

Distribution Costs

Alessandra Peter
NYU

Cian Ruane
IMF *

May 11, 2022

Abstract

We provide the first direct estimates of distribution expenses incurred by manufacturing plants and quantify their importance for aggregate output. Using a novel measure from the Indian Annual Survey of Industries, we document three key facts: (1) distribution expenses are large – they amount to over half of labor costs; (2) plants in the largest decile – relative to the smallest – spend over three times as much on distribution as a share of sales; and (3) between 2000 and 2010, distribution costs as a share of sales declined by one third. We develop a model of heterogeneous manufacturing firms that rely on the distribution sector to sell their goods across space. We quantify the model using the facts on size and systematic heterogeneity in distribution shares as well as newly constructed estimates of intranational trade. Accounting for firm heterogeneity in distribution requirements is important: welfare losses from low TFP in the distribution sector are amplified 1.7-fold. From 2000 to 2010, India saw an increase in intranational trade hand in hand with a decrease in the distribution share. In combination with the model, these trends suggest large decreases in variable costs of distribution, leading to welfare gains of 7.4% over this ten year period.

*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. 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. 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. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or its management.

1. Introduction

Before goods can be consumed, they not only need to 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 low-income countries which are typically characterized by inefficient distribution and logistics networks and underdeveloped transportation infrastructure.¹ 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 from India's Annual Survey of Industries, 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 over 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. The decline was driven primarily by large manufacturers, whose distribution expenses dropped disproportionately.

To quantify the welfare consequences of the level and trend in India's distribution share, we develop a model in which manufacturing firms purchase inputs from the distribution sector to sell their goods across space. Shipping further distances requires more distribution services and hence more productive plants, which are larger, 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 pin down the model's key parameters. Since the economy's intermediate share is high and distribution expenses constrain larger firms more, the 3.8% aggregate distribution share leads to large welfare losses of 21%. Accounting for firm heterogeneity in the distribution share amplifies losses from low TFP in the distribution sector 1.7-fold. We then use the model to assess the gains from the simultaneous fall in the aggregate distribution share and increase in interstate trade in India between 2000 and 2010. The inferred reductions in variable distribution costs led to total welfare gains of 7.4% during this period.

The paper's main contribution is to provide evidence on the size of distribution costs

¹See Arvis, Ojala, Wiederer, Shepherd, Raj, Dairabayeva and Kiiski (2018).

incurred by manufacturing firms as well as systematic heterogeneity in the cost burden. The fact that larger firms spend more on distributing their goods and are hence more exposed to low productivity in the distribution sector has large welfare consequences. This finding emphasizes a new channel through which intersectoral linkages of the type studied by [Jones \(2011\)](#), [Bartelme and Gorodnichenko \(2015\)](#), and [Liu \(2019\)](#) can hold back economic development. By directly measuring distribution costs, our paper also complements the literature on intranational trade which typically infers transport costs from price variation ([Atkin and Donaldson, 2015](#); [Asturias, Garcia-Santana and Ramos, 2019](#); [Donaldson, 2018](#)) or flows of goods across space ([Van Leemput, 2021](#)).

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, to the best of our knowledge, is unique in that plants directly report how much they spend on distribution.² Distribution expenses include transportation costs as well as other variable costs such as packing fees and transit insurance. Using this new data, we document several facts on the level of distribution costs and how these vary across manufacturers of different sizes. First, distribution costs are important in the aggregate, averaging 3.2% of aggregate sales between 1993 and 2013. This puts distribution expenses at the same order of magnitude as labor costs, which averaged 6% of sales for the same set of firms. Second, the distribution share of sales is substantially higher for larger plants, ranging from around 1% for the smallest decile to around 3.5% for the largest decile of plants. Finally, we document two important trends from the late 1990s to the late 2000s. The aggregate distribution share declined by a third, from around 3.8% to around 2.5% of sales, and this decline was particularly pronounced for larger firms.

We rationalize the positive relationship between distribution share and plant size by the fact that larger plants ship their goods further distances than smaller plants ([Holmes and Stevens \(2012\)](#)). Given that shipping costs tend to increase with distance, this results in the distribution share of sales being higher for larger firms. We also consider the three most likely alternative explanations and provide evidence against these. First, the fact that larger firms have large distribution shares could be driven by small plants reporting their distribution costs less accurately than larger plants that might have better bookkeeping. However, we find that inward distribution costs (the costs that plants pay to transport and acquire production materials) are decreasing in plant size, suggesting that measurement error is not driving our findings. Second, we show that the pattern holds when we exclude exporters, ruling out that costs associated with international

²In the Annual Survey of Manufacturers in the US, for instance, these are not recorded as plants are asked to report sales net of distribution expenses.

trade drive our findings. Finally, we rule out that our findings are driven by decreasing profit shares with plant size, as distribution expenses relative to costs are also increasing in plant size.

To assess the welfare impact of distribution costs and their changes over time, we develop a model in which heterogeneous manufacturing firms purchase distribution services in order to sell their products across space. The model builds on a closed-economy version of [Arkolakis \(2010\)](#), augmented to continuous locations and multiple sectors. Manufacturing firms differ in the quality of the variety they produce and choose both how much to sell in each location and how many locations to sell to. Selling to a specific location requires setting up a local distributor which is captured in the model 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 in distance.

More productive firms, i.e. those producing a higher quality variety, optimally choose to sell to more distant locations. Larger firms endogenously have a higher distribution share and distribution expenses act like an increasing marginal cost. Low TFP in the distribution sector therefore particularly affects the most productive firms in the economy, who reduce the geographical scope of their sales. In general equilibrium, the restriction on large and productive firms allows small, inefficient producers to survive, further lowering aggregate output and welfare.

We calibrate the model in 2000 by matching the properties of the distribution share as well as the firm size distribution and newly constructed estimates of interstate trade. The three key parameters are the ones governing distribution: (i) the productivity of the distribution sector, (ii) the rate at which variable distribution costs rise with distance and (iii) the rate at which fixed costs rise with distance. These parameters are primarily identified by (i) the aggregate distribution share, (ii) the ratio of the aggregate to average distribution share, and (iii) the interstate trade share. Intuitively, an increase in the productivity of the distribution sector primarily reduces the aggregate distribution share by lowering the price of distribution services. A reduction in the rate at which shipping costs increase with distance lowers the spending of large firms on distribution by more, and hence lowers the ratio of aggregate to average distribution costs. In principle, a high distribution share could be the result of high distribution costs or firms shipping long distances. The interstate trade share allows us to tell these apart and thereby primarily identifies how the fixed costs of shipping rise with distance.

