Development Office (FCDO). Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the Central Bank of Ireland, the IMF, its Executive Board, or its management.

1. Introduction

Before goods can be consumed, they must not only be produced but also distributed from the location of production to the final consumer. Distribution costs are therefore not only a resource cost for the economy but can easily become a bottleneck for the manufacturing sector. Even a very productive firm can only grow so large when it cannot profitably ship its goods into more distant markets. This is particularly relevant in developing countries, where growth of manufacturing plants over the life-cycle has been historically slow (Hsieh and Klenow, 2014; Akcigit, Alp and Peters, 2021) and where large infrastructure projects have generated significant productivity gains for manufacturing firms (Ghani, Goswami and Kerr, 2016; Asturias, Garcia-Santana and Ramos, 2019). Despite their apparent importance for the development of the manufacturing sector, there is little direct evidence on the magnitude of distribution costs.

In this paper, we provide the first direct estimates of distribution expenses incurred by manufacturing plants. Using a unique measure of distribution costs for Indian manufacturing plants, we document three key facts. First, distribution expenses are large in the aggregate. Between 1993 and 2013, manufacturing plants spent an average of 3.2% of their sales on distribution, which is more than half as large as labor costs. Second, larger plants spend relatively more on distributing their goods. Compared to the lowest size decile, the distribution share of sales is more three times larger for plants in the highest decile. Third, we document a substantial downward trend in the aggregate distribution share over the 2000s – from 3.8% to 2.5% of sales.

We rationalize the positive relationship between distribution share and plant size as resulting from larger plants shipping their goods further distances than smaller plants (Holmes and Stevens, 2012). We embed this channel into a model of heterogeneous manufacturing firms that purchase distribution services in order to sell their products across space. Shipping further distances requires more distribution services and hence larger plants have a higher distribution share of sales. The data on distribution expenses, together with newly constructed estimates of intranational trade in India, allow us to estimate the model's key parameters affecting the costs of distributing products across space. We find that the improvements in distribution from the 1990s to 2000s led to a nearly 25% increase in manufacturing consumption.

The new facts we document are based on data from the 1993-2013 waves of the Indian Annual Survey of Industries (ASI). The ASI is a representative survey of formal Indian manufacturing plants, which asks plants to directly report how much they spend on distribution.[1] Distribution expenses include transportation costs as well as other

[1] In the Annual Survey of Manufacturers in the US, for instance, these are not recorded as plants are

variable costs such as packing fees and transit insurance. Using this new data, we document that distribution costs as a share of sales are large, increasing in firm size, and have declined significantly over time. We show that these facts are not driven by systematic under-reporting by small plants or heterogeneous profit shares. While exporters have, on average, higher distribution shares, we show that the positive relationship with plant size also holds for firms that only sell domestically.

The model is designed to be as parsimonious as possible while still speaking directly to the data. It builds on a continuous-location version of Arkolakis (2010) and has three key ingredients: firm heterogeneity, space, and distribution services whose cost increases with distance. Manufacturing firms differ in the quality of the variety they produce and choose how much and where to sell domestically as well as whether or not to export. Selling to a specific location requires setting up a local distributor, which is captured as a fixed cost. We allow this fixed cost to be increasing in distance, reflecting, for example, worsening information frictions (Allen, 2014). In addition, shipping goods across space comes at a variable cost, which we again allow to increase with distance.

We calibrate the model to the Indian manufacturing sector in 2000 by matching the properties of the distribution share, the firm-size distribution, exporter moments, and newly constructed estimates of interstate trade. The key parameters governing domestic distribution are the productivity of the distribution sector, the rate at which variable distribution requirements increase with distance, and the rate at which fixed market access costs rise with distance. These are primarily identified by the aggregate distribution share, the ratio of the aggregate to the average distribution share, and the interstate trade share.

We first use the model to quantify the aggregate effects of intranational trade barriers. We do so using a stark counterfactual, akin to moving to free trade: In an economy in which no firms incur any variable costs of distributing their goods, consumption of manufacturing goods would be 21.3% higher. The reason a 3.8% aggregate distribution share can have such large consequences is due to the fact that (i) many inputs used by manufacturing firms are produced by other firms, hence compounding distribution costs, and (ii) large and productive firms are particularly constrained by high distribution costs. In fact, setting the size-dependence of distribution costs to zero raises manufacturing consumption by 9.2%.

We next use the model to evaluate improvements in India's distribution sector over time. The 2000s were a decade of rapid growth in India (World Bank, 2026b), and also a period of extensive investments in the country's road network, most notably the building of the Golden Quadrilateral highway network. Both Ghani et al. (2016) and Asturias

asked to report sales net of distribution expenses.

et al. (2019) find that the Golden Quadrilateral led to large income gains in the manufacturing sector, including through improvements in allocative efficiency. Over the same period, adjusting for industry composition, the aggregate distribution share fell from 3.8% to 2.9%, the ratio of the aggregate to average distribution share fell from 1.6 to 1.5, and the interstate trade share increased from 23% to 32%.

Through the lens of our model, these changes imply that the distribution sector became more productive, that variable distribution costs became less steep in distance, and that the fixed cost of accessing distant markets declined. Moving all three distribution parameters from their 2000 to their 2010 values raises manufacturing consumption by 24.5%. The largest component is the inferred decline in fixed market access costs: changing only this parameter raises manufacturing consumption by 13.9%. The increase in distribution-sector TFP raises manufacturing consumption by 4.6%, while the decline in the rate at which variable distribution costs increase with distance raises it by 3.8%.

The previous results compared two steady states of the model. In reality, firms cannot instantly create new distribution networks or enter new markets. We therefore simulate a transition in which firms meet distributors only gradually. Even when firms can access new markets with only a 10% probability each year, manufacturing consumption adjusts quickly. Over 62% of the long-run increase in aggregate consumption materializes in the first year. By 2010, the economy has achieved between 80% and 94% of the eventual gains, depending on the rate at which firms meet distributors. The reason is that firms immediately benefit from lower distribution costs in the markets they already serve. At the same time, the gains are unevenly distributed. Large and productive firms expand, while many small local firms shrink or exit. In the new steady state, the bottom 25% of firms exit altogether, and 79% of firms active in 2000 experience a decline in profits. For these firms, high distribution costs served as a shield against competition from more productive firms located elsewhere.

Beyond quantifying the importance of the distribution sector in India, the paper shows that distribution costs can bias standard misallocation measures in manufacturing: because these costs are size-dependent and often omitted from measured input costs, standard TFPR measures can partly mistake real distribution costs for distortions. It can also generate a spurious positive relationship between TFPR and firm size, a key moment used in inferring how distortions vary with productivity. In the ASI, accounting for distribution costs reduces measured TFPR dispersion by 5.1% and the elasticity of TFPR with respect to plant size by 7.0%. This suggests that part of what is often interpreted as misallocation may instead reflect real costs of reaching consumers.

Literature Our paper relates to a growing literature about the importance of intersectoral linkages for economic development. This literature typically evaluates how distortions or low productivity in individual sectors get amplified through the input-output matrix (Bartelme and Gorodnichenko, 2023; Liu, 2019; Fadinger, Ghiglino and Teteryatnikova, 2022). Our main contribution relative to this literature is to document systematic heterogeneity across firms in exposure to one specific key upstream sector and to quantify its importance for aggregate output. Similar to Jones (2011), there are complementarities between sectors in our model.[2] We show that the higher reliance of larger firms on distribution inputs amplifies the effect of complementarities between the production and distribution sectors.

Our paper is also closely related to the literature evaluating the consequences of intranational trade costs in developing countries (most notably, Atkin and Donaldson (2015); Donaldson (2018); Asturias et al. (2019); Asher and Novosad (2020); Van Leemput (2021)). In the absence of direct measures, these papers follow the traditional approach in international trade and infer the magnitude of internal trade barriers from a combination of data on the flow of goods and spatial price gaps. Another common approach in the international trade literature is to estimate fixed and variable costs of exporting from the nonlinear effects these costs have on market participation (Das, Roberts and Tybout, 2007). We complement this literature by providing an additional, arguably more direct measure of internal trade barriers: firms' distribution expenses. Another strength of our distribution expenses measure is that it also captures decreases in distribution costs from increases in competitiveness in the transportation sector. Arnold, Javorcik, Lipscomb and Mattoo (2016) show that such pro-competitive reforms were important in the transportation sector from 1993 to 2005 and led to gains to manufacturing firms.

Our model is based on a continuous-location version of Melitz (2003) and Arkolakis (2010). The main difference relative to the latter is that we focus on how variable distribution costs change with distance, rather than fixed costs. This is crucial for matching our empirical finding that the distribution share increases with firm size.

Our paper also relates to an important literature on differences in the firm size distribution between developed and developing economies, and the role of distortions. It is well established that developing countries have a larger mass of small firms (Hsieh and Olken, 2014), with this relationship even holding across Indian states of varying income levels (Jain and Kothari, 2025). Hsieh and Klenow (2014) document that Indian plants grow slowly with age and suggest that poor transportation infrastructure could be an important contributor. We contribute by evaluating the importance of distribu-

[2] Our assumption of complementarities is consistent with Peter and Ruane (2025), who estimate long-run elasticities of substitution between material inputs and services below one for Indian manufacturing plants.

tion costs, a concrete size-dependent component of firm costs, through the lens of a multi-sector general equilibrium model and showing that they matter both for aggregate consumption and the fat left tail of the size distribution.

Our findings also speak to the misallocation literature (Hsieh and Klenow, 2009) and the role of size-dependent distortions (Bento and Restuccia, 2017). Recent literature has highlighted how measurement error and misspecification make the identification of distortions from average revenue products challenging (Gollin and Udry, 2021; Bils, Klenow and Ruane, 2021; Carrillo, Donaldson, Pomeranz and Singhal, 2023; Kim, Rotemberg and White, 2025). We show that distribution costs, while not distortionary in our model, can generate dispersion in measured average revenue products and a spurious positive relationship between TFPR and firm size.

The rest of the paper is structured as follows. In Section 2. we present our empirical results, in Section 3. we lay out our model, in Section 4. we quantify the model, in Section 5. we perform our counterfactual exercises, study measured TFPR dispersion, and analyze the transition and incidence of gains, and Section 6. concludes.

2. Distribution Costs in Indian Manufacturing

We first discuss the dataset and main variables used, and then document the key stylized facts of the paper. As a share of sales, distribution costs are: (1) large, (2) increasing in plant size, and (3) trending downwards over time. Finally, we discuss our interpretation of these findings.

2.1. Data Description

The data come from the 1993-2013 waves of the Indian Annual Survey of Industries (ASI). The ASI is a nationally representative survey of the formal manufacturing sector in India. It covers all plants with more than 10 workers using power, and all plants with more than 20 workers not using power. There are, on average, over 30,000 plants per year in the survey. The survey contains standard information on plants' output, employment, labor costs, capital stocks, and intermediate input expenditures. The sampling and main variables are described in more detail in Data Appendix A.

In addition to the more standard variables, the ASI contains information on plants' expenditure on outward distribution expenses.[3] The measure of distribution expenses includes the value of outward transport costs, rebates, commissions, transit insurance

[3]'Outward' refers to distribution costs incurred on products sold, as opposed to intermediate inputs purchased.

of goods sold, and packing fees for goods sold.[4] It therefore comprehensively captures *variable* costs of selling goods. It does not include *fixed* distribution costs, such as advertising, finding distributors, or setting up a warehouse.

As far as we are aware, the measurement of outward distribution and transport expenses at the plant level is unique to the ASI. Typically, the value of products shipped is reported net of distribution costs, as is the case for the U.S. Annual Survey of Manufactures, or distribution costs are reported together with taxes and other fees. This is therefore an ideal setting for providing new insights into the magnitude of distribution costs.

Throughout the paper, we report distribution costs as a share of sales.[5] This is the most appropriate measure of the cost burden, since plants need to distribute actual goods (their sales), as opposed to the value-added contributed in supply chains. However, the distribution share of sales alone does not provide a full picture of the extent to which distribution costs affect aggregate output and consumption. Given that production chains consist of multiple steps, distribution costs are incurred multiple times. This is particularly relevant in India, where the average intermediate share is 77% in the manufacturing sector. The model in Section 3. features intermediate inputs, and therefore captures this amplification through the input-output network within manufacturing. For completeness, we also report distribution costs as a share of value-added in the summary statistics in Table 1.

2.2. Distribution Costs: Magnitude, Heterogeneity, and Explanations

Magnitude of distribution expenses in the manufacturing sector. We first document that distribution expenses are large in the aggregate. On average between 1993 and 2013, total distribution costs incurred by manufacturing firms amount to 3.2% of total sales. To put this number in perspective, distribution costs are more than half as high as labor costs, which averaged 6.0% of aggregate sales over the same period. The seemingly low labor and distribution shares in Indian manufacturing are driven by its relatively high intermediate share of sales, which averages 77%. Relative to value-added, the aggregate distribution share is 10.2%, and the labor share 19.1%.

[4]This is reported as '*other distributive expenses*' in Block J, column 10 of the schedule in more recent survey years. See p.37 of the 2014 Instruction Manual for the Annual Survey of Industries (Concepts, Definitions and Procedures) for details on what is included.

[5]We use sales net of taxes because we are interested in distribution costs relative to the value of sales from the selling firm's perspective. Taxes are a very small share of sales and have a minor impact on all the numbers reported in this section.

Heterogeneity in the distribution share across plants. The distribution share is not only large in the aggregate but also varies systematically across plants of different sizes. Table 1 summarizes key moments of this variable. The median distribution share is 0.7%, the average is 2.1%, and the 90th percentile is as high as 5.6%. As a share of value-added, these numbers are 3%, 8.4%, and 23.7%, respectively. [6]

Table 1: Summary Statistics

Distribution costs relative to	Aggregate	Mean	p10	p50	p90	Share = 0
Sales	3.2%	2.1%	0.0%	0.7%	5.6%	25.1%
Value-added	10.2%	8.4%	0.0%	3.0%	23.7%	25.1%

Notes: Data come from the Indian Annual Survey of Industries 1993 to 2013. The top row reports summary statistics for distribution costs as a share of sales. The bottom row reports the same statistics as a share of value-added. The statistics are calculated pooling all years from 1993 to 2013. See Figure 10 in Appendix A5. for a full histogram.

The fact that the aggregate distribution share is larger than the average suggests that bigger manufacturing firms have systematically higher distribution shares. Figure 1(a) confirms this and shows that the distribution share is monotonically increasing in plant size. The smallest decile of plants spend close to 1% of sales on distributing their products, while the largest decile of plants spend around 3.5%.[7] The relationship between distribution share and plant size is not driven by plants' location choices. We show in Appendix Figure 11 that the relationship between the distribution share and size holds when we control for industry $\times$ year $\times$ state fixed effects.

In the 1993-1997 survey years, the ASI records even more detailed information on sub-components of distribution expenses. For these years, we can decompose distribution costs into: (1) transport costs, which include the costs of freight and shipping, packing fees, and transit insurance; (2) commissions owed to intermediaries; (3) rebates and drawbacks owed to customers; and (4) other. Transport costs are by far the most important component of distribution costs, accounting for 46.4% aggregate distribution expenses. Commissions and rebates account for 17.8% and 12.1%, respectively, with the remaining 23.7% of distribution expenses categorized as 'other'. We show in Figure 1(b) that all components of distribution expenses (as a share of sales) behave similarly as a function of size. We therefore use total distribution spending for the remainder of the analysis.

[6] 25% of plants report zero distribution expenses. This may reflect measurement error in distribution expenses. We provide further details on these plants in Appendix A2. and show that removing them does not qualitatively affect any of our main empirical findings.

