Multinationals, Offshoring, and the Decline of U.S. Manufacturing*

Christoph E. Boehm[1] Aaron Flaaen[2] Nitya Pandalai-Nayar[1]

[1]University of Michigan
[2]Federal Reserve Board of Governors

December 19, 2015

Abstract

We provide three new stylized facts that characterize the role of multinationals in the U.S. manufacturing employment decline, using a novel microdata panel from 1993-2011 that augments U.S. Census data with firm ownership information and transaction-level trade. First, over this period, U.S. multinationals accounted for 41% of the aggregate manufacturing decline, disproportionate to their employment share in the sector. Second, U.S. multinational-owned establishments had lower employment growth rates than a narrowly-defined control group. Third, establishments that became part of a multinational experienced job losses, accompanied by increased foreign sourcing of intermediates by the parent firm. To establish whether imported intermediates are substitutes or complements for U.S. employment, we develop a model of input sourcing and show that the employment impact of foreign sourcing depends on a key elasticity – of firm size to production efficiency. Structural estimation of this elasticity finds that imported intermediates substitute for U.S. employment. In general equilibrium, our estimates imply a sizable manufacturing employment decline of 13%.

JEL Codes: F14, F16, F23

Keywords: Multinational Firms, Offshoring, Outsourcing, Manufacturing Employment

*University of Michigan, Department of Economics, 238 Lorch Hall, Ann Arbor, MI 48109-1220. chrisboe@umich.edu, aflaaen@umich.edu and npnayar@umich.edu. We would like to thank Andrei Levchenko, Alan Deardorff, Kyle Handley, Linda Tesar, and Dan Ackerberg for valuable comments, suggestions, and support. We also benefited from conversations with Isaac Sorkin and Sebastian Sotelo. This research has been supported by grants from the Michigan Institute for Teaching and Research in Economics (MITRE), for which we are very thankful. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau, the Board of Governors, or its research staff.

1 Introduction

One of the most contentious aspects of globalization is its impact on national labor markets. This is particularly true for advanced economies facing the emergence and integration of large, low-wage and export-driven countries into the global trading system. Contributing to this controversy, the United States has experienced steep declines in manufacturing employment in the last two decades, paired with extraordinary expansions of multinational activity by U.S. firms.

While a large body of research has studied the intersection of international integration and employment, particularly in developed countries, the results and policy prescriptions have been mixed. There are several factors underlying the conflicting results of this research, but prominent among them are gaps in the coverage and detail of the requisite firm-level data to disentangle competing views. Data constraints pertaining to multinational firms in the U.S. have been particularly severe, limiting research on their role in the manufacturing employment decline.

This paper uses a novel dataset together with a structural model to show that U.S. multinationals played a leading role in the decline in U.S. manufacturing employment. Our data from the U.S. Census Bureau cover the universe of manufacturing establishments linked to transaction-level trade data for the period 1993-2011. Using two directories of international corporate structure, we augment the Census data to include, for the first time, longitudinal information on the direction and extent of firms' multinational operations. To the best of our knowledge, this data permits the first comprehensive analysis of the role of U.S. multinationals in the aggregate manufacturing decline in the United States.

We begin by establishing three new stylized facts. First, U.S. multinationals averaged 30 percent of overall employment but accounted for 41 percent of the aggregate employment decline. Second, U.S. multinationals had a 3 percentage point per annum lower employment growth rate relative to a narrowly-defined control group sharing similar industry, size, and age characteristics. Finally, we use an event-study framework to compare the employment dynamics in plants which become part of a firm with multinational operations to a control group of non-transitioning plants. These transitioning plants experienced substantial job losses relative to the control group. Together, these three exercises show that U.S. multinationals contributed disproportionately to the manufacturing employment decline.

We next examine the trading patterns of multinational and other manufacturing firms in our data. We find that foreign sourcing of intermediate inputs is a striking characteristic of multinationals. Over 90% of overall U.S. intermediate imports in our sample are imported by multinationals. Further, the fraction of U.S. multinationals sourcing inputs from developing countries has nearly doubled from 1993 to 2011. To illustrate the link between these high and increasing intermediate imports by multinationals and the observed employment declines, we return to the event study. We show that the relative employment declines in transitioning plants are accompanied by large increases in imports of intermediates by the parent firm. The increase in imports is largest when the plant is shut down.

While suggestive, these stylized facts are not sufficient to establish whether foreign sourcing is a complement or a substitute for domestic employment. To understand the causal mechanism underlying these facts and to quantify their impact in the aggregate we present a model of firm sourcing decisions in the spirit of Antràs, Fort, and Tintelnot (2014). In the model firms choose the location (home, North and South) and mode (inter or intra-firm) through which they source their intermediate inputs for production. The firm's optimal sourcing strategy balances the gains from access to cheaper intermediate inputs against higher fixed costs.

The impact of foreign sourcing on U.S. employment is determined by two opposing forces. First, greater foreign competitiveness implies that firms sourcing from abroad have access to cheaper intermediates. As a result, their unit costs fall and their optimal scale increases. This effect raises their U.S. employment. On the other hand, firms reallocate intermediate production towards the location with increased competitiveness. This reduces U.S. employment.

We show that the value of a single structural constant—the elasticity of firm size with respect to production efficiency—completely determines which of the two forces dominates. Existing views in the literature on the value of this constant vary substantially. The existing range of estimates is large enough that foreign sourcing could be either complementary or substitutable with domestic employment. We therefore estimate this constant structurally using our data on the universe of U.S. manufacturing firms. Our data on cost shares of the firm from all locations and modes, as well as firm revenues and wage payments to labor is sufficient to identify the structural constant. The intuition behind this result is related to the finding in Blaum, LeLarge, and Peters (2015) that domestic cost shares and revenues are sufficient to identify changes in firm productivity due to imported inputs in a large class of models.

Our estimation demonstrates that increased foreign sourcing is a strong substitute for U.S. employment at the firm-level. This result is robust to a number of alternative estimation methods and subsamples. As a final step, we evaluate what the firm-level results imply for aggregate manufacturing employment. We implement a general equilibrium version of the model, and calibrate it using our structural parameter estimates and observed foreign sourcing shares. Our model implies a quantitatively significant employment decline in response to foreign sourcing. It generates an aggregate employment decline in U.S. multinationals of 28%, and an overall employment decline of 13%, which is larger than the direct contribution by multinationals. The latter result is due to general equilibrium effects: decreased demand from multinational firms for intermediates from other U.S. firms further reduces manufacturing employment.

This paper contributes to a growing literature documenting the impact of international integration on labor markets. Data constraints have limited previous work on the role of multinationals in the U.S. manufacturing decline. Some exceptions are Harrison and McMillan (2011), Ebenstein et al. (2014), Ebenstein, Harrison, and McMillan (2015) and Kovak, Oldenski, and Sly (2015) who have studied foreign sourcing by multinationals using BEA data. Since these data only include multinationals, they do not permit analysis of multinationals' behavior relative to a non-multinational control group. To study plant closure in multinationals, Bernard and Jensen (2007) made use of a temporary link between the BEA and the Census. However, they did not focus on offshoring.

In contrast to the limited studies on the impact of foreign sourcing by multinationals, a larger literature has examined the impact of international trade on labor markets more generally. In particular, a number of recent papers have studied the impact of import competition from China (Autor, Dorn, and Hanson, 2013; Autor et al., 2014; Acemoglu et al., 2014). Unlike our paper, these studies use industry-level data. In a firm-level study, Pierce and Schott (2013) find lower employment growth in industries that were most affected by the recent reduction in trade-policy uncertainty with China. Several papers have focused on the wage or inequality effects of trade. For instance, Hummels et al. (2014) find negative wage effects of offshoring for low skilled workers using firm-level data from Denmark.

Our finding of the substitutability between foreign sourcing and domestic manufacturing employment contributes to another active debate in the literature. A number of papers have found little to no employment substitution in various countries, including Desai, Foley, and Hines (2009) [U.S.A], Braconier and Ekholm (2000) [Sweden], Konings and Murphy (2006) [Europe], Slaughter (2000) [U.S.A.], Barba-Navaretti, Castellani, and Disdier (2010) [Italy and France] and Hijzen, Jean, and Mayer (2011) [France].

In contrast, and consistent with our results, several recent papers with data from other countries have found that firms treat foreign and domestic employment as substitutes in production. In particular, Muendler and Becker (2010) find evidence for substitutability between home and foreign employment using German data in a structural model. As in our paper, they emphasize the role of the extensive margin (in the case of that paper, of new foreign locations). We find it critical to account for the extensive margin of domestic plant deaths when calculating the employment effects of foreign operations. Other papers finding evidence for substitution are Simpson (2012) [United Kingdom], and Debaere, Lee, and Lee (2010) [South Korea]. Monarch, Park, and Sivadasan (2013) also find that offshoring firms in Census data experience declines in employment.

Finally, the structural model we develop in this paper draws on Antràs, Fort, and Tintelnot (2014), who develop a tractable model of foreign sourcing. Our model allows for a more general form of technology transfer between the parent firm and its suppliers. We also distinguish explicitly between inter and intra-firm imports in the model, as our focus is on multinationals. Moreover, our data shows that these firms' imports at arms-length are often accompanied by substantial imports from related-party suppliers. Whether firms source within or outside the firm has been extensively studied by a large empirical and theoretical literature, including Feenstra and Hanson (1996), Hanson, Mataloni, and Slaughter (2005), Antràs (2005), Antràs and Helpman (2004), Antràs and Chor (2013), and Costinot, Vogel, and Wang (2013).

The next section presents empirical evidence establishing the role of multinationals in the aggregate U.S. manufacturing employment decline, and linking this to their import patterns. Section 3 develops the partial equilibrium model, lays out the structural estimation and discusses the results. Section 4 implements the general equilibrium model and performs quantitative exercises. Section 5 concludes. Details of our data and various robustness exercises are contained in the Appendix.

2 Data and Stylized Facts

This section presents a set of stylized facts key to understanding the role of multinationals in the decline in U.S. manufacturing. To uncover these facts, we rely on a new dataset that contains production and trade information of the universe of U.S. manufacturing firms, augmented with multinational ownership and affiliate information. With this data, we show that:

1. U.S. multinationals were responsible for a disproportionate share of the aggregate manufacturing decline,
2. U.S. multinationals experienced lower employment growth than a narrow control group of establishments with similar characteristics,
3. establishments transitioning into U.S. multinational status experienced prolonged job losses while the parent firm increased imports of intermediates.

2.1 Data

Much of this paper relies on a number of restricted-use Census datasets that we have augmented with indicators of multinational affiliate and ownership status. Studying the manufacturing sector over a period of time that spans two distinct industrial classification systems is a challenging task. To create a consistent definition of manufacturing for the period 1993-2011, we apply a new concordance between the SIC/NAICS classification changes that is described in Fort and Klimek (2015). We supplement this concordance with our own set of fixes to account for known data issues, and apply it to the Longitudinal Business Database (LBD), a longitudinally-consistent dataset comprising the universe of all business establishments in the U.S. See Appendix A.1 for more details on the construction of the consistent manufacturing sample.

To identify multinational firms in the Census data, we use a new set of variables describing the international activity and ownership characteristics of U.S. firms. This information comes from a year-by-year link to a set of directories of international corporate structure. To ensure that the multinational identifiers are consistent across time, we develop a series of checks and corrections to minimize any spurious switching of firm status during our sample. For a description of these methods, as well as a summary of the data linking methodology, see Appendix A.[1]

[1] A growing literature has used alternative data sources to identify multinationals operating in the U.S. A number

The final core piece of our data is annual information on imports and exports at the firm level. We use the Longitudinal Foreign Trade Transactions (LFTTD) dataset, which contains the universe of U.S. trade transactions, linked to the firms engaged in such trade. Information in the LFTTD includes the date, value, quantity, and detailed product information (HS10) along with whether the particular transaction was conducted between related parties or at arms-length. To analyze the scope for U.S. firms to transfer portions of their domestic supply chain abroad, we utilize a novel procedure for classifying firm-level imports into those intended for further manufacture (intermediate goods) and those destined for consumption (final goods). See Boehm, Flaaen, and Pandalai-Nayar (2014) or Appendix A.4 for more details on this procedure.

2.2 Facts on Foreign Sourcing and Employment Decline

An aggregate picture of the decline in manufacturing emerges from basic statistics pertaining to our sample. The number of establishments we classify as manufacturing falls from nearly 355,000 in 1993 to under 259,000 in 2011. Table 1 shows that the annual rates of decline have been highest in U.S. multinationals and purely domestic, non-trading establishments. The only group to have experienced an increase in net establishments during this period is foreign multinational firms. This group serves as a reminder that supply chain restructuring could also stimulate U.S. employment.[2]

The employment counts in Table 2 show a similar picture of aggregate decline. Total manufacturing employment in our sample decreases from nearly 16 million workers in 1993 to 10.26 million in 2011. U.S. multinational establishments constituted 33.3% of the 1993 manufacturing employment but contributed 41% of the subsequent overall decline. While employment at other exporting and importing establishments grew in the first decade of the sample, U.S. multinationals have experienced a steady secular decline throughout our sample.

Concurrent with this employment decline has been a large increase in the participation of trade by U.S. firms. We document the fraction of firms participating in intermediate input sourcing,

of papers including Ebenstein, Harrison, and McMillan (2015) and Ramondo, Rappoport, and Ruhl (2014) have used data from the BEA to study multinationals. This is a survey and does not contain non-multinational firms. Most studies of offshoring in the U.S. have been at the industry level. Bernard et al. (2010) use firm level trade data from the U.S. Census Bureau and identify firms as multinationals based on their related party imports. This does not permit a distinction between U.S. and foreign multinationals, and rules out non-trading multinationals by assumption. Other approaches include using Orbis data (Cravino and Levchenko, 2014), and data from Dun and Bradstreet (Alfaro and Charlton, 2009).

[2]Table 3 shows that the decline in multinational firms has not been as severe as the decline in multinational-owned establishments. In the next section, we will show that the extensive margin of establishment shutdown plays an important role in understanding the decline of employment in U.S. multinationals.

separately based on whether it occurs at arms length or intra-firm, in Table 4. We split the firms into U.S. multinationals and other firms (this group includes the few foreign multinationals in our sample). The fraction of U.S. multinationals participating in arms-length input sourcing from developing countries has increased by nearly 30 percentage points, and the fraction sourcing related party inputs from these countries has doubled. In contrast, the share of firms sourcing from developed countries has only increased about 10 percentage points during our sample period. This fact motivates our analysis in later sections, which will look at sourcing patterns for developing and developed country groups separately.

2.2.1 Overall Employment Growth Differential of Multinationals

A number of establishment characteristics have been shown to be correlated with employment growth rates.[3] To the extent that any of these well-known characteristics are correlated with multinational status, attributing the decline in employment to the presence of offshore operations would be misleading. Therefore, to account for these establishment-level characteristics, we construct a set of dummy variables from the interactions of firm age, industry, establishment size, and year. More specifically, each dummy variable takes the value one if an establishment belongs to a cell defined by the interaction of the approximately 250 4-digit manufacturing industries in a year, 10 establishment size categories, and 4 firm-age categories. The setup implies around 16000 cells in the specifications pooling across years 1993-2011.[4] We fit the following regression:

$$e_{it} = \alpha + \beta M_{it} + \mathbf{\Gamma X}_{it} + u_{it} \tag{1}$$

where e_{it} is the establishment growth rate, M_{it} is an indicator for establishments owned by a U.S. multinational, and $\mathbf{X}_{it}$ is the vector of dummy variables identified above.[5]

Table 5 presents the results from this specification, pooled across all years of our sample. The

[3] See Haltiwanger, Jarmin, and Miranda (2013) for a recent example. In Appendix B.1 we decompose the within-group employment patterns into job creation/destruction rates, separated by intensive and extensive margins.

[4] If no multinational establishment exists in a particular cell, we drop that cell from the analysis. We also drop cells that contain only multinational establishments. Our establishment size categories are 0-4,5-9,10-24,25-49,50-99,100-249,250-499,500-999,1000-1999 and 2000 and above and the firm-age categories are 0-1,2-5,6-12 and greater than 12. We obtain firm-age from the LBD firm-age panel. The age of a firm is defined as the age of its oldest establishment.

[5] The growth rate is calculated following Davis, Haltiwanger, and Schuh (1996) and is defined as: $e_{i,t} = \frac{emp_{i,t+1} - emp_{i,t}}{0.5*(emp_{i,t+1} + emp_{i,t})}$

inclusion of records of zero employment before births and after deaths determines whether the measured effect captures the establishment level entry and exit margin. When pooling across years (1994-2011), and focusing only on the intensive margin, we find that multinational establishments have a slightly positive growth rate differential of 1.9 percentage points relative to non-multinational establishments. Once the extensive margin is accounted for, however, this differential changes sign and becomes significantly negative. This is consistent with the strong negative net job-destruction rates at the extensive margin in the analysis in Appendix B.1, and points to establishment closure as key to understanding employment declines in multinationals.

To understand the impact of this establishment-level result on overall employment within a firm, we run the same pooled specification with the firm as the unit of analysis. Here, we find coefficients that are significant and strongly negative: considering only the intensive margin, a multinational firm has a 1-2 percentage point lower employment growth rate than a non-multinational firm. This negative differential increases to 3 percentage points once the extensive margin (firm entry and exit) is included. Clearly, the effects of establishment closure within the multinational firm dominate any increases in employment at existing establishments, leading to aggregate decline.[6]

Table A6 displays results from this specification with different subsamples and additional controls for robustness. We conclude this set of stylized facts by examining employment and trade of establishments which become part of a multinational firm around the transition time.

2.2.2 Evidence Using an Event-Study Framework

While previous sections established the role of multinationals in the U.S. manufacturing decline, this section links this fact to their importing patterns. We analyze the change in outcomes (employment or trade) of establishments that transition into multinational status relative to a predefined control group. Using this event-study framework, we find new multinational plants are characterized by significantly lower employment growth and higher intermediate input imports.

We first divide establishments into four mutually exclusive groups: purely domestic and non-exporting, exporting, owned by a U.S. multinational or owned by a foreign multinational. An

[6] A simple aggregation exercise based on our employment weighted regression results tells us the number of jobs lost in U.S. multinational firms relative to the control group. The growth rates implied by the employment weighted specification can be directly applied to multinational employment in the sample year by year, to arrive at this number. Our estimates imply 2.02 million jobs were lost in these firms relative to a narrowly defined control group. Further details are provided in Appendix B.

establishment's state is then defined by the group it belongs to. We next explore whether changes in establishment state are an important feature of our data. To calculate the average transition rates of establishments, we divide the number of establishments transitioning from one state to another in year $t+1$ (including those that retain state) by the total establishments of that type in year t. Table 6 reports the results.

While infrequent, the transition of establishments into a multinational status provides an opportunity to assess the relationship between multinational structure and establishment-level employment dynamics in an event-study framework.[7] There have been several other recent papers that have analyzed such events for other countries, such as Barba-Navaretti, Castellani, and Disdier (2010) [Italy and France], Hijzen, Jean, and Mayer (2011) [France], and Debaere, Lee, and Lee (2010) [South Korea].[8][9]

Consider a set of establishments that transition into a multinational firm between y and $y+1$, and define a control group of similar establishments that do not transition into a multinational firm in that year.[10] For a transitioning establishment, this control group is defined as non-transitioning establishments within the same narrowly defined cells of firm age, establishment size, and 4-digit industry we utilized above. We then compare the time-path of employment growth rates of the transitioning establishments to their control group.[11]

As is clear from the table of transitions, we have relatively few multinational transitions in a given year. To gain statistical power, we therefore pool the available transitions across years and stack the datasets with "treatment" and control groups corresponding to each year of transition, which we refer to as the "event" year. We then run the following specification:

[7] Table 6 shows that multinationals have relatively high exit rates, and there are low transition rates into and out of multinational status overall. However, the large number of establishments per year in our sample provides sufficiently many transitions for our analysis.

