

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

Christoph E. Boehm
UT Austin

Aaron Flaaen
Federal Reserve Board

Nitya Pandalai-Nayar
UT Austin
and NBER

September 21, 2020

First Version: October 2015

Abstract

We provide new facts about the role of multinationals in the decline in U.S. manufacturing employment between 1993-2011, using a novel microdata panel with firm-level ownership and trade information. Multinational-owned establishments displayed lower employment growth than a narrow control group and accounted for 41% of the aggregate decline. Newly multinational establishments experienced job losses, while their parent firms increased foreign input imports. We develop a model that rationalizes this behavior and bound a key elasticity with our microdata. The estimates imply that multinational offshoring was responsible for a sizable reduction in U.S. manufacturing employment.

JEL Codes: F14, F16, F23

Keywords: Multinational Firms, Offshoring, Outsourcing, Manufacturing Employment

*We would like to thank the editor Andres Rodriguez-Clare, as well as Dan Akerberg, Pol Antràs, Javier Cravino, Rafael Dix-Carneiro, Kyle Handley, Brian Kovak, Andrei Levchenko, Eduardo Morales, Lindsay Oldenski, Esteban Rossi-Hansberg, Susan Houseman, Linda Tesar, Felix Tintelnot, Kei-Mu Yi and seminar participants at the NBER SI, NBER ITI spring meetings, CEPR-ESSIM, Michigan, Princeton, Virginia, Austin, CREI, UCSD, Yale, Colby, Cambridge, University of Oregon, UIUC, BU, Tufts, IIES, Uppsala, Warwick, Notre Dame, WCTW, SED-AM and UIBE/Georgetown for valuable suggestions. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau, the Board of Governors, or its research staff. All results have been reviewed to ensure that no confidential information is disclosed. Email: chris.e.boehm@gmail.com, aaron.b.flaaen@frb.gov and npnayar@utexas.edu

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 few decades, paired with extraordinary expansions of multinational activity by U.S. firms. These job losses in the manufacturing sector have received disproportionate attention in the public debate, in part due to the perception that manufacturing jobs offer relatively high skill-adjusted wages.

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

In this paper, we ask whether understanding the behavior of multinational enterprises is important for understanding the U.S. manufacturing employment decline. Since multinationals may affect U.S. employment for a number of reasons, we narrow our focus to one salient mechanism: the foreign sourcing of intermediate inputs. That is, we ask whether multinationals accounted for a substantial portion of the manufacturing employment decline, and if so, whether foreign sourcing of intermediates was important for explaining this decline.

To answer this question, this paper uses a novel dataset together with a structural model to show that U.S. multinationals contributed to 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. This data also allows for the distinction between U.S. multinationals and other, foreign, multinationals with headquarters abroad. The expansion of these foreign multinational firms in the U.S. offers one example of how multinational activity can also serve to increase U.S. employment. To the best of our knowledge, our dataset is the first to permit a comprehensive analysis of the role of U.S. multinationals in the aggregate manufacturing decline in the United States.

We begin by establishing four new stylized facts. First, U.S. multinationals represented 33 percent of 1993 aggregate manufacturing employment but accounted for 41 percent of the subsequent decline. Second, given their size, U.S. multinationals had surprisingly high job destruction rates and low job creation rates relative to non-multinationals throughout this

period. Third, 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 four exercises suggest 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. Moreover, the fraction of U.S. multinationals sourcing inputs from developing countries 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 a reduction in the costs of foreign sourcing leads firms to increase or decrease U.S. labor demand. To understand the causal mechanism and to quantify the aggregate impact, we present a model of importing. In the model, firms import intermediate inputs from abroad either at arms-length or from foreign affiliates. The firm's optimal sourcing decision 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, a reduction in the costs of foreign sourcing leads firms to have access to cheaper intermediates. As a result, their unit costs fall and their optimal scale increases. This “scale effect” raises their U.S. employment. On the other hand, firms respond by optimally reallocating intermediate production towards the location with lower costs. This “reallocation effect” reduces U.S. employment. We will refer to a positive *net* effect as the complements case, as lower costs of foreign sourcing raise domestic employment. If the net effect is negative, we refer to foreign sourcing and domestic employment as substitutes, as a convenient shorthand.¹

We show that in partial equilibrium, the value of a single structural constant—the elasticity of firm size with respect to its sourcing capability—completely determines which of the two forces dominates. The range of previous estimates of this constant in the literature is large enough that foreign sourcing could be either complementary or substitutable with domestic

¹We require no assumption on whether intermediate inputs or primary factors are substitutes or complements in production.

employment. We therefore develop a method to structurally estimate an upper bound on this constant using our data on the universe of U.S. manufacturing firms. While a high value of the upper bound leaves open the possibility that foreign sourcing and domestic employment are complements, a low value of the bound unambiguously implies that the two are substitutes. Our method builds on the insight of Blaum, LeLarge, and Peters (2018) that changes in cost shares are informative about changes in firm unit costs.

Our estimates of the bound are small. This suggests that increases in intermediate purchases from abroad reduce U.S. employment at the firm-level when triggered by a reduction in the costs of foreign sourcing. Our estimates of the upper bound are robust to a number of alternative specifications and across subsamples.

As a final step, we conduct two exercises to evaluate what our estimates imply for aggregate manufacturing employment. In the first exercise, we use the observed changes in firm cost shares together with our parameter bounds to obtain model-implied predictions of the employment loss due to foreign sourcing. This approach captures both the direct impact of foreign sourcing by existing firms as well as the first-order impact on domestic suppliers, holding all else equal. It suggests that about one fifth of the aggregate manufacturing employment decline can be attributed to offshoring by multinationals.

Second, we use a general equilibrium extension of the model which captures additional features such as firm entry and exit. We calibrate it using parameters consistent with our estimated bounds and aggregate import data. Again, the model implies a quantitatively significant employment decline in response to foreign sourcing. The magnitude is similar to the earlier approach as general equilibrium effects broadly offset one another. We note that all attempts to quantify the consequences of foreign sourcing on aggregate employment require strong assumptions and thus must be interpreted with caution.

This paper contributes to a growing literature documenting the impact of international integration on labor markets. Since many commonly used firm-level datasets do not contain ownership information, it is often difficult to identify multinationals and their headquarter location. To overcome this constraint, Harrison and McMillan (2011), Ebenstein et al. (2014), and Kovak, Oldenski, and Sly (2018) study foreign sourcing by multinationals using Bureau of Economic Analysis (BEA) data. However, as these data only include multinationals, they do not permit analysis of multinationals' behavior relative to a non-multinational control group.²

Whether foreign sourcing increases or decreases domestic employment remains an active debate in the literature. A number of papers have found little or no employment reduction in various countries, including Desai, Foley, and Hines (2009), Kovak, Oldenski, and Sly (2018),

²To study plant closure in multinationals, Bernard and Jensen (2007) made use of a temporary link between the BEA and the Census.

Magyari (2017) [U.S.A], Braconier and Ekholm (2000) [Sweden], and Konings and Murphy (2006) [Europe], among others. On the other hand, and consistent with our results, several recent papers with data from other countries have found that firms treat foreign and domestic employment as substitutes.^{3,4}

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., 2016). Unlike our paper, these studies use regional and industry-level data. In a firm-level study, Pierce and Schott (2016) 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 effects of trade, or inequality more generally. For instance, Hummels et al. (2014) find negative wage effects of offshoring for low skilled workers using firm-level data from Denmark.

Finally, the structural model we present draws on a growing literature studying models of firm imports, including Eaton and Kortum (2002), Halpern, Koren, and Szeidl (2015), Antràs, Fort, and Tintelnot (2017) and Blaum, LeLarge, and Peters (2018). Relative to these papers, we extend an otherwise standard model of importing to capture key features of our micro data with the objective of increasing its suitability for empirical analysis. For instance, we make minimal assumptions on firms' information sets at the time they choose their sourcing strategy. Further, we explicitly allow firms to source inputs from abroad inter and intra-firm.

Given its prominent role in the public debate, this paper focuses attention on the manufacturing sector. Recent work, such as Fort, Pierce, and Schott (2018) and Bloom et al. (2019), finds suggestive evidence that foreign sourcing has offsetting positive effects in non-manufacturing sectors. Our findings do not rule out such positive effects, and in fact we present evidence that multinational-owned non-manufacturing establishments experience higher employment growth during our sample period.⁵ Since cross-sector transitions of displaced workers are costly (e.g., Ebenstein et al. (2014)) our findings are nonetheless important for understanding the labor market implications of foreign sourcing.

The next section presents empirical evidence establishing the role of multinationals in the U.S. manufacturing employment decline, and links this to their import patterns. Section 3 develops the partial equilibrium model, lays out the structural estimation and discusses the

³For instance Muendler and Becker (2010) find evidence for substitutability between home and foreign employment using German data, albeit holding unit costs constant.

⁴Artuç, Chaudhuri, and McLaren (2010) estimate U.S. worker switching costs in response to trade shocks and find they are high. Monarch, Park, and Sivadasan (2017) also find that offshoring firms in Census data experience declines in employment.

⁵Magyari (2017) finds that *total* firm employment increases after the China shock. This finding is consistent with ours if the non-manufacturing employment increase offsets the manufacturing employment decline.

results. Section 4 presents estimates of the aggregate employment decline due to offshoring, including a general equilibrium extension of the model. Section 5 concludes.

2 Data and Stylized Facts

This section presents a set of stylized facts key to understanding the role of foreign sourcing of multinationals in the decline in U.S. manufacturing. The facts in this paper come from a novel firm-level dataset that contains production and trade information for the universe of U.S. manufacturing firms, augmented with multinational ownership and affiliate information. With these data, we show that:

1. U.S. multinationals as a group accounted for a disproportionate share of the aggregate manufacturing decline,
2. given their size, U.S. multinationals had surprisingly high job destruction rates and low job creation rates relative to non-multinationals throughout our sample period,
3. U.S. establishments of multinational firms experienced lower employment growth than a narrow control group of establishments with similar characteristics, and
4. establishments transitioning into U.S. multinational status experienced prolonged employment declines while the parent firm increased imports of intermediates.

The first three of these facts highlight that the behavior of multinational enterprises is quantitatively important for the manufacturing employment decline. The fourth fact makes a connection to their importing behavior and suggests that foreign input sourcing is a potential driver of multinational's impact on domestic employment. Based on a model of importing, we demonstrate in Section 3 that the relationship between rising imports and falling domestic employment can be interpreted as causal.

2.1 Data

2.1.1 Identifying U.S. Operations of Multinational Firms

Most previous research on U.S. multinational firms has used surveys administered by the Bureau of Economic Analysis (BEA).⁶ These surveys provide data on the activities of foreign affiliates of U.S. parent firms; however, information on the U.S. plants of these firms is limited. As an additional constraint, it is typically not possible to link the data from these surveys to the

⁶See for instance Ebenstein, Harrison, and McMillan (2015) and Ramondo, Rappoport, and Ruhl (2016).

universe of U.S. establishments available in the datasets maintained by the U.S. Census Bureau. Without such linked data, it is difficult to make statements about the behavior of multinationals *relative* to non-multinational firms, or to construct aggregate decompositions. For instance, it is not possible to study the behavior of newly-multinational firms or establishments relative to other firms.⁷

This paper addresses these shortcomings by merging new indicators of the international activity and ownership characteristics of U.S. firms into the otherwise comprehensive data of the U.S. Census Bureau. These new variables come from a year-by-year link to two directories of international corporate structure: the LexisNexis Directory of Corporate Affiliations and the Directories of Multinational Firms published by Uniworld Business Publications. In Appendix A.1 we describe our methodology that links these directories to the Census Bureau Business Register using a probabilistic name and address matching algorithm and provide extensive documentation on the quality of the merge. Relative to other research relying on such “fuzzy merging” methods, we achieve very high coverage of the *firm-level* variables of interest as we use establishments as the unit of matching. This increases the probability of a firm-level match. To ensure that the multinational identifiers are consistent across time and to minimize spurious switching of firm status, we develop a series of checks and cleaning procedures, which are listed in Appendix A.2.⁸ We define a U.S. multinational firm as a firm with at least one affiliate abroad whose ultimate parent company is headquartered in the United States. Similarly, we define a foreign multinational firm as a firm with at least one U.S. affiliate, with an ultimate parent company headquartered outside the United States.

Our data on multinational firms in the U.S. manufacturing sector are not easily compared to the survey data from the BEA. First, the BEA data lack detailed information on the unconsolidated U.S. operations of the multinational firms. For instance, while the BEA uses a firm-level industry classification and contains little information on establishments, the Census data we use has detailed industry information for each establishment of the firm. This information allows us to exclusively focus on manufacturing.⁹ Second, the sample restrictions we make to ensure a *longitudinally consistent* definition of manufacturing (see below) creates a wedge in the comparison. Finally, the ownership thresholds used for our classification and that

⁷A growing literature has used alternative data sources to identify multinationals operating in the U.S. Bernard et al. (2010) use the presence of related-party firm level trade to identify firms as multinationals in U.S. Census data. This approach 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). Most previous studies of offshoring in the U.S. that have not used BEA data have been at the industry level.

⁸Our multinational indicator covers 98% of related party trade—an alternative method to identify multinationals in U.S. Census data.

⁹For information on the distribution of multinational firms across manufacturing industries, see Flaaen (2013b).

of the BEA need not align. While our classification of affiliates is based on the majority-owned definition, the BEA uses thresholds that vary for the particular statistics under study.¹⁰

An alternative that provides a more comparable benchmark is the data used by Bernard and Jensen (2007), which looked specifically at the U.S. manufacturing plants of U.S. multinational firms in Census data. To identify multinational firms, this study used a special one-year bridge database linking the BEA and Census data for the year 1987. The authors identified multinational firms as those that held at least 10 percent of total assets outside of the United States, and then matched this firm-level variable to establishments operating in 1992. According to this definition, U.S. multinationals represented 6 percent of total manufacturing establishments, and 26 percent of total manufacturing employment. The corresponding values for the closest available year in our data, 1993, are 5 percent of total establishments and 33 percent of total manufacturing employment. The differences are likely due to the fact that the data from Bernard and Jensen (2007) do not include any new U.S. multinational establishments in the five years between 1987 and 1992. Nevertheless, this comparison confirms that our dataset has similar coverage compared to the closest alternative in the literature.

2.1.2 Other Data

This paper relies on a number of restricted-use Census datasets that we have augmented with indicators of multinational status. To create a consistent definition of manufacturing for the period 1993-2011, we apply a concordance between the SIC/NAICS classification changes described in Fort and Klimek (2016). 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.3 for details on the specific steps underlying the construction of the consistent manufacturing sample. We obtain annual employment and payroll information from the LBD.

A further 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.

Our focus in this paper is on the impact of firms moving portions of their supply chains

¹⁰A direct comparison will be possible in the future due to ongoing work to merge the BEA microdata with the Census data. Researchers at the Census Bureau are using the multinational indicators we developed in this study to compare their merge of the BEA and Census microdata.

abroad. Many firms in the data import both goods intended for further manufacture within the firm (intermediate goods) as well as those destined for immediate sale (final goods). Using total imports of the firm as a measure for foreign sourcing could over or understate the impact on employment because firms could purchase final goods for reasons entirely unrelated to domestic production processes. We therefore develop a novel procedure for classifying firm-level imports into intermediate inputs and final goods, and focus our analysis only on the subset of all import transactions that are intermediates.^{11,12}

The procedure uses the Census of Manufacturers Products-Trailer File, which lists the detailed set of products, with SIC/NAICS product codes, produced by U.S. establishments in an industry. With this information, we can define a set of products intended for final sale for that industry.¹³ Imported products which match those for final sale for a given firm are then classified as final goods imports, and the remaining imports as intermediate inputs. Importantly, this classification yields values of the intermediate share of trade that are consistent with prior estimates: 64 percent of manufacturing imports are classified as intermediates in 2007. See Boehm, Flaaen, and Pandalai-Nayar (2019) or Appendix A.4 for more details on this classification procedure.

To estimate our model in Section 3.3, we require data on firm revenues as well as cost shares from various locations and modes (inter and intra-firm) of sourcing. For revenues, we use the total value of shipments of the firm’s manufacturing establishments from the Census of Manufacturers (CMF). Total costs are constructed from information on the cost of materials inputs, firm inter-plant transfers and total machinery expenditures of the firm (also from the CMF). We aggregate these variables to the firm-level in each of the Census years in our sample (1997, 2002 and 2007), and combine these total costs with expenditures on intermediate input imports as identified above to construct cost shares of inputs from different locations/modes.

2.2 Facts on Foreign Sourcing and Employment Decline

The decline in U.S. manufacturing is reflected in several aggregate statistics of 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

¹¹Bernard et al. (2018a) also emphasize the importance of distinguishing intermediate imports and final goods imports in recent work studying the reorganization of the firm after offshoring.

¹²Our approach does not rule out the firm producing high-value intermediates in the U.S. and assembling the final product abroad prior to re-importing it. Since that final product will have a different HS code than the U.S. produced intermediates, we would include it in our foreign sourcing measure.

¹³We use the concordances outlined in Pierce and Schott (2012) to map these products from an SIC/NAICS basis to the HS codes found in the trade data.

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	-2.54	-0.32	-0.33	0.21	-1.32	2.08	-1.74
1993-2001	-1.98	0.55	0.91	2.16	-1.33	1.82	-1.11
2001-2011	-2.99	-1.01	-1.31	-1.32	-1.31	2.29	-2.25

Notes: The data are from the LBD, LFTTD, DCA, and UBP as explained in the text. This table reports the establishment counts pertaining to the “consistent” manufacturing sample as discussed in Section 2.1.2.

have experienced a sizable net increase in establishments during this period is foreign multinational firms, due to the extensive (firm entry) margin. This group serves as a reminder that multinational activity could also stimulate U.S. employment.¹⁴

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 1993 manufacturing employment but contributed a disproportionate share, 41%, of the subsequent 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 the sample period. Domestic-only establishments experienced the highest annual *rates* of employment declines, but since they accounted for a smaller share of total employment in 1993, they contributed less to the overall manufacturing decline. Further, when drawing conclusions about the role of offshoring in the observed employment decline, a simple comparison of growth rates of different groups can be difficult to interpret. For example, employment in domestic establishments may have been affected by the changing sourcing patterns of multinationals, both directly—as the multinationals switched to foreign suppliers instead of sourcing inputs from purely domestic firms—and additionally, through general equilibrium channels. Indeed, we show in Section 4.1 that more than half of the employment decline that can be attributed to multinationals comes from substituting away from domestic arms-length suppliers to foreign suppliers.¹⁵

¹⁴Appendix Table B1 shows that the decline in multinational firms has not been as severe as the decline in multinational-owned establishments. Hence, establishment shutdown was one of the important margins of the employment decline in U.S. multinationals.