[7] Plant size is measured as total output, which includes sales, the change in finished goods inventories, and revenues from other sources. We show in Panel (b) of Appendix Figure 11 that our findings are very

Figure 1: Distribution Share vs Plant Size

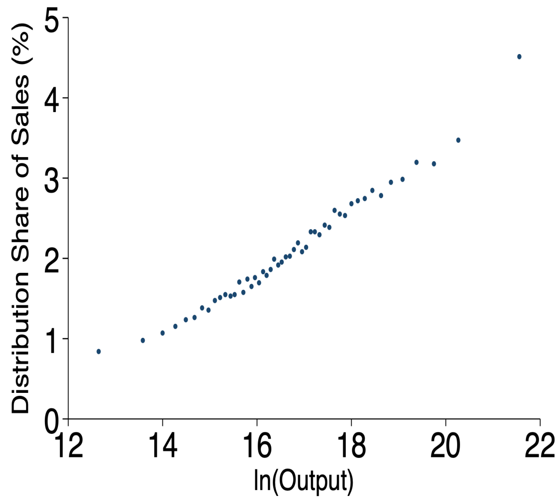

(a) Total Distribution Costs

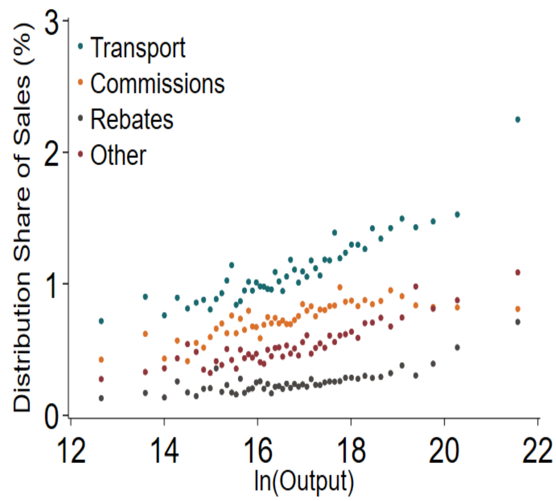

(b) Sub-Components

Notes: Data comes from Indian Annual Survey of Industries 1993 to 2013. Both subfigures show binned scatter plots of the plant-level distribution share against log(output), after residualizing on sector-year fixed effects. The left figure includes all years, while the right figure includes 1993-1997. The right figure shows a binned scatter plot of the plant-level distribution share against output, for all four subcomponents of the distribution costs: transport costs, commissions, rebates, and other. We residualize all variables on sector×year fixed effects.

Distribution costs and international exports Distribution costs, as measured in the ASI, include the costs of distributing products both domestically and internationally. While we cannot separately measure distribution costs associated with domestic and foreign sales at the plant level, we can document differences in the overall distribution share of sales between exporters and non-exporters for the subset of years in which exports are reported in the ASI – 1997 and 2008-2013. These findings are reported in Table 2.

On average, 9.2 percent of plants export, but these plants tend to be large, accounting for 25.3 percent of aggregate output. The aggregate export share of total output for exporters is just under one-third. Unsurprisingly, on account of both being larger and having international markets, exporters tend to have higher distribution costs relative to their sales than non-exporters, 3.2 percent vs. 2.5 percent. However, we show in Figure 12 that the increasing relationship between firm size and the distribution share is not solely due to a higher share of exporters among large firms, as it holds when restricting the sample to non-exporters.

similar when we use employment as a measure of plant size.

Table 2: Distribution Costs by Exporter Status

	Share of Plants	Share of Output	Export Share of Output	Distribution Share
Non-exporter	90.8	74.7	0.0	2.5
Exporter	9.2	25.3	31.6	3.2

Notes: All statistics are averages across years in which export data is available (1997, 2008–2013). The distribution share is the aggregate distribution share by firm type. The export share is the aggregate export share of total plant output. Sampling weights used for all statistics.

2.3. Interpreting the Facts

Distribution Share and Plant Size We argue that the relationship between plant size and the distribution share is the result of 1) the cost of shipping goods increasing in the distance shipped, and 2) larger plants tending to ship their goods further distances than smaller plants. These claims are well-supported by the existing evidence. A large literature in international trade has documented that transport costs increase with distance. In addition, Holmes and Stevens (2012) provide evidence that larger U.S. manufacturing plants ship their products further distances within the U.S. than smaller plants.[8] Similarly, the trade literature has extensively documented that larger plants export relatively more than smaller plants (e.g. Bernard and Jensen (1995)).[9]

One alternative explanation we address is systematic mismeasurement. If small plants fail to report a portion of their distribution expenses (because of less accurate bookkeeping, for example), we would spuriously infer that distribution costs as a share of sales are increasing in size. Smaller plants might also be more likely to do their distribution in-house, in which case some of their expenses might be reported as labor costs. While we cannot provide any direct evidence on the frequency of misreporting, we can analyze how inward distribution costs vary with plant size.[10] We would expect misreporting to affect inward distribution costs in the same way as outward distribution costs. However, we show in Appendix Figure 13 that the inward distribution share decreases with plant size, which significantly mitigates concerns about mismeasurement.

Another explanation is that the increasing relationship between distribution share and size could be driven by larger plants having lower profitability (sales relative to

[8] We cannot replicate this analysis for India as we are not aware of any comparable data.

[9] To overcome large distribution expenses, firms could open plants in additional locations. Unfortunately, we cannot directly analyze this in the data, as the ASI does not allow researchers to link plants belonging to the same firm. However, we do observe whether plants belong to multi-establishment firms and use this information in the model to gauge the importance of this issue.

[10] Inward distribution expenses are the distribution costs that plants pay to transport and acquire production materials. These are reported in the ASI for the years 1993, 1994, and 1996.

costs). If this were the case, all cost shares would be increasing in size, and the share of distribution costs in total costs would be flat. However, we show in Appendix Figure 14 that the distribution share of total costs is also increasing in plant size. The decrease in the aggregate distribution share documented in Figure 2 (a) is also similar if we measure the distribution share relative to total costs.

Distribution costs and misallocation The misallocation literature following Hsieh and Klenow (2009) typically infers distortions from measured differences in firms' revenue-to-input-cost ratios. Is heterogeneity in the distribution share of sales across plants reflective of misallocation, which lowers aggregate productivity? We argue that the positive relationship between plant size and the distribution share results from physical constraints involved in the transportation and distribution of resources across space. Given these constraints, which would also be present for a social planner, this heterogeneity in the distribution share is consistent with allocative efficiency. That said, ignoring distribution costs altogether could lead researchers to mismeasure misallocation and the relationship between distortions and firm size. We return to this issue in Section 5. after presenting the model.

2.4. The Distribution Share over Time.

So far, we have shown that distribution costs are high and increasing in plant size. In Figure 2, we document how these data patterns have changed over time. Figure 2(a) shows that aggregate distribution costs relative to sales fell considerably from 2000 to 2010. During this period, which coincided with large transport infrastructure investments, the distribution share fell from just below 4% to around 2.5%.[11] While part of this change stems from changes in sectoral composition, we find the aggregate distribution share fell from just below 4% to 2.9% over this 10-year period, even when adjusting for changes in sectoral composition.[12]

The decline in the aggregate distribution share was driven by changes in the distribution share conditional on plant size, rather than by a change in the size distribution of plants. Figure 2(b) shows that the distribution share of sales decreased for plants of all sizes. In addition, the sales share of large plants expanded during this period, and hence changes in the size distribution put upward pressure on the aggregate distribu-

[11] Major road infrastructure investments included the Golden Quadrilateral highway network, the North-South and East-West corridor, and the Pradhan Mantri Gram Sadak Yojana rural road program. We discuss these in more detail in Section 3..

[12] We adjust for industry composition by aggregating the annual sectoral distribution shares using constant sectoral sales shares from 2000.

tion share.[13] Most notably, large plants saw even larger decreases in their distribution shares than smaller plants – there was a flattening in the relationship between the distribution share and plant size. These findings suggest important potential gains for manufacturing plants from reductions in distribution costs, particularly for larger plants.

Figure 2: Distribution Share over Time

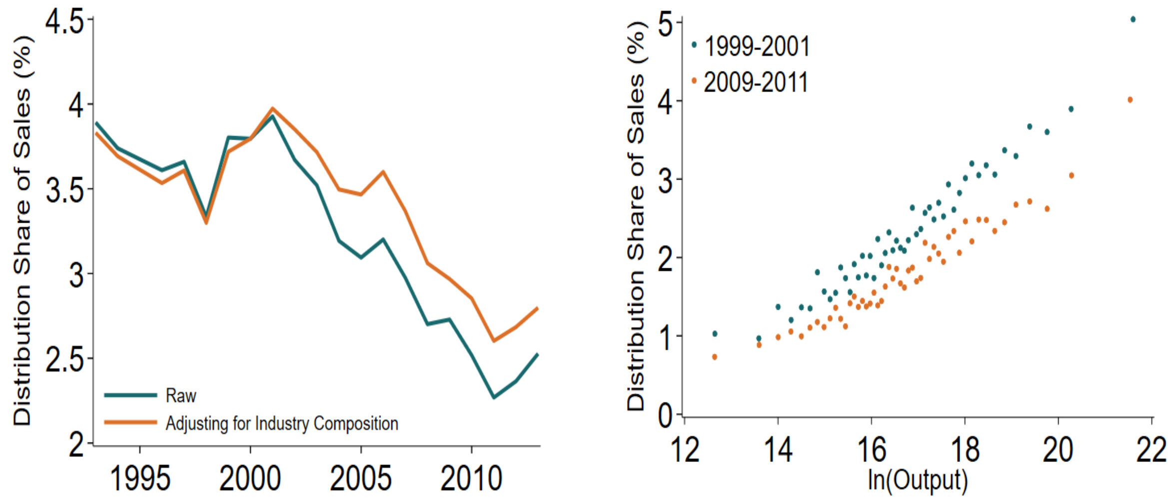

(a) Aggregate (b) Plant Heterogeneity

Notes: Data source is the Indian Annual Survey of Industries 1993 to 2013. The left figure shows the aggregate distribution share of sales over time in India, both with and without controlling for industry composition. We adjust for industry composition by holding sectoral sales shares constant at their 2000 level. The right figure shows a binned scatter plot of the plant-level distribution share against log(output), after residualizing on sector-year fixed effects. It shows 1999-2001 separately from 2009-2011. We choose 3-year windows to increase the sample size for the plot.

Infrastructure Investment and the Decline in the Distribution Share. We provide suggestive evidence that the trend in the aggregate distribution share was at least partly driven by large public infrastructure improvements in the context of the Golden Quadrilateral highway network. The GQ was launched in 2001, and 90% of the project was finished by 2006, linking together the four major economic centers of Delhi, Mumbai, Chennai, and Calcutta. In Appendix A3., we show that states in which the GQ was built saw both larger decreases in the state-level aggregate distribution share between 2000 and 2010, as well as larger decreases in the ratio of the aggregate to average distribution share. This is consistent with the GQ reducing distribution costs for firms (as shown in Asturias et al. (2019)) and suggests that some of the improvements were driven by government policy rather than technological improvements at the firm level.

[13]The sales share of the largest 10% of plants increased from 83.3% to 84.5%.

Summary We document that the aggregate distribution share in Indian manufacturing is large and that bigger firms spend more on distributing their goods. We argue that the increasing relationship between plant size and the distribution share arises from the combination of two factors: the variable cost of distributing goods is increasing in distance, and larger plants ship their goods further distances. As a result, larger plants use distribution services more intensively than smaller plants. This makes them particularly affected by the quality of Indian infrastructure and by the productivity of the distribution sector. In addition, there has not only been a significant decrease in the aggregate distribution share during the 2000s, but this decline was particularly pronounced for the largest firms.

3. Model

To rationalize the observed relationship between a firm's size and its distribution share, we develop a model in which heterogeneous manufacturing firms use inputs from the distribution sector to sell their goods across space. The model allows us to quantify the aggregate output effect of distribution costs and their systematic heterogeneity across firms, and to evaluate the gains from improvements in India's distribution sector over time.

The model is deliberately parsimonious: it includes only the minimum set of ingredients needed to match the central facts – firm heterogeneity, geography, and distribution services whose costs increase with distance. This allows us to discipline the model directly with the data, quantify how much a 3.8% aggregate distribution share matters for output, and isolate the channels through which distribution costs affect the consumption of manufactured goods. We restrict attention to manufacturing and abstract from services and agriculture, so all aggregate results should be interpreted as applying to the manufacturing sector.

3.1. Setup

The model is based on a continuous-location extension of Arkolakis (2010). Time is discrete, agents are risk-neutral and do not discount the future. The economy comprises two sectors: manufacturing and distribution. Each manufacturing firm produces a differentiated variety. Firms use distribution services to sell their products to domestic and foreign consumers, as well as to other firms as an intermediate input. Including exports ensures that distribution costs associated with reaching foreign markets are captured separately from those attributable to intranational trade. We model space in the simplest way that preserves the main mechanism: the economy lies on a circle with circumference 1, and the rest of the world is a single location outside the circle.

Final Goods Producers

At each location l along the circle, there is a final goods producer who combines all available varieties ω to produce the location's final good. Varieties are aggregated using a CES aggregator with elasticity of substitution σ.

$$Y_l = \left(\int_{\omega \in \Omega_l} \psi(\omega) c_l(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}} \tag{1}$$

Varieties differ in their quality $\psi(\omega)$. Higher-quality varieties give more utility per unit of consumption and are more efficient in producing the intermediate good.[14] The set of varieties available in location l, Ω_l, is an equilibrium object determined by the choices of all firms in the economy of which markets to sell to.

Final-goods producers take the prices of varieties as given and sell the aggregate output Y_l to both consumers and firms as intermediate inputs. The price of a unit of the final good is given by

$$P_l = \left(\int_{\omega \in \Omega_l} \psi(\omega)^{\sigma} p(\omega)^{1-\sigma} d\omega \right)^{\frac{1}{1-\sigma}} \tag{2}$$

The choice of aggregator gives rise to a standard CES demand equation for a variety with quality $\psi(\omega)$ and price $p(\omega)$:

$$y(\omega) = P_l^{\sigma-1} \; Y_l \; \psi(\omega)^{\sigma} \; p(\omega)^{-\sigma} \tag{3}$$

The final good in each location can be used either for consumption or as an intermediate input by firms producing there. Therefore,

$$Y_l = C_l + M_l \tag{4}$$

[14] Jain and Kothari (2025) document a strong empirical link between firm size and product quality in Indian manufacturing: larger plants charge higher prices, use more expensive inputs, and employ more skilled workers supporting our modeling choice of quality as the primary source of firm heterogeneity.

Domestic Consumers

At each point along the circle, there is a continuum L of workers.[15] Workers have preferences over the domestic (local) final good C_l and an imported good C_m. Agents cannot save or borrow, which gives rise to the following problem (suppressing time subscripts for ease of exposition):

$$\max_{\{C_l, C_m\}} \quad \left(\alpha C_l^{\frac{\varepsilon-1}{\varepsilon}} + (1-\alpha) C_m^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\epsilon}{\epsilon-1}}$$
$$\text{s.t} \quad P_l C_l + P_m C_m \leq y \tag{5}$$

Consumers inelastically supply their unit endowment of labor to either the manufacturing or the distribution sector for a wage w. They also receive a share $1/L$ of aggregate profits $\bar{\pi}$.

Their income y is therefore equal to $w + \bar{\pi}/L$.

Foreign Demand

In addition to selling to the final goods producer, manufacturing firms can export their goods to the rest of the world. We model India as a small-open economy, so world prices and world demand are taken as given.

Demand for each exported variety is given by

$$c_x(\omega) = C_w \psi(\omega) \left(\frac{p_x(\omega)}{P_w}\right)^{-\sigma}, \tag{6}$$

where C_w is the world-level of demand for Indian goods, and P_w is the price of the world good. The imported good is the same as the world good, so $P_m = P_w$. In what follows, we will treat the world good as a numeraire, so that $P_m = P_w = 1$.

Distribution Sector

Firms in the distribution sector are perfectly competitive. They employ labor L_D to produce distribution services d using a constant returns to scale production technology:

$$d = A_D L_D \tag{7}$$

[15] The setup is symmetric at every point on the circle.

The combination of perfect competition and constant returns to scale implies that $p_D = w/A_D$.

Manufacturing Sector

In each period t, there is an endogenously determined mass J of manufacturing firms operating along the circle. Manufacturing firms are heterogeneous along two dimensions: the quality ψ of the variety they produce and their (production) location on the circle. On the extensive margin, firms choose both the set of markets (points on the circle) in which they sell their product, and whether to export to the *rest of the world*, which we think of as a location outside of the circle. On the intensive margin, they choose the quantities to sell in each market and to export.