We first use the model to assess the welfare effects of the new facts we document. First, despite the fact that firms spend less than 4% of sales on distribution, the aggregate consequences are large: in an economy without these variable distribution re-

quirements, welfare would be 21% higher. Second, the fact that larger firms use disproportionately more distribution services is quantitatively important: in an economy in which all firms use inputs from the distribution sector equally, welfare would be 9% higher. Firm heterogeneity in exposure to the distribution sector is also important when evaluating the possibility for bottlenecks. The aggregate losses from a reduction in distribution sector TFP are amplified 1.7-fold relative to an economy where all firms spend equally on inputs from that sector.

We next recalibrate the model to 2010 to evaluate the welfare gains from changes in India's distribution sector over time. From 2000 to 2010, adjusting for industry composition, the aggregate distribution share fell from 3.8% to 2.9%, the ratio of aggregate to average distribution costs fell from 1.6 to 1.5, and the interstate trade share increased from 23% to 32%. From the concurrent reduction in expenditure on distribution and increase in interstate trade, we infer large improvements in the productivity of the distribution sector as well as a reduction in how rapidly fixed and variable costs of distribution increased with distance. Taken together, the improvements in the three parameters governing distribution led to a 28% welfare gain over the ten years. While these seem large, the 2000s were a time period of extensive investments in India's road network; most notably the building of the Golden Quadrilateral highway network that now connects India's four major economic centers.

The largest contributor to aggregate welfare gains was the reduction in the fixed cost of accessing new markets. That parameter alone, inferred mostly from the increase in interstate trade we estimate, is responsible for welfare gains of 17%. The increase in distribution sector TFP led to welfare gains of 3% and reductions in the rate at which distribution costs increase with distance increased welfare by another 5%. These 5% gains are identified by the change in the slope of the distribution share vs. plant size, one of the key new empirical findings from our data. Together, the change in the variable costs of distribution between 2000 and 2010 led to welfare gains of 7.4%.

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 gets amplified through the input-output matrix (Bartelme and Gorodnichenko, 2015; Liu, 2019; Fadinger, Ghiglini and Teteryatnikova, Forthcoming). 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 the importance of this heterogeneity for aggregate output. Similar to

Jones (2011), there are complementarities between sectors in our model.³ 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. \(2019\)](#); [Firth \(2019\)](#); [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 and arguably more direct measurement of internal trade barriers: distribution expenses incurred by firms. 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 closed-economy 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 increase with firm size.

Our paper also relates to an important literature on differences in the firm size distribution between developed and developing economies. It is well established that developing countries feature a larger mass of small firms ([Hsieh and Olken, 2014](#)). [Hsieh and Klenow \(2014\)](#) document that Indian plants grow slowly with age and suggest that poor transportation infrastructure could be an important contributor. [Bento and Restuccia \(2017\)](#) argue that distortions which are correlated with firm productivity play an important role in driving down average establishment size and output. The distribution costs we measure are increasing in firm productivity; we evaluate the importance of this concrete size-dependent distortion through the lens of a multi-sector general equilibrium model.

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 perform our quantitative exercises, and Section 5. concludes.

³Our assumption of complementarities is consistent with [Peter and Ruane \(2020\)](#), who estimate long-run elasticities of substitution between material inputs and services below one for Indian manufacturing plants.

2. Distribution Costs in Indian Manufacturing

In this section we first discuss the dataset and main variables used. We 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.⁴ The measure of distribution expenses includes the value of outward transport costs, rebates, commissions, transit insurance of goods sold, and packing fees for goods sold.⁵ 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 inclusive of taxes. This is therefore the ideal setting to provide new insights into the magnitude of distribution costs.

Throughout the paper, we report distribution costs as a share of sales.⁶ 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 on its own does not give a full picture of the extent to which distribution costs matters for aggregate GDP and welfare. Given that production

⁴'Outward' refers to distribution costs incurred on products sold, as opposed to intermediate inputs purchased.

⁵This 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.

⁶We 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.

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 Time Trends

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 large as labor costs, which averaged 6.0% of aggregate sales over the same period. The seemingly low share of labor and distribution in Indian manufacturing is driven by its relatively high intermediate share of sales, which averages 77%. Relative to value-added, the aggregate distribution share is 10.3%, and the labor share 19.1%.

Heterogeneity in the distribution share across plants. The distribution share is not only large in the aggregate but varies systematically across plants of different sizes (see Figure 5 in the Appendix). 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. All of the moments of the distribution share are driven down by the fact that 25% of plants report zero distribution expenses.⁷

Table 1: Summary Statistics

Distribution costs relative to	Aggregate	Mean	p10	p50	p90	Share = 0
Sales	3.2%	2.1%	0%	0.7%	5.6%	25.1%
Value-added	10.3%	8.4%	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.

⁷While this could be due to plants selling directly from the factory-gate, it could also partly reflect underreporting of distribution expenses. To the extent that some of the zeros in the data are due to underreporting, we are underestimating the aggregate spending on distribution and our results should be seen as a conservative lower bound.

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%.⁸ The relationship between distribution share and plant size is not driven by plants' location choices. We show in Appendix Figure 6 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 a residual 23.7% of distribution expenses being 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 the total spending on distribution for the remainder of the analysis.

Trends in the Distribution Share over Time. So far, we have shown that distribution costs are large 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%.⁹ Figure 2 also shows that, while part of this change stems from changes in sectoral composition, even when adjusting for changes in sectoral composition we find the aggregate distribution share fell from just below 4% to 2.9% over this 10-year period.¹⁰