[8] The estimated effects on employment vary across these papers, which likely reflects in part differences in data construction and sample period. For example, Hijzen, Jean, and Mayer (2011) looks at a 6 year window (t-2,t+3), forces a balanced panel (removing extensive margin effects), and constructs the control variables based on $t-2$ firm-level characteristics. Barba-Navaretti, Castellani, and Disdier (2010) look only at effects during the $t+1$ to $t+3$ period, use the Orbis dataset for the control group, and use $t-1$ for the control variables.

[9] For an application of a similar methodology to private equity transactions, see Davis et al. (2014).

[10] Note that a non-multinational establishment could either be acquired by an existing multinational firm, or the firm owning the establishment could open up operations abroad. Our results are broadly similar when considering each of these groups separately.

[11] These cells are defined in the year prior to transition, and remain constant for a given transitioning establishment across years. We drop any establishments in the control group that exit in year y, to match the implied conditioning of the survival of the treated establishments in that year. In addition, we require the establishment to have existed for at least one year prior to the potential transition, for a total minimum establishment age of 3 years.

$$e_{ik}^{y} = \mathbf{\Gamma_{ik}^{y} X_{i}^{y}} + \sum_{k=-5, k \neq 0}^{10} \delta_k T_{ik}^{y} + u_{ik}^{y}, \tag{2}$$

where the variable T_{ik}^{y} is equal to one for transitioning establishment i in year k relative to the year of transition y. (We exclude the transition year $k = 0$.) The vector $\mathbf{X_i^y}$ corresponds to the interaction of controls utilized above, and is fixed at time $k = -1$ for each event year so that the comparison groups remain the same over time. Note that the control groups are defined within an event year (i.e. differ across event years).

An establishment can appear multiple times in this specification. If the establishment exists for several years as a non-multinational until it transitions into multinational status, the establishment would show up in (potentially) several different event years: First as part of a control group for some other transitioning establishment, and then, once, as part of a "treated" group of plants in the year of its own transition. This fact has implications for the way that standard errors are calculated. It implies the need to cluster in cells that include the event year — in order to account for potentially correlated errors across the event — in addition to clustering by plant. We utilize the methodology for two-way clustering described in Cameron, Gelbach, and Miller (2011), which also allows for high-dimensional fixed effects.[12]

We use this same structure to measure the effect of multinational transitions on trading behavior; we simply replace the e_{ik}^{y} with a measure of trade: IM_{ik}^{y} or EX_{ik}^{y}. Such trade can be separately analyzed based on whether it is intra-firm, or composed of intermediate/final goods.[13]

Figure 2 shows the estimates of δ_k. Establishments that transition into multinational status experience a relative increase in their employment growth rates in the first two years. This behavior is consistent with the notion that an expansion of international activity is positively correlated with business outcomes for that firm. Subsequent years, however, show a persistently negative effect in employment growth, on the order of roughly 3-6 percentage points relative to the control group. This slight increase followed by a persistently lower employment growth rates could be explained by an initial domestic growth that coincides with (and serves to support) multinational

[12] The results are robust to clustering by firm instead of plant.

[13] A transitioning establishment associated with a complete firm identifier change could be associated with a level shift in the value of trade, a feature which would present significant challenges in interpreting the results. To prevent this complication, we restrict the sample in the analysis of trading outcomes to only those establishments that retain the same firm identifier from years $t-1$ to $t+1$. Conducting the identical employment analysis using this reduced sample yields similar results.

expansion. Such growth could include time spent by the firm learning to replicate processes within the establishment abroad. Following a successful expansion, the firm may then choose to shut down or downsize duplicated firm activities.[14] In future work, we will attempt to disentangle these competing explanations.

Our results point to the importance of studying a long horizon to understand the consequences of offshoring. We find stronger negative effects on employment than similar analyses for other countries. This discrepancy may reflect differences in the length of time under study (the papers cited above look only at the first 2-3 years following a foreign expansion), or the extent to which other studies adequately account for the extensive margin of plant/firm closings in their analysis.

To examine the role of import substitution in this decline, we estimate equation (2) after replacing the left hand side with firm-level intermediate imports (split by related party and arms-length). Figure 3 shows estimates of δ_k pertaining to imports. The figure demonstrates that transitions are associated with sizeable increases in both related-party and arms-length intermediate imports. This evidence suggests significant substitution between foreign imports and domestic employment.[15][16]

In order to attach a causal interpretation to these results, one would need to assume that the assignment to treatment (transitioning to a multinational status) is random conditional on the large set of observables we use in constructing the controls. On the one hand, after conditioning on this set of size, industry, and age categories, the residual variation may be small enough to make this assumption plausible. On the other hand, there may yet be unobserved covariates that are correlated with the treatment allocation, and thus we prefer to characterize these results as highly suggestive rather than directly causal.

[14] An alternative explanation involves transactions where the establishment is acquired by a multinational firm. Such cases often include mandatory periods where the employees cannot be laid off. In short, it might take a few years to wind down an establishment.

[15] Although the pre-transition levels are slightly higher for the arms-length imports, which suggests that the set of controls does not completely equivalize characteristics between the transitioning and control plants, the differences are small and trends are flat prior to the period of transition. As most transitions into multinational plants are by plants that belong to exporting/importing firms, and we do not condition on export status in creating a control group, the slight difference in arms-length imports is unsurprising.

[16] We demonstrate the robustness of this result to alternative specifications of firm-level trade in Appendix B.

2.2.3 Why multinationals?

The stylized facts above demonstrated that multinationals have a consistent negative employment growth rate differential, and that establishments transitioning into multinational status reduce their employment while their parent firms increase their imports. But is it the ownership (partial or total) of establishments abroad that leads to such employment declines? Or is it simply supply chain restructuring through foreign sourcing of inputs?[17] In other words, why multinationals?

To assess the role of supply chain restructuring overall relative to that occurring within multinationals, we re-run our analysis in Section 2.2.2, but consider non-multinational importer transitions instead of the multinational transitions. Intuitively, if the presence of *any* arms length imports are a sufficient indicator of significant supply chain restructuring, employment in these transitioning establishments should display a similar time-path to the multinational transitions in Figure 2. The results for employment growth differentials relative to a control group – consisting of non-multinational domestic firms based on the narrowly-defined cells of establishment characteristics as before – are shown in Figure 4. Clearly, these establishments do not display such persistent relative employment differentials. This evidence rules out the hypothesis that the presence of some arms-length imports is sufficient to predict relative employment declines. Further, we note that multinationals import the vast majority (over 90% on average) of intermediate inputs in our sample, as shown in Table 11 and discussed in the next section.[18] The importing transitions here therefore assess employment outcomes at firms that are primarily importing final goods, or importing small quantities of intermediates. This could account for the lack of employment differentials, as the degree of supply chain restructuring is minimal.[19] In fact, the tight overlap between foreign sourcing of intermediates and multinationals does not permit us to separately identify a multinational employment effect from the employment effects of foreign sourcing.

Why is foreign sourcing of intermediates concentrated in multinationals? Our data permit a closer look at whether there is a relationship between inter and intra-firm imports which lead to a greater degree of overall global production sharing in multinationals. While the share of related-party imports of multinationals is not significantly different to that of arms-length (roughly 53 vs

[17] In supply chain restructuring, we include restructuring within firms sourcing only at arms-length.

[18] This pattern is robust to excluding foreign multinationals from our sample.

[19] We do not separate the non-multinational importer transitions into new importers of intermediates and new importers of final goods, as we did not base the core analysis around multinational transitions along these lines.

47 percent on average in our sample), perhaps there exist complementarities between intra- and inter-firm imports. We explore this hypothesis by estimating the following regression for the sample period 1993-2011:

$$\log IMP^{AL}_{ijkt} = \alpha_{ijt} + \gamma_{kt} + \beta \log IMP^{RP}_{ijkt} + \epsilon_{ijkt}. \tag{3}$$

Here i is the firm, j is the partner country, k is the product code, and t is time. Hence, the α_{ijt} are firm-country-time fixed effects and the γ_{kt} are product-time fixed effects. The β coefficient then captures the extent to which a firm's AL and RP imports scale together, after absorbing common time-varying firm-by-country, or product shocks.

The results from this regression confirm that sourcing inputs within the firm in a particular foreign location induces more arms-length sourcing as well — even in narrowly defined product categories. This complementarity helps explain the concentration of imports within multinationals in our sample (see Table 7), and is presumably the reason their supply chain restructuring is large enough to show large employment effects. Underlying explanations for this finding could include network effects that enable firm sourcing closely related products from suppliers in the same countries both at arms-length or intra-firm, or lower fixed costs of joint arms-length/related-party imports than of each approach separately. We incorporate the last dimension in our structural model in the following section.

2.2.4 Linking firm-level employment to imports

In the next section, we will specify a structural model to causally link employment outcomes to foreign sourcing at the firm-level. Why do we not explore this mechanism using a reduced form specification with firm level employment regressed on imports, with appropriate instruments to capture foreign supply shocks? The reason is simple: given our data, it is difficult to construct an instrument with predictive power for firm-level imports that is also uncorrelated with firm size. For instance, a commonly used instrument is the "World Export Supply" measure, which captures supply shocks in a partner country (see Acemoglu et al. (2014) for an application to U.S. industries and Hummels et al. (2014) for a firm-level application in Denmark). Constructing this instrument with predictive power at the firm-level requires weights based on variation in the products and countries from which the firm sources. However, such weights induce a correlation with firm size because size is tightly linked to firm sourcing patterns.[20] Similar arguments apply to other commonly used instruments such as transport costs (larger firms with more sourcing destinations will source from farther away) and tariffs (larger firms are more likely to import from countries outside of a free trade agreement). One could use these instruments with the hope that firm size controls purge the instrument of this correlation, but whether this is so would remain questionable.

3 A Framework of Offshoring

We next build a structural model featuring firms' choices of supply chain structure to explore whether foreign sourcing can explain the observed changes in employment. Firms can select a sourcing *location* for their intermediates, as well as a sourcing *mode*: whether to produce intra-firm or to source from outside the firm. Intra-firm production abroad is the defining characteristic of a multinational in this model, reflecting the vertical supply chain structure of U.S. multinationals in our data. Much of the literature assumes perfect technology transfer within a firm or to its suppliers, but empirical evidence for this assumption is lacking. We therefore adopt a more general specification that allows for imperfect technology transfer across sourcing locations and modes.

We show that the model's predictions for the relationship between domestic employment and imports of intermediates depends only on a single structural constant, which is a function of two elasticities. We estimate this structural constant using the microdata at our disposal and find strong evidence that imports of intermediates substitute for domestic employment. Note that the model in this section is partial equilibrium in the sense that it describes only the manufacturing sector in the Home (U.S.) economy. The next section embeds this partial equilibrium framework in a multi-country general equilibrium model.

3.1 Demand for Manufacturing Goods

The consumer derives utility from a constant elasticity of substitution bundle of differentiated manufacturing goods and allocates a fraction of income E to the purchase of this bundle. Let $x(\omega)$ and $p(\omega)$ denote the quantity and price of a variety ω. Taking prices as given, the consumer

[20] Hummels et al. (2014) do not face this problem as they have detailed worker-level information within firms, which identifies their wage effects. Unfortunately, we do not have worker-level information.

maximizes

$$X=\left(\int_{\omega\in\Omega}\left[s\left(\omega\right)\right]^{\frac{1}{\sigma}}\left[x\left(\omega\right)\right]^{\frac{\sigma-1}{\sigma}}d\omega\right)^{\frac{\sigma}{\sigma-1}} \tag{4}$$

subject to the constraint

$$\int_{\omega\in\Omega}p\left(\omega\right)x\left(\omega\right)d\omega=E.$$

The parameter σ is the elasticity of substitution between manufacturing goods, Ω is the set of varieties produced in the country and $s\left(\omega\right)$ is a variety-specific weight. Notice that the final manufacturing varieties are not traded. From the first order conditions of this problem, we obtain the demand functions

$$x\left(\omega\right)=s\left(\omega\right)EP_X^{\sigma-1}p\left(\omega\right)^{-\sigma} \tag{5}$$

for each variety ω, where P_X is the manufacturing price index

$$P_X=\left(\int_{\omega\in\Omega}s\left(\omega\right)p\left(\omega\right)^{1-\sigma}d\omega\right)^{\frac{1}{1-\sigma}}. \tag{6}$$

3.2 Firms

There is a mass M of monopolistically competitive firms. We assume that these firms are heterogeneous along three dimensions: the weight assigned to their variety s, the vector of fixed costs $\mathbf{f}$ (discussed further below), and the scalar φ, which broadly captures the firm's productivity. We refer to φ as the firm's type, and discuss the precise mapping between φ and firm productivity below.[21] A firm is therefore fully described by the tuple $(\varphi,\mathbf{f},s)$.[22]

Each firm uses a unit continuum of intermediates, indexed ν, in the production of their unique variety. The production function is

$$x\left(\varphi,\mathbf{f},s\right)=\left(\int_0^1 x\left(\nu,\varphi,\mathbf{f},s\right)^{\frac{\rho-1}{\rho}}d\nu\right)^{\frac{\rho}{\rho-1}}. \tag{7}$$

Hence, the intermediates are imperfect substitutes with production elasticity of substitution ρ. Letting $p\left(\nu,\varphi,\mathbf{f},s\right)$ denote the price of variety ν for firm $(\varphi,\mathbf{f},s)$, cost minimization in competitive

[21] Note that type here does not refer to quality. Rather, we use it as an index of a firm that is related to firm productivity.

[22] Each firm produces a variety ω, so the tuple describes the variety ω. For brevity, we suppress the index ω for the rest of this section.

markets implies that the unit cost of $x(\varphi, \mathbf{f}, s)$ is

$$c(\varphi, \mathbf{f}, s) = \left(\int_0^1 (p(\nu, \varphi, \mathbf{f}, s))^{1-\rho} \, d\nu \right)^{\frac{1}{1-\rho}}. \tag{8}$$

The demand shifter s does not impact the firm's supply chain structure and therefore we drop this index unless it is necessary for clarity.

3.2.1 Supply chains

As we observe both significant arms-length and intra-firm intermediate input imports in the data, we allow firms the choice of integrated or arms-length sourcing within each location decision. Sourcing inside the firm is indicated by I and sourcing outside the firm by O. Consistent with our classification above, we distinguish among three possible sourcing locations, Home (H), developing (S), and developed (N). Hence, the elements of the set $\mathcal{J}$ of possible sourcing locations and modes for any variety are

1. inside the firm, at home (HI),
2. from a domestic supplier (HO),
3. at arms length from a developed country (NO),
4. inside the firm in a developed country (NI),
5. at arms length from a developing country (SO),
6. inside the firm in a developing country (SI).

We model the firm's problem as follows. First, the firm chooses its sourcing strategy $J(\varphi, \mathbf{f})$, a subset of $\mathcal{J}$. For each intermediate ν, the firm receives a price quote from each element in this set. The benefit of a larger sourcing strategy is therefore a wider range of price quotes resulting in lower input costs. On the other hand, each sourcing strategy requires an ex-ante fixed cost payment. Given their type, firms select the best option among these combinations of production efficiencies and fixed cost payments. The optimal choice of sourcing strategy will be discussed in greater detail below. For now we assume that the set $J(\varphi, \mathbf{f})$ is given.

Intermediate goods production Let j denote an element of firm $(\varphi, \mathbf{f})$'s sourcing strategy $J(\varphi, f)$. Intermediates in sourcing location/mode j are produced with production function[23]

$$x_j(\nu, \varphi) = \frac{h_j(\varphi)}{a_j(\nu)} l_j(\nu, \varphi). \tag{9}$$

The function $h_j(\varphi)$ determines the mapping from the firm's type φ to the productivity of its supplier in j. To allow for maximal generality, we initially make no assumption on the forms of h_j, $j \in J(\varphi, \mathbf{f})$, except that they are weakly increasing. We refer to h_j as the technology transfer functions. Notice that our specification nests the common assumption of perfect idiosyncratic technology transfer ($h_j(\varphi) = \varphi$), for all $j \in J$.

As in Eaton and Kortum (2002), the input efficiencies $1/a_j(\nu)$ are drawn from the Frechet distribution with location parameter T_j and dispersion parameter θ. That is, $Pr(a_j(\nu) < a) = 1 - e^{-T_j a^\theta}$. While we do not explicitly model contracting frictions or other reasons that affect whether firms integrate or source at arms-length, we allow the parameters T_j to vary across sourcing modes.[24] This assumption accommodates a number of real-world features, for instance, that arms-length suppliers in the South may have poorer quality than those that would commonly integrate with a U.S. multinational. In that case $T_{SO} < T_{SI}$, implying, on average, lower productivity draws $1/a_{SO}(\nu)$ than $1/a_{SI}(\nu)$.

Suppose the inverse productivity draws $a_j(\nu)$ have materialized. Then, taking prices as given, a potential supplier of variety ν in location/mode j maximizes

$$p_j(\nu, \varphi) \frac{x_j(\nu, \varphi)}{\tau_j} - l_j(\nu, \varphi) w_j \tag{10}$$

subject to the production function (9). Here, w_j and τ_j denote wages and iceberg transport costs. If the quantity demanded is positive and finite, optimality requires that the potential producer sets price equal to marginal cost

$$p_j(\nu, \varphi) = \frac{\tau_j a_j(\nu) w_j}{h_j(\varphi)}. \tag{11}$$

We assume that $w_{HI} = w_{HO} = w_H$ and $\tau_{HI} = \tau_{HO} = 1$.

[23] We assume that labor is the primary input into production. Assuming a Cobb-Douglas function of capital and labor would not affect our results.

[24] See for instance Antràs (2005), Antràs and Helpman (2004) and Antràs and Chor (2013) among others for theories of intra-firm production.

3.2.2 Basic model implications

Faced with price quotes from every location/mode in their sourcing strategy $J(\varphi, \mathbf{f})$, firms select the cheapest source for each intermediate ν. The distributional assumption together with basic algebra implies that the share of intermediates sourced from j is the same as the cost share of inputs from j, and equals

$$\chi_j(\varphi, \mathbf{f}) = \frac{T_j h_j(\varphi)^{\theta}(\tau_j w_j)^{-\theta}}{\sum_{k \in J(\varphi, f)} T_k h_k(\varphi)^{\theta}(\tau_k w_k)^{-\theta}}. \tag{12}$$

Clearly locations/modes with greater T_j will have larger sourcing shares. Note that $\chi_j(\varphi, \mathbf{f})$ depends on the firm's type φ as long as $h_j \neq h_k$ for some $j, k \in J(\varphi)$. We present evidence for systematic relationship between the sourcing shares and firm type below. Since the sourcing shares depend on the sourcing strategy $J(\varphi, \mathbf{f})$, they also depend on the fixed cost draws $\mathbf{f}$ that a firm must pay to set up its supply chain.

Optimal input sourcing also implies that the unit cost function (8) becomes

$$c(\varphi, \mathbf{f}) = (\gamma)^{\frac{1}{\theta}}[\Phi(\varphi, \mathbf{f})]^{-\frac{1}{\theta}} \tag{13}$$

where $\gamma = \left[\Gamma\left(\frac{\theta+1-\rho}{\theta}\right)\right]^{\frac{\theta}{1-\rho}}$ and Γ is the gamma function, and

$$\Phi(\varphi, \mathbf{f}) = \sum_{j \in J(\varphi, \mathbf{f})} T_j h_j(\varphi)^{\theta}(\tau_j w_j)^{-\theta}. \tag{14}$$

Equation 14 summarizes the firm's efficiency at producing its unique variety. We refer to this term as the firm's (overall) production efficiency. As is intuitive, firms of higher types and firms with more sourcing locations/modes have greater values of Φ and lower unit costs. Notice that neither the cost shares (12) nor the unit costs depend on the quantity the firm produces.