Domestic Production

Firms operate with a constant-returns-to-scale technology. They combine labor and an aggregate intermediate input to produce output at the factory gate. In order to sell a quantity $q(\psi, n)$ to a consumer located at a distance n along the circle a firm must not only produce the good, but also (i) pay a fixed cost $f(n)$ of accessing the market and (ii) purchase distribution services $d(n)$ to deliver the product to the consumer. The fixed cost $f(n)$ depends only on the distance n between the destination market and the firm's production location.[16] This captures costs associated with establishing and maintaining a distribution network (e.g., finding and contracting with local retailers). We follow Arkolakis (2010) in modeling the marginal cost associated with accessing a market located a distance n away:

$$f(n) = (1 - 2n)^{-\beta} \tag{8}$$

As in Arkolakis (2010), the functional form chosen captures three important dimensions of the cost of accessing a sales market:

1. $f'(n) \geq 0$. The further away a market is located, the more costly it is to find a distributor there. The parameter $\beta \geq 0$ governs how quickly costs increase with distance.

2. $\lim_{n \to 0} f'(n) > 0$. There is a positive cost of serving the home market ($n = 0$), akin to a standard fixed cost in a model of production with only one market.

[16] These fixed costs are in units of labor.

3. If $\beta > 0$, $\lim_{n \to \frac{1}{2}} f'(n) = \infty$. Accessing the entire market is prohibitively costly.[17]

Once a firm has paid the fixed cost of establishing a distribution network in a market at a distance n, it produces output for sale in that market according to the following Leontiev production function:[18]

$$q(n) = \min\left\{ M^{\theta} L_p^{1-\theta}, \frac{d(n)}{1+\epsilon n} \right\} \tag{9}$$

The physical good is produced with labor L_p and an intermediate good M. TY be sold in a market at a distance n, the physical good needs to be shipped there using distribution services $d(n)$. These distribution services are purchased from the distribution sector at unit cost p_D. In order to sell one unit of the good in a market at distance n, the firm needs to purchase $d(n) = 1 + \epsilon n$ units of distribution services. ϵ governs the rate at which the quantity of distribution services required increases with distance. This parameter captures certain aspects of the quality of a country's transportation infrastructure. For example, it could be relatively costly to ship goods over longer distances (compared to shorter distances) when highways are of poor quality. This would be captured by a high value of ϵ.

Given consumers' demand, a firm with quality ψ solves the following profit maximization in each market n:

$$\begin{aligned} \pi^*(\psi, n) \equiv \max_{q,l,m,d} \quad & p(q(\psi,n))\, q(\psi,n) - w\, l(\psi,n) - P_l m(\psi,n) - p^D\, d(\psi,n) - w\, f(n) \\ \text{s.t.} \quad & m(\psi,n)^{\theta} l(\psi,n)^{1-\theta} = q(\psi,n) \\ & d(\psi,n) = (1+\epsilon n)\, q(\psi,n) \\ & f(n) = (1-2n)^{-\beta} \end{aligned} \tag{10}$$

The second step in the firm's problem consists of choosing its optimal scale, i.e., which parts of the circle to sell to. Given that (i) production of the physical good exhibits constant returns to scale and (ii) both the fixed and variable cost of distribution are monotonically increasing in the distance to the market reached, choosing where to sell to amounts to choosing the furthest market to which the firm sells. We denote the dis-

[17]Given that the circle has circumference 1, the furthest market is located at a distance $n = \frac{1}{2}$.

[18]While output and distribution services need to be used in fixed proportions *in each market*, the elasticity of substitution between the physical good $M^{\theta} L_p^{1-\theta}$ and distribution services *at the firm-level* is larger than zero, since firms also adjust on the extensive margin of how many markets to serve.

tance between the firm and its furthest point of sale by $n^*(\psi)$, where:

$$n^*(\psi) = \underset{N\in[0,1/2]}{\text{argmax}} \int_0^N 2\pi^*(\psi, n)dn \tag{11}$$

Given that the economy is located on a circle of circumference 1, there are always two locations that are exactly distance n away from the firm, and the maximal market size is $\frac{1}{2}$. We denote by $\pi^*(\psi)$ the resulting optimal per-period profits a firm with quality ψ can achieve:

$$\pi^*(\psi) = \max_N \int_0^N 2\pi^*(\psi, n)dn = \int_0^{n^*(\psi)} 2\pi^*(\psi, n)dn \tag{12}$$

Exports

To become an exporter, the firm must pay a fixed cost f_x, denominated in labor. Upon paying the fixed cost, the firm produces output for sale to the rest of the world, using to the same production function as for domestic sales. Shipping to the rest of the world requires distribution services d_x; in order to sell one unit of good to the rest of the world, the firm has to purchase κ_x units of distribution services.

Given the demand for varieties by the rest of the world, an exporter solves the following problem:

$$\begin{aligned} \pi_x^* \equiv \max_{q_x, l_x, m_x, d_x} \quad & p_x(q_x)\, q_x - w\, l_x - P_l m_x - p^D\, d_x - w\, f_x \\ \text{s.t.} \quad & m_x^\theta l_x^{1-\theta} = q_x \\ & d_x = \kappa_x\, q_x \end{aligned} \tag{13}$$

Let J_x be the endogenously determined mass of exporting firms.

Entry and Exit

There is a large pool of potential entrants into the manufacturing sector. To enter, firms incur a sunk cost of entry f_E, measured in units of labor. Upon paying this cost, firms discover their quality ψ. The distribution from which qualities ψ are drawn is exogenous with CDF $g(\psi)$. After observing their quality, firms decide whether to remain active. If they do, they receive profits $\Pi(\psi, t) = \pi(\psi, t) + \mathbb{I}_{export}(\psi)\pi_x(t)$ in every period in which they operate. Quality is constant over time, and so $\Pi^*(\psi, t) = \Pi^*(\psi)$. In addition, firms are risk-neutral, there is no time discounting, and the probability of not receiving an exogenous exit shock is $1 - \delta$ in each period. This implies that the value of a firm with quality ψ is:

$$\sum_{s=0}^{\infty}(1-\delta)^s\Pi^*(\psi)=\frac{\Pi^*(\psi)}{\delta}$$

If $\Pi^*(\psi)\geq 0$, the firm will choose to remain active and start production. Otherwise, it will endogenously choose to exit. We denote by L^E the mass of workers used for entry and by M^E the mass of entrants each period.

3.2. Equilibrium

Given a quality distribution $g(\psi)$, an equilibrium is set of:

1. prices

$$\left\{ p^D, w, P_l, P_m, \{p(\psi,N)\}, \{p_x(\psi)\} \quad \right\}$$

2. quantities

$$\left\{ \begin{array}{l} L^D, L^E, L^P, \{c(\omega)\}, \{q(\psi)\}, \{q_x(\psi)\}, \{l(\psi)\}, \{l_x(\psi)\}, \\ \{m(\psi)\}, \{m_x(\psi)\}, \{d(\psi)\}, \{d_x(\psi)\}, J, J_x, M^E \end{array} \right\}$$

such that

1. Consumers optimize subject to their budget constraint
2. Active firms optimize subject to their constraints
3. Production = consumption + intermediate input use in each location
4. The aggregate constraints hold
5. The free entry condition holds
6. Trade is balanced in every location

We focus on equilibria that feature a stationary distribution of firms. Every period, the mass of successful entrants exactly offsets the mass of exiting firms. Given that exit is equally likely for all levels of quality and net entry is zero, the endogenous distribution of active firm qualities remains constant.

3.3. The Role of Distribution Costs

In order to illustrate the main mechanisms of the model, we start by considering the extreme case where $\beta = \epsilon = 0$. In this special case, neither fixed nor variable costs of distributing products vary with the distance between the plant and the consumer. The profit maximization problem laid out in equation (10) is now independent of the market n. This implies that all active firms sell to the entire circle. All consumers, irrespective of their location l on the circle, have access to the exact same set of varieties. The economy collapses to a standard Melitz (2003) model.

Now let $\epsilon > 0$. Selling to markets farther away requires more distribution services per unit sold. From the firms' perspective, marginal costs increase with n. This implies that they sell less and make smaller profits in markets located further away. Firms serve a market only if profits are sufficient to cover the fixed cost f.[19] This gives the following equilibrium expression for $n^*(\psi)$:[20]

$$n^*(\psi) = \frac{1}{\varepsilon}\left(\frac{A_D}{c}\left(\frac{\sigma-1}{\sigma}\right)\left(\frac{L}{\sigma}\right)^{\frac{1}{\sigma-1}}\psi^{\frac{\sigma}{\sigma-1}} - (A_D+1)\right), \tag{14}$$

where c is the unit cost of production of the firm, $c = \left(\frac{1}{\theta}\right)^{\theta}\left(\frac{w}{1-\theta}\right)^{1-\theta}$.

Equation (14) provides intuition for the main mechanism in this model. First, higher-quality firms choose to sell to more markets. Since their product is of higher quality, they make higher profits in each market. They will therefore find it worthwhile to overcome the fixed costs to reach more distant markets. This feature of the model is what allows us to replicate the main empirical finding from Figure 1 (a): distribution costs as a share of sales are increasing in firm size. The distribution share in a market of distance n is given by:

$$\frac{p_D d(n,\psi)}{p(\psi,n)q(\psi,n)} = \frac{\frac{1+\epsilon n}{A_D}}{\frac{c}{w} + \frac{1+\epsilon n}{A_D}}\,\frac{\sigma-1}{\sigma} \tag{15}$$

Conditional on distance n, the distribution share is independent of firm quality. A firm's total distribution share is then simply equal to the integral over all markets served. Since sales to more distant markets require with a higher distribution share and higher-quality firms sell to more distant markets, they spend a larger overall share of sales on distribut-

[19]Since $\beta = 0$, the fixed cost of accessing markets is constant at $f(n) = f$

[20]This formula applies to firms whose optimal choices are interior; i.e. they serve a positive fraction of consumers, but not the entire market $(n^*(\psi) \in \left[0, \frac{1}{2}\right)$. There is a positive mass of firms who sell to the whole market: $n^*(\psi) = \frac{1}{2}$

ing their products.

We now turn to discussing the role of the parameter ϵ. Consider a *decrease* in ϵ.[21] From equation (14) one can see that a reduction in ϵ will increase $n^*(\psi)$ for all active firms. This is intuitive: lowering ϵ is equivalent to a reduction in the marginal cost of selling to each market, leading to an expansion of all firms. In addition to this direct partial-equilibrium effect, firm size is affected through general-equilibrium effects on the wage. As firms expand, they demand more labor, both directly as an input into production and indirectly through their increased demand for distribution services. The wage increases to restore labor market clearing.[22] Because of the higher wage, production costs increase for all firms. This has an important effect on the minimum quality cutoff for active firms. Denote by $\underline{\psi}$ the lowest quality firm that still finds it worthwhile to remain active in equilibrium.[23] We have that:

$$\underline{\psi} = \left(c + \frac{w}{A_D}\right)^{\frac{\sigma-1}{\sigma}} \left(\frac{w}{y}\frac{\sigma}{L}\right)^{\frac{1}{\sigma}} \left(\frac{\sigma}{\sigma-1}\right)^{\frac{\sigma-1}{\sigma}} \tag{16}$$

As the wage increases, $\underline{\psi}$ also increases. The reason for this is similar to that in Melitz (2003). The expansion of high-quality (large) firms drives up the wage, making it harder for small firms to compete for workers. This drives out some small, low-quality firms. The new equilibrium distribution of firms now features larger and higher-quality firms.

Now consider $\beta > 0$, but $\epsilon = 0$. In such an environment, the distribution share of firms would be independent of productivity. While more productive firms still sell to more distant markets, the *variable* cost of distribution does not vary with distance. In terms of economic channels, β and ϵ play similar roles: the higher these are, the more firms sell locally. This affects in particular the most productive firms, which sell further, and allows small, low-productivity firms to survive. The fact that the two parameters affect firm sales similarly but have different effects on the distribution share is a relationship we use to identify the parameters of the model.

Firms enter the export market if variable profits from exporting cover the fixed cost of entry into the export market:

$$\psi^x = \left(\frac{\sigma w f_x}{Y_w}\right)^{\frac{1}{\sigma}} \left(\frac{\sigma}{\sigma-1}\left(c + \kappa_x \frac{w}{A_D}\right)\right)^{\frac{\sigma-1}{\sigma}} P_w^{-1}$$

[21] This captures any change in the technological environment that makes it relatively cheaper to ship long distances, e.g., a highway network.

[22] We have normalized the price index P to 1, which leaves the wage as the only price left to clear markets.

[23] These are firms for which $n^*(\underline{\psi}) = 0$

Provided the fixed cost of exports is high enough, $\psi^x > \underline{\psi}$. Conditional on exporting, distribution costs as a share of exports are given by:

$$\frac{p_D d_x(\psi)}{p_x(\psi) q_x(\psi)} = \frac{\frac{\kappa_x}{A_D}}{\frac{c}{w} + \frac{\kappa_x}{A_D}} \frac{\sigma - 1}{\sigma} \tag{17}$$

Equation (17) shows that distribution cost as a share of exports is independent of productivity. The positive relationship between firm productivity and distribution shares is driven by domestic distribution. It always holds conditional on the set of exporting firms, and also holds for all firms provided κ_x is large enough.

4. Quantifying the Model

In this section, we quantify the model to the Indian manufacturing sector at two points in time: 2000 and 2010. We use the facts on distribution shares from Section 2. and newly constructed estimates of interstate trade to calibrate the model's key parameters.

4.1. Interstate Trade

We construct most of the moments used to calibrate the model from the ASI. In particular, our measures of the distribution share as well as how this varies with plant size are crucial for identifying the key parameters governing firms' use of distribution services. To identify the fixed cost parameter β, we need data on the distances to which firms ship their products. However, intranational trade data for India is notoriously sparse prior to the introduction of the Goods and Services Tax (GST) in 2017.

We therefore digitize historical data on interstate shipments, which allows us to estimate the volume of trade between states in the 2000s. The raw data come from a handbook titled "Inter-State Movements / Flows of Goods by Rail, River and Air", collected and published by the Indian Directorate General of Commercial Intelligence and Statistics (DGCIS).[24] These data provide unique information on interstate trade and how it changed in India over time, yet there are two important limitations: 1) shipments are measured in quantities rather than values, and 2) they exclude shipments by road.[25] To overcome these issues and nonetheless obtain a complete estimate of the interstate trade share, we follow an approach based on Van Leemput (2021).

[24] The DGCIS publishes data based on accounting years which stretch from April to March of the following year. See Appendix A4. for a more detailed discussion of data and methodology.

[25] Capturing interstate trade by road is particularly important because of the development of large-scale road infrastructure during the 2000s, such as the Golden Quadrilateral.

The approach exploits the fact that interstate shipments for export are reported separately from those for domestic consumption. For any good, we can then construct aggregate interstate trade by rail and aggregate exports by rail and combine this with information on total aggregate exports and domestic production to infer the interstate trade share (the ratio of the value shipped interstate to the value of domestic production). The approach requires two key assumptions: 1) for each good, the share of shipments by rail vs road is the same for exports and domestic consumption, and 2) the quality of goods shipped interstate by rail is the same for export and for domestic consumption. These assumptions imply that:

$$\frac{\text{quantity exported via rail}}{\text{quantity shipped interstate via rail}} = \frac{\text{value exported via road and rail}}{\text{value shipped interstate via road and rail}} \tag{18}$$

We obtain the interstate trade share of each good by dividing its export intensity (aggregate export value/aggregate production value) by the ratio in Equation (18). We then average these shares across goods to obtain our targeted moments.