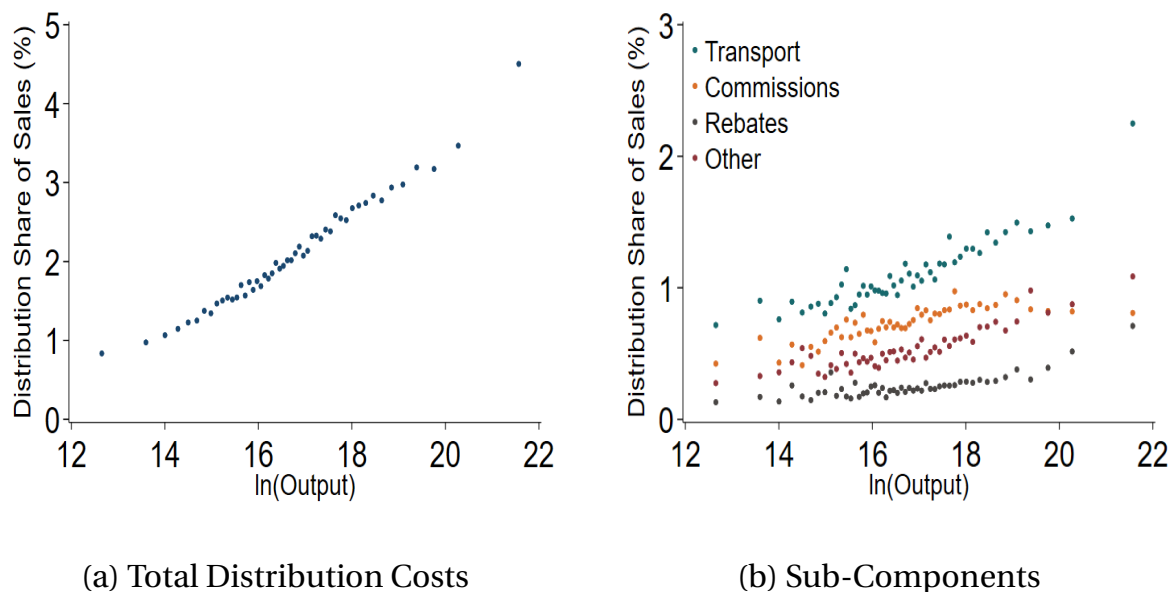
The decline in the aggregate distribution share was driven by changes in the distribution share conditional on plant size, rather than by a change in the size distribution

⁸Plant 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 6 that our findings are very similar when we use employment as a measure of plant size.

⁹Major 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..

¹⁰We adjust for industry composition by aggregating the annual sectoral distribution shares using constant sectoral sales shares from 2000.

Figure 1: Distribution Share vs Plant Size



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 $\log(\text{output})$, 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 \times year fixed effects.

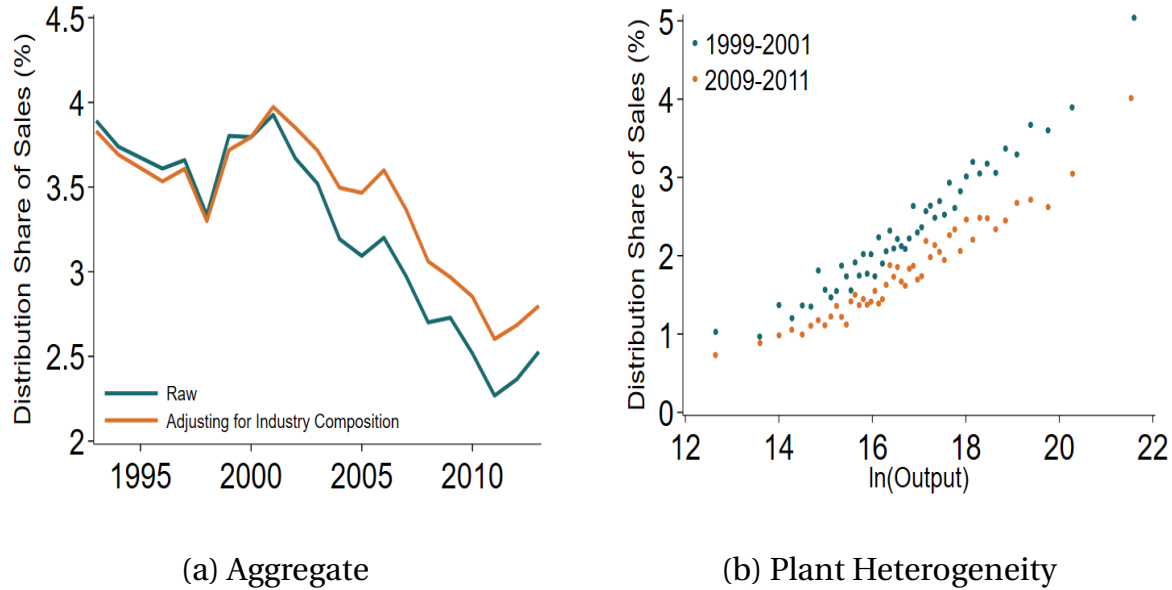
of plants. Figure 2(b) shows that the distribution share of sales decreased for plants of all sizes. In addition, the sales share of large plants expanded during this period and hence changes in the size distribution put upward pressure on the aggregate distribution share.¹¹ Most notably, large plants 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 focused among larger plants.

2.3. Potential Explanations

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, 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 are increasing in distance. In addition, [Holmes and Stevens \(2012\)](#) provide evidence that larger U.S.

¹¹The sales share of the largest 10% of plants increased from 83.3% to 84.5%.

Figure 2: Distribution Share over Time



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 plots of the plant-level distribution share against $\ln(\text{output})$, 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.

manufacturing plants ship their products further distances within the U.S. than smaller plants.¹² Similarly, the trade literature has extensively documented that larger plants export relatively more than smaller plants (e.g. [Bernard and Jensen \(1995\)](#)).

There could, of course, be alternative explanations for the patterns we document. We provide evidence against some important ones below: misreporting of distribution costs by smaller plants, larger plants charging higher markups and therefore having lower cost shares overall, and exporters driving the relationship between size and distribution costs.

One concern 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.¹³ We would expect misreporting to affect

¹²We cannot replicate this analysis for India as we are not aware of any comparable data.

¹³Inward 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.

inward distribution costs in the same way as outward distribution costs. However, show in Appendix Figure 8 that the inward distribution share is decreasing in plant size, which significantly mitigates concerns about mismeasurement.

Another concern is that the increasing relationship between distribution share and size could in principle be driven by larger plants having lower profitability (sales relative to 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 9 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.

We also investigate whether the relationship between distribution costs and size is driven by foreign exports, rather than domestic sales. We examine this in the ASI survey years in which exports are reported.¹⁴ Appendix Figure 7 shows that the positive relationship between the distribution share and size holds even when we restrict our sample to non-exporters.