We next turn to the problem determining the firm's optimal size. Given its unit costs, the firm chooses the price for its product to maximize flow profits

$$\tilde{\pi}(\varphi, \mathbf{f}) = p(\varphi, \mathbf{f})\, x(\varphi, \mathbf{f}) - c(\varphi, \mathbf{f})\, x(\varphi, \mathbf{f}) \tag{15}$$

subject to the demand function (5). The firm optimally sets its price to a constant markup over marginal cost, $p(\varphi, \mathbf{f}) = \frac{\sigma}{\sigma-1} c(\varphi, \mathbf{f})$. It is then possible to express revenues as

$$R(\varphi, \mathbf{f}) = s\Sigma P_X^{\sigma-1} \left[\Phi(\varphi, \mathbf{f})\right]^{\frac{\sigma-1}{\theta}}, \tag{16}$$

where $\Sigma = \left(\frac{\sigma-1}{\sigma}\right)^{\sigma-1} \gamma^{\frac{1-\sigma}{\theta}} E$ is a constant. In particular, the elasticity of firm revenues (a measure of firm size) with respect to production efficiency Φ is $\frac{\sigma-1}{\theta}$. As we will see below, this structural constant is critical for the employment consequences of foreign sourcing.

3.2.3 The choice of the firm's sourcing strategy

Prior to selecting its sourcing strategy the firm learns its type φ and its vector of fixed cost draws $\mathbf{f}$. In this partial equilibrium version of the model, we assume that domestic sourcing (HI and HO) does not require a fixed cost payment. In contrast, selecting a sourcing strategy $J \neq \{\text{HI,HO}\}$ requires payment of a fixed cost f_J. The vector $\mathbf{f}$ is comprised of 16 fixed cost draws, one for each J in the power set of {NO,NI,SO,SI}.

After learning φ and $\mathbf{f}$, the firm selects its sourcing strategy $J \subset \mathcal{J}$ to maximize expected profits, which can be expressed as

$$\mathbb{E}[s] \frac{\Sigma}{\sigma} P_X^{\sigma-1} \left[\Phi(\varphi, \mathbf{f})\right]^{-\frac{1}{\theta}} - w_H f_J. \tag{17}$$

Here, w_H is the wage in the Home country and fixed costs are expressed in units of labor. $\mathbb{E}$ is the expectations operator over the distribution of s. Recall that s is a firm-specific demand shifter. We assume that the realization of s is unknown at the time the firm chooses its sourcing strategy. This assumption captures the uncertainty a firm faces between setting up its production structure and selling its product to the final consumer. The demand shifter s helps interpret the structural error in the estimation below. We assume that s is independent of both the firm's type φ and its fixed costs $\mathbf{f}$. We also assume that φ and $\mathbf{f}$ are independent.

The solution to this problem is the firm's optimal sourcing strategy $J(\varphi, \mathbf{f})$ which depends on its type and its fixed cost draws. Figure 5 illustrates the stages of the firm's problem.

3.3 Implications for Domestic Dmployment

We next turn to the model's predictions for the relationship between firms' domestic employment and foreign sourcing. It is easily shown that the labor demanded by firm $(\varphi, \mathbf{f}, s)$ with sourcing strategy $J(\varphi, f)$ is

$$l_{HI}(\varphi, \mathbf{f}, s) = \Theta P_X^{\sigma-1} \frac{sE}{w_H} \underbrace{\frac{T_{HI} h_{HI}(\varphi)^{\theta} (w_H)^{-\theta}}{\Phi(\varphi, \mathbf{f})}}_{\chi_{HI},\ \text{Reallocation effect}} \underbrace{\Phi(\varphi, \mathbf{f})^{\frac{\sigma-1}{\theta}}}_{\text{Size effect}}, \tag{18}$$

where $\Theta = \left(\frac{\sigma-1}{\sigma}\right)^{\sigma} \gamma^{\frac{1-\sigma}{\theta}}$. Since the model is Ricardian in nature, intermediates that are produced at Home inside the firm reflect the firm's "comparative advantage" of intermediate production relative to other sourcing options within its sourcing strategy. The term $l_{HI}(\varphi, \mathbf{f}, s)$ is the labor required for this production.

Consider an increase in foreign competitiveness, for instance through greater values of T_j or lower wages w_j, $j \neq HI, HO$. In partial equilibrium, that is, for fixed expenditures E on manufacturing goods, a constant Home wage w_H, and a fixed manufacturing price index P_X, this increase in foreign competitiveness affects $l_{HI}(\varphi, \mathbf{f}, s)$ only through a change in Φ. Whether domestic employment rises or falls depends on the relative strength of two channels.

First, increased foreign competitiveness shifts a greater fraction of intermediate production towards that location — a reallocation effect. This decreases χ_{HI} and reduces labor demand. On the other hand, greater foreign competitiveness increases the firm's optimal size through an increase in production efficiency Φ. This has a positive effect on labor demand. While the elasticity of χ_{HI} with respect to production efficiency Φ is -1, the elasticity of firm size with respect to Φ is $\frac{\sigma-1}{\theta}$, as is also evident in the expression for revenues (equation 16). The net effect on employment therefore depends on the sign of $\frac{\sigma-1}{\theta} - 1$. If it is negative, the model implies that the reallocation effect dominates and employment declines.

Notice that the same condition characterizes the firm's labor demand after a change in its sourcing strategy, perhaps due to lower fixed costs. If the firm adds an additional location/mode to its set $J(\varphi, f)$, Φ rises and the firm's labor demand falls if and only if $\sigma - 1 - \theta < 0$.

Hence in partial equilibrium, the sign of $\sigma - 1 - \theta$ completely characterizes the within-firm domestic employment response. If $\sigma - 1 - \theta > 0$, one would expect recent productivity gains in emerging markets to increase U.S. manufacturing employment in firms that source from abroad. In contrast, if $\sigma - 1 - \theta < 0$, these same productivity gains should have led to job losses within these firms. We next estimate the value of this key structural constant using microdata on firm sourcing patterns.

3.4 Structural Estimation

Combining equations (12) and (18), the firm's labor demand at home (scaled by w_{HI}/χ_{HI} and logged) can be expressed as

$$\ln \frac{w_{HI} l_{HI} (\varphi, \mathbf{f}, s)}{\chi_{HI} (\varphi, \mathbf{f})} = \Psi_j - \frac{\sigma - 1}{\theta} \ln \chi_j (\varphi, \mathbf{f}) + (\sigma - 1) \ln h_j (\varphi) + \ln s, \; j \in J \subset \mathcal{J} \tag{19}$$

Here, Ψ_j is a fixed effect that contains only constants independent of the firm characteristics $(\varphi, \mathbf{f}, s)$.

The intuition behind the estimating equation (19) is closely related to the scale and reallocation effects discussed above. Since the model predicts that the reallocation effect is independent of parameters (recall that the elasticity of l_{HI} with respect to χ_{HI} is one), it is sufficient to estimate the scale effect. We can do so by focusing on $\frac{w_{HI} l_{HI}}{\chi_{HI}}$ rather than labor demand directly. It is easily verified that this ratio is proportional to firm revenues. Note that the intuition behind this estimating equation is closely related to the key insight of Blaum, LeLarge, and Peters (2015), who show that knowledge of firm domestic expenditure shares and revenues is sufficient to measure decreases in unit costs due to imported inputs in a large class of models of importing firms.[25] Here our equation implies that knowledge of the cost shares of the firm and firm revenues (or expenditure on domestic labor) is sufficient to estimate the scale effect.

Suppose for the moment that $h_j (\varphi)$ was observed. Then, under the assumptions made on s, (in particular, that it is independent of type φ, and fixed costs $\mathbf{f}$, and that it is revealed to the firm only after sourcing decisions are made,) the parameters in equation (19) could be consistently estimated by ordinary least squares. Controlling for the remaining variables, $-\frac{\sigma-1}{\theta}$ captures the scale effect. Intuitively, a smaller share χ_j reflects greater production efficiency resulting in greater firm scale. In contrast, large shares imply a smaller scale.

Unfortunately, $h_j (\varphi)$ is not observed and the estimation of (19) when $h_j (\varphi)$ is subsumed into

[25] Our model falls in this class, and we extend this insight further to all cost shares of a particular firm.In contrast to Blaum, LeLarge, and Peters (2015), our estimation strategy uses all the cost shares of the firm rather than only the domestic cost share, which helps us bound $\frac{\sigma-1}{\theta}$ in the absence of known firm productivity.

the error term yields a biased estimate of $-\frac{\sigma-1}{\theta}$. If $\chi_j(\varphi, \mathbf{f})$ is positively correlated with $h_j(\varphi)$, then the estimate of $-\frac{\sigma-1}{\theta}$ is biased upward. Conversely, if $\chi_j(\varphi, \mathbf{f})$ is negatively correlated with $h_j(\varphi)$, the estimate of $-\frac{\sigma-1}{\theta}$ is biased downward.

The model implies that, conditional on a particular sourcing strategy J, the terms $\chi_j(\varphi, \mathbf{f})$ and $h_j(\varphi)$, $j \in J$ cannot all be positively or all be negatively correlated. The reason is that the sum of a firm's shares over all locations/modes in its sourcing strategy must be one. Therefore, if there exists a sourcing location/mode $j \in J(\varphi, \mathbf{f})$ for which the share $\chi_j(\varphi, \mathbf{f})$ is increasing in firm type φ, some other share, say $\chi_k(\varphi, \mathbf{f})$, $k \in J$, $k \neq j$, must be decreasing in φ. The estimation of (19) by OLS therefore does deliver useful information about $\frac{\sigma-1}{\theta}$. If we condition on a particular sourcing strategy J and estimate (19) for all $j \in J$, the true value must (asymptotically) lie between the highest and the lowest estimate.

This bounding procedure can be refined further, and provides us with a range for the structural constant. We discuss the relevant details in Appendix C. Our first approach to learn about $\frac{\sigma-1}{\theta}$ is to compute the tightest possible bounds, which are reported in the next section.

While the bounds we obtain are useful, a point estimate of $\frac{\sigma-1}{\theta}$ is naturally preferable. Indeed, under certain conditions it is possible to express φ in terms of observables in equation (19) and to estimate $\frac{\sigma-1}{\theta}$ directly. By dividing two sourcing shares from m and k, $m \neq k$, as given in equation (12) by one another and rewriting the result, we obtain

$$\frac{h_m(\varphi)}{h_k(\varphi)} = \left(\frac{\chi_m(\varphi, \mathbf{f})}{\chi_k(\varphi, \mathbf{f})} \frac{T_k}{T_m}\right)^{\frac{1}{\theta}} \frac{\tau_m w_m}{\tau_k w_k}. \tag{20}$$

We next define $\eta_{m,k}(\varphi) = h_m(\varphi)/h_k(\varphi)$. If $\eta_{m,k}$ is invertible, a point we return to below, then it is possible to rewrite (19) for all j and $m \neq k$ as

$$\ln \frac{w_{HI} l_{HI}(\varphi, \mathbf{f})}{\chi_{HI}(\varphi, \mathbf{f})} = \Psi_j - \frac{\sigma - 1}{\theta} \ln \chi_j(\varphi, \mathbf{f}) + (\sigma - 1) \ln h_j \left(\eta_{m,k}^{-1}\left(\left(\frac{\chi_m(\varphi, \mathbf{f})}{\chi_k(\varphi, \mathbf{f})} \frac{T_k}{T_m}\right)^{\frac{1}{\theta}} \frac{\tau_m w_m}{\tau_k w_k}\right)\right) + \ln s. \tag{21}$$

This estimation equation is an instance of a partially linear model. It contains the linear component with regressor $\ln \chi_j(\varphi, \mathbf{f})$ and a second component of unknown functional form which only depends on the observed ratio $\chi_m(\varphi, \mathbf{f})/\chi_k(\varphi, \mathbf{f})$.

A number of semiparametric methods have been developed to consistently estimate $\frac{\sigma-1}{\theta}$ in equation (21).[26] Below we report the results for two approaches. First, we approximate the unknown function of $\chi_m(\varphi)/\chi_k(\varphi)$ by a truncated series expansion (see, e.g. Andrews, 1991), using polynomials as basis functions. Second, we approximate the unknown function of χ_m/χ_k with a step function.[27]

A necessary condition in the derivation of equation (21) is that the function $\eta_{m,k}(\varphi) = h_m(\varphi)/h_k(\varphi)$ is invertible. If this were not the case, the ratio of shares χ_m/χ_k would not provide useful information about the firm's type. It turns out that although h_m and h_k are unknown, our prior analysis allows us to tell whether $\eta_{m,k}$ is invertible. To see this, consider the following procedure. First, fix a particular sourcing strategy J. Next, estimate equation (19) separately for all $j \in J$. Denote by m the sourcing location/mode for which the largest (most upward-biased) estimate of $-\frac{\sigma-1}{\theta}$ is obtained and by k the sourcing location/mode for which the smallest (most downward-biased) estimate of $-\frac{\sigma-1}{\theta}$ is obtained. It then must be that the ratio χ_m/χ_k is strictly increasing in φ and that $\eta_{m,k}(\varphi)$ is invertible (equation 20). Notice that if the highest and lowest estimates of $-\frac{\sigma-1}{\theta}$ from the procedure above are very close together, then $\eta_{m,k}$ is not invertible, but the bias of $-\frac{\sigma-1}{\theta}$ is negligibly small.

Finally, we discuss two practical issues regarding the estimation of $-\frac{\sigma-1}{\theta}$. To allay concerns about measurement error in the shares χ_j, we estimate several specifications using the firm's shares from the previous year as instruments. For robustness, we also estimate equations (19) and (21) after replacing the left hand side variable with the log of firm revenues (recall that the model predicts that revenues are proportional to $\frac{w_{HI} l_{HI}}{\chi_{HI}}$).

Linking the Model and the Data

The structural estimation requires data on firm revenues and cost shares from the various sourcing methods $j \in J(\varphi)$. Revenues and cost share information are constructed from the Census of Manufacturers (CMF) merged with import information from the LFTTD. For revenues, we use the total value of shipments of the firm's manufacturing establishments. Total costs are constructed from information on the cost of materials inputs, firm inter-plant transfers and total machinery

[26]Notice that the constant $\frac{\sigma-1}{\theta}$ is identified when we control nonparametrically for the function $h_j(\varphi)$. Clearly, under the assumptions made above, the error term $\ln s$ is orthogonal to the regressors. Additionally, if fixed costs vary across firms, the share $\chi_j(\varphi, f)$ is not collinear with $h_j(\varphi)$.

[27]More precisely, we partition the range of χ_m/χ_k into fifty percentiles and define fifty indicator variables taking the value one if χ_m/χ_k falls between two consecutive percentiles. We then replace the unknown function of χ_m/χ_k in equation (21) by a step function based on these fifty indicator variables.

expenditures of the firm. We identify intermediate input imports of the firms using a product-level classification method based on the firm's industry. This method is discussed in detail in appendix A.4 and in Boehm, Flaaen, and Pandalai-Nayar (2014).

We use lagged values of the cost shares as instruments in robustness exercises. As the Census is quinquennial and only available in 1997, 2002 and 2007, the lagged values have to be constructed using data from the Annual Survey of Manufacturers (ASM) in 1996, 2001 and 2006. The ASM includes information on all large manufacturing establishments, but is not a complete sample of the smaller plants. Therefore we first construct a firm level total cost using establishments sampled in the ASM, and then scale up this variable using the information on total employment captured within the ASM relative to the total employment in our baseline sample. As our baseline sample is built from the LBD, it contains information on all manufacturing establishments of a firm in a year. The assumption implicit in this procedure is that the firm's cost function is the same across the establishments not captured by the ASM as it is in the surveyed portion of the firm.[28] Table 10 contains mean cost shares for multinationals sourcing from all locations for the three Census years in our sample.

3.4.1 Results and Discussion

We estimate the model in three separate cross-sections in 1997, 2002 and 2007. We find that our estimates of $\frac{\sigma-1}{\theta}$ are remarkably robust both to the method used and to the time period. Table 8 presents the bounds on $\frac{\sigma-1}{\theta}$ by year. As discussed in detail in Appendix C our procedure implies a large number of bounds. To reduce the likelihood of statistical outliers, we report the 80th percentile lower and upper bounds.[29] The widest interval for $\frac{\sigma-1}{\theta}$ is $(0, 0.86]$ in 2002, implying the true parameter value is likely in the range where foreign sourcing is a substitute for domestic employment in a firm.

Prior to estimating $\frac{\sigma-1}{\theta}$ using a semi-parametric regression, we must first show that there indeed exists a function $\eta_{m,k}$ that is invertible. In Appendix C, we show that χ_{HO} is strictly increasing and

[28] The survey methodology of the ASM assigns lower sampling weights to the smallest manufacturing plants. Therefore, if the unit costs of the firm differ across its establishments in a manner correlated with size, this assumption would be invalid. As this assumption only affects the value of instruments for cost shares, it will not bias our results as the instrument remains valid – the cost shares of the missing establishments are unlikely to be systematically correlated with the structural error in the model.

[29] This primarily affects the lower bound, as the upper bound in all estimates is always zero from the theoretical restriction on $\frac{\sigma-1}{\theta}$.

χ_{HI} is strictly decreasing in φ. This implies that $\eta_{HO,HI}$ is invertible. When estimating equation (21), we therefore control for the unknown term using polynomials or step functions in $\frac{\chi_{HO}}{\chi_{HI}}$.

Table 9 presents the baseline results for 1997, 2002 and 2007 using (a) polynomials as basis functions and (b) fifty dummy variables representing size bins. The lower panel of the table contains results for the same specifications with the cost shares instrumented by lagged values.[30] The point estimates obtained for each year and specification lie within the bounds in Table 8. The estimates range from 0.08 to 0.23, confirming that foreign sourcing is a strong substitute for domestic employment for the firms we study. The results are robust to using revenues instead of scaled payroll as a dependent variable (see Table C.3). We further estimate $\frac{\sigma-1}{\theta}$ by industry, to allow for sector-level differences in the scale effect. A kernel density of the estimates is shown in Figure 6. While the industry level estimates vary, for our sample of manufacturing industries they are never larger than one, implying that foreign sourcing is a substitute for domestic employment for all manufacturing industries.

In contrast to our estimates, Antràs, Fort, and Tintelnot (2014) find that $\frac{\sigma-1}{\theta}$ is larger than one. While closely related, their model includes a much larger set of sourcing locations, and does not distinguish between arms-length and related party imports. Further, they estimate σ and θ separately within their framework. Our model implies that estimation of the ratio $\frac{\sigma-1}{\theta}$ is sufficient for understanding the role of foreign sourcing on employment, and our more aggregated structure offers a parsimonious method to estimate this structural constant. Antràs, Fort, and Tintelnot (2014) also include data from several non-manufacturing sectors in their estimation procedure, which likely contributes to the difference in findings. Non-manufacturing sectors might have a stronger scale effect, which could result in complementarity between foreign sourcing and domestic employment.[31] We note that there is a large literature that has estimated both parameters separately in various contexts, and the range of estimates is wide. Our estimates, as well as those in Antràs, Fort, and Tintelnot (2014) are consistent with earlier findings.[32]

[30]We also include estimates for 1993, which is not a Census year. Due to incomplete LFTTD data for 1992, we are unable to instrument the 1993 regressions.

[31]In future work, we hope to explore the differences in the impact of foreign sourcing on manufacturing and services sectors.

[32]In particular, Eaton and Kortum (2002) estimate $\theta = 8.28$ as a baseline. Other estimates include Caliendo and Parro (2015), whose estimates range from 0.37 to 51.08 using sector-level data on manufacturing. In these papers, θ is also the trade elasticity. In our setup, the trade elasticity has a more complex expression. Estimates of σ as the markup in monopolistic competition models usually center around 4 (Hall, 1988). de Loecker and Warzynski (2012), who use firm-level data, find the markup in a CES framework would be 1.16.