We estimate that India's interstate trade share increased from 23% in 2000 to 32% in 2010.[26] While correctly pinning down the *level* of the interstate trade share requires the strong assumptions above, what matters for our main counterfactuals is the *change* in the interstate trade share over time. For example, while it is likely that the quality of exported goods is higher than that of those shipped domestically, what matters most for us is that relative export quality did not change over time, a much weaker assumption. In addition, assumption 1) is supported by evidence that the density of roads and rails in India is highly correlated across states (Van Leemput, 2021). We therefore view our estimates as providing strong evidence of an increase in the value of India's interstate trade over time.

4.2. Calibration

We quantify the model developed in Section 3.. Firm productivity is assumed to follow a Pareto distribution with shape parameter η. This leaves a total of ten parameters to calibrate. Two of them can be calibrated directly: the intermediate share and the exit rate of firms. We set the intermediate share θ to match the 59% material input share we

[26]Our estimates are higher than those of Van Leemput (2021) for 2011, also using the DGCIS data. While our approach is based on theirs, we make some methodological choices that differ from theirs (detailed in Appendix A4.) and restrict our attention to goods for which we can measure the interstate trade share in both 2000 and 2010. More recently, India's Ministry of Finance estimated that, in 2017, interstate trade in goods and services subject to the GST amounted to 60% of GDP (as opposed to production) (Ministry of Finance, 2018).

measure in the ASI.[27] We set the exit probability of firms to match exit rates as estimated by Peters (2020) for formal manufacturing plants in Indonesia.[28] The remaining eight parameters are calibrated to match a set of key moments in the baseline year 2000. All parameters and moments are reported in Table 3.

Table 3: Parameters and moments in 2000

Parameter	Description	Value	Moment	Data	Model
(A) Set externally					
θ	intermediate share	0.59	materials share ASI	59%	59%
δ	exit probability	0.082	exit rate (Peters (2020))	8.2%	8.2%
(B) Jointly calibrated					
A_D	TFP distrib sector	11.4	aggr distrib share	3.8%	3.9%
ϵ	slope of distrib wrt distance	43.2	aggr / avg distrib share	1.62	1.67
β	fixed cost of market access	42.4	sh sales across states	23.0%	23.7%
η	Pareto tail index	4.2	sales sh top 10%	82.8%	85.3%
f_e	entry cost	0.05	sales sh bottom 50%	1.7%	1.7%
f_x	fixed cost of exporting	0.06	sh of plants that export	12.9%	13.1%
κ_x	variable cost of exporting	2.40	distrib share exporters	4.5%	4.5%
Y_w	world demand	0.010	conditional export share	23.1%	22.9%

Notes: Table 3 reports the calibrated values for the model's key parameters, the data moments used for calibration, and their model counterparts. "Sales across states" is computed as one minus the share of sales within state. All other moments are computed using the ASI. Conditional export share is the share of plant sales that are exported conditional on export status. Parameters are chosen to match the equally weighted sum of squared deviations of model and data moments.

Identification While all eight parameters are jointly identified by all moments, they differ in their respective sensitivity to each moment. Table 3 aligns each parameter next to the moment it is most sensitive to. We now discuss identification of the three key parameters of the model – A_D, β, and ϵ – in more detail.

The productivity of the distribution sector, A_D, determines the amount of resources the economy as a whole allocates to distribution. As Figure 3 shows, the value of A_D mostly affects the aggregate distribution share, with very little impact on all other moments. In particular, an increase in A_D reduces the average and the aggregate distribution shares in similar proportions.

[27]Intermediate inputs are only purchased from other manufacturing firms in the model, and so we calibrate the intermediate share to match the material input share in the data (therefore excluding purchases of service intermediates). In Appendix B4., we re-calibrate the model for alternative values of the intermediate share and show that our conclusions are robust to higher or lower shares.

[28]The ASI is not suitable for measuring exit because plants are randomly included in the sample. We do not know of alternative estimates of exit rates for manufacturing plants in India.

The rate at which distribution requirements rise with distance, ϵ, affects the aggregate distribution share much more than the average distribution share. When $\epsilon = 0$, the average and aggregate distribution shares are identical since all firms spend the same proportion of sales distributing their products. As ϵ increases, distribution costs rise more steeply with distance, affecting larger firms more. As a result, the aggregate distribution share increases more than the average.

The rate at which the fixed cost of accessing markets increases with distance, β, is primarily identified from the average distance firms sell their products. A given level of distribution expenses could be driven by either high costs and firms selling locally, or lower costs and firms selling to more distant markets. We tease the two stories apart using the interstate trade share. In the model, "sales in state" are defined as a firm's sales within the radius corresponding to a typical state's share of Indian GDP. As a result, a segment length of approximately 7% is considered to be within-state sales. The third panel of Figure 3 illustrates that the value of β strongly affects the share of sales that cross states.

Figure 3: Identifying fixed and variable components of distribution

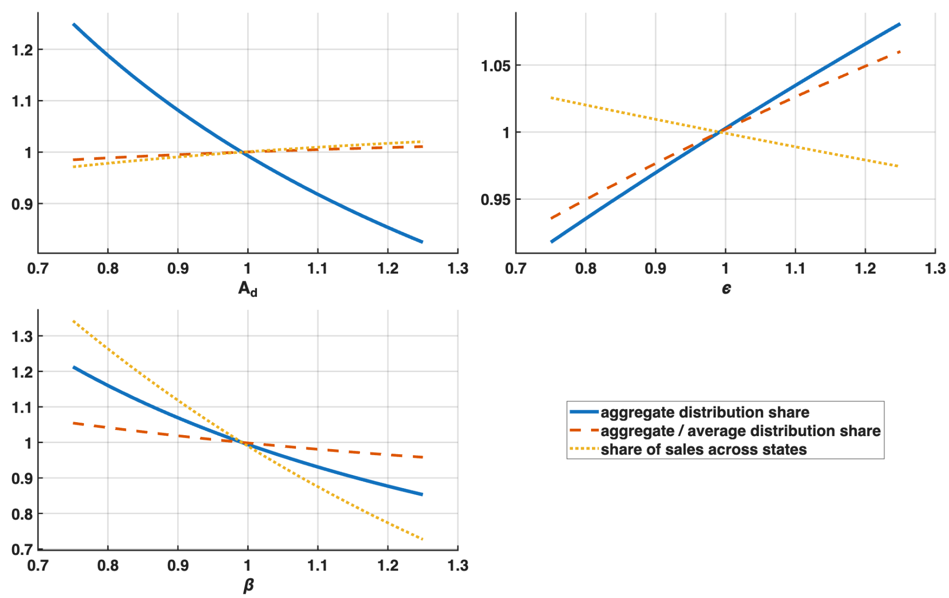

Note: Figure 3 shows sensitivity of model moments to parameters. Starting from the baseline calibration, each panel plots all three moments relative to their baseline values as a function of a single parameter, holding all other parameters constant at their baseline values. The change in each parameter is also measured relative to baseline.

We calibrate η, the shape parameter of the firm productivity distribution, and the entry cost, f_e, to match the employment distribution across firms. The lower η is, the more skewed the distribution of productivity and hence size. The larger f_e is, the more productive a firm must be in order to operate, which reduces the top 10% employment share and increases the employment share of the top 50%.

Finally, we calibrate the model's export block. Conditional on all other parameters, f_x, the fixed cost of exporting mostly affects the extensive margin, that is, the fraction of plants that choose to export. The world demand shifter Y_w governs relative sales of firms domestically vs abroad. We do not directly observe how much firms spend on distributing their products, but we infer international distribution costs from the distribution share of exporting firms (and their difference relative to the average distribution share of all firms). Figure 15 in Appendix B plots all moments against all parameters.

5. How Important are Distribution Costs?

In this section, we quantify the importance of distribution costs for aggregate consumption and for measuring misallocation. Throughout this section, we report changes in manufacturing consumption in the model economy; for simplicity, we refer to these as aggregate consumption.[29]

5.1. Aggregate Effect of Level and Heterogeneity in Distribution Costs

In Section 2., we document two key facts about distribution costs: they are large in the aggregate, and bigger firms spend more – as a share of sales – on distributing their goods. We first run a set of counterfactuals to quantify how much the level of distribution costs and the heterogeneity across firms matter for aggregate consumption.

Level of distribution costs. The Indian manufacturing sector spends around 4% of sales on distributing their products. To put that number into perspective, we run a counterfactual in which distributing goods is free ($A_D = \infty$) and hence all firms have a distribution share of zero. Comparing this economy to our baseline calibration allows us to quantify the output loss from the high level of (variable) distribution costs we find in the data. Table 4 reports all main results for this section.

In an economy in which no firms incur any variable costs of distributing their goods from the factory gate, aggregate consumption would be over 20% higher. This is not meant to be a realistically achievable counterfactual, but rather an answer to the question of how important an aggregate distribution share of 3.8% is for the economy as a whole.

[29]From 2000 to 2010, India's manufacturing share of GDP averaged 16.5% (World Bank, 2026a).

Table 4: Aggregate effects of distribution costs

	Consumption Gains
(A) Importance of the *level* of distribution costs	
No distribution costs ($A_D = \infty$)	+21.3%
(B) Importance of *heterogeneity across firms*	
Distribution costs constant wrt distance ($\epsilon = 0$)	+9.2%
Lower productivity in distribution sector $A_D \downarrow 10\%$	
$\epsilon = 43.2$	-1.7%
$\epsilon = 0$	-1.0%
Thicker tail of firm productivity ($\eta \downarrow 1\%$)	
$\epsilon = 43.2$	8.1%
$\epsilon = 0$	8.7%

Notes: Table 4 reports the change in consumption resulting from changing key parameters of the model. Part A reports the gains of eliminating distribution spending altogether by setting $A_D = \infty$. In the first row of part (B), we report the gains of setting $\epsilon = 0$. The last four rows compare the effects of changing A_D and η in the baseline model and in a model with $\epsilon = 0$ respectively. Unless otherwise specified, all parameters are set to the level calibrated in Table 3.

The consumption gains from reducing distribution costs arise through several channels. First, improving the efficiency of distribution frees up labor resources to be reallocated to the manufacturing sector. Such gains would also be present in a model without heterogeneous firms or a choice of which markets to sell to. The magnitude of this channel depends on the production structure of the economy, as produced intermediate goods also need to be shipped across space, leading to a compounding of distribution costs. In the Indian data, the materials share is high—59%—and hence a cost that accounts for 4% share of sales can have large aggregate consequences.[30]

Second, in our model, firms sell to more distant markets now that the variable cost of distributing goods is zero; consumers therefore have access to more varieties.[31] This can be seen by the fact that the interstate trade share increases by about 12%. Third, the expansion into more distant markets is most pronounced for the largest firms, which produce the highest-quality goods and spend the most on distributing them. The employment share of the top 10% of firms increases by close to 1%. Consumers, there-

[30] To illustrate this channel, we re-calibrate the model to a 10% higher and lower intermediate share in Appendix B4..

[31] Since $\beta > 0$, there is still a fixed cost of accessing markets and hence an interior solution for the choice of market access.

fore, have access to high-quality goods produced farther away. Finally, the expansion of firms drives wages up in general equilibrium. This increases the productivity threshold at which firms exit the market by about 1.5%.

Heterogeneity in distribution costs. To quantify how much the systematic size dependence of distribution costs matters for aggregate output, we run a second counterfactual in which we set ϵ to zero. In such an economy, the variable cost of selling goods locally (to a market of distance zero) is identical to the cost of selling goods to the furthest location, and all firms have equal distribution shares. We estimate large consumption gains of around 9%. The reduction in ϵ disproportionately benefits large firms, leading them to expand further and pushing out small and unproductive firms.

In addition to its direct effect on consumption by constraining larger firms, a high ϵ also makes the economy more susceptible to low TFP in the distribution sector. To quantify this complementarity, we simulate a 10% reduction in A_D in both the baseline model and a model where $\epsilon = 0$. When all firms spend equally on distribution services ($\epsilon = 0$), a decrease in A_D leads to a 1% loss in consumption, as more resources must be allocated to distribute goods across space. With our calibrated value of ϵ, larger and more productive firms are more affected by the decline in distribution TFP, leading to a more than 50% amplification of output losses.

The presence of size-dependent distribution costs also limits the gains from improving firm quality or productivity. Consider, for instance, a 1% decrease in η, the shape parameter of the Pareto distribution of firm productivity. A lower η leads to both a higher mean and a thicker right tail of firm productivity. The gains from such a productivity shift are around 6% lower in the baseline model with our calibrated ϵ, relative to a model where $\epsilon = 0$. The gains from increasing the set of high-productivity firms are dampened because these firms are disproportionately constrained by distribution costs, thereby limiting their growth potential.

Misallocation and distribution costs The market equilibrium in the model is efficient. The average revenue product (TFPR*), if measured correctly as total revenues divided by total variable costs (including variable distribution costs), is equalized across firms.

As noted in Section 2., firm revenues are typically reported net of distribution expenses in surveys. The measured average revenue product (TFPR) is therefore calculated as revenues *net of distribution costs* divided by inputs *excluding variable distribution costs*. From our model's perspective, this results in spurious TFPR dispersion that could be misinterpreted as reflecting misallocation. Moreover, it is likely to affect the relationship between measured TFPR and firm size. This is a key moment for inferring the elasticity of distortions with respect to firm productivity, which has important implications for misallocation losses (Bento and Restuccia, 2017) and for plant growth over the life cycle (Hsieh and Klenow, 2014). In the calibrated model, the slope of log(TFPR) with respect to log firm size is 0.0053, in contrast to a slope of 0 for log(TFPR*). Incorrectly accounting for distribution costs, therefore, not only generates dispersion in TFPR but also results in a spurious positive relationship between TFPR and firm size.

We estimate how important these issues are in practice in the ASI, constructing both TFPR and TFPR*.[32] In Table 5 we contrast the variances of both measures and their elasticity with respect to plant size. We find that the variance of TFPR* is 5.1% smaller than TFPR, and the slope with respect to firm size is 7% smaller. Given that these stem from mismeasurement in just one variable that accounts typically for 1-4% of sales, this is significant. To put these numbers in perspective, Bils et al. (2021) estimate that all forms of additive measurement error in the ASI (in revenues and inputs) result in a corrected variance of log(TFPR) 30% lower than the uncorrected variance.

Table 5: TFPR Dispersion and Size Dependence

	Variance	Slope on log(Employment)
TFPR	0.0235	0.0111
TFPR*	0.0223	0.0103
% Change TFPR* vs. TFPR	-5.1	-7.0

Notes: TFPR* is revenue productivity, measured as total gross output divided by total costs (labor costs + intermediates + capital costs + distribution costs). TFPR additionally nets out distribution costs from output and total costs. Both are residualized on sector-year fixed effects, and have the 1% tails trimmed. Variance is the cross-sectional variance of the trimmed residual. Slope is from a regression of the residual on log employment, controlling for sector-year fixed effects. Sampling weights used throughout.

5.2. Changes in the distribution sector over time

In this section, we recalibrate the model to the Indian data in 2010. This allows us to estimate the consumption gains from improvements in distribution during a period in which India's transportation infrastructure expanded rapidly. We recalibrate the three key parameters directly affecting the use of distribution services, A_D, ϵ, and β, targeting the 2010 values of the corresponding moments. For the aggregate and average distri-

[32]To be specific, we measure TFPR* as gross revenues divided by total costs (including distribution costs), and TFPR as revenues net of distribution costs divided by total costs excluding distribution costs. Using total costs in the denominator instead of the geometric average of inputs recovers the revenue distortion even in the presence of firm-specific production elasticities, provided that there are no relative input distortions.

bution shares, we target changes adjusted for industry composition, since our model focuses on within-sector changes. We hold all other parameters constant at their 2000 values to focus on the role of changes in distribution costs in the model.

We begin by comparing the two steady states—2000 and 2010. In the next section, we analyze transition dynamics.

As can be seen by comparing Table 3 with Table 6, there were large changes in India's distribution efficiency between 2000 and 2010. First, the aggregate distribution share fell from 3.8% to 2.9%. Second, the ratio of the aggregate to the average distribution share fell from 1.6 to 1.5. Third, the interstate trade share increased from 23% to 32%.