¹⁵The model we develop in Section 3 features such behavior. Further, general equilibrium extension in Appendix D also suggests that part of the observed manufacturing employment decline resulted from multinationals switching from purely domestic suppliers to foreign suppliers.

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

Year	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	-3.67	-2.48	-2.12	-1.86	-3.17	1.26	-2.41
1993-2001	-1.93	-0.49	0.83	1.00	-1.90	3.31	-0.55
2001-2011	-5.04	-4.04	-4.41	-4.09	-4.17	-0.35	-3.87
Net Change: 1993-2011							
Counts	-1,682,006	-775,266	-85,374	-1,048,843	-2,338,625	278,564	-5,651,550
Share	0.30	0.14	0.02	0.19	0.41	-0.05	1.00

Notes: The data are from the LBD, LFTTD, DCA, and UBP as explained in the text. This table reports the employment counts pertaining to the “consistent” manufacturing sample as discussed in Section 2.1.2.

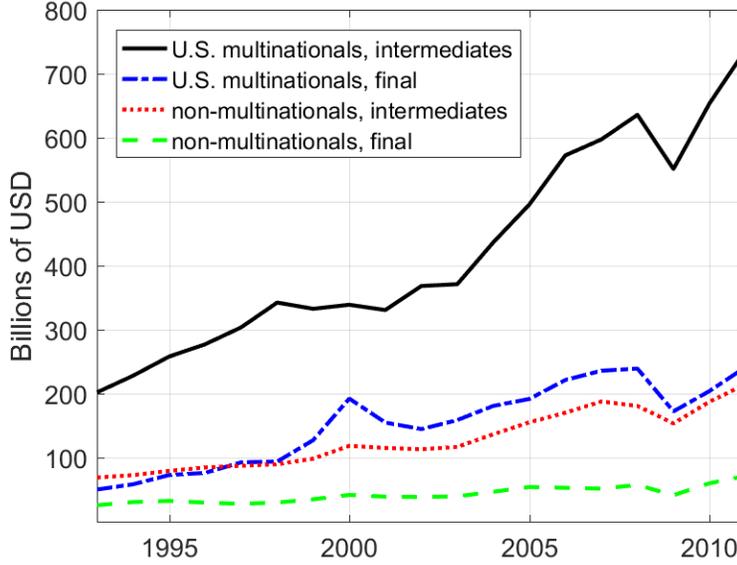
Concurrent with this employment decline has been a large increase in the participation and intensity of U.S. firms in trade. Figure 1 illustrates the growth of intermediate input imports, where we have split the sample into U.S. multinationals and other non-multinational U.S. firms. The rise in intermediate input imports of U.S. multinationals is striking. We also document the fraction of firms participating in intermediate input sourcing, separately based on whether it occurs at arms length or intra-firm, in Table 3. 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 multinational firms sourcing from developed countries has only increased about 10 percentage points during our sample period.¹⁶ Although non-multinational firms have also experienced increases in foreign input sourcing, the levels are roughly an order of magnitude smaller and they account for a small fraction of foreign input sourcing in the data.

2.2.1 Employment Growth Differential of Multinationals

We next decompose the aggregate employment growth rates of different groups of firms into job creation and job destruction rates as in, for instance, Davis and Haltiwanger (2001). We show that job creation and destruction rates vary substantially by establishment type: “domestic”

¹⁶The non U.S. multinational groups in Figure 1 and Table 3 are not directly comparable, as the group in Table 3 also includes foreign multinationals. This distinction is a result of disclosure limitations in our data as some groups have relatively few firms and thus pose limitations on the degree to which we can report splits of the sample.

Figure 1: Value of intermediate and final goods imports by firm type



Notes: The data are from the LFTTD, DCA, and UBP as explained in text. The figure shows the value of intermediate and final goods imports by firm type. The value of imports of foreign multinational firms is excluded from this figure.

non-exporting non-multinational (D), exporter (X), and U.S. multinational (MH).¹⁷ We further decompose the job creation and destruction margins into changes due to establishment births and deaths (extensive margin) or those due to employment changes at continuing establishments (intensive margin).

Formally, let employment of establishments in group $S \in \{D, X, MH\}$ at 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 calculate the job creation and destruction rates as

$$\text{Job Creation Rate: } \quad JC_{S,t} = \frac{\sum_{i \in S_{t-1}^+} \Delta E_{i,t}}{(E_{S,t} + E_{S,t-1}) / 2}, \quad (1)$$

$$\text{Job Destruction Rate: } \quad JD_{S,t} = \frac{\sum_{i \in S_{t-1}^-} |\Delta E_{i,t}|}{(E_{S,t} + E_{S,t-1}) / 2}, \quad (2)$$

where $\Delta E_{i,t} = E_{i,t} - E_{i,t-1}$ is the employment change in establishment i . To obtain rates by intensive and extensive margins, we further separate the groups into surviving establishments (existing in both $t - 1$ and t) and establishment births/deaths in a given year.

Figure 2 reports the job creation/destruction rates of the intensive margin for the three groups we study. The job creation rates in Panel A show both cyclicity and a secular decline

¹⁷Very few manufacturing firms in the data import without exporting. While Table 2 illustrated employment patterns in more disaggregate groups including importing non-exporting and exporting non-importing establishments separately, the small sample size of the group of pure importers restricted the set of facts that could be disclosed without further aggregation. For the same reason we do not report job creation and destruction rates for foreign multinationals.

Table 3: Percentage of Firms Participating in Foreign Input Sourcing

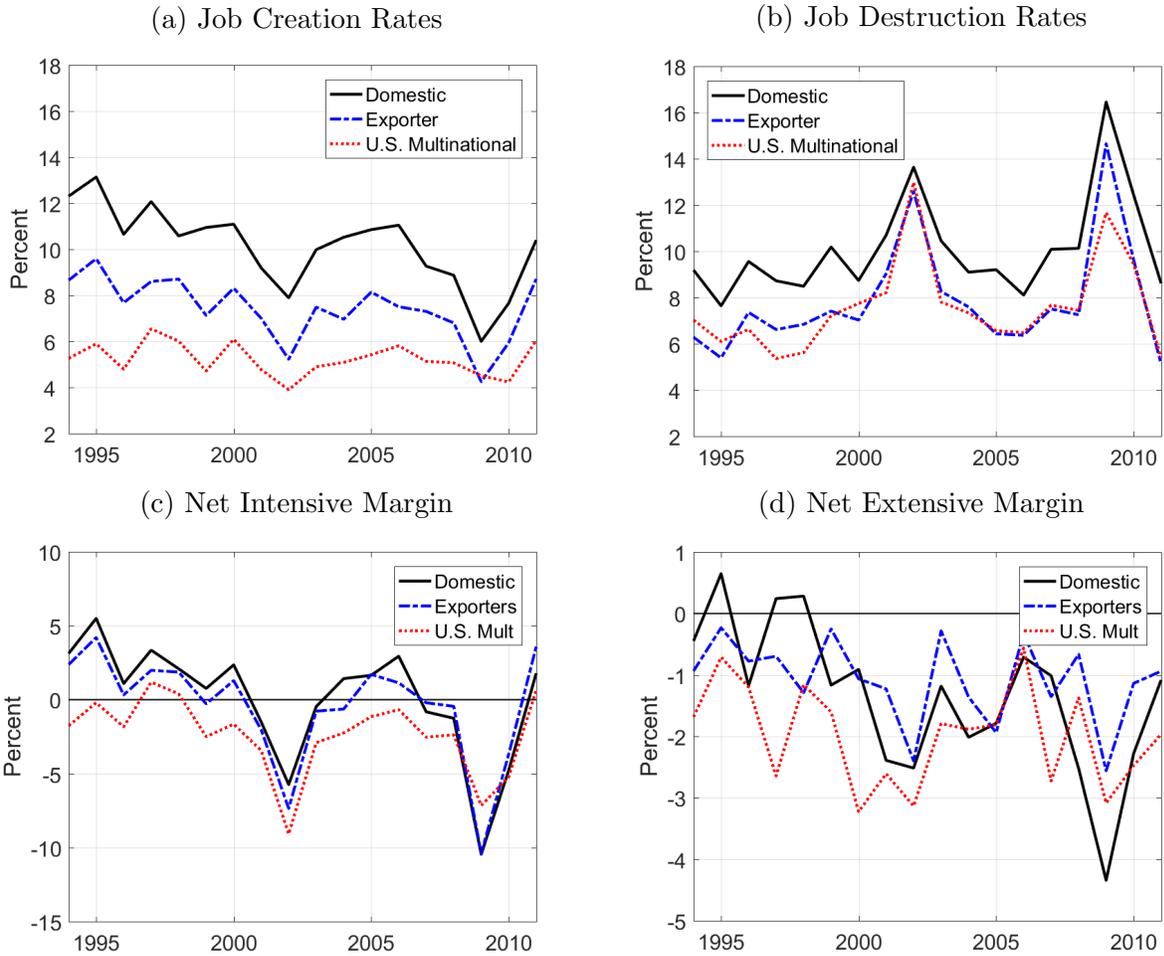
U.S. Multinationals				
Year	Arms Length Low Income	Related Party Low Income	Arms Length High Income	Related Party High Income
1993	44.35	24.62	72.63	48.17
2011	73.18	49.02	81.83	58.69
Non U.S. Multinationals				
Year	Arms Length Low Income	Related Party Low Income	Arms Length High Income	Related Party High Income
1993	1.88	0.39	5.71	1.30
2011	7.41	1.42	8.25	2.01

Notes: The data are from the LBD, LFTTD, DCA, and UBP as explained in the text. This table reports the fraction of U.S. multinationals and non U.S. multinationals that source inputs from foreign countries in percent (see also Appendix Table B1). The group non-U.S. multinationals includes foreign multinationals and other trading firms.

for domestic and exporting establishments. Multinationals' job creation rates exhibit less of a decline and less cyclicity, but their level is significantly lower in almost every year. Panel B shows the job destruction rates. These rates are higher for purely domestic firms and of comparable magnitude for exporters and multinationals. For all three groups of firms the job destruction rates are countercyclical. Previous research such as Davis and Haltiwanger (1992) has highlighted that job creation and destruction rates decrease with both firm size and firm age. Since multinationals are, on average, approximately 3 times larger than exporting establishments in our data, it is surprising that multinationals' job destruction rates are not much lower.

Panels C and D of Figure 2 combine the job creation and destruction rates at the intensive and extensive margin into net measures of employment gains by type of establishment. Panel C, which focuses on the intensive margin of establishments, shows that multinational establishments have had lower net growth rates in almost every year of our sample—with the gap narrowing only during recessions when clearly other forces are at work. These net growth rates were negative in 15 of the 18 years of the panel for multinationals which contrasts to 7(9) out of 18 for domestic (exporting) establishments. Panel D of Figure 2 shows that the relative employment loss of multinational establishments is also apparent in the extensive margin. Until 2003 multinationals had the lowest net rate of these three groups of firms in every year with one exception. From 2004 onwards the net rate of the extensive margin of multinationals was typically lower than exporters but higher than domestic firms. In summary, the picture that emerges is one where multinational firms shed more jobs than any other group of firms along

Figure 2: Job Creation and Destruction Rates: Intensive Margins and Net Job Creation



Notes: The data are from the LFTTD, DCA, and UBP as explained in text. These figures show the job creation and destruction rates at the intensive margin, and the net job creation at the intensive margin and extensive margin. as defined in equations (1) and (2).

both the intensive and the extensive margin and for the large majority of years for which we have data.

The analysis this far shows that the unconditional growth rates of multinational establishments differ systematically from both comparison groups. However, it is well-known that a variety of observable characteristics are systematically related to establishment employment growth. If any of these characteristics are correlated with multinational status, attributing the decline in employment to offshoring operations of multinationals would be misleading. To control for these establishment-level characteristics, we construct a set of indicator variables from the interactions of firm age, industry, establishment size, and year. More specifically, each indicator variable takes the value one if an establishment belongs to a cell defined by the interaction of the approximately 250 4-digit manufacturing industries, 10 establishment size categories, and 4 firm-age categories in a given year. This setup implies around 16000 cells

which we use as controls in the subsequent specifications, pooling across years 1993-2011.¹⁸ We then fit the regression

$$e_{it} = \alpha + \beta M_{it} + \mathbf{X}_{it} + u_{it} \quad (3)$$

where the dependent variable is the establishment employment growth rate, M_{it} is an indicator for establishments owned by a U.S. multinational, and \mathbf{X}_{it} is the vector of indicator variables identified above. Note that the employment growth rate is calculated following Davis, Haltiwanger, and Schuh (1996) (DHS), that is, $e_{i,t} = \frac{E_{i,t} - E_{i,t-1}}{0.5 \cdot (E_{i,t} + E_{i,t-1})}$. This definition allows us to estimate equation (3) on a sample which includes the extensive margin (with records of zero for $E_{i,t}$ in the year before entry and after exit) and on a sample which isolates variation along the intensive margin (without zeros).¹⁹

Table 4 presents the results from this specification. When pooling across years (1994-2011), and focusing only on the intensive margin, we find that multinational establishments have a *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. The importance of the extensive margin for net job destruction rates is consistent with the results from Figure 2d above. Taken together, and consistent with Bernard and Jensen (2007), this evidence points to establishment closure as important for understanding the employment decline in multinationals.

To assess the impact of this establishment-level result on overall employment within a firm, we estimate the same pooled specification with the firm as the unit of analysis.²⁰ 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 is consistent with the establishment-level estimates, since the firm-level intensive margin accounts for extensive margin changes at the establishment level (plant closings). The negative differential increases to 3 percentage points once the extensive margin (firm entry and exit) is included. The effects of establishment closure within the multinational firm dominate any increases in employment at existing establishments, leading to aggregate

¹⁸Our 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. 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. This strategy strikes a balance between high detail for a narrow comparison and broad coverage for an accurate characterization; the 4-digit NAICS is broad enough such that we do not eliminate many multinationals.

¹⁹While our preferred specification uses fully-interacted cells as flexible non-parametric controls, the results are nearly identical when using an alternative specification with polynomials in age and size as well as industry-time fixed effects as controls.

²⁰The firm here is defined as all manufacturing establishments of the same firm.

Table 4: Manufacturing Employment Growth Relative to a Control Group

	Establishment Level			
	Unweighted	Intensive Employment Weighted	Unweighted	Extensive and Intensive Employment Weighted
$\hat{\beta}$	0.019*** (0.001)	0.007*** (0.001)	-0.03*** (0.002)	-0.03*** (0.002)
Clusters	16,616	16,616	17,528	15,606
Observations	2784500	2784500	3204600	2765500
	Firm Level			
	Unweighted	Intensive Employment Weighted	Unweighted	Extensive and Intensive Employment Weighted
$\hat{\beta}$	-0.01*** (0.003)	-0.02*** (0.004)	-0.03*** (0.005)	-0.03*** (0.006)
Clusters	8,028	8,028	9,118	9,118
Observations	529100	529100	777300	777300

Notes: The data are from the LBD, LFTTD, DCA, and UBP as explained in the text. This table reports the pooled regression results from estimating equation (3) at the establishment and firm level. Standard errors are reported in parentheses. *** denotes significance at the 1 percent level.

decline.²¹ Appendix Table B5 displays results from this specification with different subsamples and additional controls for robustness.

Implications for Non-Manufacturing Employment Although this paper studies the implications of trade and multinational structure on U.S. *manufacturing* employment, it is important to keep in mind that there could be offsetting gains in other sectors of the U.S. economy. For example, the complex supply chains managed by multinational firms may require large and increasing administrative and support services by non-manufacturing establishments and workers. To evaluate whether these effects are present and offset the losses in manufacturing employment, we re-run specification (3) on the non-manufacturing establishments of the manufacturing firms in our baseline dataset. The results are shown in Table 5. We find modest evidence for positive employment effects on non-manufacturing employment of multinational firms, but only in the latter period of our sample. See Fort, Pierce, and Schott (2018) for a related analysis.

²¹A simple aggregation exercise based on our employment weighted regression results illustrates the number of jobs lost in U.S. multinational firms relative to the control group. To arrive at this number, we take the growth rates implied by the employment weighted specification and apply that to multinational employment in the sample year by year. 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.2.

Table 5: Non-Manufacturing Employment Growth Relative to a Control Group

		Establishment Level			
		Unweighted	Intensive Employment Weighted	Unweighted	Extensive and Intensive Employment Weighted
1993 - 2000	$\hat{\beta}$	-0.017** (0.007)	-0.008* (0.004)	-0.06*** (0.018)	0.001 (0.014)
	Clusters	14500	14500	15500	13000
	Observations	2,036,000	2,036,000	2,635,000	1,971,000
2001 - 2011	$\hat{\beta}$	-0.012** (0.005)	0.000 (0.004)	0.032** (0.015)	0.035*** (0.006)
	Clusters	14500	14500	15500	15000
	Observations	2,158,000	2,158,000	2,766,000	2,408,000

Notes: The data are from the LBD, DCA, and UBP. The table reports pooled regression results, where the sample is split into subsamples from 1993-2000 and 2001-2011. Standard errors are reported in parentheses. *, **, and *** denotes significance at the 10, 5, and 1 percent level.

We conclude this set of stylized facts by examining the transition dynamics of the employment and trade of establishments which become part of a multinational firm.

2.2.2 Evidence Using an Event-Study Framework

While previous sections studied the role of multinationals in the U.S. manufacturing employment decline, this section links this fact to their importing patterns. We analyze the change in outcomes (employment and trade) of establishments that transition into multinational status relative to a predefined control group. Using this event-study framework, we find that transitioning plants experience significantly lower employment growth, while their parent firms increase imports of intermediate inputs.²²

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 establishment's state is defined by the group it belongs to. We then explore whether changes in establishment state are an important feature of our data. Table 6 reports the average annual transition rates between states. As expected, most establishments maintain their status between years—as shown by the large diagonal entries—and transitions between states are quite infrequent. Only 0.03 percent of domestic establishments become part of a multinational every year. Establishments of exporting firms have a somewhat higher transition probability of 0.84 percent. We will next compare establishments which transition into multinational status to all

²²There have been several other recent papers that have analyzed such events for other countries, including Hijzen, Jean, and Mayer (2011) [France] and Debaere, Lee, and Lee (2010) [South Korea].

remaining establishments.

Establishment transitions into multinational status are endogenous outcomes and occur either when an entire firm becomes a multinational by acquiring an establishment abroad or when a multinational purchases an establishment from a non-multinational. Standard models of multinational production such as Helpman, Melitz, and Yeaple (2004) and Contessi (2015) suggest that transitions occur after positive idiosyncratic productivity shocks. Greater productivity (and similarly demand) makes it worthwhile for firms to pay the fixed costs of multinational production. This theoretical prediction is important for articulating a prior on how transitioning establishments compare to non-transitioning establishments. In particular, if transitioning plants are those whose parents experience positive productivity or demand shocks, one would expect that these plants should grow *faster* than non-transitioning plants. We will show momentarily that this positive growth differential turns negative surprisingly quickly after transitions.