3.5 Aggregation

We briefly explore the aggregate implications of our empirical model (21) in partial equilibrium. To do so, we consider the population of U.S. manufacturing firms in 1997 and predict the aggregate employment decline implied by the difference in their sourcing shares between 1997 and 2007 and our estimates of $\frac{\sigma-1}{\theta}$. We use an estimate of 0.2, which is at the upper end of our range of estimates. In this exercise we predict employment changes within firms sourcing from abroad, and first-order effects on their U.S. arms-length suppliers (sourcing from HO).

This procedure requires that we observe firms in both 1997 and 2007. It therefore cannot account for the declines in employment due to offshoring in firms that exited before 2007. Further, it underestimates the intensive margin effect in continuing firms, as some firm identifiers in the data change even though these firms continue to exist.

All else equal, this exercise suggests that 1.3 million jobs were lost due to foreign sourcing. Of this, 0.55 million jobs were lost within multinationals, and the remainder of the losses are due to declines in multinational demand for arms-length sourcing in the U.S. (0.58 million), as well as foreign sourcing by non-multinational firms. To account for general equilibrium effects such as firm entry and changes in aggregate demand, we next turn to a simple general equilibrium extension of our model.

4 General equilibrium

4.1 A general equilibrium extension

The simple model in this section aims to capture general equilibrium features such as firm entry and exit as well as adjustments in aggregate demand driven by changes in the price index of manufactured goods. We should note that we do not present a quantitative trade model that can capture several real-world features such as increases in foreign demand for U.S. manufactured goods. Given the data availability on foreign multinationals, particularly in developing countries, such a model would be extremely hard to calibrate. Rather than fitting the aggregate employment decline, our goal in this section is therefore to explore the quantitative importance of foreign sourcing of intermediates using the simplest possible framework.

Consistent with the sourcing structure in partial equilibrium, we assume a three country world, with the countries labeled Home (H), North (N) and South (S). Since the North and the South have the same economic structure, we only present the optimization problems for the North. In addition to the manufacturing sector X, there is a large absorbing sector which produces a freely traded good Z in each country. We normalize its price to unity.

Home country

Households The representative household in the Home country derives utility from the consumption of X and Z. It supplies L_H units of labor inelastically. The household maximizes utility $Z_H^{\beta} X^{1-\beta}$ subject to its budget constraint $w_H L_H = P_X X + Z_H$. Here, Z_H is Home consumption of the numeraire good. The Cobb-Douglas utility function implies that the Home consumer spends $E = (1-\beta)\, w_H L_H$ on the manufacturing good X and the remainder on Z.

Firms in the Z sector Firms in the freely-traded sector produce with linear technology $Q_H^Z = A_H L_H^Z$. Profit maximization in competitive markets implies that $w_H = A_H$ as long as Q_H^Z is strictly positive and finite.

Firms in the X sector Firms in the X sector set up supply chains and produce as described in Section 3. In this general equilibrium extension we assume that the number of firms is endogenous and determined by the following entry problem which has three stages. In the first stage, there is an unbounded mass of potential entrants who can pay fixed costs f_E to learn their type φ. In equilibrium, the number of entrants M is determined by a zero expected profit condition. Second, after learning their types, entrants must pay an additional fixed cost f_H to set up production in the Home country. Only firms with sufficiently high types φ find it profitable to do so. The lowest type that enters is φ_{LB}. Finally, those firms that produce in the Home country face the problem discussed in Section 3.

Market clearing Labor market clearing in the Home country requires that

$$\begin{aligned} L_H = L_H^Z \quad &+ \quad M\left[\int_0^{\infty}\int_{\varphi_{LB}}^{\infty}\left[l_{HI}(\varphi,s)+l_{HO}(\varphi,s)\right]dG_{\varphi}(\varphi)\,dG_s(s)\right. \\ &+ \quad \left. f_E + f_H\left(1-G_{\varphi}(\varphi_{LB})\right)+\int_{\varphi_{LB}}^{\infty} f_{J(\varphi)}dG_{\varphi}(\varphi)\right]. \qquad (22) \end{aligned}$$

Labor demand on the right hand side consists of demand from the Z sector, demand from the X sector (the first integral) and the labor demand stemming from the various fixed costs. This notation assumes that $f_{\{\text{HI,HO}\}} = 0$.

North

The representative household in the North derives utility only from the freely traded good Z, and supplies L_N units of labor inelastically.[33] Its budget constraint is $w_N L_N = Z_N$. As in the Home country, the production function for good Z is linear, $Q_N^Z = A_N L_N^Z$. Labor market clearing in the North requires that

$$L_N = L_N^Z + M \int_0^\infty \int_{\varphi_{LB}}^\infty \left[l_{NI}(\varphi, s) + l_{NO}(\varphi, s) \right] dG_\varphi(\varphi)\, dG_s(s). \tag{23}$$

We close the model with the market clearing condition for good Z,

$$Z_H + Z_N + Z_S = Q_H^Z + Q_N^Z + Q_S^Z.$$

4.2 Calibration

While the model permits very general sourcing patterns across locations/modes, we find that only a few of these are prevalent in the data. In fact, similar to Antràs, Fort, and Tintelnot (2014), there are regularities in sourcing locations/modes of the following form. First, very few firms source from abroad. Of the ones that do, most firms only import from the North at arms-length. Second, if a firm sources intra-firm from the North, then it is likely to also source from the North at arms-length. A similar pattern can be observed for imports from the South. Firms that source from all locations are typically the largest in terms of revenues.

Given these regularities and the fact that we are interested in sourcing decisions of multinationals, we restrict the equilibrium sourcing strategies to the set

$$\tilde{\mathcal{J}} = \{(\text{HO,HI}), (\text{HO,HI,NO}), (\text{HO,HI,NO,NI}), (\text{HO,HI,NO,NI,SO}), (\text{HO,HI,NO,NI,SO,SI})\}.$$

Table 11 shows the fraction of firms in the data that source according to each of these strategies.

[33] We make this assumption for simplicity as we are primarily interested in manufacturing employment in the Home country. There is no final goods trade in the X sector.

Despite this restriction we capture almost all firms (92.6% in 1997) and almost all trade (95.3% in 1997).[34]

Our calibration procedure proceeds in two steps. We first set a number of parameters equal to their direct analogues in the data or to conventional values in the literature. Second, we choose the remaining parameters to match key features of employment and imports in the manufacturing sector.

The productivity parameters A_H, A_N, and A_S are chosen to match skill-adjusted wages for the U.S., the average country in the North, and the average country in the South. Wage data are obtained from the ILO and skill adjusted using the method in Eaton and Kortum (2002). We define the South as countries with GDP per capita of less than 10 percent of the U.S. in 2000. This threshold implies that China, India, and Brazil belong to the South. The labor endowment in all three countries are set to match the skill-adjusted labor force, taken from the same source.

We next assume that firm types have a Pareto distribution with a lower bound of unity and curvature parameter α_φ. The demand elasticity σ is set to 2.3 and the dispersion parameter θ to 6. These values imply that $(\sigma - 1)/\theta$ is 0.217, roughly consistent with the upper end of our point estimates. We also set $\tau_H = 1$ and $\tau_N = \tau_S = 1.15$. Although these parameters are not important for any of the model's predictions, we note that ρ is set to 1.5 and $\mathbb{E}[s]$ to 1.0025. Finally, we must assume a functional form for $h_j(\varphi)$. We choose a simple exponential, $h_j(\varphi) = \varphi^{\kappa_j}$. We note that this choice provides a reasonable fit to the mid-range of the firm size distribution (see Appendix C for a discussion). We set κ_{HI} to one, and choose values for $\kappa_j, j \neq HI$ that are close to κ_{HI}.[35] Table 12 summarizes the values of the preset parameters.

The remaining parameters of the model are chosen to match key features of our data in 1997. These parameters are T_j, $j \in \mathcal{J}$, $j \neq$ HI, the fixed cost parameters f_J, $J \subset \tilde{\mathcal{J}}$, f_E, and f_H, the Pareto curvature parameter α_φ as well as the expenditure share β. The targets and the fit of the model in equilibrium are summarized in Table 13.

[34] This restriction of sourcing strategies greatly facilitates the numerical solution and calibration of the model because it implies a complete ordering of the sourcing strategies and that higher types choose more complex sourcing strategies.

[35] The exponential assumption is only an approximation to the true functional form of h_j. We therefore have minimal guidance on calibrating these parameters. Choosing values close to 1 ensures these parameters will not significantly influence the quantitative analysis. For robustness, Appendix D, presents the results from the quantitative exercises with alternate choices for κ_j.

4.3 Quantitative exercises

We consider two types of quantitative exercises. First, we compute the employment changes in the model when we change various productivity and fixed cost parameters individually, by an infinitesimal amount. Second, we fit the model to aggregate trade patterns and firm sourcing strategies in 2007, and compute the implied change in the size of the manufacturing sector relative to 1997.

The first panel of Table 14 reports the percent change of manufacturing employment and multinational employment when the technology parameters $T_j, j \in \{NO, NI, SO, SI\}$ are changed by one percent, one at a time. In response to changes in each of these parameters, aggregate manufacturing employment falls.

The general equilibrium effects of these parameter changes are evident in the response of the manufacturing price index, the changes in the mass of firms M and the movement of firms between different sourcing strategies (not shown). As expected, the price index always falls in response to a technological improvement that lowers the unit costs of firms whose sourcing strategy includes that location/mode. Better technology in one particular sourcing location/mode also induces transitions of firms into sourcing strategies that include that location/mode. Similar to the standard Melitz (2003) model, firms of the lowest types face lower demand as a result of the lowered cost for higher type firms. Therefore, the net effect on the mass of firms M is ambiguous. Importantly, for our calibration the general equilibrium effects do not overturn the partial equilibrium result that foreign sourcing substitutes for domestic employment.

Turning to multinational employment, two offsetting effects are in play. As shown in the partial equilibrium model in Section 3, production of intermediates is reallocated towards the location/mode with the technological improvement. However, this effect can be offset by the movement of firms into or out of strategies with that sourcing location/mode. The net effect is therefore ambiguous. For our calibration the second effect implies that multinational employment rises as T_{NI} increases.

The second panel of Table 14 shows the employment change in response to lowered fixed costs for each sourcing strategy f_J. With the exception of f_{NI}, the sign of the response of employment is the same as in the case of a technological improvement. Here, the result is driven by the extensive margin: firms enter the sourcing strategy that has lower fixed costs, reducing domestic employment in response. As above, in most cases general equilibrium effects do not change the predictions in partial equilibrium.[36] The changes in multinational employment are governed by the transition of firms between sourcing strategies. In the case of a decrease in f_{NI} non-multinational firms enter multinational status, leading to an increase in multinational employment. In contrast, decreases in f_{SO} and f_{SI} largely induce firms to switch between sourcing strategies while maintaining their multinational status.

In our second exercise, we first fix the parameters α_φ and β. We then choose the remaining parameters $T_j, j \in \{HO, HI, NO, NI, SO, SI\}$, f_J, $J \in \mathcal{J}$, and f_H to match 2007 import patterns, firm shares and the mean share of intermediates sourced from HO for the group of multinationals sourcing from all locations. Notice, we do not include any employment targets – our goal is to understand the decline in manufacturing generated by the model simply by matching observed import patterns. Table 13 illustrates the fit of our model to our calibration targets for 1997 and 2007.[37]

To match the observed trade patterns in 2007 the technology parameters $T_j, j \in \{HO, HI, NO, NI, SO, SI\}$ uniformly increase. This is shown in Table 15. T_{SO} and T_{SI} increase the most, reflecting the fact that imports from the South grew rapidly over the period 1997 - 2007. Although the fraction of multinationals increased over this time period, fixed costs of sourcing strategies increase between our two calibrations. In this model, firms respond to better technology abroad by entering the sourcing strategies that include foreign sourcing. To match the data – where the fraction of firms in these sourcing strategies has only shown small increases – the model has one counterbalancing force. Fixed costs increase to prevent the fraction of multinationals from rising beyond what is observed in the data. While this might appear counterintuitive, we note that this rise in fixed costs might reflect more complex production structures, which are harder to initially offshore. Further, in this model an increase in T_j is not separable from a decrease in τ_j or w_j. Therefore, the calibrated technology increases reflect a composite change in foreign wages and the variable costs of offshoring.

Targeting 2007 trade patterns results in an employment loss within multinational firms by 28%, slightly larger than that observed in the data (see Table 16). Total manufacturing employment

[36] In response to a decrease in f_{NI}, the fraction of firms with $J = \{\text{HO,HI,NO,NI}\}$ increases substantially. However, the mass of firms also increases in equilibrium, and the lower bound for entry falls (so entering firms are less productive). The net result is an increase in overall employment.

[37] Note that in the base year 1997, T_{HI} is not chosen and normalized to 1. In the second calibration to meet the 2007 targets, T_{HI} is also allowed to increase.

falls by 13% which accounts for roughly half of the observed decline between 1997 and 2007.[38] In addition to the direct employment loss within multinationals, increased foreign sourcing reduces the demand for intermediates from domestic suppliers. Confronted with less demand for their products these suppliers scale down production and thereby contribute to the employment decline as well.

We advice some caution should be taken in the application of these general equilibrium results. This exercise quantifies the decline in manufacturing employment due to increased foreign sourcing alone, and does not account for other factors – such as increased foreign demand for U.S. goods, or foreign multinationals – that could serve to mitigate the overall negative employment results. Indeed as we have shown in 2.2, both the number and employment of foreign multinationals in the U.S. increased during our sample. Finally, our analysis has focused on the effects on manufacturing, and it is important to note that one might suspect U.S. multinationals have increased their non-manufacturing employment. We hope to explore these effects in future research. In Appendix D, we discuss some counterfactual exercises where we only allow techonology parameters or fixed costs to change.

5 Conclusion

We present new stylized facts showing that a disproportionately large share of the manufacturing employment decline in the U.S. can be attributed to U.S. multinationals. Moreover, we find evidence that supply chain fragmentation and offshoring of intermediate input production to developing countries has played an important role in this decline. To closely examine this channel, we illustrate a tight link between domestic employment and firm-level foreign sourcing in a model of endogenous firm sourcing decisions. A key elasticity – of firm size with respect to production efficiency – governs the employment impact of changes in foreign sourcing in this framework. Structural estimation of this elasticity shows that offshoring is a strong substitute for domestic employment.

In our data, offshoring is concentrated within multinational firms, so our finding helps explain the role of these firms in the aggregate manufacturing decline. In general equilibrium, our estimates generate a quantitatively significant decline of the U.S. manufacturing sector. Note that this does not imply aggregate U.S. welfare decreases, as the gains from cheaper manufacturing goods accrue to

[38]We present the declines in the data both for our full sample and for the period 1997 - 2007. As some of the parameters in our model are calibrated using data available only in census years (years ending in 2 or 7), we present the 1997 - 2007 decline and the 1993 - 2011 decline as an additional point of comparison.

the consumer. Further, our focus is on manufacturing alone – it is possible (and indeed, likely) that foreign sourcing is complementary to employment in services in the U.S. This finding has several policy implications. In particular, it emphasizes that policy changes encouraging globalization and integration should take into account the differential impact on manufacturing workers, other workers and the consumer. Such policies can be designed to smooth the transitions for displaced manufacturing workers.

The observed concentration of both arms-length and related-party sourcing of inputs within multinationals could be attributed to several competing channels. In future work, we will assess the underlying reasons behind this strong empirical finding.

References

Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price. 2014. "Import Competition and the Great U.S. Employment Sag of the 2000s." NBER Working Papers 20395, National Bureau of Economic Research, Inc.

Alfaro, Laura and Andrew Charlton. 2009. "Intra-Industry Foreign Direct Investment." *American Economic Review* 99 (5):2096–2119.

Andrews, Donald W K. 1991. "Asymptotic Normality of Series Estimators for Nonparametric and Semiparametric Regression Models." *Econometrica* 59 (2):307–45.

Antràs, Pol. 2005. "Property Rights and the International Organization of Production." *American Economic Review* 95 (2):25–32.

Antràs, Pol and Davin Chor. 2013. "Organizing the Global Value Chain." *Econometrica* 81 (6):2127–2204.

Antràs, Pol, Teresa C. Fort, and Felix Tintelnot. 2014. "The Margins of Global Sourcing: Theory and Evidence from U.S. Firms." Working Paper 154856, Harvard University OpenScholar.

Antràs, Pol and Elhanan Helpman. 2004. "Global Sourcing." *Journal of Political Economy* 112 (3):552–580.

Autor, David, David Dorn, Gordon Hanson, and Jae Song. 2014. "Trade Adjustment: Worker-Level Evidence." *American Economic Review* 129 (4):1799–1860.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6):2121–68.

Barba-Navaretti, Giorgio, Davide Castellani, and Anne-Cèlia Disdier. 2010. "How Does Investing in Cheap Labour Countries Affect Performance at Home? Firm-Level Evidence from France and Italy." *Oxford Economic Papers* 62:234–260.

Bernard, Andrew and J. Bradford Jensen. 2007. "Firm Structure, Multinationals, and Manufacturing Plant Deaths." *Review of Economics and Statistics* 89 (2):193–204.

Bernard, Andrew, J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott. 2010. "Intra-Firm Trade and Product Contractibility." Working Paper 15881, NBER.

Blasnik, Michael. 2010. "RECLINK: A Stata Module to Probabilistically Match Records." Working paper.

Blaum, Joaquin, Claire LeLarge, and Michael Peters. 2015. "The Gains from Input Trade in Firm-Based Models of Importing." NBER Working Papers 21504, National Bureau of Economic Research, Inc.

Boehm, Christoph, Aaron Flaaen, and Nitya Pandalai-Nayar. 2014. "Input Linkages and the Transmission of Shocks: Firm Level Evidence from the 2011 Tōhoku Earthquake." Working paper, University of Michigan.

Braconier, Henrik and Karolina Ekholm. 2000. "Swedish Multinationals and Competition from High and Low-Wage Locations." *Review of International Economics* 8 (3):448–461.

Caliendo, Lorenzo and Fernando Parro. 2015. "Estimates of the Trade and Welfare Effects of NAFTA." Tech. Rep. 1.

Cameron, Colin, Jonah Gelbach, and Douglas Miller. 2011. "Robust Inference with Multiway Clustering." *Journal of Business and Economic Statistics* 29 (2):238–249.

Costinot, Arnaud, Jonathan Vogel, and Su Wang. 2013. "An Elementary Theory of Global Supply Chains." *Review of Economic Studies* 80 (1):109–144.

Cravino, Javier and Andrei Levchenko. 2014. "Multinational Firms and International Business Cycle Transmission." Working paper, University of Michigan.

Davis, Steven J. and John Haltiwanger. 1992. "Gross Job Creation, Gross Job Destruction and Employment Reallocation." *Quarterly Journal of Economics* 107 (3):819–863.

———. 2001. "Sectoral job creation and destruction responses to oil price changes." *Journal of Monetary Economics* 48 (3):465–512.

Davis, Steven J., John Haltiwanger, Kyle Handley, Ron Jarmin, Josh Lerner, and Javier Miranda. 2014. "Private Equity, Jobs, and Productivity." *American Economic Review* 104 (12):3956–90.

Davis, Steven J, John Haltiwanger, and Scott Schuh. 1996. " Small Business and Job Creation: Dissecting the Myth and Reassessing the Facts." *Small Business Economics* 8 (4):297–315.

de Loecker, Jan and Frederic Warzynski. 2012. "Markups and Firm-Level Export Status." *American Economic Review* 102 (6):2437–71.

Debaere, Peter, Hongshik Lee, and Joonhyung Lee. 2010. "It Matters Where You Go: Outward Foreign Direct Investment and Multinational Employment Growth at Home." *Journal of Development Economics* 91 (2):301–309.

Decker, Ryan, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2014. "The Secular Decline in Business Dynamism in the U.S." Tech. rep.

Desai, Mihir A., C. Fritz Foley, and James R. Hines. 2009. "Domestic Effects of the Foreign Activities of US Multinationals." *American Economic Journal: Economic Policy* 1 (1):181–203.

Eaton, Jonathan and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica* 70 (5):1741–1779.