The identification of the three parameters governing distribution costs is similar to that for the baseline year. Through the lens of the model, the changes in the first two moments imply both an increase in the productivity of the distribution sector (higher A_D) and a reduction in the rate at which distribution costs increase with distance (lower ϵ). We infer this reduction from the fact that distribution expenses fell relatively more for larger firms. While a higher A_D and lower ϵ imply an increase in the interstate trade share, we infer from the model that fixed costs of distribution must also have declined substantially in order to quantitatively explain the increase in the interstate trade share. Table 6 reports all calibrated parameters and moments for 2010.

Table 6: Parameters and moments in 2010

Parameter	Value	Moment	Data	Model
A_D	15.6	aggregate distribution share	2.9%	2.9%
ϵ	23.5	aggregate / average distribution share	1.52	1.52
β	34.6	share of sales across states	31.7%	31.7%

Notes: Table 6 reports the calibrated values for the three recalibrated parameters as well as the data moments used for calibration and their model counterpart. See Section 4.1. for a description of the share of sales across state. All other moments are computed using the ASI.

As previously described, one potential driver of the dramatic decline in the costs of distribution we infer is the massive investment in India's road network during the 2000s. India's National Highways Development Project was launched in 2001 to improve the country's underdeveloped highway network. A particular well-known investment was the Golden Quadrilateral highway network (GQ).[33] The GQ was launched in 2001 and completed in 2012, though more than 90% of the project was finished by 2006. It linked together the four distant major economic centers of Delhi, Mumbai, Chennai, and Calcutta. Asturias et al. (2019) find that the GQ led to large income gains in the manufactur-

[33]This was followed by the construction of the North-South and East-West corridor. Around the same time, the Prime Minister's Village Road Program was also begun to improve rural road access.

ing sector, including through improvements in allocative efficiency. From our model's perspective, the GQ affects A_D, ϵ, and β. A decline in the marginal cost of shipping to any location along the GQ can be mapped into an increase in A_D within our framework. The GQ was also particularly important in connecting very distant locations. This reduced the relative cost of shipping longer distances, thereby lowering ϵ. Finally, the GQ may have facilitated the flow of information and made it easier for managers to locate distributors in new markets, thereby reducing β.

Another potential driver of the decline in distribution costs is an increase in competition in the transportation and distribution sectors. Arnold et al. (2016) show that pro-competitive reforms were important in the transportation sector from 1993 to 2005 and led to productivity gains for manufacturing firms. Increases in competition in the distribution sector could lead to higher productivity and lower markups, both of which would translate into higher sectoral productivity A_D within our framework. To the extent that such increases in competition particularly affected long-distance transport and marketing, they could also have led to decreases in ϵ and β.

Aggregate Output Effects. We find that the combined improvements in India's distribution sector led to an increase in manufacturing consumption of close to 25%. That is, in 2010, consumption of the representative household that saw changes in the three parameters governing distribution (ϵ, β, A_D) would have been 25% higher than in 2000.

Table 7: Consumption Gains From Improvements in Distribution Sector

All three	A_D only	ϵ only	β only
24.5%	4.6%	3.8%	13.9%

Notes: Table 7 reports the aggregate consumption gains of the calibrated changes in distribution costs. Column 1 sets all three parameters to their 2010 values; columns 2-4 move only one parameter at a time, keeping all others at their calibrated 2000 values.

We then decompose the gains into the contribution of each of the three components. We do so by moving one parameter at a time to their 2010 value while leaving all other parameters at their 2000 values. Two key results emerge.

First, the fixed cost of accessing more distant markets played an important role. The reduction in β alone (the rate at which fixed cost increases with distance) leads to consumption gains of 14%. The main data moments that allowed us to infer the reduction in β are the increase in interstate trade over this decade. The improvement in distribution sector TF,P as well as the reduction in the rate at which *variable* costs increase with distance, also had sizeable aggregate effects on the order of 5% and 4%, respectively.

In Section 2., we found suggestive evidence that variable distribution costs declined more in states where transportation infrastructure improved as a result of the Golden Quadrilateral highway network. Our findings that aggregate consumption increased substantially as a result of declining distribution costs therefore point towards a potentially important role of infrastructure investments and other government policies for promoting the development of the manufacturing sector.

5.3. Dynamics of adjustment and incidence of gains

In reality, it takes firms time to find distributors to expand into new markets. Similarly, new entrants may need time to set up their business. To analyze the dynamics of improvements in distribution sector productivity, we simulate a modified economy in which the transition to the new steady state takes time. Starting from the steady state in the 2000s, we implement the one-time change in $\{A_d, \epsilon, \beta\}$ identified in Table 6. In each period t, a firm meets a distributor with probability ξ and can, from now on, choose N^* as well as its export status optimally. A firm that has not yet met a distributor, which happens with probability $(1-\xi)^t$ may only adjust downwards. In particular, it may stop exporting, sell to fewer markets, or exit, but not access new markets.

Similarly, we allow the mass of potential entrants to increase only gradually to the new level prevailing in the 2010 steady state. In the transition period, it adjusts by the same increment ξ each period. That is, the fraction of firms that draw a quality level ψ each period and then decide whether to remain active is a weighted average of the entry rates in the old and new states, with the weight on the new value increasing by ξ each year. Any aggregate profits present during the transition period are rebated to households.

For this section, we set $\xi = 0.1$, implying it takes on average 10 years for a firm to get the chance to expand to its optimal number of markets. In Appendix B2., we show results for a lower ($\xi = 0.05$) and higher ($\xi = 0.2$) speed of adjustment.

Figure 4 plots the evolution of manufacturing consumption (top panel) as well as the mass of active firms (bottom panel) in the transition. Even though firms may expand to new markets with only a 10% probability each year, aggregate consumption transitions to the steady state quite rapidly. In the first year, 62.3% of the increase in aggregate consumption has materialized. Note that in this first year, no firm has been able to update the number of markets it sells to. Hence, the year 1 response quantifies the direct effect of lower distribution costs, holding distance shipped fixed. Ten years after the reform, aggregate consumption reached 88% of its final level.

When the probability of meeting a new distributor is higher ($\xi = 0.2$), the transition

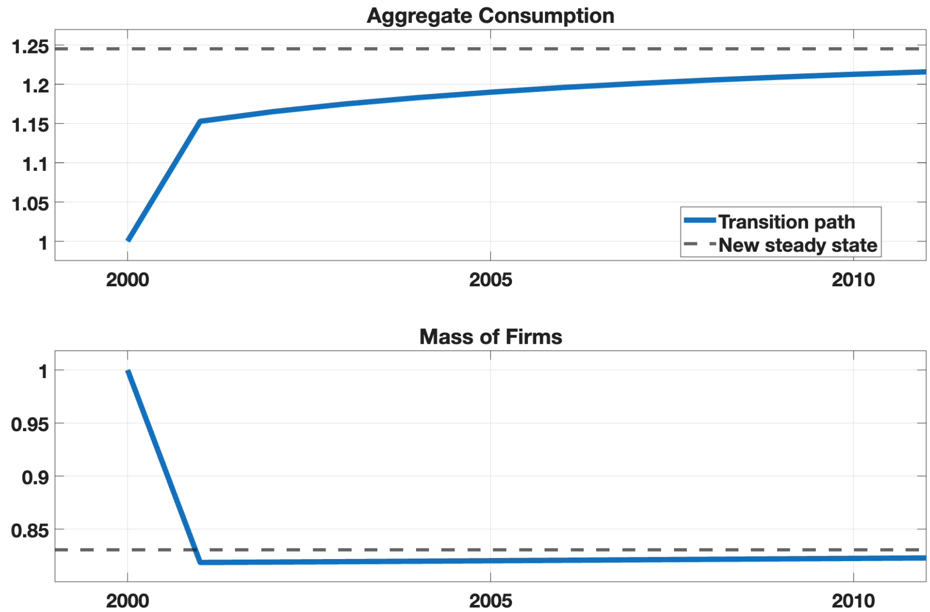

Figure 4: Consumption and the Mass of Firms during the Transition

Notes: The top panel plots the % change in aggregate consumption relative to the initial steady state in 2000, while the bottom panel plots the % change in the mass of firms. In both panels, the dashed horizontal line corresponds to the new steady state, while the solid blue line corresponds to the transition path for $\xi = 0.1$.

is even faster, reaching 94.4% of the new steady state value by 2010. Even when the probability that firms are allowed to expand into new markets is as low as 5% per year, over 80% of the gains in aggregate consumption have materialized within the first 10 years.

The reason why the transition is so rapid is two-fold. First, firms immediately benefit from the reduction in distribution costs - both variable and fixed - in all markets they serve in 2000. Even though firms cannot access new locations within India or start exporting, they do expand sales to their existing customers. Second, the expansion of sales puts upward pressure on the wage, which in turn triggers the exit of unproductive firms. In fact, in the first period, the exit rate is nearly 25%.

The lower panel of Figure 4 illustrates the exit and entry dynamics. Initially, high exit rates among unproductive firms reduce the mass of firms below the 2010 steady-state level. Gradually, the mass of firms increases thereafter. This process takes time, since entry is becoming increasingly selective. Faced with a higher wage, firms must be more productive to justify paying the entry cost.

Selection of Entrants and Sales Concentration in the data We now turn to check whether relevant moments from the data behave as predicted by our model. Of course, this evidence is only suggestive of the validity of the mechanisms in our model, as many features of Indian manufacturing changed from the late 1990s to the early 2010s.

Because the ASI is a survey with sampling weights, it is not possible to construct a good measure of exit from the data. We can, however, identify entrants based on the reported age of the plant, and group together all plants under the age of 5. We look at how the selection of entrants changed across three-year windows by plotting their average employment relative to 1999-2001 (the left panel of Figure 5). As predicted by the model, the average size of entrants increased gradually over time (more so than the average plant size of incumbents), consistent with increasingly selective entry.

Another telling moment from the data is the evolution of sales concentration. Recall that sales concentration in 2000 was used as a moment to estimate the model. We plot in the right panel of Figure 5 the sales share of the largest 10% of firms by 3-year window. Sales concentration increased by around 2 percentage points from 1999-2001 to 2002-2004, and remained at this higher level for the next decade. The rising sales concentration of larger firms is again consistent with the exit of plants and a shrinking concentration of sales in the left tail of the distribution, with larger firms in particular benefiting from the reduction in distribution costs.

Figure 5: Entrants and Sales Concentration in the Data

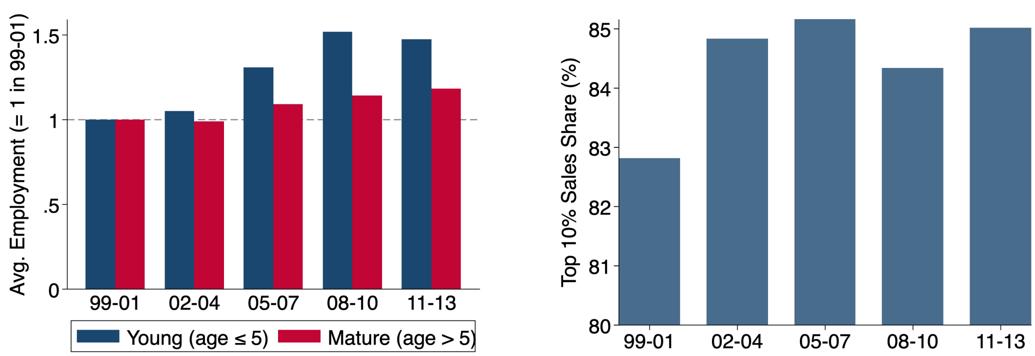

Average Size of Entrants | Top 10% Sales Share

Notes: Data source is the Indian Annual Survey of Industries 1999 to 2013. The left figure shows average employment in young (under or equal to 5 years) and mature (over 5 years) firms by 3-year window, normalized relative to the 1999-2001 window. The right figure shows the share of sales of the largest 10% of firms by 3-year window.

Winners and losers The gains from improvements in India's distribution sector were not equally distributed across existing firms but concentrated among those that were large to begin with. Figure 6 plots the change in variable profits against their value in 2000.

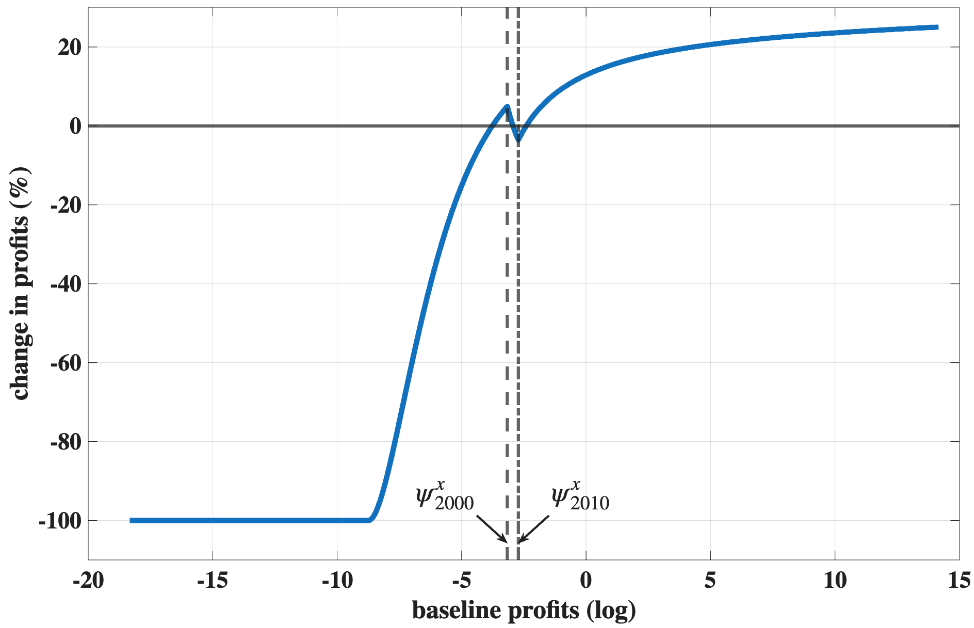

Figure 6: Changes in Profits between 2000 and 2010

Notes: Figure 6 plots the % change in firm profits between the old and the new steady states against the firms' (log) profits in the old steady state. The two vertical lines, labeled ψ^x_{2000} and ψ^x_{2010}, show the profit of the firm that is indifferent between exporting and not in the two steady states.

The bottom 25% of firms exited altogether, and 79% of firms active in 2000 saw their profits *decrease* in the new steady state. These are the small, low-productivity firms that serve mostly local markets. For them, high transportation costs served as a shield against competition from more productive and larger businesses. With a more integrated domestic market, they either shrink or go out of business.

The fact that most firms saw their bottom lines worsen after reductions in distribution costs may offer a glimpse into the political economy of public infrastructure. When many firms (by number) benefit from being shielded by competition, any measure that improves market access of productive firms may struggle to find a majority, even when the aggregate gains from such a policy would be large.

Improvements in India's domestic economy also lead to a deterioration of the terms of trade. This is most visible in the middle of the distribution, where some firms used to export but no longer find it profitable to do so in the new steady state.

5.4. Extension: Multi-Plant Firms

In the baseline model, firms are restricted to operating in a single location. This modeling choice is largely driven by data limitations: the unit of analysis in the ASI is a plant, and plants of the same firm cannot be linked. If firms open new plants to help mitigate high distribution costs, our model in the previous section and the counterfactuals we run risk overstating the importance of variable distribution costs.

To gauge the potential importance of this channel, we calibrate an extended version of the model (see Appendix B3. for the detailed setup). Upon entering and discovering their productivity, firms may open a second location along the circle for an additional fixed cost. We calibrate this fixed cost to match the fact that multi-plant firms account for 46% of sales in the data. In the model as well as the data, large and productive firms are the ones that choose to open a second plant.

If firms can overcome large distribution costs by opening a plant in a second location, the detrimental effects of low TFP in the distribution sector are naturally smaller. However, quantitatively, the two models yield very similar results. Eliminating variable distribution costs raises manufacturing consumption by 20.6%, compared to 21.3% in the baseline model, and setting $\epsilon = 0$ raises manufacturing consumption by 8.3%, compared to 9.2% in the baseline. Finally, applying the 2000–2010 changes in distribution parameters increases manufacturing consumption by 24.5%, which is essentially identical to the baseline result.