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. 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 (all defined in period $y - 1$) we utilized above. We then compare the time path of employment growth rates of the transitioning establishments to their control group.²³ In our data, a new multinational establishment could result either from being acquired by a multinational firm or from the parent firm expanding operations abroad. As the latter case better approximates the experiment of adding a new foreign establishment to a firm, we restrict the sample for our baseline results to the set of transitioning establishments where the firm identifier remains the same between period $y - 1$ and period $y + 1$. The same restriction is applied to the control group. Finally, we highlight that while the control variables include a rich set of observables, they will not capture idiosyncratic shocks, which drive these transitions in standard models. Hence, transition should not be interpreted as an exogenous treatment. We will return to this discussion below.

As is clear from Table 6, 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

²³These 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.

Table 6: Establishment-Level Transition Probabilities

(percent)

t\ t+1	Dom	Exp	U.S. Mult	For Mult	Exit
Dom	84.59	5.41	0.03	0.04	9.93
Exp	13.56	79.80	0.84	0.52	5.29
U.S. Mult	0.27	1.85	90.95	0.87	6.06
For Mult	0.45	1.94	1.60	90.36	5.70
Entry	84.35	12.65	1.16	1.85	

Notes: The data are from the LBD, DCA, and UBP. The 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 sample ranges from 1993 to 2011.

we refer to as the “event” year. We then estimate the specification

$$e_{it}^y = \Gamma^y \mathbf{X}_i^y + \sum_{k=-5, k \neq 0}^{10} \delta_k T_{ik}^y + u_{it}^y, \quad (4)$$

where the variable T_{ik}^y is equal to one for transitioning establishment i in year k relative to the year of transition y . That is, calendar time t is $t = y + k$. The coefficients of interest δ_k are identified as the average effect of transitioning establishments relative to their control group, k periods after the transition year y .

The excluded group in the coefficients of interest δ_k is the transition year $k = 0$. The vector \mathbf{X}_i^y contains the indicators we utilized as controls above, and is fixed at time $k = -1$ for each event year, so that the comparison groups remain the same over time. The coefficients Γ^y on these indicators are allowed to differ by event year y . Note that the control groups are defined separately for each event year (and thus 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 the sample as follows: First as part of a control group for other transitioning establishments, and then, once, as part of a “treated” group of plants in the year of its own transition. Standard errors are clustered by plant and cell using the two-way clustering method by Cameron, Gelbach, and Miller (2011).²⁴

To study the behavior of imports around multinational transitions, we replace e_{ik}^y with imports (IM_{ik}^y) in equation (4). Such trade can be separately analyzed based on whether it is at arms-length or between related parties. Note that we observe trade at the *firm* level, unlike employment, which is measured at the *plant* level. Hence, this exercise compares the imports of the parent firm of transitioning establishments to the imports of the parent firms of the control

²⁴The standard errors change little when we alternatively cluster by firm or plant or plant and year.

group. As decisions to offshore or shut down plants are likely made at the firm level, this level of aggregation is preferable to using plant-level imports (which are also difficult to measure). Note that, as above, we focus our analysis on transitions associated with a firm acquiring an establishment abroad, as opposed to domestic mergers/acquisitions. This avoids changes in measured imports arising from changes in the transitioning establishment’s parent firm.

Panel A of Figure 3 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 transitions are driven by positive idiosyncratic productivity or demand shocks. Subsequent years, however, show a persistent negative effect on employment growth, in the order of 2 to 5 percentage points relative to the control group. The initial period of relative employment growth could reflect time spent by the firm replicating production processes abroad. Following a successful expansion, the firm may then choose to shut down or downsize duplicated firm activities.²⁵ Panel B of Figure 3 reports the implied cumulative effects. 10 years post transition, employment in transitioning plants falls approximately 10 percentage points relative to the control group.

Our results point to the importance of studying a long horizon to understand the consequences of offshoring. While a number of studies found weak positive or no effects on domestic employment, our evidence points to large but delayed negative effects. This discrepancy could, in part, reflect differences in the length of time under study. Finally, the exercise here does not condition on survival and fully captures the extensive margins of plant and firm closings as we use DHS growth rates on the left hand side of equation (4).

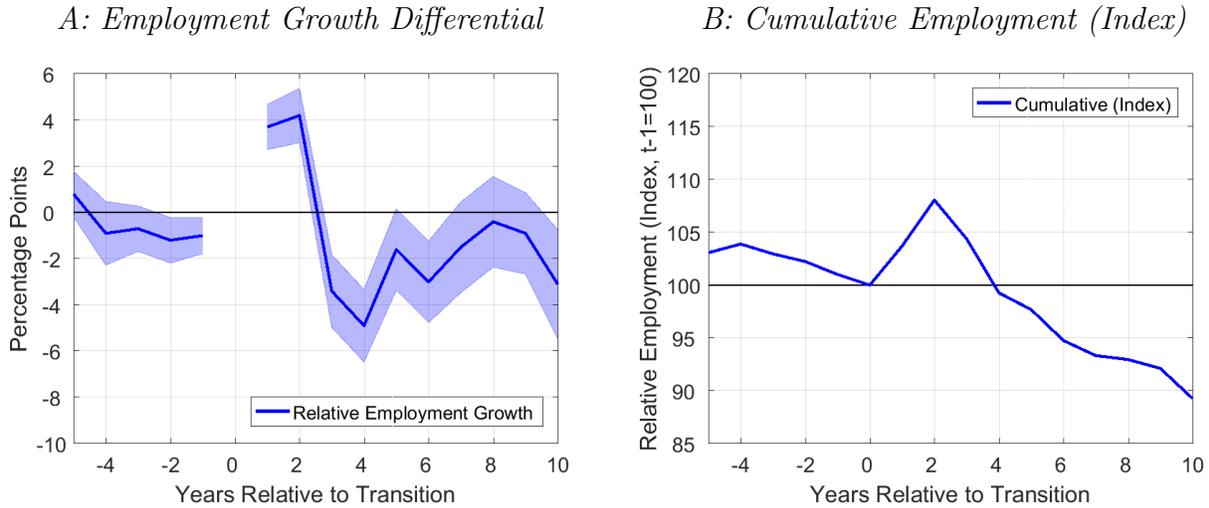
To examine the role of substitution towards imports in this decline, we estimate equation (4) after replacing the left hand side with firm-level intermediate imports (split by related party and arms-length).²⁶ Panel C of Figure 3 shows estimates of δ_k pertaining to imports. The figure demonstrates that transitions are associated with sizable increases in both related-party and arms-length intermediate imports. This suggests that firms may replace domestic production with intermediates imported from abroad. We find no evidence of pre-trends in any of the variables under consideration.²⁷

²⁵The precise timing of employment changes presented here should be interpreted with some caution, due—at least in part—to data reporting. Employment (a stock variable), is measured as of March in any given year, and hence reported employment changes following a multinational transition could mis-align with other variables (such as imports) by nearly a year. Further, the data we rely on to identify multinational operations may also imperfectly time the actual opening of foreign operations.

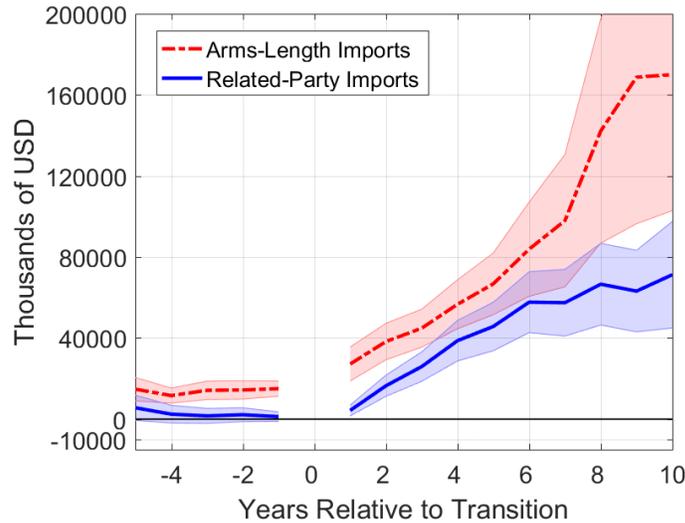
²⁶We use the level of imports as many firms either do not import or have consecutive periods of zero imports, ruling out using growth rates.

²⁷As most transitions into multinational status are by plants that belong to exporting/importing firms, and we do not condition on export status when creating a control group, the level difference in arms-length imports is unsurprising. We have found no difference in the results using alternative control groups that explicitly condition on the level of imports.

Figure 3: Evidence from Multinational Transitions



C: Importing Differentials of Multinational Transitions



Notes: The data are from the LFTTD, DCA, and UBP as explained in text. Panel A plots the annual employment growth rate differential of establishments that transition into a multinational firm in year ($t = 0$), relative to a control group with similar firm age, establishment size, and industry (in year $t = -1$). Panel B reports the cumulative coefficient estimates from Panel A, on an index scale. Panel C plots the related-party and arms-length intermediate input imports of the parent firm a transitioning establishment, relative to a control group which consists of establishments that are not part of a multinational firm in year $t = 0$. See equation (4) and text for details. The shaded areas corresponds to a 95 percent confidence interval.

The results in Figure 3 are robust to various alternative specifications. For instance, we also study transitions where an existing multinational firm acquires a non-multinational establishment (a merger/acquisition), and group all transitions (new multinationals and mergers/acquisitions) together. We find that our main results are robust to either form of transition (see Appendix B.1 for details.) We also report alternative specifications for our results on firm-level trade following multinational transitions in Appendix B. For instance, following research by Autor, Dorn, and Hanson (2013) and Pierce and Schott (2016), Appendix B.1.3 studies inter-

mediate input imports by multinationals from China. Consistent with these papers, Appendix Figure B5 shows large increases in imports from China post-transition.

2.2.3 Discussion

The evidence thus far suggests that U.S. multinationals contributed disproportionately to the manufacturing employment decline, with multinational establishments having lower employment growth rates than a narrow control group. We further documented that foreign sourcing of multinationals in the aggregate increased dramatically between 1993 and 2011, and that establishment transitions into multinational status were associated with establishment-level employment loss as well as increases in firm-level imports. We provide a number of additional facts in Appendix B. For instance, we show that arms-length *exports* also rise after transitions into multinational status, which is consistent with the view that transitions occur after idiosyncratic productivity or demand shocks, and that either these shocks or the increase in imported intermediates reduce unit costs and increase the optimal firm size—a prediction of the model below. We briefly discuss three related issues which help clarify the role of these facts.

Why multinationals? The evidence on lower employment growth after transitions into multinational status raises the question of the precise role of ownership of foreign affiliates. In particular, it is possible that foreign sourcing alone, perhaps exclusively at arms-length, is responsible for such employment declines. Since multinationals account for the vast majority of intermediate input imports (see Figure 1), these two channels are statistically difficult to distinguish. For the remainder of the paper we treat intermediate input sourcing as a key characteristic of multinational firms.²⁸

Alternative/additional explanations This section presented a series of facts to provide guidance for theories, which aim to explain the manufacturing employment decline. These facts are consistent with the view that foreign sourcing by multinationals *caused* the decline, but neither establish this causality nor do they rule out additional or alternative explanations. Besides trade-related explanations for the employment decline, other work has emphasized the role of technology. It is possible, for instance, that multinationals substituted more than other groups of firms towards labor-saving technology. While the facts we established above are

²⁸Bernard and Fort (2015) demonstrate that a nontrivial fraction of U.S. imports are imported by wholesale importers that are not in our sample. We do not observe many firms with only direct arms-length imports of intermediates, which could be because they purchase intermediates produced abroad that have been imported by an intermediary. Our data would incorrectly identify these as domestic input purchases, and as a result we would understate the economy-wide foreign sourcing of inputs. If this was true, however, these firms would appear in our control group of the analysis in Section 2.2.2, and thus bias the estimates toward zero.

not directly informative about these hypotheses, explanations which argue that the manufacturing employment decline is *exclusively* caused by technology have difficulties replicating the substantial increase in imported intermediates over the relevant time horizon.²⁹ In Section 3 we develop a model of foreign sourcing and confront this model with the observed increase in intermediate input imports. After estimating a key elasticity, the model implies that foreign sourcing substitutes for U.S. manufacturing employment.

IV-based strategies An alternative approach to get at the underlying causal relationship is to rely on IV-based strategies. In Appendix B.5 we present the results from a firm-level IV specification in the spirit of Acemoglu et al. (2016) and Hummels et al. (2014) where the instruments are constructed from changes in tariffs and exchange rates, using firms’ imports as weights.³⁰ While the OLS estimates reflect a strong positive correlation between employment growth and import growth, the coefficient becomes negative, albeit insignificantly so, when estimated using these instruments. Further, there is reason to believe that the instruments do not fully remove all confounding variation. Our approach of addressing this identification problem using a structural model is driven by the notorious difficulty in constructing convincing instruments with sufficient power at the firm level.

3 A Framework of Offshoring

We next develop a model of foreign sourcing that can explain the observed changes in employment. Our model builds on Eaton and Kortum (2002), Halpern, Koren, and Szeidl (2015), Antràs, Fort, and Tintelnot (2017), Tintelnot (2017), Bernard et al. (2018b) and Blaum, LeLarge, and Peters (2018), but relative to these contributions we adapt the model to the precise features of our micro data with the objective of increasing its suitability for empirical analysis. In particular, we extend an otherwise standard model of importing in three dimensions. First, firms can select a sourcing *location* for their intermediates, as well as a sourcing *mode*: whether to produce intra-firm or purchase goods from a different firm in a foreign location. Intra-firm production abroad is the defining characteristic of a multinational in the model, reflecting the vertical supply chain structure of U.S. multinationals in our data.

Second, we augment the model with a large number of structural shocks. Rather than

²⁹Fort, Pierce, and Schott (2018) discuss both trade and technology-related explanations and find that the data broadly support both views.

³⁰We have considered various other instruments based on transport costs, foreign GDP growth and exports to other countries (the “world export supply” instrument). Neither of these were relevant in our data. We also found that this class of instruments tends to exhibit a trade-off between exogeneity and relevance. A reasonable first stage typically requires firm-specific and time-varying weights which are lagged by only a few years—raising concerns about the exclusion restriction.

being characterized by a single productivity parameter, we increase flexibility in the model by allowing firms' productivity to vary by production location and mode. For instance, a firm can be highly productive at producing inside its own plants within the U.S., but its suppliers in an emerging market can be relatively unproductive. We also allow firms to differ along an unobserved demand component and regarding their fixed costs of foreign sourcing.

Third, we impose minimal assumptions on firms' information sets prior to choosing their sourcing strategy. Since it is unclear whether firms have perfect knowledge about their productivity in a foreign location or about their demand prior to starting production, we prefer to be as agnostic as possible on this dimension.

The model does not impose that an increase in imports following a reduction in the costs of foreign sourcing leads to a contraction in domestic employment. Instead, we show that whether imported inputs and domestic employment are complements or substitutes depends on a single structural constant. A high value of this constant implies that input imports and domestic employment are complementary, while a low value implies they are substitutes. We develop a method allowing to estimate an upper bound on this structural constant under fairly weak and conventional assumptions. Since our estimate of the upper bound is relatively small, we conclude that foreign sourcing and domestic employment are substitutes. While the estimation is based on the same microdata we used in the previous section, the unit of analysis in the model is best interpreted as a *firm*, as the decision to source intermediate inputs from abroad is presumably made at the firm-level. This contrasts to the previous section where we also analyzed the employment growth rates of *plants*.³¹

The model in this section is in partial equilibrium in the sense that it only describes firms in the manufacturing sector in the Home (U.S.) economy and takes sector-level aggregates as fixed. In the next section, we develop a simple general equilibrium extension, which requires stronger assumptions to illustrate how foreign sourcing might affect aggregate manufacturing employment.

3.1 Firms

There is a mass M of monopolistically competitive firms, each producing a unique variety. Firms are indexed by i , sourcing locations by ℓ , and sourcing location/mode *pairs* by j and k . We assume that firms are heterogeneous along three dimensions: a set of location/mode-specific productivity shocks $\zeta_{i,j}$ which we summarize in vector ζ_i , a demand shock δ_i , and a vector of fixed costs \mathbf{f}_i . We discuss all sources of heterogeneity in greater detail below. For now it is

³¹Neither the model nor the data we presented above are informative about the impact of foreign sourcing on workers of different skill levels or the wage premium of workers in global firms. A recent literature has focused on these questions using alternative models and data (see, e.g., Helpman, Itskhoki, and Redding (2010)).

sufficient to note that a firm is fully described by the tuple $(\zeta_i, \delta_i, \mathbf{f}_i)$, which we will refer to as the firm's type. Firms' shocks may be arbitrarily correlated. For instance, it is possible that a firm's productivity shocks $\zeta_{i,j}$ are correlated across locations/modes j , or that firms' demand shocks δ_i are correlated with their productivity shocks.

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

$$x_i = \left(\int_0^1 x_i(\nu)^{\frac{\rho-1}{\rho}} d\nu \right)^{\frac{\rho}{\rho-1}}. \quad (5)$$

Hence, the intermediates are imperfect substitutes with elasticity of substitution ρ . Letting $p_i(\nu)$ denote firm i 's price of variety ν , cost minimization in competitive factor markets implies that the unit cost of x_i is

$$c_i = \left(\int_0^1 p_i(\nu)^{1-\rho} d\nu \right)^{\frac{1}{1-\rho}}. \quad (6)$$

3.1.1 Supply chains

As we observe significant arms-length and intra-firm imported inputs in the data, we allow firms the choice of integrated or arms-length sourcing in each location. Sourcing inside the firm is indicated by I and sourcing outside the firm by O . Firms can source from $\ell = 1, \dots, N$ locations in addition to the home country H . The elements of the set $\mathcal{J} = \{H, 1, \dots, N\} \times \{I, O\}$ represent all possible sourcing locations and modes for varieties.

We model the firm's problem as follows. First, the firm chooses its sourcing strategy $J_i = J(\zeta_i, \delta_i, \mathbf{f}_i)$, 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 $(\zeta_i, \delta_i, \mathbf{f}_i)$, firms select the best option among these combinations of unit costs and fixed cost payments. We will return to the choice of firms' sourcing strategy below. For now we assume that the set $J(\zeta_i, \delta_i, \mathbf{f}_i)$ is given.