Ebenstein, Avraham, Ann Harrison, and Margaret McMillan. 2015. "Why are American Workers getting Poorer? China, Trade and Offshoring." NBER Working Papers 21027, National Bureau of Economic Research, Inc.

Ebenstein, Avraham, Ann Harrison, Margaret McMillan, and Shannon Phillips. 2014. "Estimating the Impact of Trade and Offshoring on American Workers using the Current Population Surveys." *The Review of Economics and Statistics* 96 (3):581–595.

Feenstra, Robert and Gordon Hanson. 1996. "Foreign Investment, Outsourcing and Relative Wages." In *The Political Economy of Trade Policy: Papers in Honor of Jagdish Bhagwati*, edited by Gene Grossman Robert Feenstra and Douglas Irwin. MIT Press.

Flaaen, Aaron. 2013a. "Constructing a Bridge Between the LexisNexis Directory of Corporate Affiliations and the U.S. Business Register." Working paper, University of Michigan.

———. 2013b. "Multinational Firms in Context." Working paper, University of Michigan.

Fort, Teresa and Shawn Klimek. 2015. "What's in a Name? The Effect of Industry Classification Changes on U.S. Employment Composition." Working paper, U.S. Census Bureau.

Hall, Robert E. 1988. "The Relation between Price and Marginal Cost in U.S. Industry." *Journal of Political Economy* 96 (5):921–47.

Haltiwanger, John C., Ron S. Jarmin, and Javier Miranda. 2013. "Who Creates Jobs? Small vs. Large vs. Young." *Review of Economics and Statistics* 95 (2):347–361.

Hanson, Gordon H., Raymond J. Mataloni, and Matthew J. Slaughter. 2005. "Vertical Production Networks in Multinational Firms." *The Review of Economics and Statistics* 87 (4):664–678.

Harrison, Ann and Margaret McMillan. 2011. "Offshoring Jobs? Multinationals and U.S. Manufacturing Employment." *The Review of Economics and Statistics* 93 (3):857–875.

Helpman, Elhanan, Marc Melitz, and Stephen Yeaple. 2004. "Exports vs FDI with Heterogeneous Firms." *American Economic Review* 94:300–316.

Hijzen, Alexander, Sebastien Jean, and Thierry Mayer. 2011. "The Effects at Home of Initiating Production Abroad: Evidence from Matched French Firms." *Review of World Economics* 147 (3):457–483.

Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang. 2014. "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data." *American Economic Review* 104 (6):1597–1629.

Konings, Jozef and Alan Patrick Murphy. 2006. "Do Multinational Enterprises Relocate Employment to Low-Wage Regions." *Weltwirtschaftliches Archiv (Review of World Economics)* 142 (2):267–286.

Kovak, Brian, Lindsey Oldenski, and Nicholas Sly. 2015. "The Labor Market Effects of Offshoring by U.S. Multinational Firms: Evidence from Changes in Global Tax Policies." Working paper.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6):1695–1725.

Monarch, Ryan, Jooyoun Park, and Jagadeesh Sivadasan. 2013. "Gains from offshoring? Evidence from US Microdata." *US Census Bureau Center for Economic Studies Paper, No. CES-WP-13-20* .

Muendler, Marc-Andreas and Sascha O. Becker. 2010. "Margins of Multinational Labor Substitution." *American Economic Review* 100 (5):1999–2030.

Pierce, Justin and Peter Schott. 2012. "A Concordance Between U.S. Harmonized System Codes and SIC/NAICS Product Classes and Industries." *Journal of Economic and Social Measurement* 37 (1-2):61–96.

———. 2013. "The Surprisingly Swift Decline of U.S. Manufacturing Employment." Working Papers 13-59, Center for Economic Studies, U.S. Census Bureau.

Ramondo, Natalia, Veronica Rappoport, and Kim J. Ruhl. 2014. "Horizontal vs Vertical FDI: Revisiting Evidence from U.S. Multinationals." Working paper, Columbia Business School.

Simpson, Helen. 2012. "Investment Abroad and Labour Adjustment at Home: Evidence from UK Multinational Firms." *Canadian Journal of Economics* 45 (2):698–731.

Slaughter, Matthew. 2000. "Production Transfer within Multinational Enterprises and American Wages." *Journal of International Economics* 50 (2):449–472.

Wasi, Nada and Aaron Flaaen. 2014. "Record Linkage Using Stata: Pre-Processing, Linking and Reviewing Utilities." Working paper, University of Michigan.

Table 1: Summary Statistics: Establishment Counts by Type: 1993-2011

year	Domestic Only	Exporter Only	Importer Only	Exporter & Importer	U.S. Multinational	Foreign Multinational	Total
1993	252,965	41,353	6,911	30,237	17,119	6,178	354,763
2011	159,133	39,034	6,513	31,391	13,488	8,952	258,511
				Average Annual Percent Change			
1993-2011	-2.41	-0.30	-0.31	0.20	-1.25	1.97	-1.65
1993-2001	-1.76	0.49	0.81	1.92	-1.18	1.62	-0.98
2002-2011	-2.97	-0.70	-1.87	-0.87	-1.12	2.84	-2.13

Source: LBD-LFTTD-DCA-UBP as explained in text.
This table reports the establishment counts pertaining to the "constant" manufacturing sample used in section 2.

Table 2: Summary Statistics: Employment Counts by Type: 1993-2011

	Domestic Only	Exporter Only	Importer Only	Exporter & Importer	U.S. Multinational	Foreign Multinational	Total
1993	3,433,510	2,133,327	267,090	3,663,103	5,314,411	1,102,240	15,913,681
2011	1,751,504	1,358,061	181,716	2,614,260	2,975,786	1,380,804	10,262,131
				Average Annual Percent Change			
1993-2011	-3.48	-2.35	-2.01	-1.76	-3.01	1.19	-2.28
1993-2001	-1.72	-0.44	0.74	0.89	-1.69	2.93	-0.49
2002-2011	-4.68	-3.19	-3.89	-3.22	-3.67	0.80	-3.19
				Net Change: 1993-2011			
Counts	-1,682,006	-775,266	-85,374	-1,048,843	-2,338,625	278,564	-5,651,550
Percent Contribution	0.30	0.14	0.02	0.19	0.41	-0.05	1.00

Source: LBD-LFTTD-DCA-UBP as explained in text.
This table reports the employment counts pertaining to the "constant" manufacturing sample used in section 2.

Table 3: Summary Statistics: Firm Counts by Type: 1993-2011

year	Non U.S. Multinationals	U.S. Multinationals	Total
1993	302,669	2,539	305,208
2011	218,572	2,036	220,608
	Average Annual Percent Change		
1993-2011	-1.54	-1.10	-1.54
1993-2001	-1.17	-0.17	-1.16
2002-2011	-2.02	-2.06	-2.02

Source: LBD-LFTTD-DCA-UBP as explained in text.
This table reports the firm counts pertaining to the "constant" manufacturing sample used in section 2.

Table 4: Percentage of Firms Participating in Foreign Input Sourcing: 1993-2011

	Non U.S. Multinationals				U.S. Multinationals			
year	Arms Length Low Income	Related Party Low Income	Arms Length High Income	Related Party High Income	Arms Length Low Income	Related Party Low Income	Arms Length High Income	Related Party High Income
1993	1.88	0.39	5.71	1.30	44.35	24.62	72.63	48.17
2011	7.41	1.42	8.25	2.01	73.18	49.02	81.83	58.69
				Percent Change				
1993-2011	294	264	44.5	54.6	65	99	12.7	21.8

Source: LBD-LFTTD-DCA-UBP as explained in text.

This table reports the fraction of U.S. multinationals and non U.S. multinationals that sourced inputs from foreign countries. These are non-exclusive shares of the total number of firms in 3. Non U.S. multinationals includes foreign multinationals and other trading firms.

Table 5: Pooled Regression Results

	Establishment Level			
	Intensive		Extensive and Intensive	
	Unweighted	Employment Weighted	Unweighted	Employment Weighted
β	0.019***	0.007***	-0.03***	-0.03***
S.E.	(0.001)	(0.001)	(0.002)	(0.002)
Clusters	16,616	16,616	17,528	15,606

	Firm Level			
	Intensive		Extensive and Intensive	
	Unweighted	Employment Weighted	Unweighted	Employment Weighted
β	-0.01***	-0.02***	-0.03***	-0.03***
S.E.	(0.002)	(0.004)	(0.005)	(0.006)
Clusters	8,028	8,028	9,118	9,118

Source: LBD-LFTTD-DCA-UBP as explained in text.
This table reports the pooled regression results from 1 at the establishment and firm level.

Table 6: Average Establishment-Level Transition Probabilities: 1993-2011

t\t+1	Dom	Exp	U.S. Mult	For Mult	Exit
Dom	85%	5%	0%	0%	10%
Exp	13%	80%	1%	1%	5%
U.S. Mult	0%	2%	91%	1%	6%
For Mult	0%	2%	2%	90%	6%
Entry	84%	13%	1%	2%	

Source: LBD, DCA, and UBP
This table reports average probability of transition from state i in t to j in $t+1$ where $\{i, j \in D, X, MH, MF, Entry, Exit\}$. The average number of establishments corresponding to each type is in Table 1.

Table 7: Inter-Firm and Intra-Firm Sourcing

	Country Level		Industry & Country Level	
	RP Indicator	Log RP Imports	RP Indicator	Log RP Imports
Coef.	1.84***	0.39***	1.765***	0.49***
Std. Err.	(0.006)	(0.002)	(0.004)	(0.001)
Fixed Effects				
Firm X time	Yes	Yes	No	No
Country X Time	Yes	Yes	No	No
Industry X Time	No	No	Yes	Yes
Firm X Country X Time	No	No	Yes	Yes
R^2	0.51	0.61	0.52	0.64
Observations	1,776,800	380,400	5,012,000	1,033,000

Source: LFTTD
This table reports the results from equation 3. The dependent variable is the log of a firm's inter-firm imports from a particular country or industry within a country.

Table 8: Estimation Results: Bounding

Year	Upper Bound	Lower Bound
1997	0.61***	0
	(0.10)	
2002	0.86***	0
	(0.10)	
2007	0.79***	0
	(0.07)	

Source: LBD,LFTTD, CMF and ASM
This table reports bounds on $\frac{\sigma-1}{\theta}$ implied by the bounding procedure in section 3.4. The upper bound is the 80th percentile of all lower bounds calculated by applying the procedure to different sourcing strategies, as discussed in appendix C.

Table 9: Estimation Results: Semiparametric Regressions

Year	1993		1997		2002		2007	
$\frac{\sigma-1}{\theta}$	0.16***	0.11***	0.16***	0.12***	0.11***	0.08***	0.11***	0.06***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)
Higher order F.E.	YES	NO	YES	NO	YES	NO	YES	NO
Size percentiles	NO	YES	NO	YES	NO	YES	NO	YES
Instrumented	NO	NO	NO	NO	NO	NO	NO	NO
Observations	72,700	72,700	79,600	79,600	67,400	67,400	71,800	71,800
R-squared	0.96	0.96	0.95	0.95	0.96	0.97	0.97	0.97
$\frac{\sigma-1}{\theta}$			0.17**	0.23***	0.10	0.22***	0.19*	0.14***
			(0.068)	(0.011)	(2.718)	(0.010)	(0.095)	(0.009)
Higher order F.E			YES	NO	YES	NO	YES	NO
Size percentiles			NO	YES	NO	YES	NO	YES
Instrumented			YES	YES	YES	YES	YES	YES
Observations			76,000	76,000	64,000	64,000	67,400	67,400

Source: LBD,LFTTD, CMF and ASM

This table reports point estimates for $\frac{\sigma-1}{\theta}$ from the polynomial approximation and size bin approaches discussed in 3.4.We use fifty size bins for the approximation. The lower panel displays results where the cost shares are instrumented with lagged values.

Table 10: Cost Shares for Firms with $J = \{HO, HI, NO, NI, SO, SI\}$

year	χ_{HO}	χ_{HI}	χ_{NO}	χ_{NI}	χ_{SO}	χ_{SI}
1997	0.51	0.32	0.05	0.06	0.04	0.03
2002	0.48	0.34	0.05	0.07	0.05	0.04
2007	0.50	0.29	0.06	0.07	0.05	0.05

Source: LBD, LFTTD and CMF
This table reports the average cost shares from different sourcing locations/modes for firms that source from all possible locations and modes.

Table 11: Firm Sourcing Patterns

Year	{HO,HI}	{HO,HI, NO}	{HO,HI, NO,NI}	{HO,HI,NO, NI,SO}	{HO,HI,NO, NI,SO,SI}	Other
	Fraction of firms with sourcing strategy:					
1997	74.5%	9.9%	2.6%	2.6%	3.1%	7.4%
2002	66.8%	11%	2.9%	3.4%	4.6%	11.3%
2007	61.6%	9.6%	2.1%	3.5%	6.3%	16.9%
	Fraction of imports in sourcing strategy:					
1997	0%	0.6%	1.4%	4.2%	89.1%	4.7%
2002	0%	0.5%	1.6%	3.3%	91.0%	3.7%
2007	0%	0.2%	0.8%	3.4%	92.0%	3.5%

Source: LBD, LFTTD and CMF
This table reports the fraction of firms sourcing from five of the most prominent sourcing strategies, as well as the fraction of imports accounted for by firms in each of these sourcing strategies. "Other" includes sourcing strategies $J \in \{\{HO, HI, SO\}, \{HO, HI, SI\}, \{HO, HI, NI\}, \{HO, HI, NO, SO\}, \{HO, HI, NO, SI\}, \{HO, HI, NI, SI\}, \{HO, HI, NI, SO\}, \{HO, HI, SI, SO\}, \{HO, HI, NO, NI, SI\}, \{HO, HI, NO, SO, SI\}, \{HO, HI, NI, SO, SI\}\}$.

Table 12: Calibration Stage 1

Parameter	Value	Note
σ	2.3	Demand elasticity
θ	6	Frechet shape parameter
b_ϕ	1	Lower bound of the Pareto distribution
A_H	14.32	Skill-adjusted wages in Home, from the ILO
A_N	8.29	Average skill-adjusted wages in North from the ILO
A_S	1.02	Average skill-adjusted wages in South from the ILO
L_H	0.301	Skill-adjusted labor force in Home, from the ILO
L_N	0.822	Total skill-adjusted labor force in North from the ILO
L_S	2.35	Total skill-adjusted labor force in South from the ILO
τ_H	1	Domestic transport costs
τ_N	1.15	Transport costs from North
τ_S	1.15	Transport costs from South
ρ	1.5	Elasticity of substitution of tasks
$\mathbb{E}[s]$	1.0025	Expected value of demand shifter
κ_{HI}	1	Home within firm technology transfer parameter
κ_{HO}	1.1	Home outside supplier technology transfer parameter
κ_{NI}	0.98	North within firm technology transfer parameter
κ_{NO}	0.95	North outside supplier technology transfer parameter
κ_{SI}	0.97	South within firm technology transfer parameter
κ_{SO}	0.93	South outside supplier technology transfer parameter

This table summarizes the first stage of the baseline calibration.

Table 13: Quantitative Exercises: Model Fit

	1997		2007	
	Targets	Model	Targets	Model
NO imports/Manuf sector sales	0.020	0.011	0.024	0.017
NI imports/Manufacturing sector sales	0.035	0.035	0.051	0.051
SO imports/Manufacturing sector sales	0.017	0.019	0.044	0.045
SI imports/Manufacturing sector sales	0.015	0.017	0.030	0.033
Fraction of trade with $J = \{HO, HI, NO, NI\}$	0.015	0.024	0.008	0.016
Fraction of trade with $J = \{HO, HI, NO, NI, SO\}$	0.044	0.023	0.035	0.013
Fraction of trade with $J = \{HO, HI, NO, NI, SO, SI\}$	0.935	0.902	0.954	0.917
Fraction of firms with $J = \{HO, HI\}$	0.804	0.814	0.741	0.786
Fraction of firms with $J = \{HO, HI, NO\}$	0.107	0.088	0.116	0.108
Fraction of firms with $J = \{HO, HI, NO, NI\}$	0.028	0.008	0.025	0.006
Fraction of firms with $J = \{HO, HI, NO, NI, SO\}$	0.028	0.005	0.042	0.003
Fraction of firms with $J = \{HO, HI, NO, NI, SO, SI\}$	0.033	0.085	0.076	0.097
Mean χ_{HO} with $J = \{HO, HI, NO, NI, SO, SI\}$	0.514	0.514	0.500	0.500
Home multinational/total manufacturing employment	0.307	0.305	-	0.317
Manufacturing employment share	0.168	0.169	-	0.130

This table summarizes the fit of the model to calibration targets in 1997 and 2007.

Table 14: Quantitative Exercises: Local Effects

1 % change in:	manufacturing employment (in percent)	multinational employment (in percent)
T_{NO}	-0.05	-0.07
T_{NI}	-0.10	0.67
T_{SO}	-0.07	-0.18
T_{SI}	-0.04	-0.12
f_{NO}	-0.02	-0.18
f_{NI}	0.01	1.25
f_{SO}	-0.04	-0.12
f_{SI}	-0.11	-0.28

This table summarizes the responses of key variables to one percent increases in foreign technology parameters or one percent decreases in fixed costs of foreign sourcing.

Table 15: Quantitative Exercises: Parameter changes

Technology	T_{HI}	T_{HO}	T_{NI}	T_{NO}	T_{SI}	T_{SO}
Change	273 %	275 %	555%	495 %	712 %	923 %
Fixed Costs	{HO,HI}	{HO,HI, NO}	{HO,HI, NO,NI}	{HO,HI,NO, NI,SO}	{HO,HI,NO, NI,SO,SI}	
Change	112 %	184 %	193 %	222 %	216 %	

This table summarizes changes in the technology and fixed cost parameters between the baseline calibration (1997) and the final calibration (2007).

Table 16: Quantitative Exercises: Manufacturing Decline

	Data (1993- 2011)	Data (1997- 2007)	Model
Manufacturing Employment	-0.36 %	-0.25 %	-0.13 %
Multinational Employment	-0.44 %	-0.27 %	-0.28 %
Non-MN Employment Employment	-0.31 %	-0.24 %	- 0.07 %

This table summarizes the decline in aggregate manufacturing employment within the model. We show the declines in the data over two periods – the full sample and a shorter period between the census years 1997 and 2007, as some of our calibration targets are only available in census years and have been chosen to match data in 1997 and 2007.

Figure 1: Share of Trade and Firm Participation in Trade, by Type

Value of Trade by Firm Type

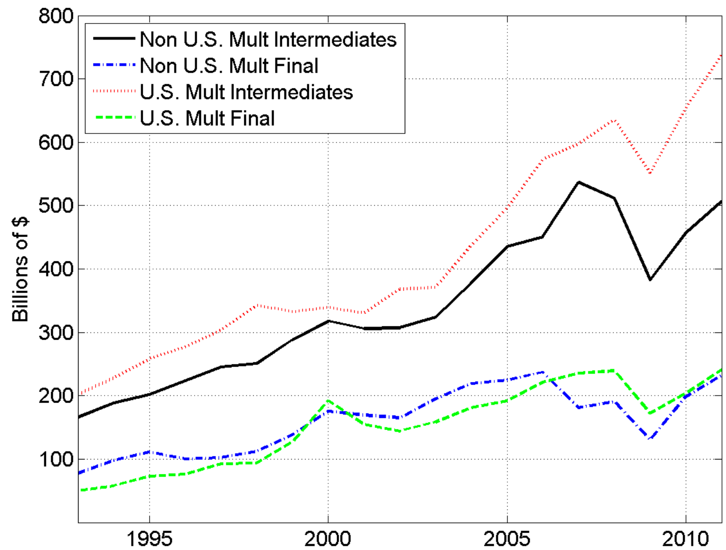

Source: LFTTD-DCA-UBP as explained in text.
These figures report the value of intermediate and final goods trade by firm type, as well as the share of intermediate inputs imported from low income countries by U.S. multinationals.