We therefore conclude that our baseline model of one-location plants is an appropriate modeling assumption for the focus of this paper. A full treatment of endogenous location choice, which is beyond the scope of our data, remains a very promising area for future research.

6. Conclusion

In this paper, we examine the importance of a well-functioning distribution sector for manufacturing firms and aggregate output. Using a unique measurement of outward distribution expenses for Indian manufacturing plants, we document that distribution expenses are not only large, but also – as a share of sales – over three times larger for big relative to small plants. These empirical findings suggest that a poorly functioning distribution sector can easily become a bottleneck: as firms grow, they require more and more distribution services, and marginal costs rise.

To quantify the importance of this new finding, we build a model of heterogeneous manufacturing firms that purchase inputs from the distribution sector in order to sell their goods across space. We pin down the rate at which variable shipping costs increase with distance by matching the observed relationship between plant distribution shares and size. Accounting for differential exposure of manufacturing firms to the distribution sector is important: aggregate consumption losses from low TFP in distribution are 1.7

times larger than in a world with equal distribution shares across firms.

We then evaluate the consolidated consumption effects of improvements in India's distribution sector between 2000 and 2010. Rather than analyzing individual policies, we infer changes in fixed and variable costs of accessing markets from three trends: the aggregate distribution share declined from 3.8% to 2.9%, the relationship between size and the distribution share flattened, and intranational trade increased from 23% to 32%. Overall, this led to consumption gains of 24.5% over the ten years. While these gains materialized quickly, they were unevenly distributed. A quarter of plants are predicted to exit in the first few years alone: these are small, low-productivity plants that were previously shielded from competition by large intranational trade costs.

Our findings reinforce the importance of policies affecting competition in the distribution and transportation sectors in developing countries. As of 2020, India has the second-highest restrictions on foreign entry in distribution services among OECD countries and among the highest barriers to entry in the transport and distribution sectors (OECD, 2020). An additional implication of our findings is that such frictions may dampen firms' incentives to innovate or invest in cost-reducing measures. The private benefits of such investments are substantially reduced if the firm's growth potential is constrained by high costs of expanding into new domestic markets. This could therefore contribute to the weak life-cycle growth of plants in India and other developing markets documented in Hsieh and Klenow (2014).

References

Akcigit, Ufuk, Harun Alp, and Michael Peters, "Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries," *American Economic Review*, 2021, *111* (1), 231–75.

Allen, Treb, "Information Frictions in Trade," *Econometrica*, 2014, *82* (6), 2041–2083.

Arkolakis, Costas, "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of Political Economy*, 2010, *118* (6), 1151–1199.

Arnold, Jens Matthias, Beata Javorcik, Molly Lipscomb, and Aaditya Mattoo, "Services Reform and Manufacturing Performance: Evidence from India," *Economic Journal*, 2016, *126* (590), 1–39.

Asher, Sam and Paul Novosad, "Rural Roads and Local Economic Development," *American Economic Review*, 2020, *110* (3), 797–823.

Asturias, Jose, Manuel Garcia-Santana, and Roberto Ramos, "Competition and the welfare gains from transportation infrastructure: Evidence from the Golden Quadrilateral of India," *Journal of the European Economic Association*, 2019, *17* (6), 1881–1940.

Atkin, David and Dave Donaldson, "The Size and Implications of Intra-national Trade Costs," July 2015. NBER Working Paper No. 21439.

Bartelme, Dominick and Yuriy Gorodnichenko, "Linkages and Economic Development," March 2023. Working paper.

Bento, Pedro and Diego Restuccia, "Misallocation, Establishment Size, and Productivity," *American Economic Journal: Macroeconomics*, 2017.

Bernard, Andrew and J. Bradford Jensen, "Exporters, Jobs and Wages in U.S. Manufacturing, 1976-87," *Brookings Papers on Economic Activity: Microeconomics*, 1995.

Bils, Mark, Peter J. Klenow, and Cian Ruane, "Misallocation or Mismeasurement?," *Journal of Monetary Economics*, 2021, *124* (S), S39–S56.

Carrillo, Paul, Dave Donaldson, Dina Pomeranz, and Monica Singhal, "Misallocation in Firm Production: A Nonparametric Analysis Using Procurement Lotteries," November 2023. Working paper.

Das, Sanghamitra, Mark J. Roberts, and James R. Tybout, "Market Entry Costs, Producer Heterogeneity, and Export Dynamics," *Econometrica*, 2007, *75* (3), 837–873.

Donaldson, Dave, ""Railroads of the Raj: Estimating the Impact of Transportation Infrastructure.," *American Economic Review*, 2018, *108* (4-5), 899–934.

Fadinger, Harald, Christian Ghiglino, and Mariya Teteryatnikova, "Income Differences, Productivity, and Input-Output Networks," *American Economic Journal: Macroeconomics*, 2022, *14* (2), 367–415.

Ghani, Ejaz, Arti Grover Goswami, and William R. Kerr, "Highway to Success: The Impact of the Golden Quadrilateral Project for the Location and Performance of Indian Manufacturing," *Economic Journal*, 2016, *126* (591), 317–357.

Gollin, Douglas and Christopher Udry, "Heterogeneity, Measurement Error, and Misallocation: Evidence from African Agriculture," *Journal of Political Economy*, 2021, *129* (1), 1–80.

Holmes, Thomas J. and J. J. Stevens, "Exports, Borders, Distance, and Plant Size," *Journal of International Economics,*, 2012, *88* (1), 91–103.

Hsieh, Chang-Tai and Ben Olken, """The Missing Missing Middle," *Journal of Economic Perspectives*, 2014.

_ and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 2009, *124* (4), 1403–1448.

_ and _ , "The Life Cycle of Plants in India and Mexico," *Quarterly Journal of Economics*, 2014, *129* (3), 1035–1084.

Jain, Anil K. and Siddharth Kothari, "Why are Manufacturing Plants Smaller in Developing Countries? Theory and Evidence from India," International Finance Discussion Paper 1417, Board of Governors of the Federal Reserve System August 2025.

Jones, Charles I., "Intermediate Goods, Weak Links, And Superstars: A Theory of Economic Development," *American Economic Journal: Macroeconomics*, 2011, *3* (2), 1–28.

Kim, Hang, Martin Rotemberg, and T. Kirk White, "Manufacturing Dispersion: How Data Cleaning Choices Affect Measured Misallocation and Productivity Growth in the Annual Survey of Manufactures," Center for Economic Studies Working Paper CES-25-67, U.S. Census Bureau September 2025.

Leemput, Eva Van, "A passage to India: Quantifying internal and external barriers to trade," *Journal of International Economics*, 2021, *131.*

Liu, Ernest, "Industrial Policies in Production Networks," *Quarterly Journal of Economics*, 2019, *134* (4), 1883–1948.

Melitz, Mark, "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.," *Econometrica*, 2003, *71*, 1695–1725.

Ministry of Finance, "A New, Exciting Birds-Eye View of the Indian Economy Through the GST," 2018. Chapter 2, Economic Survey 2017-18, Economic Division, Ministry of Finance, Government of India.

OECD, "OECD Services Trade Restrictiveness Index: Policy trends up to 2020," 2020.

Peter, Alessandra and Cian Ruane, "Intermediate Input Substitutability," 2025.

Peters, Michael, "Heterogeneous Mark-Ups and Endogenous Misallocation," *Econometrica*, 2020, *88* (5).

World Bank, "Manufacturing, value added (% of GDP) – India," World Development Indicators, series NV.IND.MANF.ZS 2026. Accessed May 5, 2026.

—, "World Development Indicators: GDP growth (annual %) – India," https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?locations=IN 2026. Accessed April 27, 2026.

A Data Appendix

A1. Annual Survey of Industries

We use the Indian Annual Survey of Industries (ASI) for the years 1993 to 2013. The microdata is publicly available through India's Ministry of Statistics and Programme Implementation (MOSPI). The reference period of the survey is the accounting year, which in India begins on the 1st of April and ends on the 31st of March the following year. We reference the surveys by the earlier of the two years covered. The ASI is a representative sample of plants with at least 10 workers, though the threshold is 20 workers for plants that don't use power. ASI sampled plants fall into two schemes: Census and Sample. Census plants are surveyed every year. They consist of plants with more than 100 workers and plants in 12 of the industrially 'backwards' states.[34] The remaining plants fall into the Sample scheme and are sampled at random within state 3-digit industry category. Sampling weights are provided with the data.

We construct plant output as the sum of the gross value of product sales, changes in inventories and other sources of revenue. Employment includes both paid and unpaid labor. Labor costs include wages and salaries, bonuses, contributions to the firms' provident (pension) fund, and other welfare expenses. The capital stock is measured as the average of beginning and end of year book value of the net fixed capital stock. Intermediate inputs include materials and fuels consumed as well as other expenditures. We construct a harmonized sectoral classification consisting of 50 manufacturing sectors that are consistently defined throughout our time period.[35]

A2. Plants with Zero Distribution Expenses

In this appendix we provide some additional details on plants that report zero distribution costs.

On average over the sample period, 25% of plants report zero distribution expenses, but this share rises from 22 percent to around 27 percent over time (Figure 7 (a)). Small plants are far more likely to report zero distribution expenses (Figure 7 (b)). Among plants with 1-9 employees, almost half report zero distribution expenses. In contrast, for plants with over 250 employees, this share falls to below 10%.

While the aggregate and average distribution shares are mechanically higher when plants with zero distribution costs are dropped, Figure 8 shows that our two main em-

[34] The size threshold for plants to be included in the Census sample is 200 in some years during the 1990s.

[35] The official sectoral classification (NIC) changed in 1998, 2004 and 2008. We use official NIC concordances to construct our harmonized classification.

Figure 7: Firms with Zero Distribution Costs

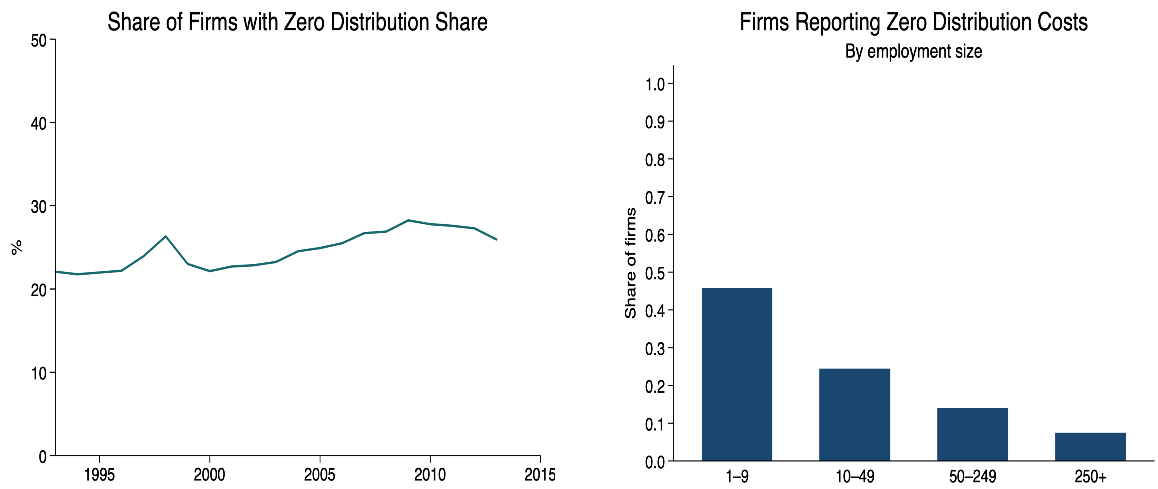

(a) Over Time (b) By Firm Size Category

Notes: Data source is the Indian Annual Survey of Industries 1993 to 2013. The left figure shows the share of firms over time whose reported distribution costs are zero. The right figure shows the share of firms by employment bin reporting zero distribution costs, pooling all years.

pirical findings, 1) the decline in the aggregate distribution share, and 2) the increasing plant-level relationship between the distribution share and size, still look very similar even with this sampling restriction.

Table 8 reports some additional summary statistics comparing plants that report strictly positive vs. zero distribution costs. While plants with zero distribution costs are smaller, they also tend to have slightly higher labor and capital shares, and slightly lower material and fuel shares, than firms reporting positive distribution costs. Table 9 additionally lists the sectors with the highest and lowest share of plants reporting zero distribution costs. We do not see any systematic patterns that lead us to be concerned about how we interpret this measure of distribution costs.

A3. The Golden Quadrilateral Highway Network

In this appendix we do a simple exercise to see whether trends in the aggregate distribution share and the relationship between the plant-level distribution share and size are correlated with a measurable large improvement in infrastructure in India: the Golden Quadrilateral highway network.

India's National Highways Development Project was launched in 2001 with the goal of improving the country's underdeveloped highway network. One of the main invest-

Figure 8: Main Findings for Firms with Strictly Positive Distribution Costs

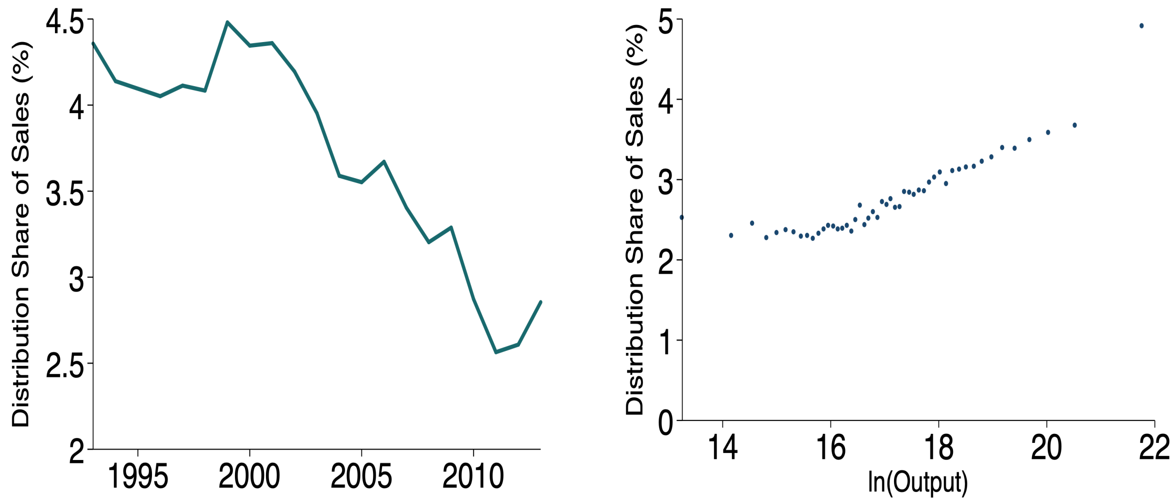

(a) Aggregate Share (b) Relationship with Size

Notes: Data source is the Indian Annual Survey of Industries 1993 to 2013. The left figure shows the aggregate distribution share constructed for the sample of plants with strictly positive distribution costs. The right figure shows a binned scatter plot of the plant-level distribution share of sales against ln(output), residualized on sector-year fixed effects, for plants with strictly positive distribution costs.

Table 8: Summary Statistics for Plants with Zero Distribution Costs

	Distribution Costs >0	Distribution Costs =0
Labor	96.12	38.49
Cost Shares		
Labor	0.106	0.157
Materials+Fuels	0.736	0.688
Capital	0.059	0.078

Notes: The table reports averages of several variables for plants with zero distribution costs and those with strictly positive distribution costs. Within-year averages are taken using sampling weights, which are then averaged across years. Labor is measured in number of workers. The labor, material+fuels and capital share are all measured as a share of total costs (the residual is non-material and fuel intermediate inputs).