Intermediate goods production Intermediates in sourcing location/mode j are produced with production function

$$x_{i,j}(\nu) = \frac{\zeta_{i,j}}{a_{i,j}(\nu)} l_{i,j}(\nu). \quad (7)$$

While a common assumption in the literature is that firms can fully transfer their inherent productivity to their suppliers ($\zeta_{i,j} = \zeta_i$ for all j), our framework permits *any* form of technology transfer across firms' suppliers. For instance, it is conceivable that firms can transfer their technology to suppliers, but that production in location/mode j is also subject to idiosyncratic

shocks. In that case $\zeta_{i,j} = \bar{\zeta}_i \cdot \tilde{\zeta}_{i,j}$ where $\mathbb{E}[\tilde{\zeta}_{i,j}] = 1$. We emphasize this generality since a sufficiently flexible model is critical for taking the model to the data.

As in Eaton and Kortum (2002), the input efficiencies $1/a_{i,j}(\nu)$ are drawn from a Fréchet distribution with location parameter T_j and dispersion parameter θ . That is, $\Pr(a_{i,j}(\nu) < a) = 1 - e^{-T_j a^\theta}$. 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.³² This assumption accommodates a number of real-world features, for instance, that arms-length suppliers in a developing country may, on average, have lower productivity than those that would commonly integrate with a U.S. multinational. In that case $T_{\ell O} < T_{\ell I}$ for some country ℓ , implying, on average, lower productivity draws $1/a_{i,\ell O}(\nu)$ than $1/a_{i,\ell I}(\nu)$. Unlike $\zeta_{i,j}$, T_j is common across firms. Distinguishing these two productivity variables is convenient for the exposition below.

Suppose the inverse productivity draws $a_{i,j}(\nu)$ have materialized. Then, taking prices as given, a potential supplier of variety ν in location/mode j sets her price equal to marginal cost

$$p_{i,j}(\nu) = \frac{\tau_j a_{i,j}(\nu) w_j}{\zeta_{i,j}}, \quad (8)$$

where w_j and τ_j denote wages and iceberg transport costs. As is clear from this expression, the firm-destination/mode specific productivity shock $\zeta_{i,j}$ can equivalently represent idiosyncratic variation in transport costs or wages.

3.1.2 Basic model implications

Faced with price quotes from every location/mode in their sourcing strategy $J(\zeta_i, \delta_i, \mathbf{f}_i)$, 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

$$s_{i,j} = \frac{T_j \zeta_{i,j}^\theta [\tau_j w_j]^{-\theta}}{\sum_{k \in J(\zeta_i, \delta_i, \mathbf{f}_i)} T_k \zeta_{i,k}^\theta [\tau_k w_k]^{-\theta}}. \quad (9)$$

Clearly, for fixed $J(\zeta_i, \delta_i, \mathbf{f}_i)$, locations/modes with greater T_j , greater $\zeta_{i,j}$ and lower $\tau_j w_j$ have larger cost shares. Further, since $\zeta_{i,j}$ varies across firms, the model can accommodate cost shares of any size—which is critical for matching the data. Finally, the cost shares depend on the sourcing strategy $J(\zeta_i, \delta_i, \mathbf{f}_i)$. This implies that they also depend on the fixed cost draws \mathbf{f}_i that a firm must pay to set up its supply chain, and the demand shock δ_i .

³²See for instance Antràs (2005), Antràs and Helpman (2004) and Antràs and Chor (2013) among others for theories of intra-firm production.

Optimal input sourcing implies that the unit cost function (6) becomes

$$c(\boldsymbol{\zeta}_i, \delta_i, \mathbf{f}_i) = (\gamma)^{\frac{1}{\theta}} [\Phi(\boldsymbol{\zeta}_i, \delta_i, \mathbf{f}_i)]^{-\frac{1}{\theta}} \quad (10)$$

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

$$\Phi(\boldsymbol{\zeta}_i, \delta_i, \mathbf{f}_i) = \sum_{j \in J(\boldsymbol{\zeta}_i, \delta_i, \mathbf{f}_i)} T_j \zeta_{i,j}^{\theta} [\tau_j w_j]^{-\theta}. \quad (11)$$

Equation (11) summarizes the firm's efficiency at producing its unique variety. Following Antràs, Fort, and Tintelnot (2017) we refer to this term as the firm's (overall) sourcing capability. As is intuitive, firms with greater productivity in location/mode j (greater $\zeta_{i,j}$) and firms with more sourcing locations/modes have, all else equal, greater values of Φ_i and lower unit costs. Neither the cost shares (9) nor the unit costs depend on the quantity the firm produces.

3.1.3 Optimal firm size and the scale elasticity

We next turn to the problem determining the firm's optimal size. We assume that firm i faces an iso-elastic demand curve of the form

$$x_i = \delta_i E P_X^{\sigma-1} p_i^{-\sigma}, \quad (12)$$

where E is the household's expenditure on the manufacturing bundle, $P_X = (\int_{i \in \mathcal{I}} \delta_i p_i^{1-\sigma} di)^{\frac{1}{1-\sigma}}$ is the manufacturing price index, δ_i the firm's demand shock, and σ the demand elasticity. Given its unit costs, the firm chooses the price for its product to maximize flow profits, $p_i x_i - c_i x_i$, subject to the demand function (12). The firm optimally sets its price to a constant markup over marginal cost, $p_i = \frac{\sigma}{\sigma-1} c_i$. It is then possible to express revenues as

$$R(\boldsymbol{\zeta}_i, \delta_i, \mathbf{f}_i) \propto \delta_i E P_X^{\sigma-1} \cdot \Phi(\boldsymbol{\zeta}_i, \delta_i, \mathbf{f}_i)^{\frac{\sigma-1}{\theta}}. \quad (13)$$

In our framework, the elasticity of firm revenues (a standard measure of firm size) with respect to its sourcing capability Φ is

$$\varepsilon_{R,\Phi} := \frac{\partial \ln R}{\partial \ln \Phi} = \frac{\sigma-1}{\theta}. \quad (14)$$

As we will discuss momentarily, this *scale elasticity* is critical for the employment consequences of foreign sourcing.

3.1.4 The choice of the firm's sourcing strategy

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 $J \neq \{\text{HI}, \text{HO}\}$ necessitates the payment of fixed cost $f_{i,J}$. The vector \mathbf{f}_i is comprised of 2^{2N} fixed cost draws, one for each J in the power set of $\{1, \dots, N\} \times \{I, O\}$.

Prior to selecting its sourcing strategy, the firm learns its vector of fixed cost draws \mathbf{f}_i . Since it is not clear to what extent firms know about their productivity in all possible sourcing locations (as summarized in vector ζ_i) and its demand shock δ_i when they choose their sourcing strategy, we are agnostic on this point and simply assume that firms' information set is ι_i . This set contains the firms' fixed cost draws \mathbf{f}_i , but it may also contain signals on or actual values of the productivity draws ζ_i and demand δ_i . Our estimation strategy below will not depend on any timing assumptions regarding the revelation of shocks to the firm.

Upon learning ι_i , a firm selects its sourcing strategy $J_i \in \mathcal{J}$ to maximize expected profits

$$\mathbb{E} \left[\delta_i E P_X^{\sigma-1} \Sigma [\Phi(\zeta_i, \delta_i, \mathbf{f}_i)]^{\frac{\sigma-1}{\theta}} \mid \iota_i \right] - w_H f_{i,J}. \quad (15)$$

In this expression, Σ is a constant, w_H is the wage in the Home country, and the fixed costs $f_{i,J}$ are denominated in units of labor. $\mathbb{E}[\cdot \mid \iota_i]$ denotes the expectations operator conditional on information set ι_i . The solution to this problem is the firm's optimal sourcing strategy J_i which depends on all available information at the time of the choice.

3.2 Implications for Domestic Employment

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 $(\zeta_i, \delta_i, \mathbf{f}_i)$ with sourcing strategy $J(\zeta_i, \delta_i, \mathbf{f}_i)$ is

$$l_{HI}(\zeta_i, \delta_i, \mathbf{f}_i) \propto \frac{\delta_i E P_X^{\sigma-1}}{w_H} \cdot \underbrace{\frac{T_{HI} \zeta_{i,HI}^\theta [\tau_H w_H]^{-\theta}}{\Phi(\zeta_i, \delta_i, \mathbf{f}_i)}}_{s_{i,HI}, \text{ Reallocation effect}} \cdot \underbrace{\Phi(\zeta_i, \delta_i, \mathbf{f}_i)^{\varepsilon_{R,\Phi}}}_{\text{Scale effect}}. \quad (16)$$

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 $l_{HI}(\zeta_i, \delta_i, \mathbf{f}_i)$ is the labor required for this production.

Consider a reduction in the costs of foreign sourcing, 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 , lower costs of foreign sourcing affect $l_{i,HI}$ only through a change in Φ_i . Whether domestic employment rises or falls depends on the relative strength of two channels.

First, lower costs of foreign sourcing lead firms to shift a greater fraction of intermediate production towards the location with lower costs—a reallocation effect. This decreases $s_{i,HI}$ and thereby reduces domestic labor demand. On the other hand, lower costs of sourcing from abroad reduce the firm’s unit costs and increase sourcing capability Φ_i . This effect increases the firm’s optimal size and its domestic labor demand with elasticity $\varepsilon_{R,\Phi}$ —the scale elasticity we identified in equations (13) and (14). The net effect on employment is determined by the sign of $\varepsilon_{R,\Phi} - 1$. If negative, the model implies that the reallocation effect dominates and employment declines after a reduction in the costs of foreign sourcing. In this case we refer to foreign and domestic employment as substitutes. Conversely, if $\varepsilon_{R,\Phi} - 1$ is positive, foreign and domestic employment are complements. This notion of substitutability/complementarity is consistent with the terminology in Antràs, Fort, and Tintelnot (2017) and independent of the elasticity of substitution between intermediates ρ .

The same condition characterizes the firm’s change in labor demand after a change in its sourcing strategy, for instance due to lower fixed costs. If the firm adds an additional location/mode to its set J_i (e.g., it engages in offshore activities), Φ_i rises and the firm’s labor demand falls if and only if $\varepsilon_{R,\Phi} - 1 < 0$.

Hence in partial equilibrium, the sign of $\varepsilon_{R,\Phi} - 1$ completely characterizes the within-firm domestic employment response. If $\varepsilon_{R,\Phi} - 1 > 0$, one would expect recent productivity gains in emerging markets to increase U.S. manufacturing employment in firms that source from abroad. In contrast, if $\varepsilon_{R,\Phi} - 1 < 0$, 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.

3.2.1 Discussion

Absent general equilibrium effects, it is sufficient to know the scale elasticity $\varepsilon_{R,\Phi} := \frac{\partial \ln R}{\partial \ln \Phi}$ to characterize the consequences of foreign sourcing. As noted above, in this simple model of importing this scale elasticity is $\varepsilon_{R,\Phi} = \frac{\sigma-1}{\theta}$. Although this fact implies that knowledge of the scale elasticity provides information on σ and θ , and vice-versa, we do not emphasize this connection as our estimation method is more general than this particular mapping.

We illustrate in Appendix C.1 that an alternative framework with non-constant returns to scale changes the structural parameters underlying the scale elasticity. In that alternative model the estimation procedure we turn to momentarily will still correctly bound the scale elasticity

$\varepsilon_{R,\Phi}$ and this elasticity will still fully characterize the effect of foreign sourcing on domestic employment. However, without knowledge of the degree of returns to scale, the estimated scale elasticity $\varepsilon_{R,\Phi}$ is uninformative about σ and/or θ . Similarly, knowledge of σ and θ are only informative of the scale elasticity if the degree of returns to scale is known. Introducing love of variety for intermediates as in Benassy (1998) has similar implications as increasing returns to scale.

3.3 Structural Estimation

3.3.1 Estimating equation

Combining the cost share (equation 9) with the relationship for revenues (equation 13) or domestic labor demand (equation 16) yields

$$\ln y_i = \alpha_j - \varepsilon_{R,\Phi} \ln s_{i,j} + (\sigma - 1) \ln \zeta_{i,j} + \ln \delta_i, \quad j \in J, \quad (17)$$

where y_i is either revenues R_i or scaled firm payroll $\frac{w_{i,HI} l_{i,HI}}{s_{i,HI}}$. In the model the two variables are proportional to one another and both serve as measures of firm size. α_j is a location/mode-specific constant independent of firm characteristics. Equation (17) implies that holding the firm-location/mode-specific supply shock $\zeta_{i,j}$ and the demand shock δ_i constant, a reduction in the log cost share $\ln s_{i,j}$ by one infinitesimal unit raises the log firm size by $\varepsilon_{R,\Phi}$, which is precisely the elasticity which determines whether foreign sourcing and domestic employment are complements or substitutes as defined above.

Equation (17) holds for all sourcing locations/modes j in a firm's sourcing strategy J_i . For instance, for a purely domestic firm, it holds for $j = HI$ and $j = HO$. If the firm also sources at arms-length from country ℓ , then the equation also holds for $j = \ell O$, and so forth. Hence, and letting n_i denote the number of location/modes in firm i 's sourcing strategy, we have n_i equations per firm.

To obtain a single estimating equation per firm, we take the simple average of (17) over all locations/modes j in firm i 's sourcing strategy. This gives

$$\ln y_i = \frac{1}{n_i} \sum_{j \in J_i} \alpha_j - \varepsilon_{R,\Phi} \overline{\ln s_i} + u_i, \quad (18)$$

where $\overline{\ln s_i} = \frac{1}{n_i} \sum_{j \in J_i} \ln s_{i,j}$. The error u_i now contains averages of the firm's shocks,

$$u_i = (\sigma - 1) \overline{\ln \zeta_i} + \ln \delta_i, \quad (19)$$

where $\overline{\ln \zeta_i} = \frac{1}{n_i} \sum_{j \in J_i} \ln \zeta_{i,j}$. The location/mode-specific constants α_j are fixed effects scaled by $1/n_i$. It turns out that relative to equation (17), equation (18) has a key advantage. As we will show momentarily, the variation identifying $\varepsilon_{R,\Phi}$ in equation (18) is—to a first order—only the *number* of sourcing location/modes. In contrast, the variation identifying $\varepsilon_{R,\Phi}$ in equation (17) is firm i 's cost share from location/mode j , which is affected by the productivity and wages in location/mode j —in addition to the number of sourcing locations/modes (see equation 9). We will demonstrate that the former source of identifying variation is desirable in our context, because it allows us to bound elasticity $\varepsilon_{R,\Phi}$. We will therefore use equation (18) as our estimating equation. Of course, if the error u_i was uncorrelated with the average share $\overline{\ln s_i}$, it would be possible to estimate the scale elasticity $\varepsilon_{R,\Phi}$ consistently by Ordinary Least Squares (OLS).

3.3.2 Intuition

To understand the intuition connecting our estimating equation (18) to the scale effect, it is useful to consider the following approximation.

Lemma 3.1. *To a first order, $\overline{\ln s_i} = \alpha - \ln n_i$, for some constant α .*

Proof. See Appendix C.2. □

Lemma 3.1 states that a first order Taylor approximation of the average log cost share is decreasing in the number of sourcing locations/modes. It follows that, up to a first order, estimating equation (18) can be viewed as a regression of the log firm size on the log number of sourcing locations. To a first order, the log number of sourcing locations captures the scale effect.

To see why, note that the ideal experiment to estimate elasticity $\varepsilon_{R,\Phi}$ is to take two otherwise identical firms and to add one foreign location/mode to the sourcing strategy of the first of these two firms. (This can be done by reducing the fixed costs of sourcing from that location/mode for firm one.) After this addition, the first firm sources a positive fraction of intermediates from the newly added location/mode. This has two implications. First, the average cost share falls as it is decreasing in the number of sourcing locations/modes (Lemma 3.1). Second, because the firm only purchases intermediates from the new location/mode which are cheaper there than in any other location/mode of its sourcing strategy, its unit costs fall and its optimal scale increases. Thus, *ceteris paribus*, a lower average log cost share is associated with greater size.³³

In practice, the first order approximation of the average log cost share is surprisingly precise. In Table 7 below, we show that the R-squared of a regression of the average log cost share on

³³This intuition is closely related to Blaum, LeLarge, and Peters (2018), who show that changes in firm's domestic cost shares are informative about changes in unit costs.

the log number of sourcing locations/modes (and a set of fixed effects) is 92 percent. Of course, both the average log cost share and the log number of sourcing locations are choices of the firm, and hence endogenous. We address this issue below.

3.3.3 An upper bound for the scale elasticity

The estimation of the scale elasticity $\varepsilon_{R,\Phi}$ on the basis of equation (18), is complicated by the fact that $\overline{\ln \zeta_i}$ and δ_i are unobserved and thus enter the error term. We will next demonstrate that under plausible assumptions we can sign this omitted variable bias, leading to an upper bound of the scale elasticity $\varepsilon_{R,\Phi}$.

Lemma 3.1 implies that conditional on the firm's endogenous but observed sourcing strategy J_i , which determines n_i , $\overline{\ln s_i}$ is *uncorrelated* with the structural shocks in the error term—again up to a first order. This statement is obvious for the demand shock δ_i , since this shock affects the shares $s_{i,j}$ only through sourcing strategy J_i (see equation 9). However, it is also true for the productivity shocks $\zeta_{i,j}$ in location/mode j . Conditional on J_i , an increase in $\zeta_{i,j}$ raises the cost share from location/mode j , but decreases the cost shares from all other locations/modes $k \neq j$, because shares sum to one. Up to a first order, these effects cancel, leaving the average log share $\overline{\ln s_i}$ unchanged.

An implication of this fact is that the omitted variable bias that would arise if equation (18) was estimated by OLS, is solely determined by the correlation between the number of sourcing locations/modes n_i and the shocks $\overline{\ln \zeta_i}$ and δ_i . This feature is the critical advantage of estimating the *average* specification (18) rather than specifications which are based on individual shares as equation (17).³⁴

Given that the bias is solely determined by the covariance between $\ln n_i$ and u_i allows us to succinctly state the assumption for obtaining an upper bound for the scale elasticity.

Assumption 1. *Firms choose their sourcing strategy J_i , and hence the number of sourcing locations/modes n_i , such that*

$$\text{Cov} [\ln n_i, u_i] = \text{Cov} [\ln n_i, (\sigma - 1) \overline{\ln \zeta_i} + \ln \delta_i] \geq 0.$$

Assumption 1 states that firms with greater demand shocks δ_i and greater average productivity $\overline{\ln \zeta_i}$ are weakly more likely to source from more locations/modes. We view this assumption as

³⁴For the *average* specification (18) the condition for an upward-biased scale elasticity is stated in Assumption 1. An analogous assumption for the specification based on individual shares as in equation (17) is unlikely to hold. The reason is that the error contains the location/mode-specific component $\zeta_{i,j}$ which can be positively correlated with the share $s_{i,j}$, because greater productivity in destination/mode j lead to a greater cost share from that destination/mode. Such a positive covariance would bias the scale elasticity downwards. Hence, specification (17) cannot easily be used to bound the scale elasticity.

plausible and consistent with standard models in the literature.