Figure 2: Employment Growth Differential of Multinational Transitions

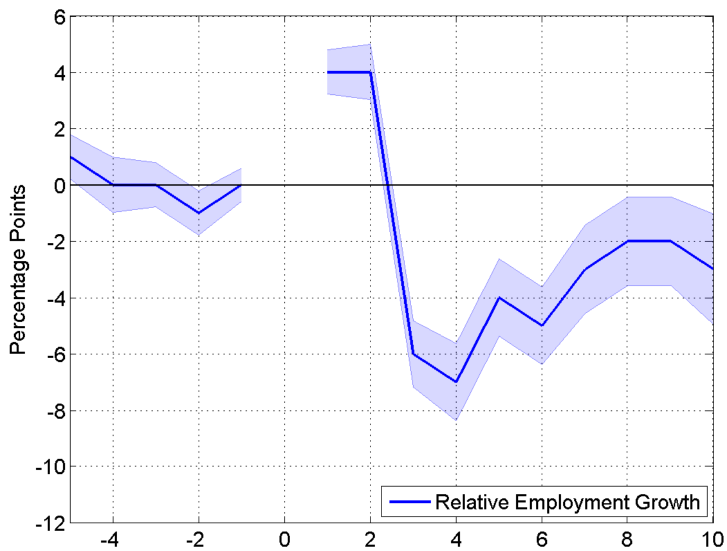

Source: LFTTD-DCA-UBP as explained in text.
This figure plots the pre and post annual deviations in the employment growth rate of establishments that transition into part of a multinational firm in year ($t = 0$), relative to a control group based on interacted effects of firm age, establishment size, and industry (in year $t = -1$). The control group consists of establishments that are not part of a multinational firm in year $t = 0$. See equation 2. The shaded area corresponds to a 95 percent confidence interval.

Figure 3: Importing Differentials of Multinational Transitions

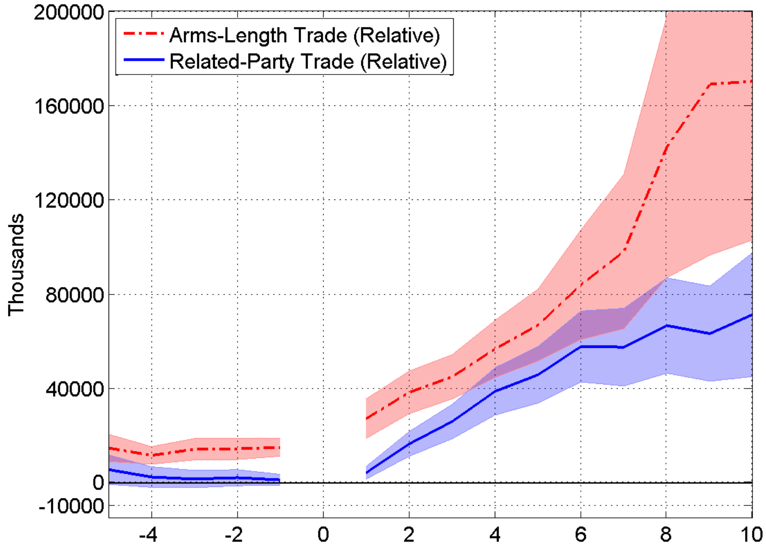

Source: LFTTD-DCA-UBP as explained in text.
This figure reports the related-party and arms-length intermediate input imports of the parent firm of an establishment that transitions into part of a multinational firm in year ($t = 0$), relative to a control group based on interacted effects of firm age, establishment size, and industry (in year $t = -1$). See equation 2, modified to reflect firm-level imports as dependent variables. The shaded area corresponds to a 95 percent confidence interval.

Figure 4: Employment Growth Differential of Importer Transitions

Source: LFTTD-DCA-UBP as explained in text.
This figure plots the pre and post annual deviations in the employment growth rate of establishments that begin importing from abroad year ($t = 0$), relative to a control group based on interacted effects of firm age, establishment size, and industry (in year $t = -1$). The control group consists of establishments that are not part of a multinational firm in year $t = 0$, nor have recorded positive imports in the period ($t - 3, t = 0$). See equation 2. The shaded area corresponds to a 95 percent confidence interval.

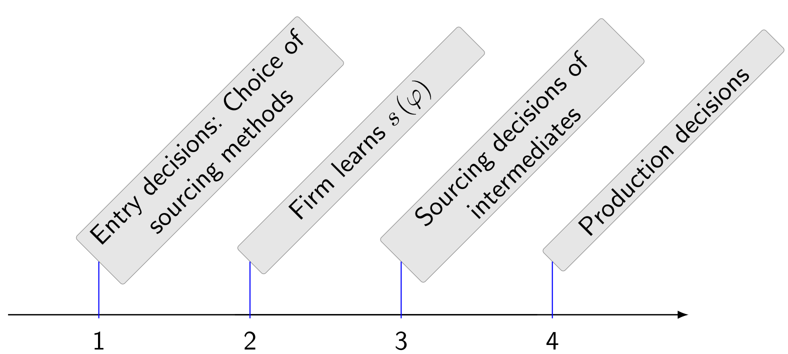

Figure 5: Stages of the Firm's Problem

Figure 6: Estimation results by industry

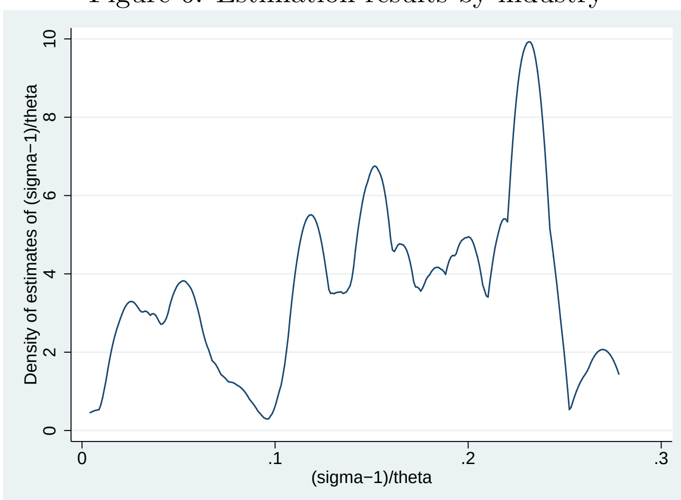

Source: CMF-LFTTD as explained in text.
This figure plots the kernel density of the results of the estimation of $\frac{\sigma-1}{\theta}$ using equation 21 by industry in 1997. The results are similar for other estimation years 2002 and 2007.

A Data Appendix

A.1 Creating Constant Manufacturing Sample

An important challenge for our analysis of U.S. manufacturing employment over such an extended period of time is defining exactly what plant-level operations constitute manufacturing. This task is complicated by the fact that our sample coincides with two distinct industry classification systems (SIC and NAICS) as well as periodic revisions to these systems.

To construct a constant manufacturing sample, we begin with the Longitudinal Business Database (LBD), an assembly of the Standard Statistical Establishment List (SSEL) that has been augmented with longitudinal identifiers and standardized across years. We drop establishments listed as government, and establishments listed as "dead". Next, we utilize a new concordance of manufacturing classification systems outlined in Fort and Klimek (2015) for smoothing out discrepancies between industries defined as manufacturing between SIC and NAICS. There remain several acknowledged data issues of the Fort and Klimek (2015) manufacturing definition, principally related to manufacturing establishments that are re-coded into NAICS 55 - "Management of Companies and Enterprises" in 2002. We set up the following two rules to broadly account for establishments that transition into and out of a FK-manufacturing industry during our sample. First, we drop establishments (in all years) that are re-classified out of manufacturing during our sample; and second, we retain establishments (in all years) that are ever reclassified into manufacturing during our sample. This system prevents the possibility of spurious establishment "births" or "deaths" being recorded as a consequence of a classification change.

Figure A1 illustrates how our constant manufacturing sample compares to manufacturing employment from two other sources: published totals from the Current Employment Survey and Pierce and Schott (2013).

A.2 Identifying Plants Owned by Multinationals

The discussion that follows is an abbreviated form of the full technical note (see Flaaen (2013a)) documenting the bridge between the DCA and the Business Register.

A.2.1 External Sources of Information

Identification of foreign ownership and affiliate information comes from two external sources, the LexisNexis Directory of Corporate Affiliations (DCA) and Uniworld Business Publications.

The LexisNexis DCA is the primary source of information on the ownership and locations of U.S. and foreign affiliates. This directory describes the organization and hierarchy of public and private firms, and consists of three separate databases: U.S. Public Companies, U.S. Private Companies, and International – those parent companies with headquarters located outside the United States. The U.S. Public database contains all firms traded on the major U.S. exchanges, as well as major firms traded on smaller U.S. exchanges. To be included in the U.S. Private database, a firm must demonstrate revenues in excess of $1 million, 300 or more employees, or substantial assets. Those firms included in the International database, which include both public and private companies, generally have revenues greater than $10 million. Each database contains information on all parent company subsidiaries, regardless of the location of the subsidiary in relation to the parent.

Uniworld Business Publications (UBP) provides a secondary source used to identify multinational structure, and serves to increase the coverage and reliability of these measures. UBP has produced periodic volumes documenting the locations and international scope of i) American firms operating in foreign countries; and ii) foreign firms with operations in the United States. Although only published biennially, these directories benefit from a focus on multinational firms, and from no sales threshold for inclusion.

Because there exist no common identifiers between these directories and Census Bureau data infrastructure, we rely on probabilistic name and address matching — so-called "fuzzy merging" — to link the directories to the Census data infrastructure.

A.2.2 The Matching Procedure: An Overview

The matching procedure uses a set of record linking utilities described in Wasi and Flaaen (2014). This program uses a bigram string comparator algorithm on multiple variables with differing user-specified weights.[39] The primary variables for matching include the establishment name along with geographic indicators of street, city ,zip code, and state.

Recognizing the potential for false-positive matches, w use a relatively conservative criteria for identifying matches between the directories and the Census Bureau data. In practice, the procedure generally requires a match score exceeding 95 percent, except in those cases where ancillary evidence provides increased confidence in the match.[40] This matching proceeds in an iterative fashion, in which a series of matching procedures are applied with decreasingly restrictive sets of matching requirements. In other words, the initial matching attempt uses the most stringent standards possible, after which the non-matching records proceed to a further matching iteration, often with less stringent standards. In each iteration, the matching records are assigned a flag that indicates the standard associated with the match.

See Table A1 for a summary of the establishment-level match rate statistics by year and type of firm. Table A2 lists the corresponding information for the Uniworld data.

A.3 Creating Panel of Multinational Plants

The external directories allow for relatively easy categorization of the multinational status of U.S. plants. If the parent firm contains addresses outside of the United States, but is headquartered within the U.S., we designate this establishment as part of a U.S. multinational firm. If the parent firm is headquartered outside of the United States, we designate

[39]The term bigram refers to two consecutive characters within a string (the word *bigram* contains 5 possible bigrams: "bi", "ig", "gr", "ra", and "am"). The program is a modified version of Blasnik (2010), and assigns a score for each variable between the two datasets based on the percentage of matching bigrams. See Flaaen (2013a) or Wasi and Flaaen (2014) for more information.

[40]The primary sources of such ancillary evidence are clerical review of the matches, and additional parent identifier matching evidence.

this establishment as part of a Foreign multinational firm.

This paper seeks to understand how changes in multinational status affect labor market outcomes in the United States. To achieve this end, we must take the yearly multinational identifiers and construct a panel across many years. The challenge with this exercise comes from the fact that the directories are matched year-by-year, utilizing little longitudinal information.[41] This implies the possibility that a multinational plant may not be successfully matched every year, and our data could have spurious entries and exits from multinational status throughout the panel.

To mitigate this concern, we develop a series of checks and rule-based procedures to correct and smooth out any unlikely firm switching. These steps can be classified as those accounting for changes within a year across plants of a given firm, and those correcting for multinational status across years for a particular plant.

A.3.1 Within-Year Rules

First, we apply our multinational indicators to all establishments within a firm provided there are no disagreements in the DCA/UBP information among the establishments. This is an attractive feature of our methodology as the researcher must only successfully match one plant of a given firm to apply that information throughout the firm. To resolve any conflicting information within a year, we first attempt to use corroborating evidence from the secondary source (typically Uniworld), and then turn to the maximum employment share of a particular type of match. Finally, we conduct manual checks on the data, particularly on those firms that demonstrate very large amounts of related-party trade but have not been captured by our matching procedure.

A.3.2 Checks and Rules for Across Years

Another important step in creating a panel of establishment information on the scope of international operations is to check and correct for any potentially spurious transitions of establishment type over time. First, if there is only one missing year of a multinational indicator in the establishment's history, we fill it in manually. Second, if there is a gap of two years in this indicator that corresponds to gap years in the Uniworld coverage, we also fill it manually. Similarly, if an establishment is identified as a multinational in only one year in it's history, we remove the flag. Finally, we fill in 2 year gaps provided that in the intervening period the share of related party trade remains high.

A.4 Classification of Intermediate/Final Goods Trade

Firm-level data on imports available in the LFTTD do not contain information on the intended use of the goods.[42] Disentangling whether an imported product is used as an intermediate input for further processing — rather than for final sale in the U.S. — has

[41] The only longitudinal information used is by applying prior clerical edits forward in time for a particular establishment, provided that the name and address information remains unchanged.

[42] This is one advantage of the survey data on multinational firms available from the Bureau of Economic Analysis. There are, however, a number of critical disadvantages of this data source, as outlined in Flaaen (2013b).

important implications for the effect of offshoring on U.S. employment. Fortunately, the Census Bureau data contains other information that can be used to distinguish intermediate input imports from final goods imports. In brief, identifying the principal products produced by U.S. establishments in a given detailed industry should indicate the types of products that, when imported, should be classified as a "final" good – that is, intended for final sale without further processing. The products imported outside of this set, then, would be classified as intermediate goods.[43] Such product-level production data exists as part of the "Products" trailer file of the Census of Manufacturers. As detailed in Pierce and Schott (2012) (see page 11), combining import, export, and production information at a product-level is useful for just such a purpose.

It is important to acknowledge that the Census data on trade exists at the firm level, while the other information used in this paper is , principally, at the plant level. Utilizing the establishment industry information, however, will allow us to parse a firm's trade based on the intermeidate/final distinction for a given establishment, thereby generating some heterogeneity in firm trade across establishments.[44]

A.4.1 Creating a NAICS-Based set of Final/Intermediate Products

As part of the quinquennial Census of Manufacturers (CM), the Census Bureau surveys establishments on their total shipments broken down into a set of NAICS-based (6 digit) product categories. Each establishment is given a form particular to its industry with a list of pre-specified products, with additional space to record other product shipments not included in the form. The resulting product trailer file to the CM allows the researcher to understand the principal products produced at each manufacturing establishment during a census year.

There are several data issues that must be addressed before using the CM-Products file to infer information about the relative value of product-level shipments by a particular firm. First, the trailer file contains product-codes that are used to "balance" the aggregated product-level value of shipments with the total value of shipments reported on the base CM survey form. We drop these product codes from the dataset. Second, there are often codes that do not correspond to any official 7-digit product code identified by Census. (These are typically products that are self-identified by the firm but do not match any of the pre-specified products identified for that industry by Census.) Rather than ignoring the value of shipments corresponding to these codes, we attempt to match at a more aggregated level. Specifically, we iteratively try to find a product code match at the 6, 5, and 4 digit product code level, and use the existing set of 7-digit matches as weights to allocate the product value among the 7-digit product codes encompassing the more aggregated level.

We now discuss how this file can be used to assemble a set of NAICS product codes that are the predominant output (final goods) for a given NAICS industry. Let x_{pij} denote the shipments of product p by establishment i in industry j during a census year. Then the total

[43]To be more precise, this set will include a combination of intermediate and capital goods.

[44]To be more precise, the total trade at each establishment of a firm must be identical. The shares of intermediate/final goods will vary.

output of product p in industry j can be written as:

$$X_{pj} = \sum_{i=1}^{I_j} x_{pij},$$

where I_j is the number of firms in industry j. Total output of industry j is then:

$$X_j = \sum_{p=1}^{P_j} X_{pj}.$$

The share of industry output accounted for by a given product p is therefore:

$$S_{pj} = \frac{X_{pj}}{X_j}.$$

One might argue that the set of final goods products for a given industry should be defined as the set of products where $S_{pj} > 0$. That is, a product is designated as a "final good" for that industry if *any establishment* recorded positive shipments of the product. The obvious disadvantage of employing such a zero threshold is that small degrees of within-industry heterogeneity will have oversized effects on the classification.

Acknowledging this concern, we set an exogenous threshold level W such that any p in a given j with $S_{pj} > W$ is classified as a final good product for that industry. The upper portion of Table A3 documents the number of final goods products and the share of intermediate input imports based on several candidate threshold levels. The issues of a zero threshold are quite clear in the table; a small but positive threshold value (0.1) will have a large effect on the number of products designated as final goods. This shows indirectly that there are a large number of products produced by establishments in a given industry, but a much smaller number that comprise the bulk of total value.

There are several advantages to using the CM-Products file rather than using an input-output table.[45] First, within a given CM year, the classification can be done at the firm or establishment level rather than aggregating to a particular industry. This reflects the fact that the same imported product may be used as an input by one firm and sold to consumers as a final product by another. Second, the CM-Products file is one of the principal data inputs into making the input-output tables, and thus represents more finely detailed information. Related to this point, the input-output tables are produced with a significant delay – the most recent available for the U.S. is for year 2002. Third, the input-output tables for the U.S. are based on BEA industry classifications, which imply an additional concordance (see below) to map into the NAICS-based industries present in the Census data.

We now turn to the procedure to map firm-level trade into intermediate and final goods using the industry-level product classifications calculated above.

[45] Another option is to use the CM-Materials file, the flip side of the CM-Products file. Unfortunately, the CM-Materials file contains significantly more problematic product codes than the Products file, and so concording to the trade data is considerably more difficult.

A.4.2 Mapping HS Trade Transactions to the Product Classification

The LFTTD classifies products according to the U.S. Harmonized Codes (HS), which must be concorded to the NAICS-based product system in order to utilize the classification scheme from the CM-Products file. Thankfully, a recent concordance created by Pierce and Schott (2012) can be used to map the firm-HS codes present in the LFTTD data with the firm-NAICS product codes present in the CM-Products data.

A challenge of this strategy is that the LFTTD exists at a firm-level, while the most natural construction of the industry-level classification scheme is by establishment. More concretely, for multi-unit, multi-industry firms, the LFTTD is unable to decompose an import shipment into the precise establishment-industry of its U.S. destination. By using the industry of each establishment to classify the firm's imports, we generate heterogeneity in the intermediate/final goods trade across the establishments of the firm.

Once the firm-level trade data is in the same product classification as the industry-level filter created from the CM-Products file, all that is left is to match the trade data with the filter by NAICS industry. Thus, letting M_{ij} denote total imports from a firm i (firm i is classified as being in industry j), we can then categorize the firm's trade according to:

$$\left.\begin{aligned} M_{ij}^{\text{int}} &= \sum_{p \notin P_j} M_{ipj} \\ M_{ij}^{\text{fin}} &= \sum_{p \in P_j} M_{ipj} \end{aligned}\right\} \qquad \text{where} \qquad P_j = \{p \mid S_{pj} \geq W\}\,. \tag{A1}$$

The bottom section of Table A3 shows some summary statistics of the intermediate share of trade according to this classification system, by several values of the product-threshold W. There are at least two important takeaways from these numbers. First, the share of intermediates in total imports is roughly what is reported in the literature using IO Tables. Second, the share of total trade occupied by intermediate products is not particularly sensitive to the exogenous threshold level. While there is a small increase in the share when raising the threshold from 0 to 0.1 (about 3 percentage points), the number is essentially unchanged when raising it further to 0.2.