Firms vs Plants The unit of analysis in the ASI is a plant, rather than a firm.¹⁵ This is potentially important for the interpretation of the fact that larger plants spend more on distribution. Suppose that more productive firms choose to overcome high distribution costs by opening up plants in various locations. If this is the case, an increasing relationship between the distribution share and size at the plant-level could go hand in hand with the opposite relationship at the firm-level.¹⁶ While the ASI does not allow linking plants that belong to the same firm, it does have an indicator for whether a plant belongs to a multi-plant firm.¹⁷

If it were the case that firms open new plants to reduce their variable distribution expenses, we should see that multi-plant establishments have lower distribution shares. However, the opposite relationship holds in the data. On average, the distribution share of multi-plant establishments is 1.8 percent higher than that of single-plant firms. To test this more formally and account for the fact establishments belonging to multi-plant firms are larger, we regress the plant-level distribution share on a set of industry-year fixed effects, a dummy for whether a plant is part of a multi-plant firm and $\log(\text{output})$. Even controlling for firm size, plants that are part of a multi-plant firm have a 0.4 percent *higher* distribution share than single-plant firms. Given that multi-plant firms do not

¹⁴These are 1997, 2008, 2009, 2011, 2012 and 2013.

¹⁵While this is the case for most surveyed units, some firms can choose to file consolidated filings across different plants in the same state and industry.

¹⁶We thank an anonymous referee for pointing out this possibility.

¹⁷Plants that belong to multi-plant firms account for 13% of establishments and 46% of revenues.

have systematically lower distribution shares, we proceed with the rest of the analysis at the plant-level without distinguishing between plants and firms.

Decline in Distribution Share over Time Lastly, 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 A2., 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 a 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 as opposed to technological improvements at the firm-level.

In 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 been not only a significant decrease in the aggregate distribution share during the 2000s, but this decline was particularly pronounced for the largest firms.

3. Model

In order 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 welfare consequences of accounting for the magnitude of distribution costs and the systematic heterogeneity across firms as well as evaluate the gains from improvements in India's distribution sector over time.

3.1. Setup

The model is based on a closed-economy, continuous-location extension of [Arkolakis \(2010\)](#). Time is discrete, agents are risk-neutral and do not discount the future. The economy consists of two sectors, manufacturing and distribution. Each manufacturing firm produces a differentiated variety. Firms purchase distribution services in order to sell their products to both consumers and as intermediate input to other firms across space. We model space in the simplest way that preserves the main mechanism: the economy lies on a circle with circumference 1.

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 .

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

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

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

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

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

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

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

$$Y_i = C_i + M_i \quad (4)$$

Consumers

At each point along the circle, there is a continuum L of workers.¹⁸ Agents cannot save or borrow, which gives rise to the following problem (suppressing time subscripts for ease of exposition):

$$\begin{aligned} \max_{C_i} \quad & u(C_i) \\ \text{s.t} \quad & P_i C_i \leq y \end{aligned} \quad (5)$$

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 $w + \bar{\pi}/L$.

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:

$$d = A_D L_D \quad (6)$$

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. Firms choose both the set of markets (points on the circle) in which they sell their product (the extensive margin problem) and also the quantity to sell in each market (the intensive margin problem).

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

¹⁸The setup is symmetric at every point on the circle.

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 production location of the firm.¹⁹ 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:

$$f(n) = (1 - 2n)^{-\beta} \quad (7)$$

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 in distance.
2. $\lim_{n \rightarrow 0} f'(n) > 0$. 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.
3. If $\beta > 0$, $\lim_{n \rightarrow \frac{1}{2}} f'(n) = \infty$. Accessing the entire market is prohibitively costly.²⁰

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:²¹

$$q(n) = \min \left\{ M^\theta L_p^{1-\theta}, \frac{d(n)}{1 + \epsilon n} \right\} \quad (8)$$

The physical good is produced with labor L_p and an intermediate good M . In order to 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 infrastruc-

¹⁹These fixed costs are in units of labor.

²⁰Given that the circle has circumference 1, the furthest market is located at a distance $n = \frac{1}{2}$.

²¹While 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.

ture. For example, it could be relatively costly to ship goods 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 :

$$\begin{aligned} \pi^*(\psi, n) &\equiv \max_{q,l,m,d} p(q(\psi, n)) q(\psi, n) - w l(\psi, n) - P_I m(\psi, n) - p^D d(\psi, n) - w f(n) & (9) \\ \text{s.t. } & m(\psi, n)^\theta l(\psi, n)^{1-\theta} = q(\psi, n) \\ & d(\psi, n) = (1 + \epsilon n) q(\psi, n) \\ & f(n) = (1 - 2n)^{-\beta} \end{aligned}$$

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 amounts to choosing the furthest market to which the firm sells. We denote the distance between the firm and its furthest point of sale by $n^*(\psi)$, where:

$$n^*(\psi) = \operatorname{argmax}_{N \in [0, 1/2]} \int_0^N 2\pi^*(\psi, n) dn \quad (10)$$

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:

$$\pi^*(\psi) = \max_N \int_0^N 2\pi^*(\psi, n) dn = \int_0^{n^*(\psi)} 2\pi^*(\psi, n) dn \quad (11)$$

Entry and Exit

There is a large pool of potential entrants into the manufacturing sector. In order to enter, firms pay a sunk cost of entry equal to f_E 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 choose whether or not to remain active. If they do, they receive profits $\pi(\psi, t)$ 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 $\{p^D, w, P, \{p(\psi, N)\}\}$
2. quantities $\{L^D, L^E, L^P, \{c(\omega)\}, \{q(\psi)\}, \{l(\psi)\}, \{m(\psi)\}, \{d(\psi)\}, J, M^E\}$

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

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 (9) 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 one-location closed-economy version of Melitz (2003).

Now let $\epsilon > 0$. Selling to markets at a further distance requires more distribution services per unit sold. From the firms' perspective, marginal costs are increasing in n . This implies that they sell less and make smaller profits in markets located further away. Firms only serve a market if profits are sufficient to cover the fixed costs f .²² This gives the following equilibrium expression for $n^*(\psi)$:²³

$$n^*(\psi) = \frac{1}{\epsilon} \left(\frac{A_D}{c} \left(\frac{\sigma - 1}{\sigma} \right) \left(\frac{L}{\sigma} \right)^{\frac{1}{\sigma-1}} \psi^{\frac{\sigma}{\sigma-1}} - (A_D + 1) \right), \quad (12)$$

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 (12) 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:

$$\frac{p_D d(n, \psi)}{p(\psi, n) q(\psi, n)} = \frac{\frac{1+\epsilon n}{A_D}}{\frac{c}{w} + \frac{1+\epsilon n}{A_D}} \frac{\sigma - 1}{\sigma} \quad (13)$$

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 come with a higher distribution share and higher quality firms sell to more distant markets, they spend a larger overall share of sales on distributing their products.