Under Assumption 1, Lemma 3.1 immediately implies that OLS estimation of the scale elasticity $\varepsilon_{R,\Phi}$ from equation (18) delivers an estimate which is biased upwards. Since this result only holds up to a first order due to the approximation in Lemma 3.1, we use an additional, two stage least squares (2SLS) estimator using $\ln n_i$ as an “instrument” for $\overline{\ln s_i}$. Notice that $\ln n_i$ is not a true instrument, as under Assumption 1, the exclusion restriction is violated — the “instrument” is positively correlated with the error term. The resulting estimate is again upward biased, however in this case without an approximation. We summarize our bounding strategy in the following proposition.

Proposition 3.2. *Under Assumption 1, the following statements hold.*

1. *Up to a first order, the OLS estimate of $\varepsilon_{R,\Phi}$ based on equation (18) is biased upwards.*
2. *Without approximation, the 2SLS estimate of $\varepsilon_{R,\Phi}$ based on equation (18), using $\ln n_i$ as an “instrument” for $\overline{\ln s_i}$, is biased upwards.*

Proof. Under Assumption 1, $\text{Cov}[\overline{\ln s_i}, u_i] \leq 0$ up to a first order. The assumption $\text{Cov}[\ln n_i, u_i] \geq 0$ immediately determines the asymptotic bias of the 2SLS estimator. \square

3.4 Results

3.4.1 Baseline estimates

Table 7 presents the baseline results from a sample pooled across the Census years 1997, 2002 and 2007, and using $\ln R_i$ (revenues) as the dependent variable. We begin by estimating (18) without controls for productivity or demand. Under Assumption 1, the estimate for the scale elasticity $\varepsilon_{R,\Phi}$ will be biased upwards. As specification (1) of Table 7 shows, the estimate of $\varepsilon_{R,\Phi}$ is 1.1. If the bias is sufficiently large so that the true scale elasticity is below 1, the reallocation effect dominates the scale effect in our sample and foreign sourcing is a substitute for domestic employment.

We confirm this intuition by showing that adding productivity controls reduces the estimate of the scale elasticity. In specification (2) of Table 7 we add a third order polynomial in unit labor productivity. The estimate drops from 1.1 to 0.97. Adding industry-year fixed effects to control for common industry-specific demand and supply shocks further reduces the estimate of the scale elasticity to 0.8 and significantly below one. This estimate should still be interpreted as an upper bound on the elasticity, since the control variables imperfectly soak up the confounding variation. For instance, idiosyncratic demand disturbances remain in the error term. Under the assumption that the demand shock $\ln \delta_i$ is positively correlated with the number of sourcing

Table 7: Baseline estimates: pooled sample

Specification	Dependent variable: $\ln R_i$ (Revenues)			
	(1)	(2)	(3)	(4)
Scale Elasticity $\varepsilon_{R,\Phi}$	1.101*** (0.0506)	0.965*** (0.0457)	0.803*** (0.0175)	0.498*** (0.0213)
Polynomial in Unit Labor Productivity	No	Yes	Yes	Yes
Industry-Year FE	No	No	Yes	Yes
Year FE	Yes	Yes	No	No
Observations	726000	726000	726000	726000
R-squared	0.227	0.286	0.405	0.403
2SLS (instrument $\ln n_i$)	No	No	No	Yes
First Stage F				6895
First Stage R2				0.9208

Notes: The data are from the LBD, the CMF, and the LFTTD. The estimates are based on equation (18). As suggested by the model, we control for location/mode fixed effects interacted with $1/n_i$ in all specifications. In the pooled sample, the α_j are allowed to vary by year. To comply with Census disclosure requirements, the sample sizes are rounded. Heteroscedasticity robust standard errors are reported in parentheses. *** indicates that the estimate is significantly different from zero at the 1 percent level.

locations $\ln n_i$, a simple extension of Proposition 3.2 implies that the estimate of the scale elasticity is still biased upwards. Similarly, if our productivity controls imperfectly capture supply disturbances and the residual confounding variation is positively correlated with the number of sourcing locations, the estimate of the scale elasticity is still biased upwards.³⁵

The final column of Table 7 uses the number of sourcing location/modes n_i as an “instrument” for the average log cost share in estimating equation (18). As noted above, the 2SLS estimate is biased upwards even without requiring the first order approximation. The estimate of roughly 0.5 suggests that the true scale elasticity is likely very small. Hence, all our estimates thus far indicate that foreign sourcing is a substitute for domestic employment.

3.4.2 Robustness and extensions

We next discuss a number of extensions and robustness checks for the estimation of $\varepsilon_{R,\Phi}$.

Estimates by Census year Table 8 presents the estimates of the scale elasticity separately for each Census year in our data. As above, we estimate specification (18) first without productivity controls (specification 1), and then sequentially add a third order polynomial in unit

³⁵We have also estimated specification (18) without unit labor productivity controls, but with industry-year fixed effects. The estimates (not reported) are very similar.

Table 8: Baseline estimates: by Census year

Year	Specification	Dependent variable: $\ln R_i$ (Revenues)			
		(1)	(2)	(3)	(4)
1997	Scale Elasticity $\varepsilon_{R,\Phi}$	1.293*** (0.0945)	1.039*** (0.0743)	0.842*** (0.0283)	0.346*** (0.0348)
	Observations	270000	270000	270000	270000
	R-Squared	0.196	0.307	0.410	0.406
	First Stage F				1951
	First Stage R2				0.92
2002	Scale Elasticity $\varepsilon_{R,\Phi}$	0.838*** (0.0992)	0.748*** (0.0847)	0.673*** (0.0302)	0.486*** (0.0348)
	Observations	249000	249000	249000	249000
	R-Squared	0.196	0.275	0.392	0.391
	First Stage F				2454
	First Stage R2				0.92
2007	Scale Elasticity $\varepsilon_{R,\Phi}$	1.176*** (0.0526)	1.054*** (0.0475)	0.848*** (0.0222)	0.639*** (0.0375)
	Observations	207000	207000	207000	207000
	R-Squared	0.252	0.296	0.417	0.416
	First Stage F				2901
	First Stage R2				0.92
	Polynomial in Labor Productivity	No	Yes	Yes	Yes
	Industry FE	No	No	Yes	Yes
	2SLS (Instrument $\ln n_i$)	No	No	No	Yes

Notes: The data are from the LBD, the CMF, and the LFTTD. The estimates are based on equation (18). As suggested by the model, we control for location/mode fixed effects interacted with $1/n_i$ in all specifications. To comply with Census disclosure requirements, the sample sizes are rounded. Heteroscedasticity robust standard errors are reported in parentheses. *** indicates that the estimate is significantly different from zero at the 1 percent level.

labor productivity (specification 2), and industry-year fixed effects (specification 3). Specification (4) reports the 2SLS estimate when the number of sourcing locations/modes is used as an “instrument”. The upper bound of the scale elasticity is significantly below one in every Census year and relatively stable.

Alternative dependent variables We next replace the dependent variable with the log of scaled payroll, $\ln \frac{w_{i,HI} l_{i,HI}}{s_{i,HI}}$. The model predicts that scaled payroll is proportional to firm revenues and payroll may better capture the employment consequences of foreign sourcing.

Table 9: Robustness: alternative dependent variables

Dependent Variable Specification	Scaled Payroll		Domestic Revenues	
	(1)	(2)	(3)	(4)
Scale Elasticity $\varepsilon_{R,\Phi}$	1.051*** (0.0480)	0.484*** (0.0208)	1.109*** (0.0506)	0.489*** (0.0221)
Polynomial in Labor Productivity	No	Yes	No	Yes
Industry-Year FE	No	Yes	No	Yes
Year FE	Yes	No	Yes	No
2SLS (Instrument $\ln n_i$)	No	Yes	No	Yes
Observations	726000	726000	726000	726000
R-squared	0.224	0.382	0.217	0.394
First Stage F		6895		6895
First Stage R2		0.921		0.921

Notes: The data are from the LBD, the CMF, and the LFTTD. The estimates are based on equation (18) and the sample is pooled across years 1997, 2002, and 2007. As suggested by the model, we control for location/mode fixed effects interacted with $1/n_i$ in all specifications. In the pooled sample, the α_j are allowed to vary by year. To comply with Census disclosure requirements, the sample sizes are rounded. The dependent variable scaled payroll is calculated as $\ln \frac{w_{i,HI} l_{i,HI}}{s_{i,HI}}$. Domestic revenues refer to the log of total revenues minus exports. Heteroscedasticity robust standard errors are reported in parentheses. *** indicates that the estimate is significantly different from zero at the 1 percent level.

Table 9 reports the OLS estimates of $\varepsilon_{R,\Phi}$ without controls (specification 1) and the 2SLS estimate with controls (specification 2). Both estimates are almost unchanged. We also consider a measure of domestic revenues by removing exports from total revenues. Since many importing firms also export, one may be concerned that a simultaneous choice of the extensive margins of importing and exporting affects our results. The estimates, however, barely change when we consider domestic revenues instead of total revenues (specifications (3) and (4) of Table 9).

Measurement error A further concern is that the estimates are small due to measurement error in the average log cost share. This is unlikely as one can show that measurement error in individual cost shares leaves the average log cost share unaffected up to a first order. If, for instance, one cost share is measured too high, then other cost shares are measured too low as the shares sum to one. To a first order, this type of measurement error cancels out and leaves the average log cost share unchanged. Of course, the 2SLS estimates, which are smaller than the OLS estimates, also alleviate concerns about measurement error.

A related concern is that logarithms of small *individual* cost shares could become arbitrarily negative because the logarithm is unbounded. Such small values could unduely affect the *average* log cost share—our right hand side variable of interest. To rule out that this is the

Table 10: Robustness: Removing small shares and inter/intra firm distinction

Specification	Dependent variable: $\ln R_i$ (Revenues)			
	(1)	(2)	(3)	(4)
Scale Elasticity $\varepsilon_{R,\Phi}$	0.840*** (0.0228)	0.508*** (0.0356)	0.729*** (0.0101)	0.394*** (0.0505)
Polynomial in Labor Productivity	No	Yes	Yes	Yes
Industry-Year FE	No	Yes	Yes	Yes
2SLS (Instrument $\ln n_i$)	No	Yes	No	Yes
Remove small shares [†]	Yes	Yes	No	No
Distinction between AL and RP shares	Yes	Yes	No	No
Observations	726000	726000	726000	726000
R-squared	0.396	0.395	0.407	0.403
First Stage F		5757		6610
First Stage R2		0.922		0.927

Notes: The data are from the LBD, the CMF, and the LFTTD. The estimates are based on equation (18) and the sample is pooled across years 1997, 2002, and 2007. As suggested by the model, we control for location/mode fixed effects interacted with $1/n_i$ in all specifications. In the pooled sample, the α_j are allowed to vary by year. AL abbreviates to arms-length and RP related-party. To comply with Census disclosure requirements, the sample sizes are rounded. Heteroscedasticity robust standard errors are reported in parentheses. *** indicates that the estimate is significantly different from zero at the 1 percent level.

†: The average log cost share is calculated after reassigning cost shares from locations/modes below 1 basis point to the other share from the same country if available, and otherwise to the domestic arms-length (HO) share.

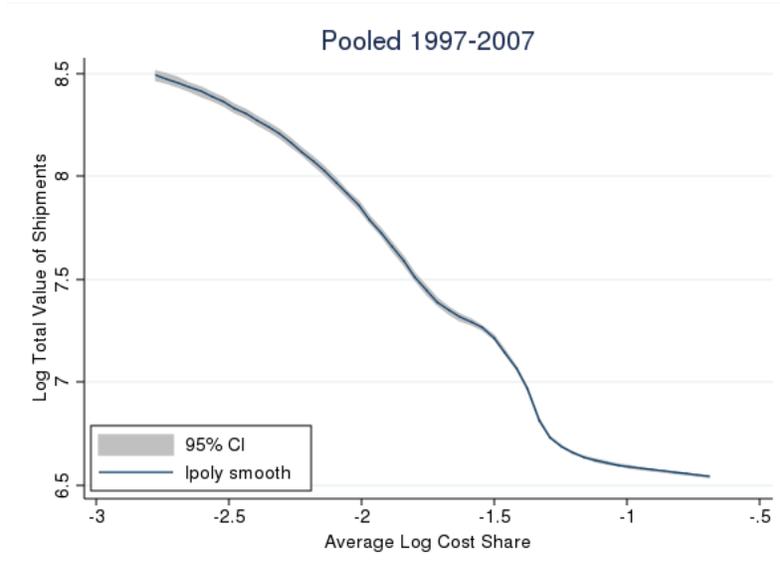
case, we reassign cost shares from locations/modes below 1 basis point to the other cost shares of the same firm.³⁶ As Table 10 shows in specifications (1) and (2), this correction has essentially no effect on the estimates.

Further robustness To show that our results are not driven by the distinction between arms-length and related-party sourcing we calculate cost shares exclusively by sourcing location (and not mode) and estimate the scale elasticity. The point estimate for the upper bound, reported in Table 10 specifications (3) and (4), are even lower than before.

Finally, we show in Figure 4 that the scale elasticity does not vary substantially with the number of sourcing locations/modes. The figure plots the log firm size as a function of the average log cost share, estimated using a kernel-weighted local polynomial. Since the slope is

³⁶More precisely, we reassign cost shares from locations/modes below 1 basis point to the other share from the same country if available. For instance, if a firm sources from a particular country at arms length less than 1 basis point of its total costs, then we reassign this value to the related-party cost share from the same country if the firm also sources from an affiliate in the same country. If the firm does not source from the same country, we reassign the value to the domestic arms-length (HO) share.

Figure 4: Elasticity of Size with Respect to Average Cost Shares



Notes: The data are from the CMF, the CMF, and the LFTTD as explained in the text. The figure plots the log firm size as a function of the average log cost share, estimated using a kernel-weighted local polynomial.

relatively constant, there is little evidence that the scale elasticity varies with firm size.³⁷

Instrumental variables estimation An alternate approach to estimating equation (17) would be to use an instrument which isolates variation in the cost shares orthogonal to the above confounders. In the model, the *only* source of such variation is heterogeneous fixed cost draws which lead otherwise identical firms to choose different sourcing strategies. We do not have an instrument that isolates such variation. One approach which has been taken in the literature is to difference the estimating equation and to use the Hummels et al. (2014) World Export Supply (WEX) instrument to estimate $\varepsilon_{R,\Phi}$. Unfortunately, this instrument is weak in our sample. As discussed in Section 2, alternative instruments based on exchange rate shocks or tariff/transport cost changes also have little to no relevance for the imports of the full population of U.S. firms in our 19-year panel.

3.4.3 Relationship to other estimates

In contrast to our estimates, Antràs, Fort, and Tintelnot (2017) find that $\varepsilon_{R,\Phi}$ is larger than one. This difference could be explained by a number of factors. Most importantly, our data and approach to creating a time-consistent manufacturing sample is very different from that

³⁷For small firms with average log cost shares between -1.5 and -0.5 the elasticity is slightly larger. However, this fact is driven by the approximate discreteness of the average log cost share. Using Lemma 3.1, firms with two, three, and four sourcing locations/modes have approximate average log costs shares of $-\ln 2 \approx -0.69$, $-\ln 3 \approx -1.10$, and $-\ln 4 \approx -1.39$.

used by Antràs, Fort, and Tintelnot (2017) as we drop from our sample any establishments that are headquarter establishments, likely to be primarily headquarter establishments, or non-manufacturing establishments. It is possible that non-manufacturing sectors and/or the non-manufacturing components of firms have a stronger scale effect. This would lead to larger estimates of $\varepsilon_{R,\Phi}$.³⁸ An additional difference is that we focus on imports of intermediates rather than all imports.

Kovak, Oldenski, and Sly (2018) find that foreign sourcing has a modestly positive effect for continuing multinationals. Their estimate can be interpreted as the net scale and reallocation effect for the firms in their data. In our model, this would imply a value of $\varepsilon_{R,\Phi}$ that is greater than one. There are several possible reasons for this difference. First, their employment measure explicitly includes U.S. headquarters (which likely include some non-manufacturing employment). Second, since their data only covers existing multinationals, changes in employment (and foreign sourcing) are not observable in the first year in which a firm becomes a multinational. Third, our data includes detailed information on arms-length input sourcing by multinationals, in contrast to Kovak, Oldenski, and Sly (2018). We therefore capture a broader range of activities of foreign sourcing. Finally, the methodologies differ. Kovak, Oldenski, and Sly (2018) also find a negative net effect when they examine the extensive margin.

4 Aggregation

4.1 A naive policy counterfactual

We briefly explore the aggregate implications of our estimates. Under the assumption that U.S. aggregates and firms' idiosyncratic shocks remain unchanged, our model in Section 3 implies that $d \ln l_{i,j} = (1 - \varepsilon_{R,\Phi}) \cdot d \ln s_{i,j}$. Hence, taking the changes in firm cost shares $d \ln s_{i,j}$ from the data, and aggregating across firms i , we can predict the aggregate employment decline implied by the changes in firm-level sourcing shares between 1997 and 2007 using our estimates of $\varepsilon_{R,\Phi}$. We use the tightest upper bound estimate of 0.5, which likely understates the within-firm job loss due to foreign sourcing. In this exercise we predict employment changes within firms sourcing from abroad (changes in $s_{i,HI}$), and first-order effects on their U.S. arms-length

³⁸Since the restricted-access microdata available for this paper do not include information on the cost structure of non-manufacturing/headquarter establishments, we are unable to fully assess the reasons behind the discrepancy between our results and Antràs, Fort, and Tintelnot (2017). We perform one exercise that supports the interpretation in the text. In Appendix B we replicate the transitions exercise from Section 2.2.2 on the non-manufacturing establishments of multinationals. We find no evidence of a relative drop in employment growth in the non-manufacturing sector.

suppliers (changes in $s_{i,HO}$).³⁹

This aggregation approach requires that we observe firms’ cost shares $s_{i,j}$ in both 1997 and 2007. It therefore cannot account for the declines in employment due to offshoring in firms that exited before 2007. It also underestimates the intensive margin effect in continuing firms, as some firm identifiers in the data change even though these firms continue to exist. Hence, both these factors, coupled with the fact that our estimate of 0.5 is an upper bound, suggest this procedure will underestimate the job-loss due to foreign sourcing. Note also that the results of this exercise are relative to the existing macroeconomic environment in 1997 and do not capture factors that might have changed between 1997 and 2007, such as new policies in other countries or structural change, but are not reflected in firms’ sourcing shares.

All else equal, this exercise suggests that for a scale elasticity of 0.5, 0.81 million lost jobs can be attributed to foreign sourcing, approximately 22 percent of the overall decline between 1997 and 2007. Of these 0.81 million, 0.34 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. (0.36 million), as well as foreign sourcing by non-multinational firms (0.11 million).