A.5 Creating the Firm-Level Sample

Much of our analysis is at the firm level, so we build a sample of U.S. multinational firms from the panel of multinational plants (constructed as detailed in Section A.3). As the Corporate Directories are matched at the establishment level, when aggregating up to the firm, there are occasional conflicts in the definition of a firm between the Census and the Directories. We rely on the Census definition of a firm. Conflicts are resolved as follows:

- We define a firm in the panel as a U.S. multinational in a particular year if our matches are completely consistent in that year, and there are no conflicts.

- In the special case of a conflict where the Census classifies a firm as a set of establishments, but our matches to the Directories indicate a subset of those establishments belongs to a foreign multinational and a subset to a U.S. multinational, we classify the firm as a U.S. multinational if the employment share of the firm in the matched U.S. multinational sample is larger than that matched as a Foreign multinational.

Note, firm identifiers in the Census are sometimes problematic longitudinally. An example is that the firm identifier changes when the firm goes from being a single unit to a multi-unit establishment. Further, mergers and acquisitions can lead in some cases to the birth of a new firm identifier, and in others to the continuation of one of the merged identifiers. As such, results pertaining to the extensive margin that use the firm identifier as the basis of analysis will be overstated. This is a problem faced by all longitudinal firm-level analysis using Census Bureau data. We do not use longitudinal information in classifying U.S. and foreign multinationals, or non multinational firms. However, some of our analysis in 2 uses the growth rates of employment in the firm. In these cases, we use establishment level outcomes as the baseline (as these identifiers are longitudinally consistent), and present the firm-level results for robustness. The structural estimation relies on repeated cross-sections of the firm-level data and does not suffer from this issue.

Table A1: DCA Establishments and Match Rates, by Firm Type

	Panel A: Total DCA			Panel B: U.S. Multinationals			Panel C: Foreign Multinationals		
	DCA (Total)	Matched to BR	Match Rate	DCA (Total)	Matched to BR	Match Rate	DCA (Total)	Matched to BR	Match Rate
1993	61,646	43,190	0.70	21,482	14,387	0.67	8,270	5,810	0.70
1994	64,090	44,904	0.70	22,396	15,110	0.67	9,326	6,437	0.69
1995	65,223	45,743	0.70	22,952	15,448	0.67	9,365	6,414	0.68
1996	64,152	41,713	0.65	22,353	13,806	0.62	10,057	6,331	0.63
1997	60,884	41,290	0.68	20,962	13,583	0.65	9,556	6,328	0.66
1998	59,043	40,854	0.69	20,012	13,218	0.66	9,416	6,282	0.67
1999	58,509	40,697	0.70	20,157	13,408	0.67	9,218	6,054	0.66
2000	68,672	48,875	0.71	18,728	12,631	0.67	9,900	6,755	0.68
2001	70,522	50,105	0.71	18,516	12,477	0.67	10,089	6,864	0.68
2002	97,551	66,665	0.68	31,260	21,004	0.67	13,168	8,483	0.64
2003	123,553	86,838	0.70	25,905	17,465	0.67	11,101	7,398	0.67
2004	117,639	84,450	0.72	24,028	16,923	0.70	10,152	7,156	0.70
2005	110,106	80,245	0.73	20,870	15,191	0.73	9,409	6,865	0.73
2006	110,826	79,275	0.72	21,335	15,539	0.73	9,981	7,243	0.73
2007	112,346	81,656	0.73	22,500	16,396	0.73	10,331	7,555	0.73
2008	111,935	81,535	0.73	23,090	16,910	0.73	9,351	6,880	0.74
2009	111,953	81,112	0.72	22,076	16,085	0.73	11,142	8,193	0.74
2010	111,998	79,661	0.71	21,667	15,785	0.73	11,308	8,181	0.72
2011	113,334	79,516	0.70	21,721	15,557	0.72	11,619	8,357	0.72

[1]Notes: U.S. multinationals are defined as establishments whose parents are U.S. firms that have a foreign affiliate in the DCA. Foreign multinationals are defined as establishments owned by firms whose headquarters are in a foreign location.

Figure A1: Comparison of Constant Manufacturing Employment Samples: 1993-2011

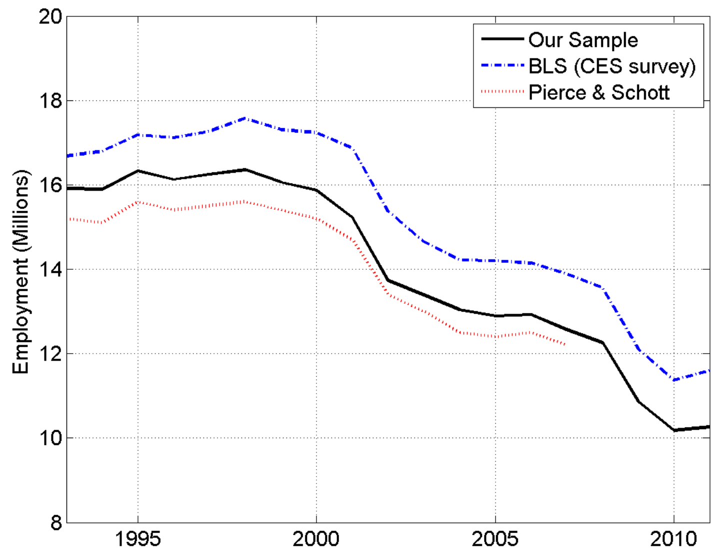

Source: BLS, Pierce and Schott (2013) and the LBD.

Table A2: Uniworld Match Statistics: 2006-2011

	# of Uniworld Establishments	Matched to B.R.	Percent Matched
Foreign Multinationals			
1992	1,597	1,223	0.77
1995	1,625	1,213	0.75
1998	2,020	1,555	0.77
2000	2,371	1,862	0.79
2002	2,780	2,154	0.77
2004	3,220	2,347	0.73
2006	3,495	2,590	0.74
2008	3,683	2,818	0.76
2011	6,188	4,017	0.65
U.S. Multinationals[1]			
1993	2,553	1,746	0.68
1996	2,502	1,819	0.73
1999	2,438	1,942	0.80
2001	2,586	2,046	0.79
2004	3,001	2,403	0.80
2005	2,951	2,489	0.84
2007	4,043	3,236	0.80
2009	4,293	3,422	0.80

[1]U.S. multinationals include only the establishments identified as the U.S. headquarters.

Table A3: Appendix Table Comparing the Results from Threshold Values W

	Threshold Values		
	$W = 0$	$W = 0.1$	$W = 0.2$
Number of Final Good Products per Industry			
Median	19	1	1
Mean	25	1.52	1.14
Min	1	1	0
Max	154	6	3
Implied Share of Intermediate Inputs			
Imports	60.9	63.90	63.97
Exports	52.0	54.96	55.04

This table is applicable to the year 2007.

B Additional Results

B.1 Within-Group Decompositions

In a first level of disaggregation, we show that job creation and destruction rates vary substantially by establishment type: U.S. multinational, exporter or purely domestic. U.S. multinationals have had persistently high job destruction rates and low job creation rates. In contrast, exporting and domestic establishments have higher job creation rates than destruction rates during business cycle expansions.

Employment growth is affected jointly by the rate of job creation and that of job destruction, and further by the extent to which this pertains to establishment births and deaths rather than employment changes at continuing establishments. Following the common practice exemplified by Davis and Haltiwanger (2001) we decompose the changes in within-group employment into job creation/destruction rates, separated by intensive and extension margins. Formally, let employment at establishments in group $S \in \{D, X, MH, MF\}$ in time t be denoted as $E_{S,t}$. Defining S^+_{t-1} and S^-_{t-1} as the set of establishments in S that increase (decrease) employment between t-1 and t, we can then define the job creation ($JC_{S,t}$) and destruction ($JD_{S,t}$) rates as:

$$\text{Job Creation Rate:} \qquad JC_{S,t} = \frac{\sum_{i \in S^+_{t-1}} \Delta e_{i,S,t}}{\left(E_{S,t} + E_{S,t-1}\right)/2} \tag{A2}$$

$$\text{Job Destruction Rate:} \qquad JD_{S,t} = \frac{\sum_{i \in S^-_{t-1}} |\Delta e_{i,S,t}|}{\left(E_{S,t} + E_{S,t-1}\right)/2} \tag{A3}$$

Separating these groups further into those surviving establishments (existing in both $t-1$ and t) will yield intensive margin growth rates, while focusing on establishment births/deaths in a given year will yield rates corresponding to the extensive margin.

Figures A2 report the intensive job creation/destruction rates of the three relevant groups we study. In Panel A, the job creation rates show both cyclicality and a secular decline for both domestic and exporting establishments.[46] The job creation rate for multinational firms is lower and slightly less cyclical than the other groups. The high cyclicality of job destruction rates is very much evident in Panel B of Figure A2. Taking into account that both JC and JD rates are known to decrease with both firm size and firm age (see Davis and Haltiwanger (1992) and Haltiwanger, Jarmin, and Miranda (2013)), and that multinationals are 3 times (20 times) larger than exporting (domestic) establishments, it is striking how similar the job destruction rates for multinationals are to the other two groups. With this in mind, it appears that job destruction plays a more important role for multinationals relative to non-multinational establishments, and has been an important driver of the observed aggregate decline in employment in this group.

Figure A3 translates the job creation and destruction rates into a net measure of employment gains by type of establishment. Panel A shows that multinational establishments

[46] This decline in job creation rates is consistent with other evidence on the decline in the overall dynamism of U.S. businesses, as documented in Decker et al. (2014).

have had lower net growth rates than the domestic/exporting groups in nearly every year of our sample. While domestic/exporting firms were on net adding jobs following the 2001 recession in the U.S., the multinational establishments continued to shed jobs through the 2008/2009 financial crisis. In this way multinationals are shown to be a contributor to the "jobless recovery" of the 2003-2007 expansion.

B.2 Other Results on Transitions

B.2.1 Assumptions of Firm-Level trade Following an Establishment Death

There are at least two distinct approaches to account for the role of establishment death on the import activity at the *firm*-level. The estimates in Figure 3 fill in the post-death values for a given establishment with the actual imports of the firm associated with that establishment. [47] This approach better captures the import substitution that may occur if a plant is closed in response to offshore activities. If this was the case, we would see a larger import differential relative to the benchmark calculation. On the other hand, if establishment deaths are associated with broad firm decline, then this differential import measure would be smaller relative to the benchmark.

An alternative approach would be to fill in a value of zero trade for all years following an establishment death. If transitioning establishments are dying at a higher rate than non-transitioning establishments, this would reduce the differential importing patterns following the transition. A final approach would be to ignore the extensive-margin effects and simply allow the observations to be dropped upon an establishment death.

Below we demonstrate the effects of these assumptions on our estimates of import behavior surrounding the event study . In our baseline sample underlying Figure 3, we create a balanced panel and fill the pre-birth or post-death observations with the value at the firm immediately following preceding its birth/death. To assess the alternative approach we fill the pre-birth and post-death trade values with zero (which we call the "zeros-fill" results). Finally, the "no-ext margin" results demonstrate our estimates when completely ignoring these extensive margin effects.

Figure A4 reports the coefficient estimates from the baseline, zero-fill, and no-ext margin samples corresponding to related-party imports before and after the transition to multinational status. The evidence points to transitioning plants with a higher death rate than the control group, an effect which pulls the differential import behavior down relative to the baseline. On the other hand, filling in the firm imports after death actually increases the importing differential. This evidence further supports the hypothesis of employment substitution of these firms.

B.2.2 Other Trade Effects Following Multinational Transitions

We estimate equation (2) using various types of firm-level trade corresponding to establishments that transition into part of a multinational firm. The results pertaining to related-party and arms-length intermediate imports are shown in Figure 3. New U.S. multinationals may also begin importing final goods from an arms-length or intra-firm supplier abroad. The

[47] If the entire firm disappears, we then record zeros in that period and all future periods.

results that show the differential imports of final goods of new multinationals are shown in Figure A5. Perhaps more surprisingly, we also find strong growth in export volumes in the years following a multinational transition. The increase in exports (shown in Figure A6), together with the broad increase in importing activity, demonstrates the overall modifications of the production structure of these firms that accompany expansions abroad.

What do our results imply in the aggregate? To convert the estimates from Figure 2 into a measure of total job gain/loss from new multinational activities.

Further details are available in Appendix B.3.1.

B.3 Quantifying Job Loss: Back-of-the-Envelope Calculations

B.3.1 Job Loss from Multinational Transitions

This section describes how we convert the estimates on relative employment growth rates of new multinational plants into a measure of the aggregate net gains of employment. The coefficients from Figure 2 represent relative employment effects, expressed in percentage points, of a transitioning plant. These effects represent averages that span the entire period (1993-2011) for which plants may be transitioning into a multinational firm. To translate these percentage points into jobs, one challenge is to identify the appropriate base on which to apply the relative percentage differentials. Unfortunately, the average size of transitioning plants is not currently available. However, using the productivity/size ordering of firms implied by models such as Helpman, Melitz, and Yeaple (2004), and confirmed using similar data sources in Flaaen (2013b), we assign these transitioning plants an average size that is between that of exporters and multinational plants.

Another challenge comes from what to assume when the time-path of a given transitioning plant extends beyond our estimates (which currently end at $t = 10$ years post transition). While we could extrapolate our estimates in the later years of the estimation in , we instead follow the more conservative assumption and terminate the counterfactual time path once the estimates from equation (2) run out. (Essentially, we assume that the growth rate differentials in all years $t > 10$ are zero.) Of course, extrapolating the estimates beyond year 10 would magnify the job losses – adding an additional percentage point or two in accounting for the total job loss – resulting from multinational transitions.

Formally, we compute the job loss as

$$\sum_{t=1994}^{2010} T_t E_t \sum_{i=1}^{\min\{10,2010-t\}} \delta_i \prod_{j=1}^{i-1} (1+\delta_j) \tag{A4}$$

where T_t is the number of transitioning plants in event year t, E_t is the average size of transitioning plants in event year t, and δ are the coefficient estimates from equation (2). Table A4 provides further details. The result is an estimate of approximately 400,000 jobs lost due to these transitioning plants, roughly 7 percent of the total 5.65 million decline in manufacturing employment in our sample.

B.3.2 Job Loss from all Multinational Activity: Total

A similar exercise can be done using the coefficient estimates from Table 5. This calculation is somewhat easier in that we simply apply the employment growth rate differential to the average establishment size of multinationals, and then multiply by the total number of multinational establishments in each year. Table A5 shows the results. The first set of calculations uses the weighted regression coefficient pertaining to the intensive/extensive establishment growth rate, whereas the second set of calculations uses the unweighted regression coefficient. The numbers are large: between 2.02 and 2.45 million manufacturing jobs over our full sample.

B.4 Regression Evidence: Robustness

Table A6 presents results from running the specification in equation 1 for various subsamples of our data. The results are also robust to including lagged establishment or firm employment growth rates as controls (available upon request).

Figure A2: Job Creation and Destruction Rates by Group: Intensive Margin

A. Job Creation Rates

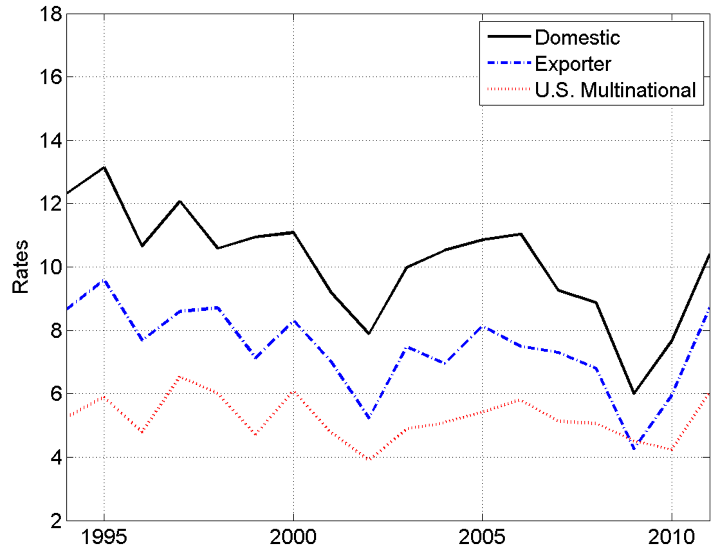

B. Job Destruction Rates

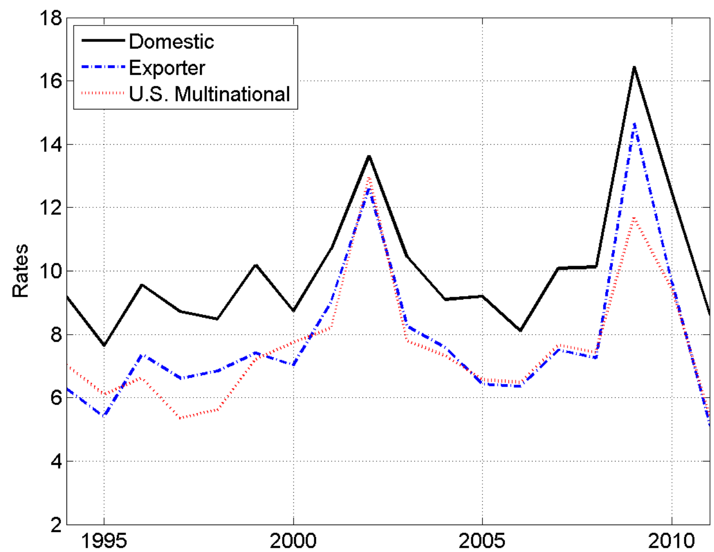

Source: LFTTD-DCA-UBP as explained in text.
These figures report the decomposition of within-group growth rates of employment at the intensive margin. See equation A2 in the text.

Figure A3: Net Growth Rates by Group:

A. Intensive Margin

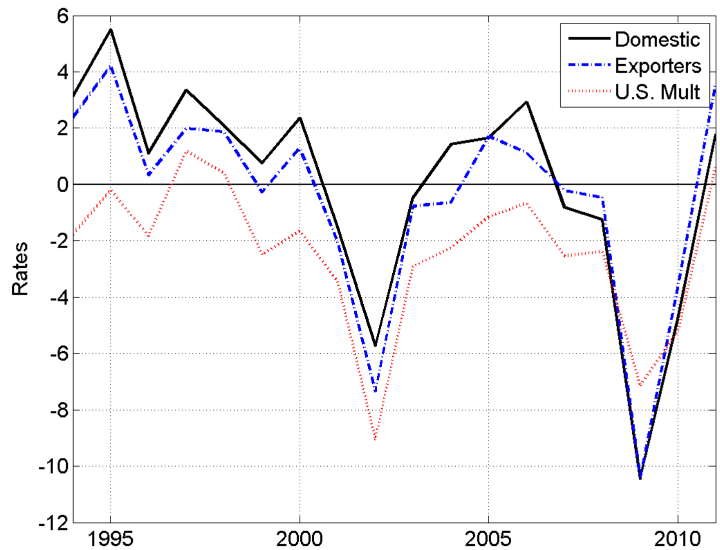

B. Extensive Margin

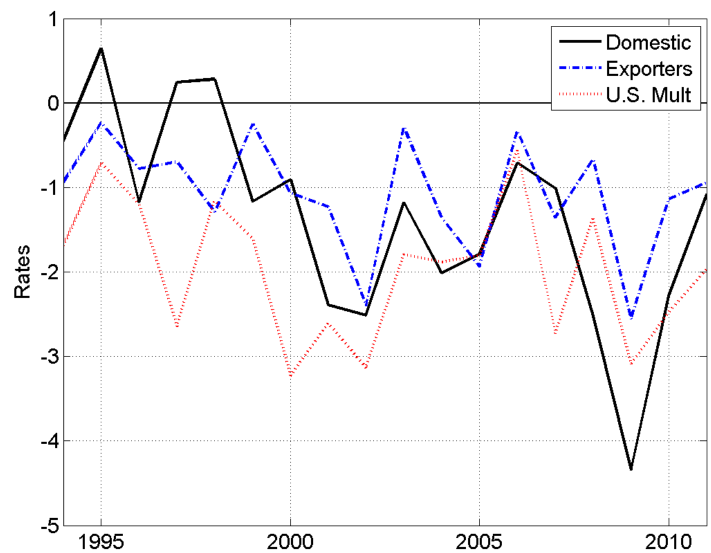

Source: LFTTD-DCA-UBP as explained in text.
These figures report the decomposition of within-group growth rates of employment at the intensive margin. See equation A2 in the text.