We now turn to discussing the role of the parameter ϵ . Consider a *decrease* in ϵ .²⁴ From equation (12) 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

²²Since $\beta = 0$, the fixed cost of accessing markets is constant at $f(n) = f$

²³This formula applies to firms whose optimal choices are interior; i.e. they serve a positive fraction of consumers, but not the entire market ($n^*(\psi) \in [0, \frac{1}{2}]$). There is a positive mass of firms who sell to the whole market: $n^*(\psi) = \frac{1}{2}$

²⁴This captures any change in the technological environment that makes it relatively cheaper to ship long distances, e.g. a highway network.

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.²⁵ 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.²⁶ We have that:

$$\underline{\psi} = \left(c + \frac{w}{A_D} \right)^{\frac{\sigma-1}{\sigma}} \left(\frac{w \sigma}{y L} \right)^{\frac{1}{\sigma}} \left(\frac{\sigma}{\sigma-1} \right)^{\frac{\sigma-1}{\sigma}} \quad (14)$$

As the wage increases, $\underline{\psi}$ increases also. The reasoning 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, who sell further, and allows small low productivity firms to survive. The fact that the two parameters affect firm sales in a similar way but have different effects on the distribution share is a relationship we rely on to identify the parameters of the model.

4. Quantitative Analysis

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, as well as newly constructed estimates of interstate trade to calibrate the model's key parameters. We find that there were significant improvements in the functioning of the distribution sector with large aggregate welfare gains. The combination of improvements in TFP of the distribution sector, reductions in the fixed cost of accessing markets as well as a flattening of the distribution cost curve across space led to cumulative welfare gains of around 28% over 10 years.

²⁵We have normalized the price index P to 1, which leaves the wage as the only price left to clear markets.

²⁶These are firms for which $n^*(\underline{\psi}) = 0$

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 is crucial for identifying the key parameters governing firms' use of distribution services. In order to identify the fixed cost parameter β , we need data on how far 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 that allow us to estimate the amount of trade between states in the 2000s. The raw data comes from a handbook called the "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).²⁷ 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.²⁸ To overcome these issues and nonetheless obtain a complete estimate of the interstate trade share, we follow an approach based on [Van Leemput \(2021\)](#).

The approach takes advantage of the fact that interstate shipments for export are reported separately from interstate shipments 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, 2) the quality of goods shipped interstate by rail is the same for export and for domestic consumption. These assumptions imply that:

$$\frac{\text{quantity exported via rail}}{\text{quantity shipped interstate via rail}} = \frac{\text{value exported via road and rail}}{\text{value shipped interstate via road and rail}} \quad (15)$$

We obtain the interstate trade share of each good by dividing its export intensity (aggregate export value / aggregation production value) by the ratio in Equation (15). 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 2010.²⁹ While correctly pinning down the *level* of the interstate trade share requires

²⁷The DGCIS publishes data based on accounting years which stretch from April to March of the following year. See Appendix A3. for a more detailed discussion of data and methodology.

²⁸Capturing interstate trade by road is particularly important because of the development of large-scale road infrastructure during the 2000s such as the Golden Quadrilateral.

²⁹Our estimates are higher than those of [Van Leemput \(2021\)](#) for 2011, also using the DGCIS data. While

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 road and rail in India are highly correlated across states (Van Leemput, 2021). We therefore view our estimates as providing strong evidence of an increase in the value of interstate trade in India 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 seven 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 measure in the ASI and set the exit probability of firms to match exit rates as estimated by Peters (2020) for formal manufacturing plants in Indonesia.³⁰ The remaining five parameters are calibrated to match a set of key moments in the baseline year 2000. All parameters and moments are reported in Table 2.³¹

Identification While all five main parameters are jointly identified by all moments, they differ in their respective sensitivity to each moment. Table 2 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 how many resources the economy as a whole spends on 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 by similar proportions.

our approach is based on theirs, we make some different methodological choices (detailed in Appendix A3.) 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 interstate trade in goods and services subject to the GST amounted to 60% of GDP (as opposed to production) in 2017 (Ministry of Finance, 2018).

³⁰Intermediate 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). 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.

³¹Parameters are chosen to match the (equally weighted) sum of squared deviations of model and data moments.

Table 2: Parameters and moments in 2000

Parameter	Value	Moment	Data	Model
(A) Set externally				
θ	0.59	intermediate share	59%	59%
δ	0.082	exit rate (Peters (2020))	8.2%	8.2%
(B) Jointly calibrated				
A_D	5.4	aggregate distribution share	3.8%	3.8%
ϵ	38.7	aggr / avg distribution share	1.62	1.62
β	47.8	share of sales across states	23.0%	23.0%
η	4.1	sales share top 10%	85.9%	82.8%
f_e	0.25	sales share bottom 50%	1.7%	1.7%

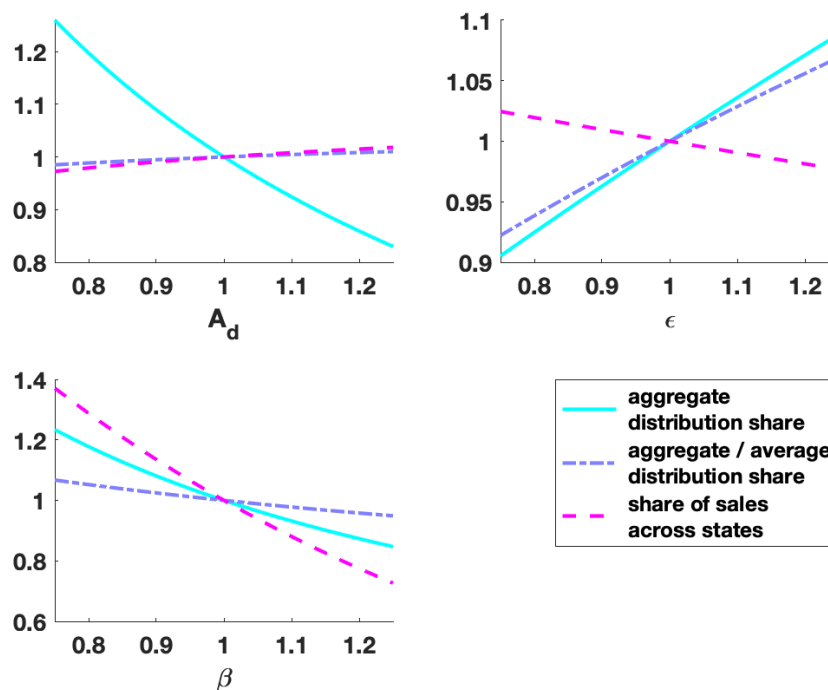
Notes: Table 2 reports the calibrated values for the model's six key parameters as well as the data moments used for calibration and their model counterpart. See Section 4.1. for a description of the share of sales across state. All other moments are computed using the ASI.