4.2 General Equilibrium

We next consider a general equilibrium extension of the model in order to capture several equilibrium features that were absent in the above “naive” counterfactual. These features include firm entry and exit, changes in the manufacturing price index, and substitution of manufacturing demand due to relative price changes, and as such are dependent on additional assumptions, such as preferences and labor mobility. To conserve on space we only provide a rough outline in this section and relegate the full exposition of the model to Appendix D.

The model is similar to that in Antràs, Fort, and Tintelnot (2017). We assume a three country world, composed of Home, North, and South. We distinguish between the North and South in this model for two reasons. First, we observe increased U.S.-based production by firms from predominantly developed countries (Tables 1 and 2). Second, while U.S. imports from developing countries grew rapidly over our sample period, imports from developed countries also increased (Table 3). By distinguishing the North from the South, our counterfactuals can capture both of these facts.

³⁹In particular, we compute these employment changes as

$$\sum_i (l_{i,j,2007} - l_{i,j,1997}) = (1 - \varepsilon_{R,\Phi}) \cdot \sum_i l_{i,j,1997} \frac{s_{i,j,2007} - s_{i,j,1997}}{s_{i,j,1997}},$$

for location/mode $j = HI$ (job loss within offshoring firms) and for $j = HO$ (first-order effects on U.S. arms-length suppliers). As before, firms are indexed by i . The sums are taken over all firms which are present in 1997 and 2007.

All countries produce a differentiated and traded non-manufacturing good, but only Home and North produce differentiated manufacturing goods. Similar to above, intermediates in manufacturing can be sourced from other countries, including the South. The households in all three countries consume manufacturing goods from Home and the North, as well as the differentiated non-manufacturing goods. Manufacturing firms are heterogeneous and only the most productive firms will pay the fixed costs of offshoring to obtain access to cheaper intermediates from abroad. In line with the estimate in specification (4) of Table 7, we assume a scale elasticity of approximately 0.5.

In our main counterfactual, we reduce the fixed costs of offshoring to match observed changes in trade flows over our sample period. All other parameters remain the same. Falling fixed costs of foreign sourcing imply more intermediate input sourcing from abroad. As a result, U.S. manufacturing employment falls. The increase in intermediate input imports explains approximately one-fifth of the aggregate manufacturing employment decline, similar to the naive counterfactual in Section 4.1. Note that the increase in foreign sourcing leads to small welfare gains in the Home country. For a full description of the model, its calibration, and the counterfactual exercises, see Appendix D. All attempts to assess the role of foreign sourcing in general equilibrium require strong assumptions and so should be interpreted with caution.

5 Conclusion

We present new facts showing that U.S. multinationals played a key role in the U.S. manufacturing employment decline. These firms accounted for 41% of the aggregate manufacturing employment decline between 1993-2011, and, given their size, displayed low job creation rates and high job destruction rates relative to non-multinationals throughout this period. U.S. multinational firms also had lower employment growth rates than a narrow control group of similar non-multinational firms. In our data, multinationals account for over 90% of all arms-length and related-party input imports and the fraction of U.S. multinationals which source from developing countries nearly doubled between 1993-2011. Using an event-study approach, we find that newly multinational establishments decrease employment while their parent firms increase imports of intermediate inputs. In summary, these facts suggest that U.S. multinationals reduced employment by offshoring the production of intermediate inputs to developing countries.

We then study a model of foreign input sourcing in which a single key elasticity—of firm size with respect to its sourcing capability—governs the employment impact of offshoring. Estimates of this elasticity imply that foreign sourcing is a substitute for domestic employment. Finally,

we present two approaches to quantify the aggregate impact on manufacturing employment. These suggest that the role of offshoring is sizable. An important caveat is that our paper exclusively studies the manufacturing sector. It is possible (and indeed, likely) that foreign sourcing is complementary to employment in the U.S. service sector (see Bloom et al. (2019)).

As in many models, the model we present in this paper generally implies U.S. welfare gains, rather than losses, from foreign sourcing. These welfare gains reflect consumers' access to cheaper manufacturing goods and imply that globalization is generally beneficial. However, our results add to the mounting evidence that international trade differentially exposes manufacturing workers to competition from abroad. This suggests that policy interventions that assist displaced workers may be desirable.

References

- Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price. 2016. "Import Competition and the Great U.S. Employment Sag of the 2000s." *Journal of Labor Economics* 34 (S1):S141–S198.
- Alfaro, Laura and Andrew Charlton. 2009. "Intra-Industry Foreign Direct Investment." *American Economic Review* 99 (5):2096–2119.
- 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. 2017. "The Margins of Global Sourcing: Theory and Evidence from US Firms." *American Economic Review* 107 (9):2514–64.
- Antràs, Pol and Elhanan Helpman. 2004. "Global Sourcing." *Journal of Political Economy* 112 (3):552–580.
- Arkolakis, Costas and Fabian Eckert. 2017. "Combinatorial Discrete Choice." Working paper.
- Artuç, Erhan, Shubham Chaudhuri, and John McLaren. 2010. "Trade Shocks and Labor Adjustment: A Structural Empirical Approach." *American Economic Review* 100 (3):1008–45.
- 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.
- Benassy, Jean-Pascal. 1998. "Is there always too little research in endogenous growth with expanding product variety?" *European Economic Review* 42 (1):61–69.

- Bernard, Andrew, Teresa Fort, Valerie Smeets, and Frederic Warzynski. 2018a. “Heterogeneous Globalization: Offshoring and Reorganization.” Working paper.
- 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. “Intrafirm Trade and Product Contractibility.” *American Economic Review* 100 (2):444–48.
- Bernard, Andrew B. and Teresa C. Fort. 2015. “Factoryless Goods Producing Firms.” *American Economic Review Papers and Proceedings* 105 (5):518–523.
- Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott. 2018b. “Global Firms.” *Journal of Economic Literature* 56 (2):565–619.
- Blasnik, Michael. 2010. “RECLINK: A Stata Module to Probabilistically Match Records.” Working paper.
- Blaum, Joaquin, Claire LeLarge, and Michael Peters. 2018. “The Gains from Input Trade in Firm-Based Models of Importing.” *American Economic Journal: Macroeconomics* .
- Bloom, Nicholas, Kyle Handley, Andre Kurmann, and Philip Luck. 2019. “The Impact of Chinese Trade on U.S. Employment: The Good, the Bad and the Apocryphal.” Working paper.
- Boehm, Christoph E., Aaron Flaaen, and Nitya Pandalai-Nayar. 2019. “Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake.” *The Review of Economics and Statistics* 101 (1):60–75.
- Braconier, Henrik and Karolina Ekholm. 2000. “Swedish Multinationals and Competition from High and Low-Wage Locations.” *Review of International Economics* 8 (3):448–461.
- Cameron, Colin, Jonah Gelbach, and Douglas Miller. 2011. “Robust Inference with Multiway Clustering.” *Journal of Business and Economic Statistics* 29 (2):238–249.
- Contessi, Silvio. 2015. “Multinational firms’ entry and productivity: Some aggregate implications of firm-level heterogeneity.” *Journal of Economic Dynamics and Control* 61:61 – 80.
- 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, and Scott Schuh. 1996. “Small Business and Job Creation: Dissecting the Myth and Reassessing the Facts.” *Small Business Economics* 8 (4):297–315.
- 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.

- 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.
- Flaen, 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. 2016. “The Effect of Industry Classification Changes on U.S. Employment Composition.” Working paper, U.S. Census Bureau.
- Fort, Teresa, Justin Pierce, and Peter Schott. 2018. “New Perspectives on the Decline of U.S. Manufacturing Employment.” Working Paper 24490, NBER.
- Halpern, László, Miklós Koren, and Adam Szeidl. 2015. “Imported Inputs and Productivity.” *American Economic Review* 105 (12):3660–3703.
- 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, Oleg Itskhoki, and Stephen Redding. 2010. “Inequality and Unemployment in a Global Economy.” *Econometrica* 78 (4):1239–1283.
- 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.
- Johnson, Robert C and Guillermo Noguera. 2016. “A Portrait of Trade in Value Added over Four Decades.” Working Paper 22974, National Bureau of Economic Research.
- 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. 2018. “The Labor Market Effects of Offshoring by U.S. Multinational Firms: Evidence from Changes in Global Tax Policies.” Working paper.
- Magyari, Ildiko. 2017. “Firm Reorganization, Chinese Imports, and US Manufacturing Employment.” Working paper.
- Melitz, Marc J. and Stephen J. Redding. 2015. “New Trade Models, New Welfare Implications.” *American Economic Review* 105 (3):1105–1146.
- Monarch, Ryan, Jooyoun Park, and Jagadeesh Sivadasan. 2017. “Domestic gains from offshoring? Evidence from TAA-linked U.S. microdata.” *Journal of International Economics* 105:150 – 173.
- 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.
- Pierce, Justin R. and Peter K. Schott. 2016. “The Surprisingly Swift Decline of US Manufacturing Employment.” *American Economic Review* 106 (7):1632–62.
- Ramondo, Natalia, Veronica Rappoport, and Kim J. Ruhl. 2016. “Intrafirm trade and vertical fragmentation in U.S. multinational corporations.” *Journal of International Economics* 98:51 – 59.
- Tintelnot, Felix. 2017. “Global Production with Export Platforms.” *The Quarterly Journal of Economics* .
- Wasi, Nada and Aaron Flaaen. 2015. “Record Linkage Using Stata: Pre-Processing, Linking and Reviewing Utilities.” *The Stata Journal* 15 (3):672–697.

ONLINE APPENDIX
(NOT FOR PUBLICATION)

A Data Appendix

A.1 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.1.1 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.1.2 The Matching Procedure: An Overview

The matching procedure uses a set of record linking utilities described in Wasi and Flaaen (2015). This program uses a bigram string comparator algorithm on multiple variables with differing user-specified weights.⁴⁰ 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, we use a relatively conservative criterion 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.⁴¹ 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

⁴⁰The 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 (2015) for more information.

⁴¹The primary sources of such ancillary evidence are clerical review of the matches, and additional parent identifier matching evidence.

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.

A.1.3 Matching Procedures: Yearly Steps

The following list describes the specific steps in the routine used to match the DCA and Census data (the UBP data matching used a similar routine).

1. Match DCA to Compustat (and then to Compustat-Bridge) for those DCA observations with a Compustat Identifier.
 - (a) The DCA to Compustat bridge was accomplished external to the Census data architecture, using similar name and address matching procedures. The percentage of DCA firms matched to Compustat firms was in line with our other match rates (≈ 70 percent).
2. Apply the standardization routines to the name and address variables of both DCA and the Business Register (BR). Remove DCA observations
3. Tier 1 Matching
 - (a) Restrict BR to LBD observations (save non-matching observations for Tier 2)
 - (b) Apply *relink2* of Compustat-linked DCA observations to BR (using name, street address, zip code, and requiring city, state, and firmid to match exactly)
 - (c) Apply *relink2* of non Compustat-linked DCA observations to BR (using name, street address, zip code, and requiring city and state to match exactly)
 - (d) Apply *relink2* of non-matching DCA to BR (using name, street address, zip code, and city, but now only requiring state to match exactly)
 - (e) Evaluate matches
 - if “match score” is above 0.95, classify as a match⁴².
 - if “match score” is between 0.80 and 0.95, evaluate manually⁴³
 - if “match score” is below 0.80, classify as a non-match
 - (f) Append evaluated-as-match DCA observations to the other matched observations, and send non-matching DCA observations to Tier 2 matching
4. Tier 2 Matching
 - (a) For the non-matching DCA observations, try to find an existing match with the same (DCA) parent identifier. Take the corresponding BR firm identifier (alpha or ein) for this match, and search for match over BR observations with identical alpha/ein
 - apply *relink2* of DCA to BR (using name, street address, and city, and requiring state and alpha/ein to match exactly)

⁴²We have looked at a several thousand of these potential matches and see that the false-positive rate for these is very small (i.e. less than 0.5 percent)

⁴³The manual evaluation of matches is the one step in which we utilize some longitudinal information. (Without this, the set of potential matches to evaluate was too large – in the range of 5-6 thousand per year.) Rather than continue to manually review common matches (and non-matches) from year to year, we keep the pool of manually evaluated matches from previous years and automatically accept as a match any potential match that exactly aligns with a match evaluated in a previous year. The same is true for previously-evaluated non-matches.

- if “match score” is above 0.70 classify as as match – spot checks have shown no false positives when requiring the alpha/ein to match
- (b) Implement additional name standardization routines
- (c) Apply *reclink2* of DCA to non-LBD-matched BR observations (using name, street address, and city, and requiring state to match exactly)
- if “match score” is above 0.95, classify as a 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. It is important to note that we implement the matching at the establishment level, whereas the variables we are using from these external directories are firm-level by their nature. Hence, the true degree of firm-level matching is in practice much higher than the establishment match rates from Table A1.

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

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.

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			
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

Notes: U.S. multinationals include only the establishments identified as the U.S. headquarters.

A.1.4 Within-Year Rules

To apply the matched establishment-level information to other, non-matched establishments within a given year, we apply the following steps. 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. As indicated above, this attractive feature of our methodology requires that the researcher only successfully match one plant of a given firm to apply that information throughout the firm. If there is 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.2 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 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.⁴⁴ 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.2.1 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.

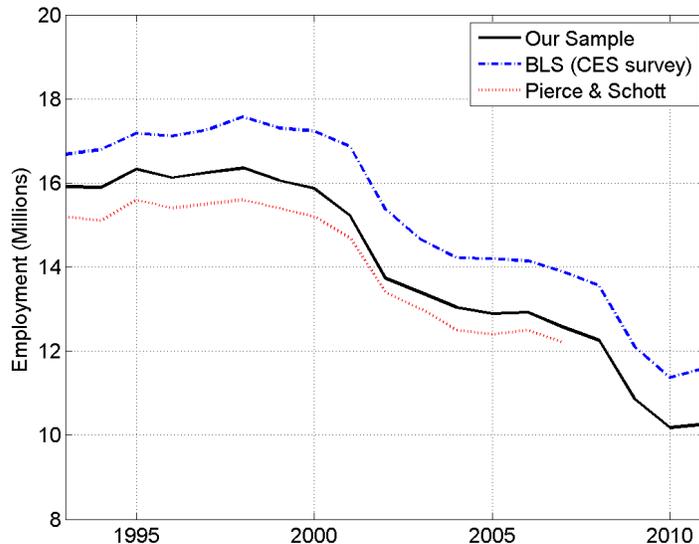
A.3 Creating Consistent Manufacturing Sample

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 consistent 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 (2016) for smoothing out discrepancies between industries defined as manufacturing between SIC and NAICS. There remain several acknowledged data issues of the Fort and Klimek (2016) manufacturing definition, principally related to manufacturing establishments that are re-coded into a non-manufacturing industry in 2002, specifically, NAICS 55 - “Management of Companies and Enterprises”. To deal with establishment transitions between manufacturing and non-manufacturing industries, we set up the following two rules. 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.

⁴⁴The 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.

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



Notes: The data are from the BLS, Pierce and Schott (2016) and the LBD.

Figure A1 illustrates how our consistent manufacturing sample compares to manufacturing employment from two other sources: published totals from the Current Employment Survey and Pierce and Schott (2016).

Finally, to evaluate whether there were offsetting gains to the non-manufacturing employment component of a manufacturing firm, we create a non-manufacturing sample that is the complement set of establishments of these manufacturing firms to what we identify above (hence, non-manufacturing establishments). This sample is used for results in Figure B6 and Table (??).

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.⁴⁵ Disentangling whether an imported product is used as an intermediate input for further processing — rather than for final sale in the U.S. — has 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.⁴⁶ 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

⁴⁵This 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 Section 2.

⁴⁶To be more precise, this set will include a combination of intermediate and capital goods.

establishment industry information, however, will allow us to parse a firm’s trade based on the intermediate/final distinction for a given establishment, thereby generating some heterogeneity in firm trade across establishments.⁴⁷

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 output (in U.S. \$) 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.

⁴⁷To be more precise, total trade at each establishment of a firm must be identical. The shares of intermediate/final goods will vary.

Acknowledging this concern, we set a 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.⁴⁸ 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.

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:

$$\left. \begin{aligned} M_{ij}^{\text{int}} &= \sum_{p \notin P_j} M_{ipj} \\ M_{ij}^{\text{fn}} &= \sum_{p \in P_j} M_{ipj} \end{aligned} \right\} \quad \text{where} \quad P_j = \{p \mid S_{pj} \geq W\}. \quad (\text{A1})$$

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

⁴⁸Another 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 concurring to the trade data is considerably more difficult.

Table A3: Comparing the Results from Threshold Values W

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

Notes: This table is applicable to the year 2007.

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 threshold level W . 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.2). 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. While we did use longitudinal information in classifying establishments by multinational status, we do not this information in aggregating up to the firm-level. However, some of our analysis in Section 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 B1: 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

Notes: The data are from the LBD, LFTTD, DCA, and UBP as explained in text. This table reports the firm counts pertaining to the “consistent” manufacturing sample used in section 2.

B Additional Results

B.1 Other Results on Transitions

B.1.1 Alternative Control and Treatment Groups for Employment Results

The baseline results on the relative employment growth rates of plants that transition to become part of a multinational firm shown in Figure 3 combines two cases: 1) where the non-multinational plant is acquired by a multinational firm, and 2) where the non-multinational plant’s firm becomes a multinational. These cases may offer distinct implications for employment growth for a number of reasons. For example, plants involved in a merger/acquisition may experience employment declines apart from whether or not the acquiring firm was also a multinational. We can identify these cases based on whether the firm identifier of the plant changes in the pre and post periods of the setup described in Section 2.2.2.

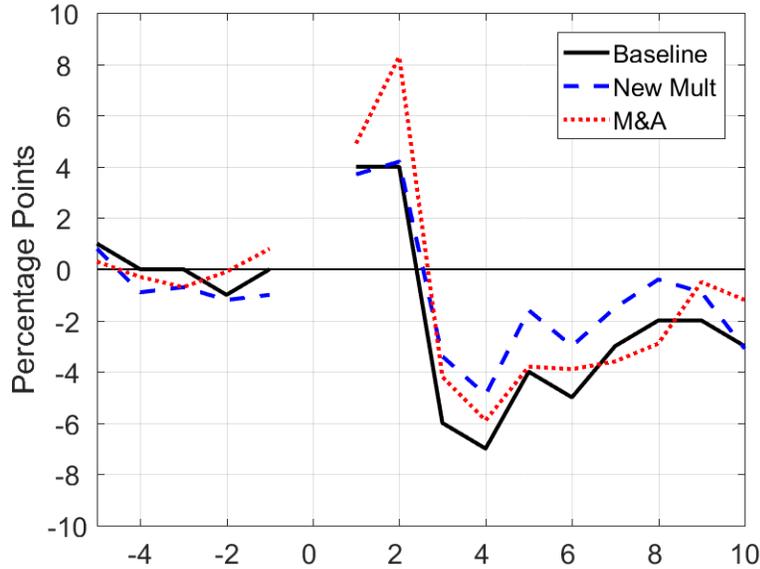
To evaluate these cases separately, we re-run specification (4) for a sample of establishments that retain the same firm identifier in periods -1 and +1 (the “New Mult” sample) and separately for a sample of establishments that change firm identifiers between periods -1 and +1 (the “M&A” sample). Since these sample restrictions apply to both the control group and the treated group (new multinational plants), we will not conflate standard M&A effects with those we attribute to foreign sourcing. The results are shown in Figure B1. Although, as expected, there are differences in the magnitudes of employment effects based on the type of plant transition into a firm with foreign sourcing, the overall message of large and persistent employment declines remains.