Figure A4: Importing Differentials of Multinational Transitions, Balanced Panel

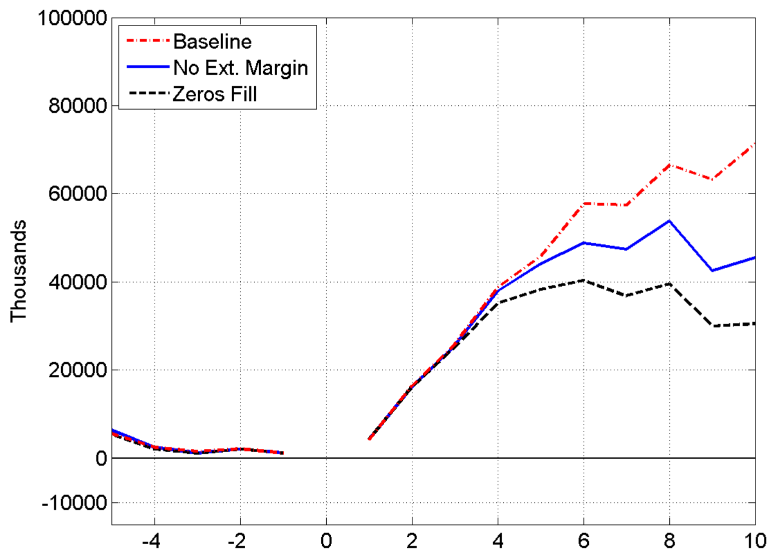

Source: LFTTD-DCA-UBP as explained in text.
This figure reports the related-party intermediate input imports of the parent firm of the transitioning establishment relative to a control group, as outlined in equation 2. *Zero Fill* refers to a balanced panel with zeros for trade after an establishment death. *No Ext. Margin* refers to the sample with no extensive margin effects following the establishment death.

Figure A5: Final Goods Importing Differentials of Multinational Transitions

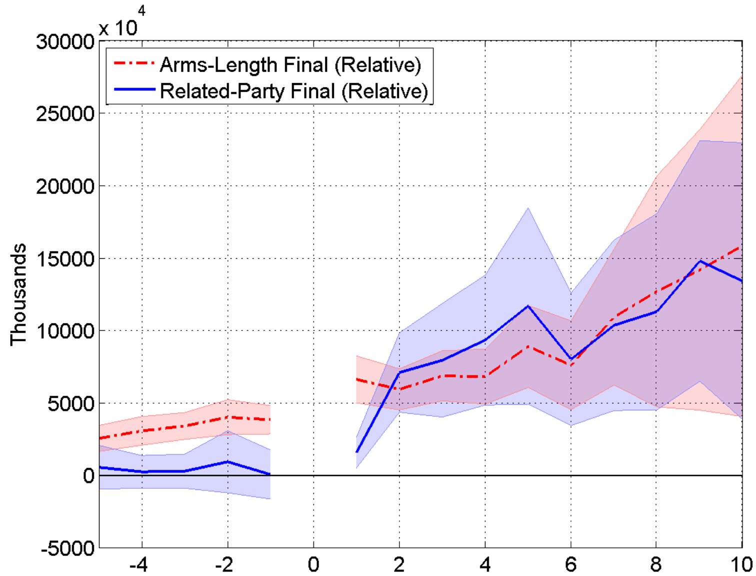

Source: LFTTD-DCA-UBP as explained in text.
This figure reports the related-party and arms-length final goods imports of the parent firm of an establishment that transitions into part of a multinational firm in year ($t = 0$), relative to a control group based on interacted effects of firm age, establishment size, and industry (in year $t = -1$). See equation 2, modified to reflect firm-level imports as dependent variables. The shaded area corresponds to a 95 percent confidence interval.

Figure A6: Exporting Differentials of Multinational Transitions

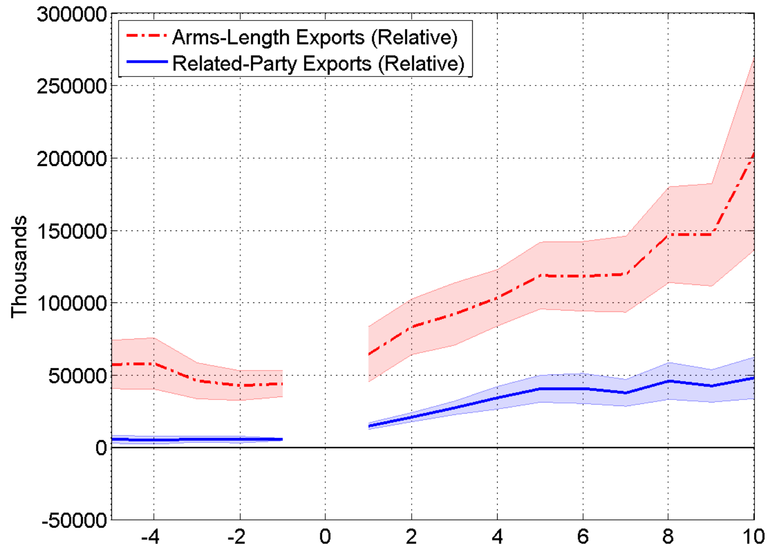

Source: LFTTD-DCA-UBP as explained in text.
This figure reports the related-party and arms-length exports of the parent firm of an establishment that transitions into part of a multinational firm in year ($t = 0$), relative to a control group based on interacted effects of firm age, establishment size, and industry (in year $t = -1$). See equation 2, modified to reflect firm-level imports as dependent variables. The shaded area corresponds to a 95 percent confidence interval.

Table A4: Appendix Table Detailing Aggregate Job Loss from New Multinational Plants

Year	Average Size	# of Transitions	Cumul. Jobs per Estab.	Total Job Gains
1994	203	344	-45	-15,424
1995	204	498	-45	-22,436
1996	205	915	-45	-41,344
1997	202	762	-45	-33,977
1998	205	851	-45	-38,590
1999	208	994	-46	-45,593
2000	197	962	-43	-41,774
2001	195	699	-43	-30,048
2002	193	1,060	-43	-45,062
2003	181	623	-36	-22,185
2004	178	723	-32	-23,204
2005	175	539	-29	-15,401
2006	174	535	-24	-12,799
2007	174	837	-16	-13,428
2008	169	679	-9	-6,255
2009	164	352	3	964
2010	152	465	12	5,759
		Total		**-400,796**
		Share of 5.65 million lost		**0.07**

Source: Estimates based on Table 1, Table 2, and Figure 2.

Table A5: Appendix Table Detailing Aggregate Job Loss from All Multinational Plants

	Average Size	# of Mult Establishments	Extensive, Weighted: Avg. Differential Employment per Establishment[1]	Extensive, Weighted: Total per year	Extensive, Unweighted: Avg. Differential Employment per Establishment[2]	Extensive, Unweighted: Total per year
1994	310	17,119	-8.0	137,112	-9.7	166,341
1995	311	16,269	-8.0	130,612	-9.7	158,456
1996	309	16,316	-8.0	129,956	-9.7	157,660
1997	306	16,365	-7.9	129,359	-9.6	156,935
1998	313	15,950	-8.1	128,823	-9.8	156,285
1999	312	16,084	-8.0	129,307	-9.8	156,872
2000	299	16,466	-7.7	127,067	-9.4	154,155
2001	297	15,886	-7.7	121,800	-9.3	147,766
2002	296	15,386	-7.6	117,568	-9.3	142,631
2003	279	14,930	-7.2	107,524	-8.7	130,446
2004	275	14,823	-7.1	105,186	-8.6	127,609
2005	270	14,692	-7.0	102,480	-8.5	124,326
2006	270	14,534	-7.0	101,095	-8.4	122,646
2007	269	14,482	-6.9	100,475	-8.4	121,894
2008	261	14,641	-6.7	98,763	-8.2	119,817
2009	254	14,456	-6.5	94,562	-7.9	114,721
2010	235	13,865	-6.1	83,888	-7.3	101,771
2011	222	13,562	-5.7	77,721	-7.0	94,290
			Total	**2,023,296**		**2,454,619**
			Share of 5.65 million lost	**0.36**		**0.43**

Source: Estimates based on Table 1, Table 2, and Table 5.
[1]This column applies the coefficient estimates from the intensive/extensive and weighted estimates from Table 5.
[2]This column applies the coefficient estimates from the intensive/extensive and unweighted estimates from Table 5.

Table A6: Regression Results: Subsamples

		Establishment Level			
		Intensive		Extensive and Intensive	
		Unweighted	Employment Weighted	Unweighted	Employment Weighted
1993 - 2000	β	0.02***	0.01***	-0.04***	-0.03***
	S.E.	(0.001)	(0.001)	(0.003)	(0.002)
	Clusters	8179	8179	8606	7081
2001 - 2011	β	0.02***	0.004***	-0.03***	-0.03***
	S.E.	(0.001)	(0.001)	(0.002)	(0.002)
	Clusters	8437	8437	8922	8922
		Firm Level			
		Intensive		Extensive and Intensive	
		Unweighted	Employment Weighted	Unweighted	Employment Weighted
1993 - 2000	β	-0.01***	-0.03***	-0.06***	-0.04***
	S.E.	(0.004)	(0.006)	(0.01)	(0.01)
	Clusters	3481	3481	3931	3931
2001 - 2011	β	-0.01***	-0.01***	-0.01***	-0.01***
	S.E.	(0.003)	(0.005)	(0.006)	(0.008)
	Clusters	4547	4547	5187	5187

Source: LBD, DCA, and UBP. The table reports pooled regression results, where the sample is split into subsamples from 1993-2000 and 2001-2011.

C Appendix: Structural estimation

C.1 Estimation

This appendix lays out the procedure we use to find bounds of the constant $(\sigma - 1)/\theta$. The model predicts that

$$R(\varphi) = \frac{\sigma}{\sigma - 1} \frac{w_j l_j(\varphi)}{\chi_j(\varphi)}$$

so the results present here apply whether we use revenues $R(\varphi)$ or $\frac{w_j l_j(\varphi)}{\chi_j(\varphi)}$ as the dependent variable.

Revenues of a firm of type φ are given by

$$R(\varphi) = \left(\frac{\sigma}{\sigma - 1}\right)^{1-\sigma} (\gamma)^{\frac{1-\sigma}{\theta}} E P_X^{\sigma-1} [\Phi(\varphi)]^{\frac{\sigma-1}{\theta}} s(\varphi)$$

and the sourcing share from location/mode j is

$$\chi_j(\varphi) = \frac{T_j [h_j(\varphi)]^{\theta} (\tau_j w_j)^{-\theta}}{\Phi(\varphi)}$$

Next we construct the sum of shares over some strict subset I of J.

$$\sum_{j \in I} \chi_j(\varphi) = \frac{\sum_{j \in I} T_j [h_j(\varphi)]^{\theta} (\tau_j w_j)^{-\theta}}{\Phi(\varphi)}$$

Solving for $\Phi(\varphi)$, substituting into the expression for revenues, and taking logs gives

$$\ln R(\varphi) = \Psi_I - \frac{\sigma - 1}{\theta} \ln \sum_{j \in I} \chi_j(\varphi) + \frac{\sigma - 1}{\theta} \ln \left(\sum_{j \in I} T_j [h_j(\varphi)]^{\theta} (\tau_j w_j)^{-\theta} \right) + \ln s(\varphi) \quad \text{(A5)}$$

where Ψ_I is a fixed effect. Strictly speaking Ψ_I does actually not depend on the set I. However, since the nonparametric term does depend on I, we always allow the constant to depend on I.

Next, we fix a particular sourcing strategy J and partition it into the strict subsets $I_1, ..., I_S$. We then estimate equation (A5) for all $I_1, ..., I_S$ and obtain S estimates of $-\frac{\sigma-1}{\theta}$. Now the same logic as described in the text applies. As the sample size tends to infinity, the true value of $-\frac{\sigma-1}{\theta}$ must lie between the smallest and the largest estimates we obtain. Of course, in practice these bounds are estimated with error.

C.2 Invertibility of $\eta_{m,k}$

To show that $\eta_{m,k}$ is invertible for some m and k, we estimate specification 19 separately for all $j \in \mathcal{J}$. Note, the technology transfer function is not conditional on a sourcing strategy, but only on a location/mode j. Unlike the bounding procedure, therefore we do not condition on a particular sourcing strategy $J \subset \mathcal{J}$, but pool all observations that source from a given location/mode. The results of the estimation are shown in Table A7.

The table shows that $\frac{\sigma-1}{\theta}$ is severely upward biased when $j = HO$. In contrast, the estimate is most downward biased when $j = HI$. For all other $j \in \mathcal{J}$, the estimates are quite close together and lie between these two extremes. The structure of our model now suggests that χ_{HO} is strictly increasing in φ while χ_{HI} is strictly decreasing in φ. This implies that $\eta_{HO,HI}$ is strictly increasing and therefore invertible.

We next estimate 21 for $j = HO$ and plot the semi-parametric component $(\sigma - 1)$ .. as a function of $\frac{chi_{HO}}{\chi_{HI}}$. The result is shown in Figure C.2. As expected, the technology transfer function is increasing in the share ratio.

Table A7: Bias in single share estimates of $\frac{\sigma-1}{\theta}$

	χ_{HI}	χ_{HO}	χ_{NI}	χ_{NO}	χ_{SI}	χ_{SO}
$\frac{\sigma-1}{\theta}$	1.798***	-1.263***	0.095***	0.156***	0.137***	0.254***
	(0.0259)	(0.0571)	(0.0184)	(0.0158)	(0.0308)	(0.0188)
Observations	32,000	32,000	2,100	6,000	1000	2900
R2	0.168	0.054	0.024	0.014	0.051	0.070

Source: LBD, LFTTD and CMF

This table reports the results from estimating 19 for all firms in 1997. The single shares are instrumented with lagged shares. F statistics for the first stage are significant at conventional levels. The results for years 2002 and 2007 (not shown) are similar.

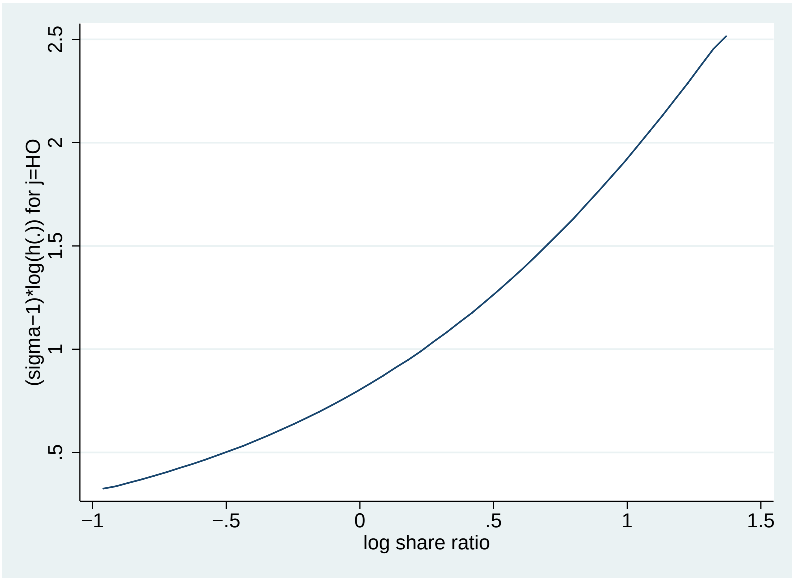

Source: LBD, LFTTD and CMF

This figure displays the results of plotting $h_j(\varphi)$ on $\frac{\chi_{HO}}{\chi_{HI}}$ as discussed in the text. The size distribution of $\frac{\chi_{HO}}{\chi_{HI}}$ is truncated at the 15th and 85th percentiles.

C.3 Estimation results: robustness

Table A8: Estimation Results: Semiparametric Regressions (Robustness)

Year	1993		1997		2002		2007	
$\frac{\sigma-1}{\theta}$	0.16***	0.12***	0.16***	0.12***	0.16***	0.08***	0.11***	0.06***
(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	
Higher order F.E.	YES	NO	YES	NO	YES	NO	YES	NO
Size percentiles	NO	YES	NO	YES	NO	YES	NO	YES
Instrumented	NO	NO	NO	NO	NO	NO	NO	NO
Observations	72,700	72,700	79,500	79,500	67,200	67,200	71,800	71,800
R-squared	0.96	0.96	0.95	0.96	0.96	0.97	0.97	0.97
$\frac{\sigma-1}{\theta}$			0.14	0.24***	0.18***	0.23***	0.20*	0.15***
			(0.200)	(0.011)	(0.024)	(0.011)	(0.118)	(0.010)
Higher order F.E.			YES	NO	YES	NO	YES	NO
Size percentiles			NO	YES	NO	YES	NO	YES
Instrumented			YES	YES	YES	YES	YES	YES
Observations			76,000	76,000	63,800	63,800	67,400	67,400

Source: LBD,LFTTD, CMF and ASM
This table reports point estimates for $\frac{\sigma-1}{\theta}$ from the polynomial approximation and size bin approaches discussed in 3.4, where the dependent variable is firm revenues. The lower panel displays results where the cost shares are instrumented with lagged values.

Table A9: Robustness to κ_j

	Baseline	All $\kappa_j = 1, j \neq HO$	$\kappa_j = 1, \in \{NI, SI\}$
Manufacturing Employment	-0.23	-0.14	-0.13
Multinational Employment	-0.20	-0.27	-0.27
Non-MN Employment Employment	-0.24	-0.08	-0.07

This table summarizes the decline in aggregate manufacturing employment within the model under alternative assumptions on which κ_j.

D Appendix: Quantitative Exercises

D.1 Robustness to choices of κ_j

This section presents results of fitting the model in Section 4 to the calibration targets in Table 13 with alternative choices for $\kappa_j, j \in \{HO, NO, NI, SO, SI\}$. We present the declines in employment between 1997 and 2007 implied by the model when (a) all technology transfer parameters with the exception of κ_{HO} are set to 1, and (b) all within firm technology transfer parameters are set to 1 ($\kappa_j = 1, \in \{NI, SI\}$).

Table A10: Robustness: Quantitative Exercises

	Data (1997- 2007)	Baseline	Only T_j changes	Only fixed costs change
Manufacturing Employment	-0.25	-0.23	-0.06	-0.03
Multinational Employment	-0.27	-0.20	0.06	0.13
Non-MN Employment Employment	-0.24	-0.24	-0.12	-0.10

This table summarizes the decline in aggregate manufacturing employment within the model under alternative assumptions on which parameters change between 1997 - 2007.

D.2 Counterfactual exercises

We next discuss the changes in employment implied in our baseline model if we (a) only allow the technology parameters $T_j, j \in HO, HI, NO, NI, SO, SI$ to change between 1997 and 2007 or (b) only allow the fixed costs of each sourcing strategy $f_j, J \in \mathcal{J}$ to change.

Table A10 presents the results of these alternative calibrations. Notice that in both counterfactual exercises, we do not change any of the calibration targets, so we have more targets than parameters to fit the model. Manufacturing employment falls in aggregate in both cases, but by a smaller amount than in the baseline. Further, multinational employment actually increases, with the largest effect in the calibration where only fixed costs fall to match observed importing patterns. The large declines in fixed costs in this case result in entry into multinational activity, which dominates the within-firm effect of declining domestic employment due to import substitution. Similar reasoning applies to the case with only technological improvements, but the effect is smaller.[48]

[48] We note that as we do not have many parameters to fit our calibration targets in these exercises, the fit of the model is not as close as in the baseline, which also affects outcomes.