The rate at which distribution requirements rise with distance, ϵ , affects the aggregate distribution share much more than the average. 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 expenses rise more and more with distance which affects larger firms more. As a result, the aggregate distribution share increases more than the average. As the second panel of Figure 3 shows, ϵ has a much stronger effect on the ratio of the aggregate to average distribution share than A_D does.

The rate at which the fixed cost of accessing markets rises with distance, β , is mostly identified by how far firms sell their products on average. 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 to be a firm’s sales within the radius that corresponds to a typical state’s share of Indian GDP. As a result, a segment length of around 7% is considered within-state sales. The third panel of Figure 3 illustrates that the value of β strongly affects the share of sales that cross states.

Last, 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 is η , the more skewed is the distribution of productivity and hence size. The larger is f_e , the more productive a firm must be in order to operate, which reduces the top 10% employment

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 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.

share and increases the employment share of the top 50%. Figure 10 in Appendix B plots all moments against all parameters.

4.3. Welfare Impacts 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 welfare. In the next subsection, we evaluate the welfare gains from observed changes in distribution costs in the 2000s.

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 welfare consequences of the high level of (variable) distribution costs we find in the data. Table 3 reports all main results for this section.

Table 3: Welfare effects of distribution costs

	Welfare Gains
<u>(A) Importance of the <i>level</i> of distribution costs</u>	
No distribution costs ($A_D = \infty$)	+20.9%
<u>(B) Importance of <i>heterogeneity across firms</i></u>	
Distribution costs constant wrt distance ($\epsilon = 0$)	+8.8%
Lower productivity in distribution sector $A_D \downarrow 10\%$	
$\epsilon = 38.7$	-1.7%
$\epsilon = 0$	-1.0%
Thicker tail of firm productivity ($\eta \downarrow 1\%$)	
$\epsilon = 38.7$	7.4%
$\epsilon = 0$	7.9%

Notes: Table 3 reports the welfare effects of changing key parameters of the model. Part A reports the welfare gains of eliminating distribution spending altogether by setting $A_D = 0$. In the first row of part (B), we report the welfare gains of setting $\epsilon = 0$. The last four rows compare the welfare 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 2.

In an economy in which no firms incur any variable costs of distributing their goods from the factory gate, aggregate welfare would be 21% higher. The fact that these welfare gains are an order of magnitude larger than the distribution share in the data stems in part from the fact that the intermediate share is high. To a first order, the effect of an increase in distribution sector TFP is equal to its Domar weight. With an aggregate distribution share of manufacturing sales of 3.8% and an intermediate share of 59%, the distribution sector's Domar weight is 9.3%. We find that this is roughly in line with our model predictions: a 1% increase in distribution sector productivity leads to a 0.16% increase in welfare. Because distribution services are Leontiev in production, the Domar weight will overstate the increase in GDP for large increases in distribution sector productivity, and the gains from making distribution costless (as in Table 3) are bounded.

The welfare gains from reducing the cost of distribution come from 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. 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 of goods.³² This can be seen by the fact that the interstate trade share increases by 11%. Third, the expansion into more distant markets is most pronounced for the largest firms, who produce the highest quality goods and spend the most distributing them. The employment share of the top 10% of firms increases by around 1%. Consumers therefore have access to high-quality goods produced further away. Finally, the expansion of firms drives wages up in general equilibrium. This increases the threshold level of productivity under which firms exit the market by 2%.

Heterogeneity in distribution costs. To quantify how much the systematic size dependence of distribution costs matters for welfare, 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 welfare gains of 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 welfare 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 ($\epsilon = 0$), the decrease in A_D leads to a 2% welfare loss, as more resources need to be spent to distribute goods across space. With our calibrated value of ϵ , larger and more productive firms are more affected by this decline in distribution TFP, leading to a nearly 50% amplification of welfare 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 7% lower in the baseline model with our calibrated ϵ , relative to a model where $\epsilon = 0$. The welfare gains from increasing the set of high-productivity firms are dampened because these firms are disproportionately constrained by distribution inefficiencies, thereby limiting their growth potential.

³²Since $\beta > 0$, there is still a fixed cost of accessing markets and hence an interior solution for the choice of market access.

4.4. Changes in the distribution sector over time

In this section, we re-calibrate the model to the Indian data in 2010. This allows us to estimate the welfare gains from improvements in distribution during a period in which India's transportation infrastructure expanded rapidly. We re-calibrate 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 distribution share, we target the changes adjusted for industry composition, given our model is focused on within-sector changes. We hold the entry costs and the Pareto shape parameter constant at their value in 2000 in order to focus on the role of changes in distribution costs in the model. As previously shown in Figure 2 Panel (b), the decrease in the aggregate distribution share occurred entirely within size groups, rather than due to changes in the size distribution of firms.³³

As can be seen when comparing Table 2 to Table 4, there were large changes in India's distribution efficiency from 2000 to 2010. First, the aggregate distribution share fell from 3.8% to 2.9%. Second, the ratio of the aggregate to 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.³⁴ Through the lens of the model, the changes in the first two moments imply that there was 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 4 reports all calibrated parameters and moments for 2010.

Table 4: Parameters and moments in 2010

Parameter	Value	Moment	Data	Model
A_D	7.6	aggregate distribution share	2.9%	2.9%
ϵ	24.6	aggregate / average distribution share	1.52	1.52
β	37.8	share of sales across states	31.7%	31.7%

Notes: Table 4 reports the calibrated valued for the three re-calibrated 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.

³³The sales share of the top 10% of firms very modestly increased during the period we consider, which would be a force pushing the aggregate distribution share to increase between 2000 and 2010.

³⁴Figure 11 in Appendix B shows the corresponding identification figure.

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 with the goal of improving the country's underdeveloped highway network. A particular well-known investment was the Golden Quadrilateral highway network (GQ).³⁵ 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 manufacturing sector, including through improvements in allocative efficiency. From the perspective of our model, the GQ has effects on 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 lowered the relative cost of shipping longer distances - a decline in ϵ . Finally, the GQ may have also enabled easier flows of information and made it easier for managers to locate distributors in new markets - a decline in β .