B.1.2 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 Panel C of Figure 3 fill in the post-death values for a given establishment with the actual imports of the firm associated with that establishment.⁴⁹ 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 than otherwise. On the other hand, if establishment deaths are associated

⁴⁹If the entire firm disappears, we then record zeros in that period and all future periods.

Figure B1: Employment Growth Differential of Multinational Transitions: Transition Types



Notes: The data are from the LFTTD, DCA, and UBP as explained in text. This figure plots the annual employment growth rate differential of establishments that transition into a multinational firm in year ($t = 0$), relative to a control group with similar firm age, establishment size, and industry (in year $t = -1$). The sample is split between those establishments that also experience switches in firm identifiers (M&A Sample) and those that do not (new Mult Sample). This split also applies to the control group, where the control group is not part of a multinational firm in year $t = 0$. See equation (4) and text for details. Standard errors are suppressed, but the post-transition results are significant at traditional levels.

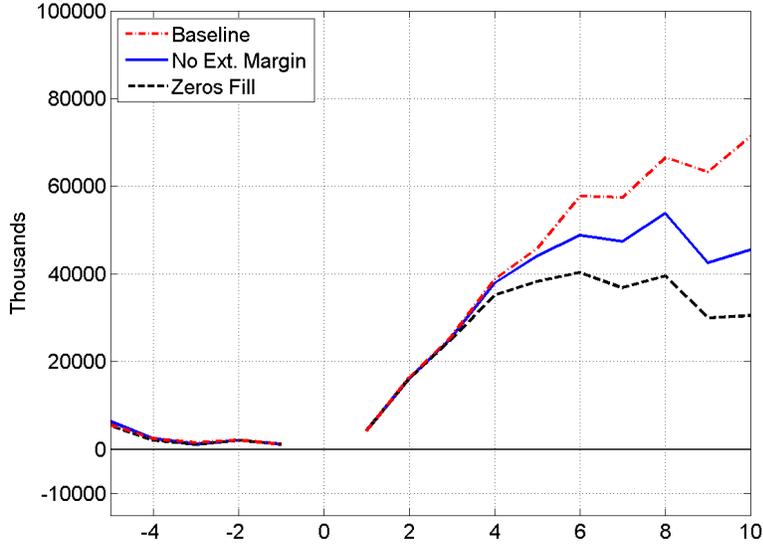
with broad firm decline, then this differential import measure would be smaller relative to the benchmark.

An alternative approach is 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 Panel C of 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 B2 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.

Figure B2: Importing Differentials of Multinational Transitions, Balanced Panel



Notes: The data are from the LFTTD, DCA, and 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, see equation (4). *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.

B.1.3 Other Trade Effects Following Multinational Transitions

We estimate equation (4) 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 results that show the differential imports of final goods of new multinationals are shown in Figure B3. We also find strong growth in export volumes in the years following a multinational transition, both to foreign affiliates and unaffiliated parties. The increase in exports (shown in Figure B4), is consistent with the interpretation that transitions occur after positive idiosyncratic shocks; alternatively these results may indicate that the multinational reorganization allows the firm to focus on the highest value products or transaction. A third interpretation involves greater interaction with other markets that naturally leads to sales opportunities abroad (both internally and externally).

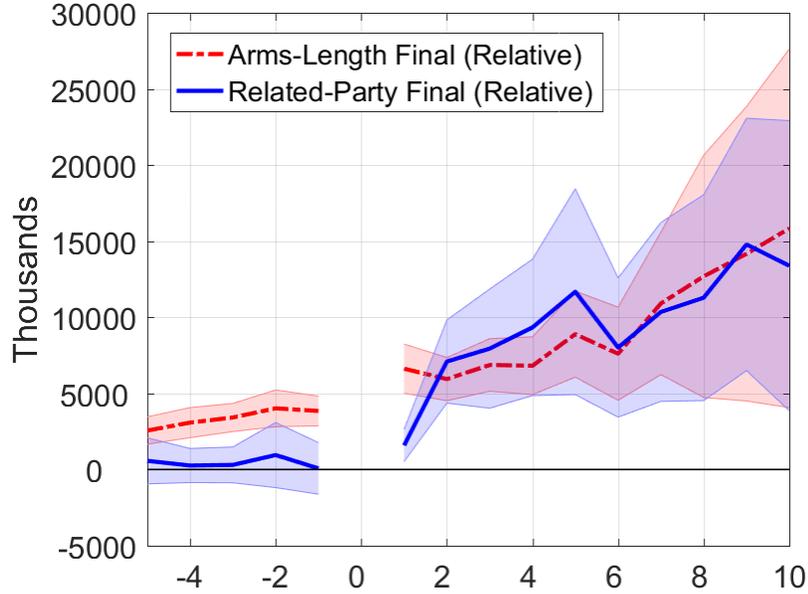
Finally, Figure B5 plots firm-level imports from China corresponding to establishments that transitions into part of a multinational firm. As would be expected, the imports from China are a large contributor to the overall increase in imports following a firm transition.

B.2 Quantifying Job Loss: Back-of-the-Envelope Calculations

B.2.1 Job Loss from Multinational Transitions

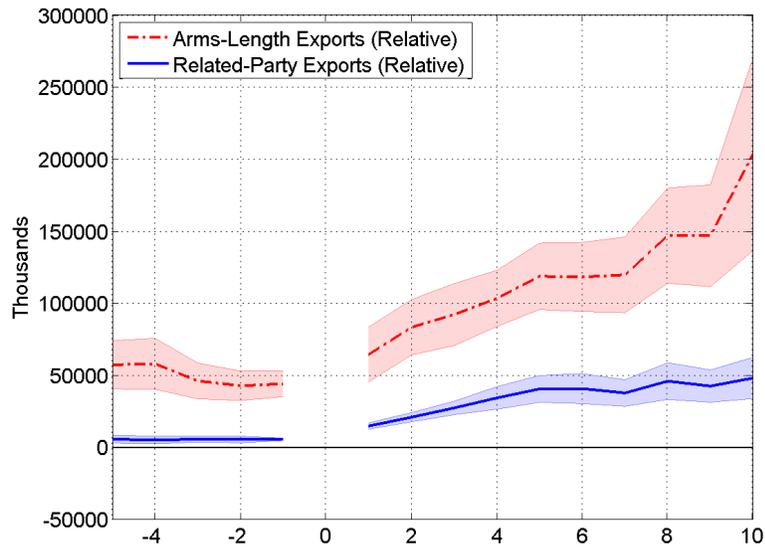
This section describes how we convert the estimates on relative employment growth rates of new multinational plants into a measure of the aggregate net gains of employment. The coefficients from Figure 3 represent relative employment effects, expressed in percentage points, of a transitioning plant. These effects represent averages that span the entire period (1993-2011)

Figure B3: Final Goods Importing Differentials of Multinational Transitions



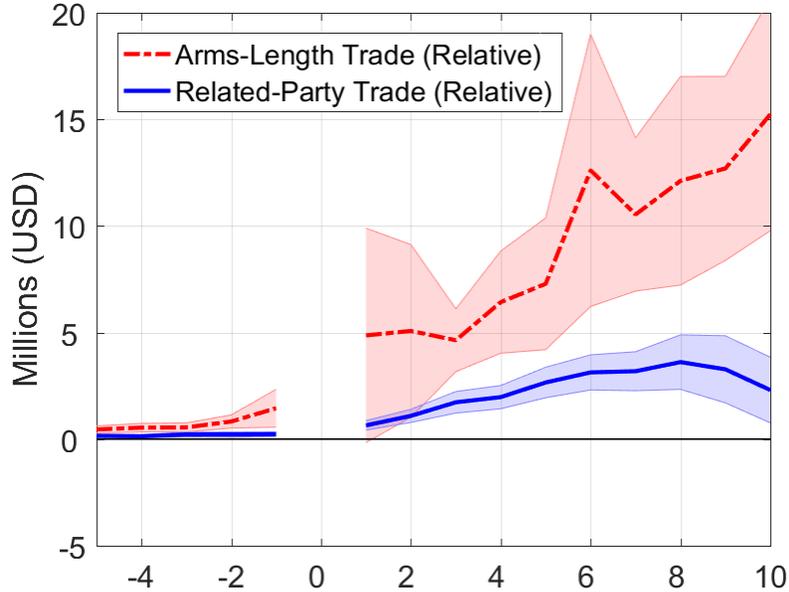
Notes: The data are from the LFTTD, DCA, and UBP as explained in text. This figure reports the related-party and arms-length final goods imports of the parent firm of an establishment that transitions into a multinational firm in year ($t = 0$), relative to a control group with similar firm age, establishment size, and industry (in year $t = -1$). See equation (4), modified to have firm-level imports as the dependent variable. The shaded area corresponds to a 95 percent confidence interval.

Figure B4: Exporting Differentials of Multinational Transitions



Notes: The data are from the LFTTD, DCA, and 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 a multinational firm in year ($t = 0$), relative to a control group with similar firm age, establishment size, and industry (in year $t = -1$). See equation (4), modified to have firm-level imports as the dependent variable. The shaded area corresponds to a 95 percent confidence interval.

Figure B5: Importing Differentials from China



Notes: The data are from the LFTTD, DCA, and UBP as explained in text. This figure reports the related-party and arms-length imports from China of the parent firm of an establishment that transitions into a multinational firm in year ($t = 0$), relative to a control group with similar firm age, establishment size, and industry (in year $t = -1$). See equation (4), modified to have firm-level imports (from China) as the dependent variable. The shaded area corresponds to a 95 percent confidence interval.

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), 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 (4) 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 resulting from multinational transitions.

Formally, we compute the job loss as

$$\sum_{t=1994}^{2010} T_t E_t \sum_{i=1}^{\min\{10, 2010-t\}} \delta_i \prod_{j=1}^{i-1} (1 + \delta_j) \quad (\text{B1})$$

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 (4). Table B2 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.

Table B2: 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

Notes: Estimates based on Table 1, Table 2, and Figure 3.

B.2.2 Job Loss from all Multinational Activity: Total

A similar exercise can be done using the coefficient estimates from Table 4. 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 B3 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.3 Other Results

B.3.1 Impact of Multinational Transitions on Non-Manufacturing Plants

To assess whether the non-manufacturing establishments in new multinationals also experience relatively negative employment growth rates, we re-estimate equation (4) for this group of establishments. We construct control groups from all non-manufacturing establishments in a manner similar to Section 2.2.2. Unlike in the manufacturing sector, however, the non-manufacturing establishments in our data encompass a very broad swathe of the economy and come from sectors as diverse as wholesale or technology. Observing a common pattern after transition for this group, relative to controls, is therefore very unlikely. The results from

Table B3: Aggregate Job Loss from All Multinational Plants

	Average Size	# of Mult Establishments	Extensive, Weighted		Extensive, Unweighted	
			Avg. Differential Employment per Establishment ¹	Total per year	Avg. Differential Employment per Establishment ²	Total per year
1994	310	17,119	-8.0	137,194	-9.7	166,846
1995	311	16,269	-8.0	130,690	-9.7	158,937
1996	309	16,316	-8.0	130,034	-9.7	158,138
1997	306	16,365	-7.9	129,436	-9.6	157,412
1998	313	15,950	-8.1	128,900	-9.8	156,760
1999	312	16,084	-8.0	129,384	-9.8	157,348
2000	299	16,466	-7.7	127,143	-9.4	154,622
2001	297	15,886	-7.7	121,873	-9.3	148,214
2002	296	15,386	-7.6	117,638	-9.3	143,063
2003	279	14,930	-7.2	107,588	-8.7	130,842
2004	275	14,823	-7.1	105,249	-8.6	127,997
2005	270	14,692	-7.0	102,541	-8.5	124,704
2006	270	14,534	-7.0	101,155	-8.4	123,018
2007	269	14,482	-6.9	100,535	-8.4	122,264
2008	261	14,641	-6.7	98,821	-8.2	120,180
2009	254	14,456	-6.5	94,619	-7.9	115,169
2010	235	13,865	-6.1	83,938	-7.3	102,079
2011	222	13,562	-5.7	77,768	-7.0	94,576
Total				2,024,504	2,462,070	
Share of 5.65 million lost				0.36	0.44	

Notes: Estimates based on Table 1, Table 2, and Table 4.

¹This column applies the coefficient estimates from the intensive/extensive and weighted estimates from Table 4. ²This column applies the coefficient estimates from the intensive/extensive and unweighted estimates from Table 4.

this exercise are in Figure B6. New non-manufacturing multinational establishments do not experience job losses relative to the control group. However, the results are noisy, and do not allow us to identify whether the scale effect for this group of establishments as a whole outweighs the reallocation effect.

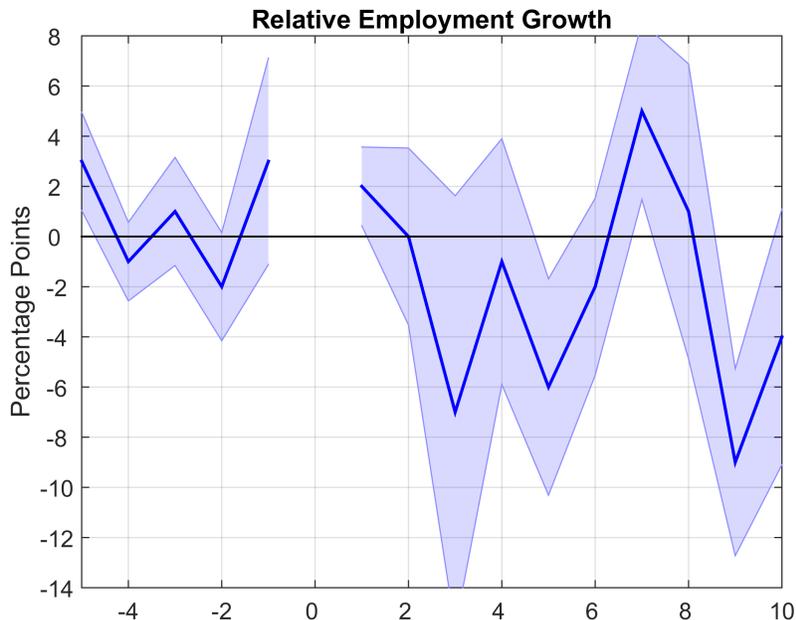
B.3.2 Location Complementarities

Why is direct 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 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:

$$\log IM_{ijkt}^{AL} = \alpha_{ijt} + \gamma_{kt} + \beta \log IM_{ijkt}^{RP} + \epsilon_{ijkt}. \quad (\text{B2})$$

Here i is the firm, j is the partner country, k is the product code, and t is time. Hence,

Figure B6: Transitions: Non-manufacturing



Notes: The data are from the LFTTD, DCA, and UBP as explained in text. This figure reports the employment growth rates of non-manufacturing plants that transition into a multinational firm in year ($t = 0$), relative to a control group with similar firm age, establishment size, and industry (in year $t = -1$). The shaded area corresponds to a 95 percent confidence interval.

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 B4), 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.

B.4 Regression Evidence: Robustness

Table B5 presents results from running the specification in equation (3) 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).

B.5 IV

Table B6 presents results from a specification where we regress firm-level employment growth rates on import growth, instrumenting for the import growth with changes in tariffs and exchange rates. Specifically, we instrument for import growth $\Delta \ln IM_{i,t}$ of firm i in year t with

Table B4: Inter-Firm and Intra-Firm Sourcing

	Country Level		Industry & Country Level	
	RP Indicator	Log RP Imports	RP Indicator	Log RP Imports
Coef.	1.84*** (0.006)	0.39*** (0.002)	1.765*** (0.004)	0.49*** (0.001)
Fixed Effects				
Firm \times time	Yes	Yes	No	No
Country \times Time	Yes	Yes	No	No
Industry \times Time	No	No	Yes	Yes
Firm \times Country \times 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

Notes: The data are from the LFTTD. This table reports the results from equation (B2). The dependent variable is the log of a firm’s inter-firm imports from a particular country and industry. Standard errors are reported in parentheses. *** denotes significance at the 1 percent level.

$\sum_{j,h} w_{i,j,h,t-5} \Delta \tau_{i,j,h,t}$ and $\sum_j w_{i,j,t-5} \Delta \ln Q_{j,t}$, where $\tau_{i,j,h,t}$ is the tariff paid by firm i to import product h from country j , and $Q_{j,t}$ is the bilateral nominal exchange rate between the U.S. and country j . We use firm-specific weights constructed from the share of the firm’s imports of a product from a particular location in period $t - 5$, $w_{i,j,h,t-5} = \frac{IM_{i,j,h,t-5}}{IM_{i,t-5}}$ and $w_{i,j,t-5} = \sum_h \frac{IM_{i,j,h,t-5}}{IM_{i,t-5}}$. Unfortunately, with the exception of this specification, we are largely unable to construct instruments with predictive power for firm-level imports. We have found that other commonly used instruments such as the “World Export Supply” measure, which captures supply shocks in a partner country (see Acemoglu et al. (2016) or Hummels et al. (2014)), transport costs and GDP growth rates are not relevant in our data.

We include firm fixed effects in our IV estimation. The OLS specification demonstrates a strong positive correlation between employment growth and import growth. With the instruments, the correlation is negative and insignificant.

B.6 Patterns in Firm Sourcing

Table B7 illustrates the fraction of firms in each of the five most common broad sourcing categories together with the fraction of input imports by firms in each sourcing category. We also include a residual “Other” category that encompasses sourcing strategies not explicitly in the table. The broad categories are combinations of the two home sourcing location/modes HI and HO as well as sourcing from developed countries inter or intrafirm (NO and NI) or developing countries inter or intrafirm (SO and SI). Multinationals are firms that have any NI or SI sourcing, which is the sum of the third, fourth and fifth columns of the table.