Table 9: Sector Shares of Plants, Output, and Labor with Zero Distribution Costs

Sector	Plants	Output	Labor
Bottom 10 Sectors by Share of Plants			
Wearing apparel	0.073	0.022	0.039
Knitted and crocheted fabrics	0.076	0.021	0.035
Office, accounting and computing machinery	0.087	0.072	0.056
Medical appliances and instruments	0.099	0.085	0.069
Man-made fibers	0.103	0.012	0.023
Insulated wire and cable	0.109	0.039	0.061
Railway and tramway locomotives	0.115	0.184	0.209
Electric motors, generators and transformers	0.119	0.062	0.067
Electronic valves, tubes and components	0.123	0.073	0.067
Transport equipment n.e.c.	0.129	0.036	0.059
Top 10 Sectors by Share of Plants			
Basic iron and steel	0.278	0.073	0.095
Refined petroleum products	0.284	0.495	0.393
Glass and glass products	0.284	0.045	0.127
Furniture	0.321	0.043	0.102
Beverages	0.359	0.081	0.102
Tobacco products	0.460	0.316	0.188
Wood, cork, straw and plaiting materials	0.473	0.160	0.271
Coke oven products	0.494	0.500	0.537
Sawmilling and planing of wood	0.505	0.305	0.447
Non-metallic mineral products n.e.c.	0.549	0.080	0.340

Notes: The table reports by sector the share of firms, output and labor with zero distribution costs. Only the 10 sectors with the highest and lowest share of plants are included.

ments was the Golden Quadrilateral highway network (GQ). The GQ was launched in 2001 and completed in 2012, though more than 90% of the project was finished by 2006. It linked together the four distant major economic centers of Delhi, Mumbai, Chennai, and Calcutta, and linked together roughly half of India's states. Asturias et al. (2019) find that the GQ led to large income gains in the manufacturing sector through a reduction in transport costs.

We first consider whether changes in the aggregate distribution share were larger in states through which the GQ passed than other states. In the top left panel of Figure 9, we show the state by state levels of the aggregate distribution share for a time window of 1999-2000 and 2009-2010, while in the top right panel we plot the changes between the two windows.[36] We separate states in which the GQ was built and states through which the GQ did not pass. It is clear that state distribution shares fell more on average in states through which the GQ passed than states through which it did not pass. The change averaged -1.2 percent for GQ states, and -0.7 percent for non-GQ states.[37]

Similarly, the bottom left panel shows the ratio of the aggregate to average distribution shares state-by-state separately for the two windows, while the bottom right panel shows how these ratios have changed over time. Once again, we find that states through which the GQ passed saw large declines in this ratio than states through which the GQ did not pass. For the average GQ state, the aggregate to average distribution share ratio was 14% lower in 09-10 than 99-00, while it was 11% higher for non-GQ states.

While these results are only suggestive of the construction of the GQ having a measurable impact on the distribution expenses of Indian manufacturing plants, they are consistent with government policies and infrastructure investments playing a role in the distribution share trends we document in the paper.

A4. Interstate Trade

We obtain data on interstate trade from the Indian Directorate General of Commercial Intelligence and Statistics (DGCIS). DGCIS collects and publishes historical data on interstate goods flows by rail, river and air in the "Inter-State Movements / Flows of Goods by Rail, River and Air". Data are collected by DGCIS for tax purposes based on invoices of commodity consignments. These are reported for 70 different goods categories, including both agricultural commodities and manufactured goods. Shipments between

[36] We average 1999 and 2000 as well as 2009-2010 to minimize the role of transitory changes in distribution shares. Our results are very similar for different choices of windows. We drop the states Sikkam, Telangan and Meghalaya as these are very small states with very few industrial plants.

[37] We find a smaller albeit still substantial difference when we control for differences in industry composition across states.

Figure 9: The GQ and Changes in Distribution Shares

(a) State-level Distribution Shares

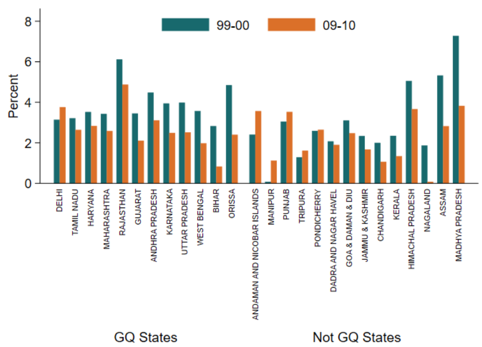

(b) 09-10 – 99-00 Distribution Share

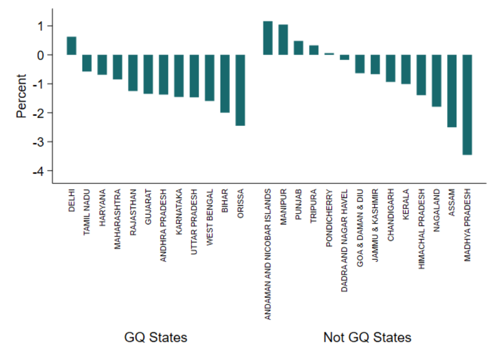

(c) Ratio of Aggregate / Average Distribution Shares

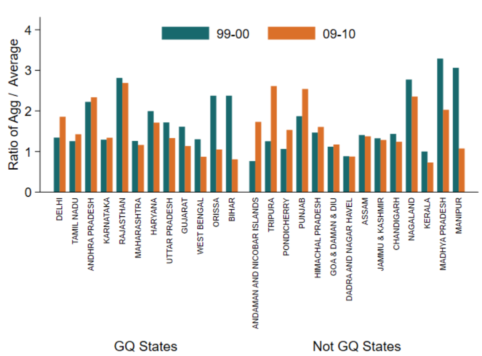

(d) 09-10 / 99-00 Ratio

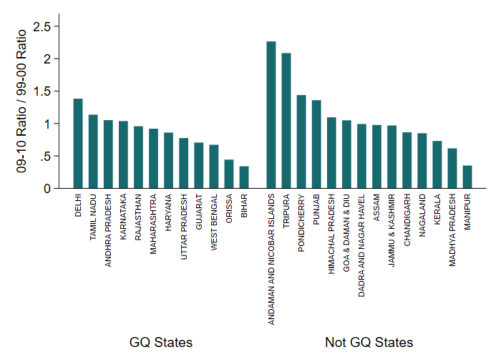

Source: Indian Annual Survey of Industries 1999,2000,2009 and 2010. The top left figure shows the state-level aggregate distribution share, separately for an average of 1999-2000 and 2009-2010. The top right panel plots the corresponding change in the aggregate distribution share state-by-state. The bottom left panel plots the state-level ratio of the aggregate to the (unweighted) average distribution share, separately for an average of 1999-2000 and 2009-2010. The bottom right panel plots the ratio of the 09-10 to the 99-00 ratios.

all Indian states are reported, with separate entries for shipments for international export.[38] For example, shipments from Delhi to 'Other ports of Karnataka' are reported, as well as shipments from 'Karnataka (excluding ports)' to 'Other ports of Karnataka'. This is the only dataset of which we are aware which contains historical information on interstate trade in India.[39]

As described in section 4.1., the approach to constructing the interstate trade share (adjusted for shipments by road) is possible because we have information both on shipments for export and for domestic consumption. For each commodity, we first construct:

$$\frac{\text{quantity exported by rail}}{\text{quantity shipped interstate by rail}} \tag{19}$$

We then make the following assumptions: 1) for each commodity, the share of shipments by rail is the same for exports and domestic consumption, 2) the quality of goods shipped interstate by rail is the same for export and for domestic consumption. It follows that:

$$\frac{\text{quantity exported by rail}}{\text{quantity shipped interstate by rail}} = \frac{\text{aggregate value exported}}{\text{aggregate value shipped interstate}} \tag{20}$$

For each commodity, we then construct it's aggregate export intensity, given by:

$$\text{export intensity} \equiv \frac{\text{aggregate value exported}}{\text{aggregate value of domestic production}} \tag{21}$$

For each commodity, we then obtain the interstate trade share by dividing Equation 21 by 20:

$$\text{interstate trade share} = \frac{\frac{\text{aggregate value exported}}{\text{aggregate value of domestic production}}}{\frac{\text{aggregate value exported}}{\text{aggregate value shipped interstate}}} \tag{22}$$

With appropriate data, the steps described above can, in principle, be implemented for each of the 70 good categories in the DGCIS data. In practice, however, we can only use a much more limited set of goods. The first issue is that we can only implement this ap-

[38]Exports are reported for the main port states of Andhra, Gujarat, Karnataka, Kerala, Maharashtra, Tamil Nadu and West Bengal. The main state excluded is Delhi, which accounts for 8% of Indian exports (Van Leemput, 2021). Given this is omitted in both 2000 and 2010, we don't expect it to bias our estimates of the change in the interstate trade share.

[39]Van Leemput (2021) is the first paper we have seen to use this data to calculate India's interstate trade share. Their focus is on measuring the level of the interstate trade share in 2011-2012, while we focus on the change from 1999-2000 to 2009-2010.

proach for commodities that were exported in both 2000 and 2010. India's trade liberalization was still underway in 2000, especially for agricultural commodities. Restricting to commodities for which international exports are positive reduces our sample to 16 goods. The second issue is that we need comparable data on aggregate export values and aggregate production for each commodity. We obtain this from a variety of sources, including the ASI, Indiastat and CEIC.[40] DGCIS uses its own commodity classification which we concord manually to other classifications for which we can find data on aggregate export values and domestic production values. Finally, we do not use the interstate trade shares for some goods because they are either greater than 100% or change unrealistically between 2000 and 2010.[41] These issues could be due to measurement error in the data or differences in the quality of exports vs. domestic shipments. After accounting for the limited concordance and other measurement issues, we end up with a set of 5 agricultural and manufacturing good categories for which we measure the interstate trade share in both 2000 and 2010: coal and coke, metal products, pulses, sugar and mollasses, and iron and steel. We obtain our data moments by taking the average interstate trade share across these 5 categories in each year.

[40]We scale up output from the ASI by 3/2 to account for informal production.

[41]For example, the interstate trade share for 'Other grains' increases from 1.1% in 2000 to 96.5% in 2010.

A5. Figures

Figure 10: Histogram of Distribution Share of Sales

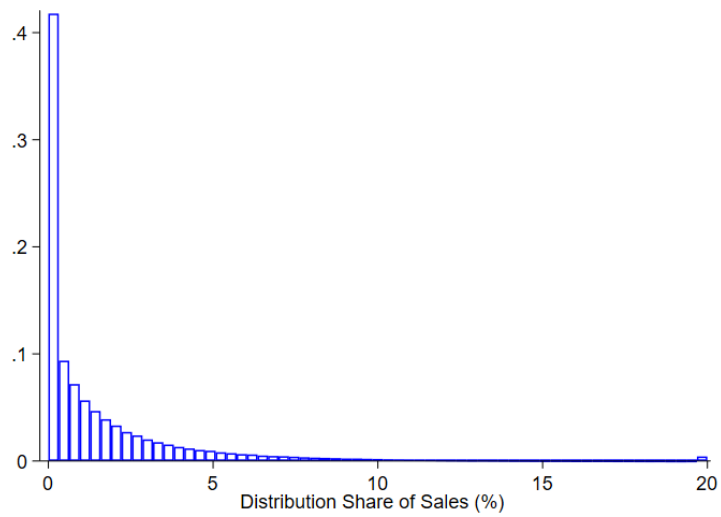

Source: Indian Annual Survey of Industries 1993 to 2013. The figure shows a histogram of the distribution share of sales in India, winsorized at 20%.

Figure 11: Distribution Share vs Plant Size: Robustness

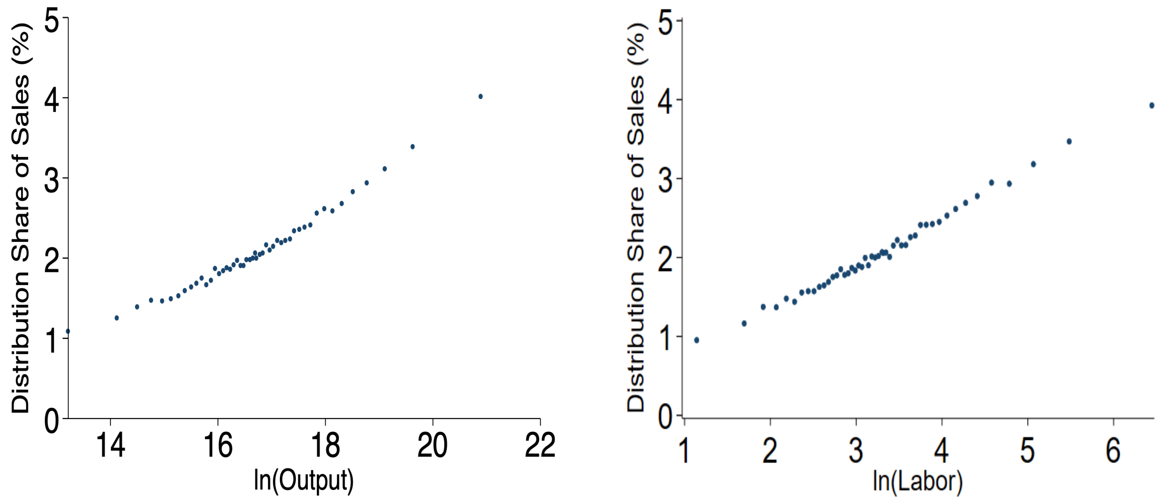

(a) vs. Output (b) vs. Labor

Source: Indian Annual Survey of Industries 1993 to 2013. Both subfigures show binned scatter plots of the plant-level distribution share against plant size, after residualizing on industry×year×state fixed effects. We use the most detailed industry classifications available in each year, 4-digit NIC87, 5-digit NIC98, 5-digit NIC04 and 5-digit NIC08. There are between 400 and 700 manufacturing industries at this level of disaggregation. Subfigure (a) uses output as a measure of plant size, and subfigure (b) uses labor.

Figure 12: Distribution Share of Sales for Non-Exporters

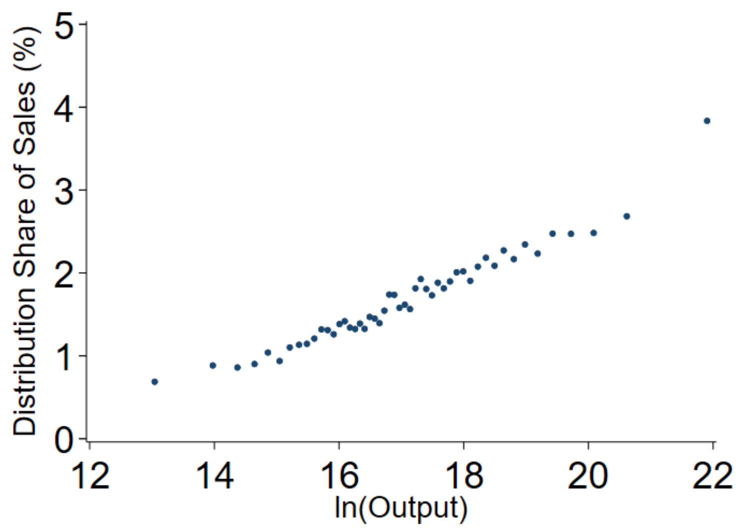

Source: Indian Annual Survey of Industries 1997, 2008, 2009, 2011, 2012 and 2013. The figure shows a binned scatter plot of the plant-level distribution share against output for plants that report no exports. We residualize both variables on sector×year fixed effects.

Figure 13: Inward Distribution Share vs. Plant Size

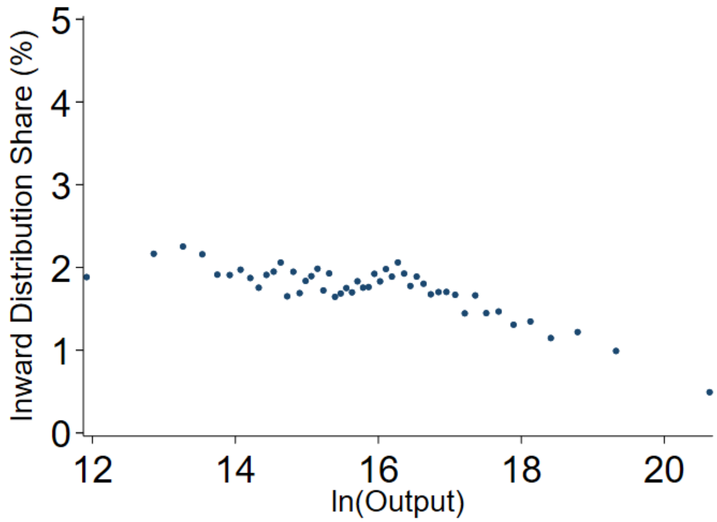

Source: Indian Annual Survey of Industries 1993, 1994 and 1996. Note: the figure shows a binned scatter plot of plants' inward distribution share of materials and fuels against log(output). We residualize both variables on sector×year fixed effects. The ASI survey years used are 1993, 1994 and 1996 as these are the only years in which inward distribution costs are reported. They include expenses incurred for acquiring and transporting production materials.