Another potential driver of the decline in distribution costs is increases 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 that these led to productivity gains for manufacturing firms. Increases in competition in the distribution sector could lead to increases in productivity and decreases in markups, both of which would map 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 .

Welfare Impacts. We find that the combined improvements in India's distribution sector led to welfare gains of 28%. That is, in 2010, welfare of the representative consumer who saw changes in the three parameters governing distribution (ϵ, β, A_D) would have been 28% higher than in 2000.

Table 5: Welfare Gains From Improvements in Distribution Sector

All three	A_D only	ϵ only	ϵ and A_D	β only
27.8%	2.9%	5.1%	7.4%	17.1%

Notes: Table 5 reports the welfare 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.

³⁵This was followed by the construction of the North-South and East-West corridor. Around the same time, the Prime Ministers Village Road Program was also begun to improve rural road access.

We then decompose the welfare 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 increase with distance) lead to welfare gains of 17%. The main data moments that allowed us to infer the reduction in β is the increase in interstate trade over this decade. The improvement in distribution sector TFP as well as the reduction in the rate at which *variable* costs increase with distance also had sizeable welfare effects on the order of 3% and 5% 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 welfare increased substantially as a result of declining distribution costs therefore points towards a potentially important role of infrastructure investments and other government policies for promoting the development of the manufacturing sector.

5. 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 in 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 welfare losses from low TFP in distribution are 1.7 times larger than in a world in which all firms have equal distribution shares.

We then evaluate the consolidated welfare 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%.

Improvements in the TFP of the distribution sector led to a 3% increase in welfare. Declining variable costs of accessing more distant markets added another 5% to welfare. The fact that these improvements happened simultaneously is crucial - improvements in the TFP of the distribution sector alone would have had much smaller aggregate effects.

Our findings reinforce the importance of policies affecting competition in distribution and transportation sectors in developing countries. As of 2020, India has the second highest restrictions to foreign entry in distribution services among OECD countries, and has among the highest barriers to entry in transport and distribution sectors (OECD, 2020). An additional implication of our findings is that such frictions may dampen the incentives of firms to innovate or invest in cost-reducing measures. The private benefits of such investments are greatly diminished if the firm's growth potential is restricted 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

- 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.
- Arvis, Jean-Francois, Lauri Ojala, Christina Wiederer, Ben Shepherd, Anasuya Raj, Karlygash Dairabayeva, and Tuomas Kiiski, “Connecting to Compete 2018 : Trade Logistics in the Global Economy,” *World Bank*, 2018.
- 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,” 2015. NBER Working Paper No. 21251.
- 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.
- 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 Economics Journal: Macroeconomics*, Forthcoming.
- Firth, John, “I’ve Been Waiting on the Railroad: The Effects of Congestion on Firm Production,” 2019.

- 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, “The Life Cycle of Plants in India and Mexico,” *Quarterly Journal of Economics*, 2014, 129 (3), 1035–1084.
- Jones, Charles I., “Intermediate Goods, Weak Links, And Superstars: A Theory of Economic Development,” *American Economic Journal: Macroeconomics*, 2011, 3 (2), 1–28.
- 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, “The Aggregate Importance of Intermediate Input Substitutability,” 2020.
- Peters, Michael, “Heterogeneous Mark-Ups and Endogenous Misallocation,” *Econometrica*, 2020, 88 (5).

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.³⁶ 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.³⁷

A2. 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 investments 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,

³⁶The size threshold for plants to be included in the Census sample is 200 in some years during the 1990s.

³⁷The official sectoral classification (NIC) changed in 1998, 2004 and 2008. We use official NIC concordances to construct our harmonized classification.

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 4, 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.³⁸ 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.³⁹

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.

A3. 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 all Indian states are reported, with separate entries for shipments for international export.⁴⁰ For example, shipments from Delhi to 'Other ports of Karnataka' are reported,


³⁸We 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.

³⁹We find a smaller albeit still substantial difference when we control for differences in industry composition across states.


⁴⁰Exports 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

Figure 4: The GQ and Changes in Distribution Shares

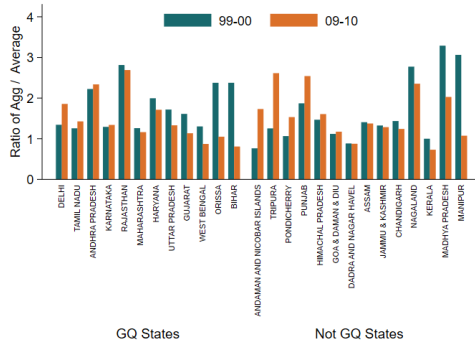
(a) State-level Distribution Shares



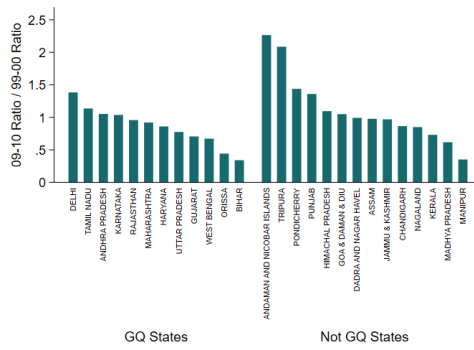
(b) 09-10 – 99-00 Distribution Share



(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.

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.⁴¹

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:

$$\frac{\text{quantity exported by rail}}{\text{quantity shipped interstate by rail}} \quad (16)$$

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:

$$\frac{\text{quantity exported by rail}}{\text{quantity shipped interstate by rail}} = \frac{\text{aggregate value exported}}{\text{aggregate value shipped interstate}} \quad (17)$$

For each commodity, we then construct its aggregate export intensity, given by:

$$\text{export intensity} \equiv \frac{\text{aggregate value exported}}{\text{aggregate value of domestic production}} \quad (18)$$

For each commodity, we then obtain the interstate trade share by dividing Equation 19 by 17:

$$\text{interstate trade share} = \frac{\frac{\text{aggregate value exported}}{\text{aggregate value of domestic production}}}{\frac{\text{aggregate value exported}}{\text{aggregate value shipped interstate}}} \quad (19)$$

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 approach 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

(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.

⁴¹Van 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.


and aggregate production for each commodity. We obtain this from a variety of sources, including the ASI, Indiastat and CEIC.⁴² 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 2010.⁴³ 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.

⁴²We scale up output from the ASI by 3/2 to account for informal production.

⁴³For example, the interstate trade share for 'Other grains' increases from 1.1% in 2000 to 96.5% in 2010.


A4. Figures

Figure 5: 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 6: Distribution Share vs Plant Size: Robustness




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 7: 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 8: 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 $\log(\text{output})$. We residualize both variables on $\text{sector} \times \text{year}$ 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 9: 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 $\text{sector} \times \text{year}$ fixed effects. Total costs include intermediate expenditures, labor costs and the imputed rental on capital assuming a rental rate of 20%.

B Model Appendix

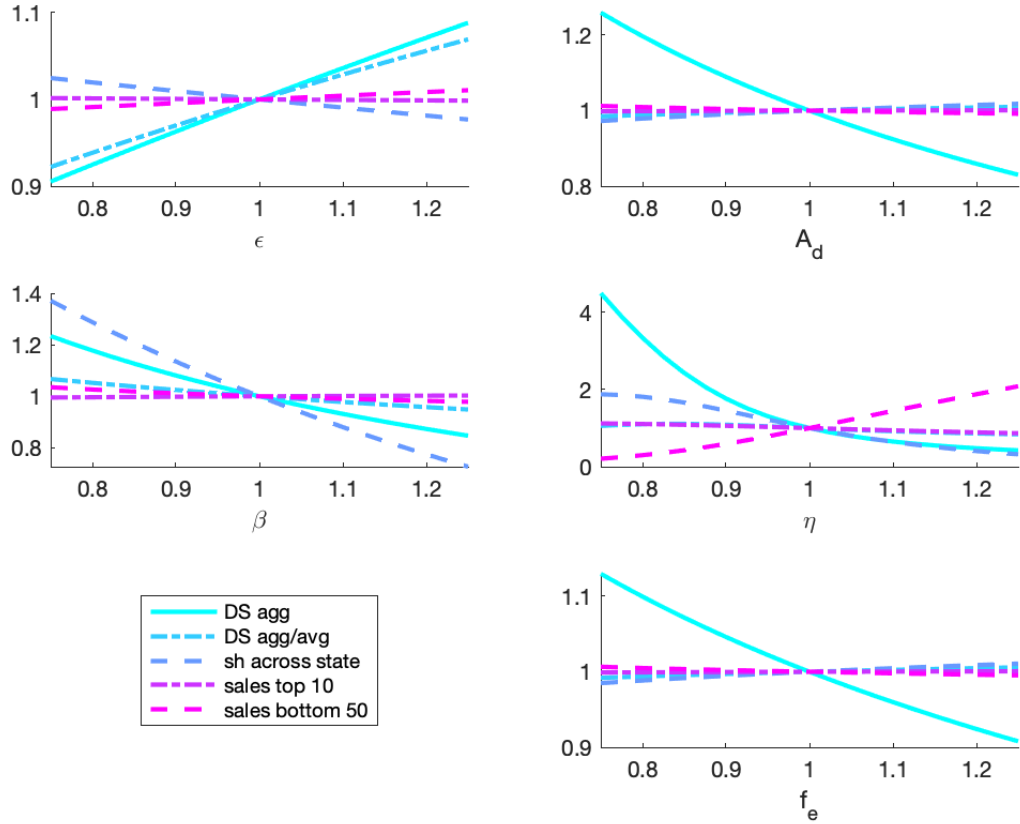


Figure 10: Identification 2000: all parameters

Note: Figure 10 shows sensitivity of model moments to parameters. Starting from the baseline calibration, each panel plots all six 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.

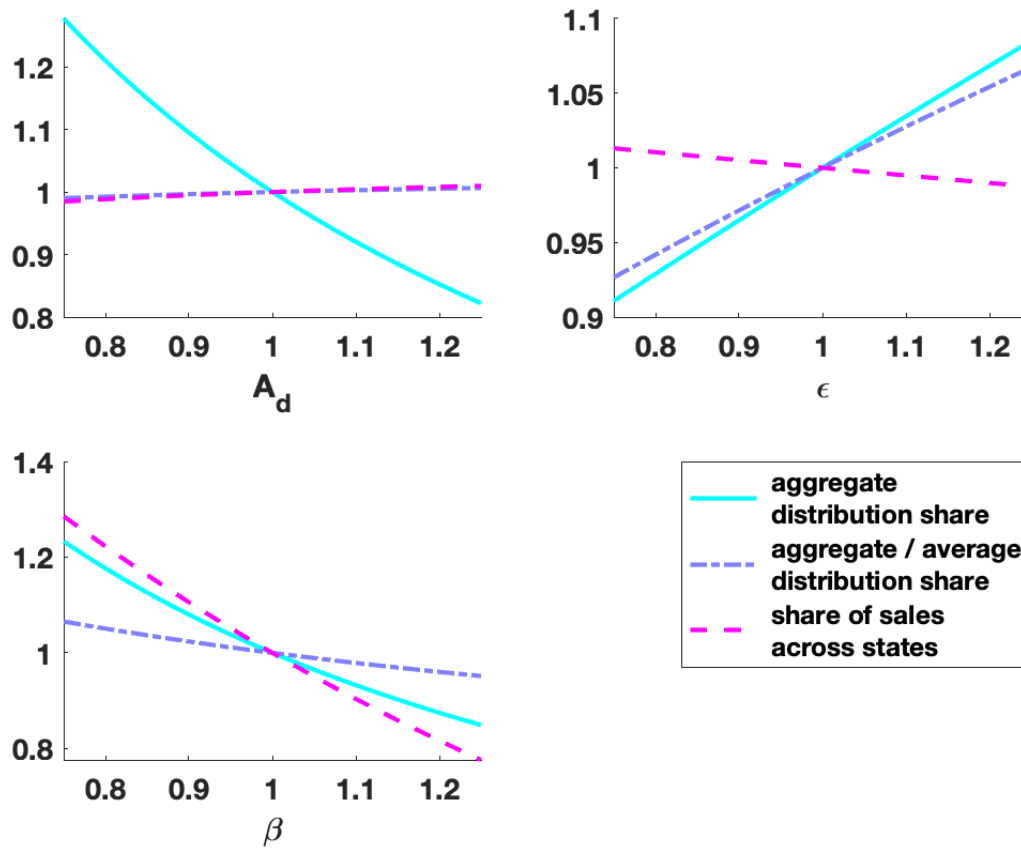


Figure 11: Identification 2010

Note: Figure 11 shows sensitivity of model moments to parameters. Starting from the baseline calibration, each panel plots all three 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.