Figure 14: Distribution Share of Total Costs vs. Plant Size

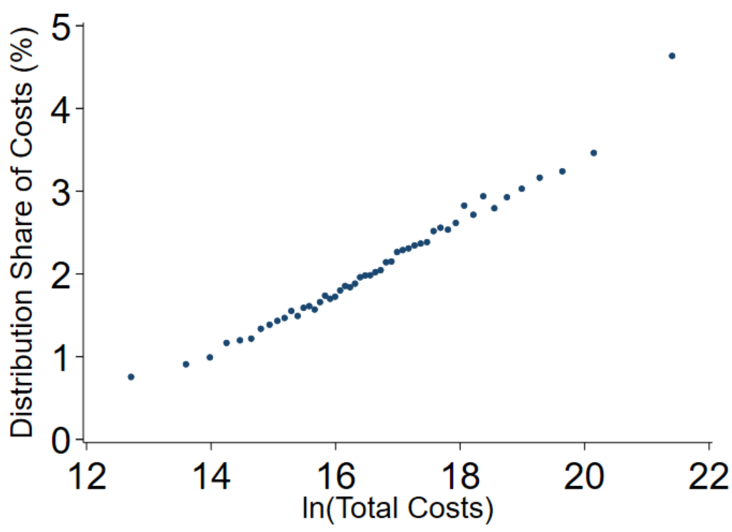

Source: Indian Annual Survey of Industries 1993-2013. The figure shows a binned scatter plot of the plant-level distribution share of total costs against output. We residualize both variables on sector×year fixed effects. Total costs include intermediate expenditures, labor costs and the imputed rental on capital assuming a rental rate of 20%.

B Model Appendix

B1. Identification of Model Parameters

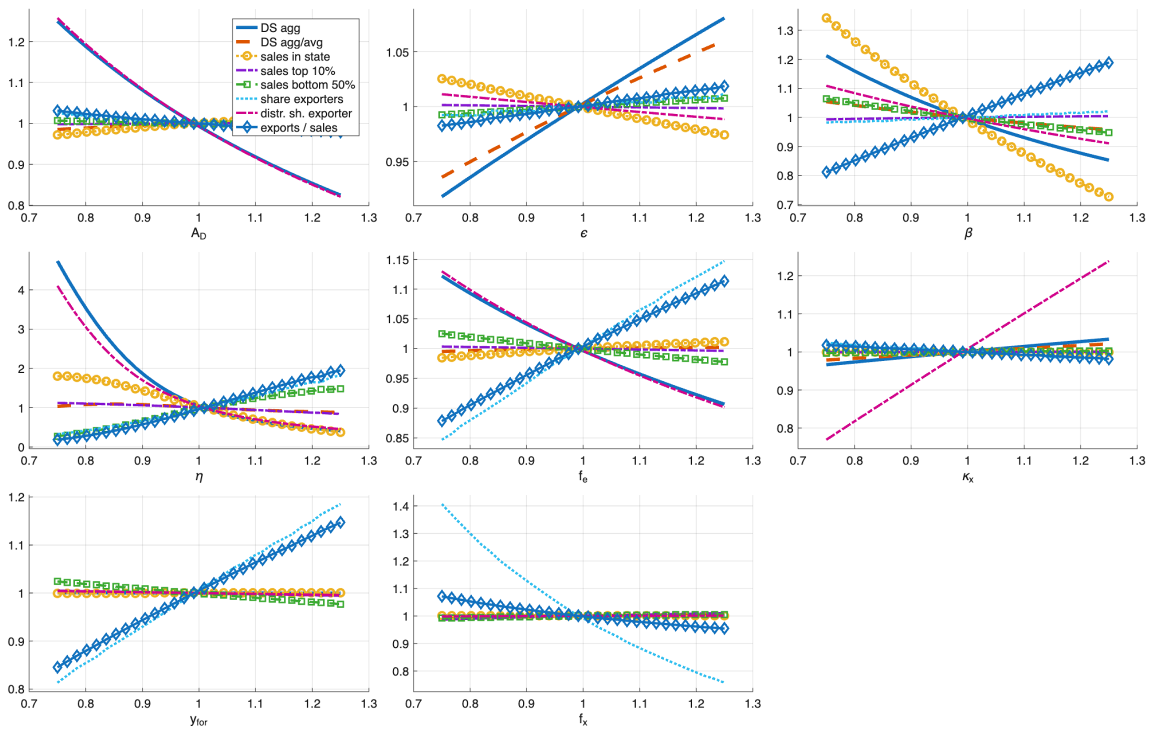

Figure 15: Identification 2000: all parameters

Note: Figure 15 shows sensitivity of model moments to parameters. Starting from the baseline calibration, each panel plots all eight moments relative to their baseline value as a function of one parameter, keeping all other parameters constant at their baseline value. The change in each parameter is also measured relative to baseline.

B2. Gradual adjustment over time for alternative values of ξ

In this section, we replicate Figure 4 for higher and lower values of xi ($\xi = 0.2$ and $x = 0.05$). The remainder of the model is identical to the one laid out in Section 5.3.

B3. Multi-Plant Extension

Consider the model laid out in Section 3. with one additional feature. After paying the entry cost and discovering their idiosyncratic quality ψ, firms have the option to open a second location. This second location operates with the same quality level. We assume

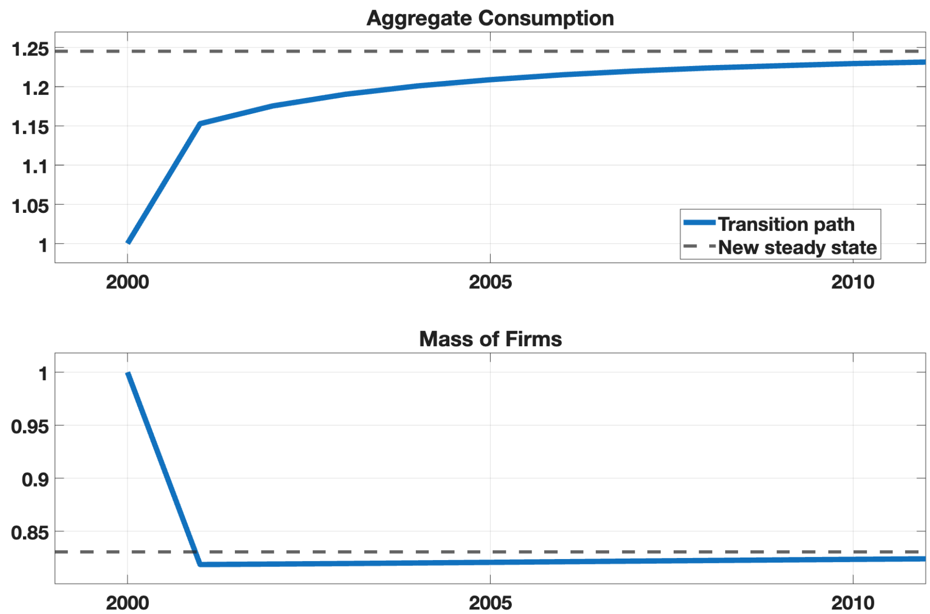

Figure 16: Consumption and the Mass of Firms during the Transition

Notes: The top panel plots the % change in aggregate consumption relative to the initial steady state in 2000, while the bottom panel plots the % change in the mass of firms. In both panels, the dashed horizontal line corresponds to the new steady state, while the solid blue line corresponds to the transition path for $\xi = 0.2$.

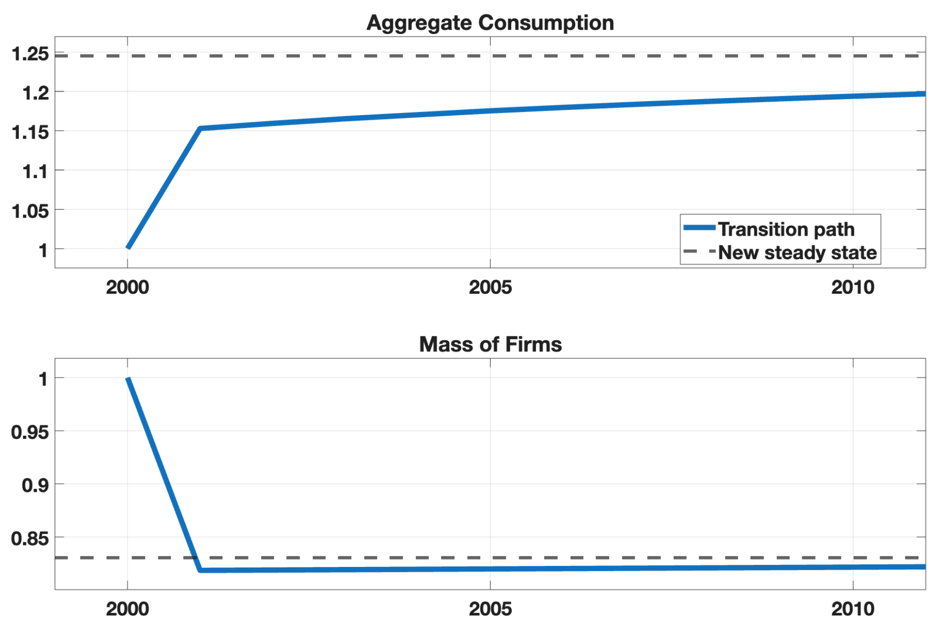

Figure 17: Consumption and the Mass of Firms during the Transition

Notes: The top panel plots the % change in aggregate consumption relative to the initial steady state in 2000, while the bottom panel plots the % change in the mass of firms. In both panels, the dashed horizontal line corresponds to the new steady state, while the solid blue line corresponds to the transition path for $\xi = 0.05$.

that firms place this second location at the opposite end of the circle from their first one—given that the equilibrium is symmetric, doing so is optimal.

Opening a second location comes at a fixed cost, f_{mp}, measured in units of labor. Since opening a second plant simply doubles a firm's variable profits, they will choose to do so if

$$\frac{\Pi * (\psi)}{\delta} \geq w f_{mp} \tag{23}$$

Equation 23 defines a threshold level of ψ above which firms choose to open a second plant.

Calibration Table 10 reports the results of the calibration of the multi-plant model. Relative to the baseline, there is one additional parameter and one additional moment: the fixed cost of opening a second plant which is chosen to match the share of aggregate sales accounted for by multiplant firms.

Table 10: Parameters and moments in 2000 – multi-plant model

Parameter	Description	Value	Moment	Data	Model
(A) Set externally					
θ	intermediate share	0.59	materials share ASI	59%	59%
δ	exit probability	0.082	exit rate (Peters (2020))	8.2%	8.2%
(B) Jointly calibrated					
A_D	TFP distrib sector	6.8240	aggr distrib share	3.8%	3.8%
ϵ	slope of distrib wrt distance	39.6507	aggr / avg distrib share	1.62	1.62
β	fixed cost of market access	44.5366	sh sales across states	23.0%	22.8%
η	Pareto tail index	4.5190	sales sh top 10%	82.8%	86.0%
f_e	entry cost	0.1014	sales sh bottom 50%	1.7%	1.7%
f_x	fixed cost of exporting	0.0471	sh of plants that export	12.9%	12.8%
κ_x	variable cost of exporting	2.3108	distrib share exporters	4.5%	4.5%
Y_w	world demand	4.1687	conditional export share	23.1%	23.0%
f_{mp}	fixed cost 2nd plant	7.4614	sales share of multiplant firms	46.0%	46.0%

Notes: Table 10 reports the calibrated values for the model's key parameters, the data moments used for calibration, and their model counterparts. "Sales across states" is computed as one minus the share of sales within state. Conditional export share is the share of plant sales that are exported conditional on export status. All other moments are computed using the ASI.

For the over-time calibration in 2010, we use the same percentage change in the three parameters, $\{A_D, \epsilon, \beta\}$, that we had inferred for the baseline model in Section 5.2.. That is, A_D increased by 36%, ϵ decreased by 46% and β by 18.4%.

Results

Table 11: Aggregate effects of distribution costs – multi-plant model

	Consumption Gains
(A) Importance of the *level* of distribution costs	
No distribution costs ($A_D = \infty$)	+20.6%
(B) Importance of *heterogeneity across firms*	
Distribution costs constant wrt distance ($\epsilon = 0$)	+8.3%
Lower productivity in distribution sector $A_D \downarrow 10\%$	
$\epsilon = 39.7$	-1.8%
$\epsilon = 0$	-1.1%
Thicker tail of firm productivity ($\eta \downarrow 1\%$)	
$\epsilon = 39.7$	7.2%
$\epsilon = 0$	7.8%

Notes: Table 11 reports the aggregate consumption effects of changing key parameters of the model. Part A reports the gains of eliminating distribution spending altogether by setting $A_D = \infty$. In the first row of part (B), we report the gains of setting $\epsilon = 0$. The last four rows compare the effects of changing A_D and η in the baseline model and in a model with $\epsilon = 0$ respectively. Unless otherwise specified, all parameters are set to the level calibrated in Table 10.

Table 12: Consumption Gains From Improvements in Distribution Sector

All three	A_D only	ϵ only	β only
24.5%	4.5%	3.4%	14.4%

Notes: Table 12 reports the gains of the calibrated changes in distribution costs. Column 1 moves all three parameters to their 2010 values, columns 2-4 move only one parameter at a time while keeping all others at their value calibrated for 2000.

B4. Robustness: Intermediate Share

To quantify the importance of an economy's intermediate share for our results, we repeat the main counterfactuals for a 10 percentage point higher and a 10 percentage point lower intermediate share. For each scenario, we re-calibrate all other parameters to fit the same set of moments as in the main text. Table 13 summarizes the results.

As for the multi-plant extension, we use the percentage change in parameters across time to quantify improvements in the distribution sector.

Table 13: Aggregate effects of distribution costs for different values of the intermediate share

	Higher θ	Lower θ
(A) Importance of the *level* of distribution costs		
No distribution costs ($A_D = \infty$)	+25.7%	+17.3%
(B) Importance of *heterogeneity across firms*		
Distribution costs constant wrt distance ($\epsilon = 0$)	+11.0%	+7.4%
Lower productivity in distribution sector $A_D \downarrow 10\%$		
$\epsilon =$ calibrated value	-2.0%	-1.5%
$\epsilon = 0$	-1.2%	-0.9%
Thicker tail of firm productivity ($\eta \downarrow 1\%$)		
$\epsilon =$ calibrated value	+8.7%	+7.5%
$\epsilon = 0$	+9.6%	+8.0%

Notes: Table 13 reports the aggregate consumption effects of changing key parameters of the model. The first column corresponds to an economy with a higher intermediate share than the baseline, $\theta = 0.65$, while the second column corresponds to an economy with a lower intermediate share than the baseline, $\theta = 0.53$. Part A reports the gains of eliminating distribution spending altogether by setting $A_D = \infty$. In the first row of part (B), we report the gains of setting $\epsilon = 0$. The last four rows compare the effects of changing A_D and η in the baseline model and in a model with $\epsilon = 0$, respectively. Unless otherwise specified, all parameters are set to the level calibrated in Table 3.

Table 14: Consumption Gains From Improvements in Distribution Sector for different values of the intermediate share

	Higher θ	Lower θ
All three changes	+28.5%	+21.1%
Only A_D	+5.3%	+3.9%
Only ϵ	+4.4%	+3.1%
β only	+15.7%	+12.5%

Notes: Table 14 reports the gains from changing distribution costs over time under two alternative values of θ. The first row reports the gain from changing all three components jointly. The remaining rows report the gains from changing each component in isolation.