Table B5: Relative Employment Growth: Subsamples

		Establishment Level			
		Intensive		Extensive and Intensive	
		Unweighted	Employment Weighted	Unweighted	Employment Weighted
1993 - 2000	$\hat{\beta}$	0.02*** (0.002)	0.01*** (0.001)	-0.04*** (0.004)	-0.03*** (0.003)
	Clusters	8179	8179	8606	7081
2001 - 2011	$\hat{\beta}$	0.02*** (0.001)	0.005*** (0.002)	-0.03*** (0.003)	-0.02*** (0.003)
	Clusters	8437	8437	8922	8525
		Firm Level			
		Intensive		Extensive and Intensive	
		Unweighted	Employment Weighted	Unweighted	Employment Weighted
1993 - 2000	$\hat{\beta}$	-0.01*** (0.004)	-0.03*** (0.007)	-0.06*** (0.01)	-0.04*** (0.01)
	Clusters	3481	3481	3931	3931
2001 - 2011	$\hat{\beta}$	-0.01*** (0.003)	-0.01*** (0.006)	-0.01*** (0.007)	-0.01*** (0.009)
	Clusters	4547	4547	5187	5187

Notes: The data are from the LBD, DCA, and UBP. The table reports pooled regression results, where the sample is split into subsamples from 1993-2000 and 2001-2011. Standard errors are reported in parentheses. *** reports significance at the 1 percent level.

Table B6: IV Specification for Firm-Level Employment Growth

	$\Delta \ln \text{Emp}$	
	OLS	IV
$\Delta \ln \text{IM}$	0.040*** (0.0048)	-0.021 (0.1519)
Fixed Effects		
Firm	Yes	Yes
Year	Yes	Yes
First stage F		155.73
R ²	0.348	
Observations	41,800	41,800

Notes: The data are from the LBD and the LFTTD. This table reports the results from an OLS regression relating changes in firm employment growth to changes in firm import growth, and the same regression using changes in lagged tariff and nominal exchange rate movements, with country-weights lagged by five years, as instruments in a first stage. Standard errors are reported in parentheses. *** denotes significance at the 1 percent level.

Table B7: 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%

Notes: The data are from the 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 all other possible sourcing strategies.

C Appendix: Model Extensions and Proofs

C.1 Model with non-constant returns to scale

In this Appendix we suppose instead that firms produce with the production function

$$x_i = \left(\int_0^1 x_i(v)^{\frac{\rho-1}{\rho}} dv \right)^{\frac{\rho}{\rho-1}(1-\lambda)}.$$

Here, the parameter $\lambda > -1$ captures the degree of decreasing returns to scale. For $\lambda = 0$ returns to scale are constant, and we return to the model in the paper. For $\lambda > 0$ returns to scale are decreasing and for $\lambda < 0$ returns to scale are increasing. With this production function, but leaving the remainder of the model unchanged, firm revenues are

$$R_i \propto (\delta_i E(P_X)^{\sigma-1})^{\frac{1}{1+(\sigma-1)\lambda}} \Phi_i^{\frac{\sigma-1}{\theta} \cdot \frac{1-\lambda}{1+(\sigma-1)\lambda}},$$

where Φ_i is defined as in the paper. The scale elasticity therefore is

$$\frac{\partial \ln R}{\partial \ln \Phi} = \frac{\sigma-1}{\theta} \cdot \frac{1-\lambda}{1+(\sigma-1)\lambda}.$$

As noted in the text, decreasing returns to scale implies that the scale elasticity is smaller than $\frac{\sigma-1}{\theta}$. Conversely, for increasing returns to scale that the scale elasticity is greater than $\frac{\sigma-1}{\theta}$.

The domestic employment response in partial equilibrium continues to be determined by the scale and reallocation effect,

$$l_{i,j} \propto \frac{(\delta_i E(P_X)^{\sigma-1})^{\frac{1}{1+(\sigma-1)\lambda}}}{w_j} \cdot s_{i,j} \cdot \Phi_i^{\frac{\sigma-1}{\theta} \cdot \frac{1-\lambda}{1+(\sigma-1)\lambda}}.$$

That is, the scale elasticity fully characterizes the domestic employment response in partial equilibrium.

The estimating equation in the model takes the form

$$\ln R_i = \alpha + \frac{1}{n_i} \sum_{j \in J_i} \alpha_j - \frac{\sigma-1}{\theta} \frac{1-\lambda}{1+(\sigma-1)\lambda} \frac{1}{n_i} \sum_{j \in J_i} \ln s_{i,j} + u_i,$$

where

$$u_i = (\sigma-1) \frac{1-\lambda}{1+(\sigma-1)\lambda} \frac{1}{n_i} \sum_{j \in J_i} \ln \zeta_{i,j} + \frac{1}{1+(\sigma-1)\lambda} \ln(\delta_i).$$

It is easy to see that our bounding strategy also applies to this modified model with non-constant returns to scale.

We finally note that introducing love of variety for intermediate inputs, as in Benassy (1998), generates predictions comparable to increasing returns to scale. Again, the scale elasticity differs from $\frac{\sigma-1}{\theta}$.

C.2 Proof of Lemma 3.1

Lemma. *To a first order, $\overline{\ln s_i} = \alpha - \ln n_i$, for some constant α .*

Proof. We begin with the linear approximation of

$$\ln \Phi_i = \ln \sum_{k \in J_i} T_k \zeta_{i,k}^\theta (\tau_k w_k)^{-\theta}$$

in $T_j (\tau_j w_j)^{-\theta}$ around $\bar{T} (\bar{\tau} \bar{w})^{-\theta} = \frac{1}{n} \sum_k T_k (\tau_k w_k)^{-\theta}$ and $\ln \zeta_{i,j}$ around $\ln \bar{\zeta}$. This yields

$$\ln \Phi_i = \ln \bar{\Phi}_i + \sum_{j \in J_i} \frac{\bar{T} (\bar{\zeta})^\theta (\bar{\tau} \bar{w})^{-\theta}}{\sum_{k \in J_i} \bar{T} (\bar{\zeta})^\theta (\bar{\tau} \bar{w})^{-\theta}} \left\{ \theta [\ln \zeta_{i,j} - \ln \bar{\zeta}] + \left(\ln \left(T_j (\tau_j w_j)^{-\theta} \right) - \ln \left(\bar{T} (\bar{\tau} \bar{w})^{-\theta} \right) \right) \right\}.$$

Note that

$$\frac{\bar{T} \bar{\zeta}^\theta (\bar{\tau} \bar{w})^{-\theta}}{\sum_{k \in J_i} \bar{T} (\bar{\zeta})^\theta (\bar{\tau} \bar{w})^{-\theta}} = \frac{1}{n_i}$$

so that

$$\ln \Phi_i = \ln \bar{\Phi}_i + \sum_{j \in J_i} \frac{1}{n_i} \left\{ \theta [\ln \zeta_{i,j} - \ln \bar{\zeta}] + \left(\ln \left(T_j (\tau_j w_j)^{-\theta} \right) - \ln \left(\bar{T} (\bar{\tau} \bar{w})^{-\theta} \right) \right) \right\}.$$

Further

$$\ln \bar{\Phi}_i = \ln \sum_{k \in J_i} \bar{T} (\bar{\zeta})^\theta (\bar{\tau} \bar{w})^{-\theta} = \ln n_i + \ln \bar{T} \bar{\zeta}^\theta (\bar{\tau} \bar{w})^{-\theta}$$

and hence

$$\ln \Phi_i = \ln n_i + \ln \bar{T} \bar{\zeta}^\theta (\bar{\tau} \bar{w})^{-\theta} + \sum_{j \in J_i} \frac{1}{n_i} \left\{ \theta [\ln \zeta_{i,j} - \ln \bar{\zeta}] + \left(\ln \left(T_j (\tau_j w_j)^{-\theta} \right) - \ln \left(\bar{T} (\bar{\tau} \bar{w})^{-\theta} \right) \right) \right\}.$$

Next write the log share as

$$\begin{aligned} \ln s_{i,j} &= \ln \frac{T_j \zeta_{i,j}^\theta (\tau_j w_j)^{-\theta}}{\sum_{k \in J_i} T_k \zeta_{i,k}^\theta (\tau_k w_k)^{-\theta}} \\ &= \theta \ln \zeta_{i,j} + \ln T_j (\tau_j w_j)^{-\theta} - \ln \Phi_i \end{aligned}$$

Plugging in the first order approximation of $\ln \Phi_i$ and again summarizing constants in α implies

$$\ln s_{i,j} = \alpha + \theta \ln \zeta_{i,j} + \ln T_j (\tau_j w_j)^{-\theta} - \ln n_i - \theta \frac{1}{n_i} \sum_{k \in J_i} \ln \zeta_{i,k} - \frac{1}{n_i} \sum_{k \in J_i} \ln \left(T_k (\tau_k w_k)^{-\theta} \right)$$

Finally, take the simple average over the firm's sourcing strategy to obtain

$$\begin{aligned} \frac{1}{n_i} \sum_{j \in J_i} \ln s_{i,j} &= \alpha + \theta \frac{1}{n_i} \sum_{j \in J_i} \ln \zeta_{i,j} + \frac{1}{n_i} \sum_{j \in J_i} \ln T_j (\tau_j w_j)^{-\theta} - \ln n_i \\ &\quad - \theta \frac{1}{n_i} \sum_{k \in J_i} \ln \zeta_{i,k} - \frac{1}{n_i} \sum_{k \in J_i} \ln T_k (\tau_k w_k)^{-\theta} \\ &= \alpha - \ln n_i. \end{aligned}$$

□

D Appendix: General Equilibrium

D.1 A general equilibrium extension

For this general equilibrium version of the model, we assume a three country world, with the countries labeled Home (H), North (N) and South (S). We distinguish between the North and South in this model for two reasons. First, we observe increased U.S.-based production by firms from predominantly developed countries (Tables 1 and 2). Second, while U.S. imports from developing countries grew rapidly over our sample period, imports from developed countries also increased (Table 3). By distinguishing the North from the South, our counterfactuals can capture both of these facts. For simplicity, we drop the distinction between intra and inter-firm sourcing. We further assume that 1) $\zeta_{i,j} = \zeta_i$, that is, firms' productivity does not vary by location/mode, and 2) fixed costs of foreign sourcing are common across firms. In our baseline version of the model, the Home and North countries are symmetric and produce differentiated final manufactured goods that are freely traded between all three countries. We model the North as symmetric to the Home country as we lack data to calibrate the foreign sourcing of other developed countries. The South does not produce and export final manufacturing goods, only intermediates. To conserve space we only present the problems of the Home and South country. In addition to the manufacturing sector in H and N , there is a nonmanufacturing sector in each country, which produces a differentiated and freely traded non-manufacturing good.

Notation In this section, a double subscript indicates first the destination and second the source. For instance, X_{HN} will denote Home's consumption of the North manufactured bundle.

Households

The representative household in the Home country derives utility from the consumption of the manufacturing aggregate Y and the nonmanufacturing aggregate Z . It supplies L_H units of labor inelastically. The household maximizes utility $(Z_H)^\beta (Y_H)^{1-\beta}$ subject to the budget constraint $w_H L_H = P_H^Y Y_H + P_H^Z Z_H$. As a result, the Home consumer spends $P_H^Y Y_H = (1 - \beta) w_H L_H$ on the manufacturing aggregate and $P_H^Z Z_H = \beta w_H L_H$ on the nonmanufacturing aggregate.

The composite manufacturing good Y_H is an aggregate of the Home manufacturing good X_{HH} and the North manufacturing good X_{HN} ,

$$Y_H = \left((a_{HH})^{\frac{1}{\varepsilon}} (X_{HH})^{\frac{\varepsilon-1}{\varepsilon}} + (a_{HN})^{\frac{1}{\varepsilon}} (X_{HN})^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

The demand functions from the Home and North manufacturing goods are given by

$$X_{Hj} = a_{Hj} Y_H \left(\frac{P_j^X}{P_H^Y} \right)^{-\varepsilon}, \quad j = H, N,$$

and the manufacturing price index is

$$P_H^Y = \left(a_{HH} (P_H^X)^{1-\varepsilon} + a_{HN} (P_N^X)^{1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}}.$$

For simplicity, we assume that final goods are freely traded, so these goods' prices are identical in each destination.

The composite nonmanufacturing good Z is an aggregate of differentiated manufacturing goods Q from Home, the North, and the South,

$$Z_H = \left((b_{HH})^{\frac{1}{\varepsilon}} (Q_{HH})^{\frac{\varepsilon-1}{\varepsilon}} + (b_{HN})^{\frac{1}{\varepsilon}} (Q_{HN})^{\frac{\varepsilon-1}{\varepsilon}} + (b_{HS})^{\frac{1}{\varepsilon}} (Q_{HS})^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

The associated demand functions are

$$Q_{Hj} = b_{Hj} Z_H \left(\frac{P_j^Q}{P_H^Z} \right)^{-\varepsilon}, \quad j = H, N,$$

and the nonmanufacturing price index is

$$P_H^Z = \left(b_{HH} (P_H^Q)^{1-\varepsilon} + b_{HN} (P_N^Q)^{1-\varepsilon} + b_{HS} (P_S^Q)^{1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}}.$$

Nonmanufacturing sector

A representative Home firm in the nonmanufacturing sector operates a linear technology $Q_H = A_H L_H^Q$, where A_H is Home's total factor productivity (TFP) in the nonmanufacturing sector, and L_H^Q denotes labor in this sector. Profit maximization in competitive markets implies that $A_H P_H^Q = w_H$ as long as Q_H is strictly positive and finite.

Manufacturing sector

A Home aggregating firm, which sells to Home, North and South now replaces the consumer as described in Section 3. The firm produces a CES aggregate

$$X_H = \left(\int_{\omega \in \Omega_H} [x(\omega)]^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}}$$

from intermediates $x(\omega)$. Note that we abstract from various sources of firm heterogeneity in this general equilibrium version of the model, including the firm-specific demand shock. The demand for variety ω is

$$x(\omega) = X_H \left(\frac{p(\omega)}{P_H^X} \right)^{-\sigma},$$

and the price index

$$P_H^X = \left(\int_{\omega \in \Omega_H} (p(\omega))^{1-\sigma} d\omega \right)^{\frac{1}{1-\sigma}}.$$

Firms in the X sector

Home firms in the X sector set up supply chains and produce intermediates $x(\omega)$ as described in Section 3. Recall that we assume here that $\zeta_{i,j} = \zeta_i$ and common fixed costs across firms. 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 productivity ζ . In equilibrium, the number of entrants M_H is determined by a zero expected profit condition:

$$\mathbb{E}_{\zeta} [\max \{ \pi(\zeta) - w_H f_H, 0 \}] - w_H f_E = 0,$$

where f_H is the fixed cost of entering into production at Home, and

$$\pi(\zeta) = \frac{1}{\sigma} \left(\frac{\sigma - 1}{\sigma} \right)^{\sigma-1} X_H (P_H^X)^\sigma \left((\gamma)^{\frac{1}{\theta}} [\Phi(\zeta)]^{-\frac{1}{\theta}} \right)^{1-\sigma} - w_H \sum_k 1\{J(\zeta) = J_k\} f_k.$$

Second, after learning their types, entrants must pay the additional fixed cost f_H to set up production in the Home country. Only firms with sufficiently high types find it profitable to do so. The lowest type that enters is ζ_H^{LB} . The equation determining entry into Home production is

$$\pi(\zeta_H^{LB}) - w_H f_H = 0.$$

Finally, those firms that produce in the Home country set up their supply chain as discussed in Section 3. As before, all fixed costs are measured in units of labor.

Labor market clearing

Labor market clearing in the Home country requires that

$$\begin{aligned} L_H = & L_H^Q + M_H \left[\int_{\zeta_H^{LB}}^{\infty} l_{HH}(\zeta) dG_\zeta(\zeta) + f_E + f_H (1 - G_\zeta(\zeta_H^{LB})) \right. \\ & \left. + \int_{\zeta_H^{LB}}^{\infty} \sum_k 1\{J(\zeta) = J_k\} f_k dG_\zeta(\zeta) \right] + M_N \int_{\zeta_N^{LB}}^{\infty} l_{NH}(\zeta) dG_\zeta(\zeta) \end{aligned} \quad (D1)$$

Labor demand on the right hand side consists of demand L_H^Q from the nonmanufacturing sector, demand from Home firms ($l_{HH}(\zeta)$) in the manufacturing sector (the first integral), the labor demand stemming from the various fixed costs and the demand from North firms ($l_{NH}(\zeta)$) in the manufacturing sector sourcing intermediates from Home (the last integral).

South

The representative household in the South has preferences analogous to Home and North. The production of the non-manufacturing good is also analogous to the Home and North. However, there is no final manufacturing sector in the South, or only to the extent that Home and North produce intermediates in the South. As a result, the labor market clearing condition requires that

$$L_S = L_S^Q + M_H \int_{\zeta_H^{LB}}^{\infty} l_{HS}(\zeta) dG_\zeta(\zeta) + M_N \int_{\zeta_N^{LB}}^{\infty} l_{NS}(\zeta) dG_\zeta(\zeta), \quad (D2)$$

where $l_{HS}(\zeta)$ and $l_{NS}(\zeta)$ denote Home's and North's labor demand in the South.

Output market clearing

The output market clears for all nonmanufacturing goods

$$Q_j = Q_{Hj} + Q_{Nj} + Q_{Sj}, \quad j \in \{H, N, S\},$$

and for all manufacturing goods

$$X_j = X_{Hj} + X_{Nj} + X_{Sj}, \quad j \in \{H, N\}.$$

D.2 Calibration

While the model does not restrict firms' sourcing patterns from Home, North, and South, we find that only three of these are prevalent in the data.⁵⁰ We therefore restrict the model to these three equilibrium sourcing strategies, so that $\tilde{\mathcal{J}}_H = \{(H), (H,N), (H,N,S)\}$. This restriction together with nonstochastic fixed costs conveniently implies a complete ordering of sourcing strategies which simplifies the numerical solution: Most manufacturing firms do not source from abroad. They only pay fixed costs f_E to learn their type and f_H to enter into domestic production. The somewhat more productive firms additionally pay a fixed cost to also source from the North. Finally, only the most productive firms source from the North and the South in addition to their production in the Home country. To do so, they pay an additional fixed cost.⁵¹

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 in a baseline year. For the baseline calibration, we choose parameters for both the Home and the North to match the U.S.

The productivity parameters A_H and A_S are chosen to match skill-adjusted wages for the U.S. 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 $\alpha_\zeta = 4.5$ similar to Melitz and Redding (2015). The demand elasticity σ is set to 4 and the dispersion parameter θ to 5.97. These values imply that $(\sigma - 1)/\theta$ is approximately 0.5, the tightest upper bound we estimate. In line with the literature, we choose ϵ to be 4. We also set $\tau_{jj} = 1$ and $\tau_{jk} = 1.15$, $j, k \in \{H, N, S\}$, $j \neq k$. Although these parameters are not important for any of the model's predictions, we note that ρ is set to 1.5. Table D1 summarizes the values of the preset parameters.

The remaining parameters of the model are chosen to match key features of our micro-data in 1997.⁵² These parameters are T_j , $j \in \{H, N, S\}$, the fixed cost parameters f_J , $J \subset \{(H,N), (H,N,S)\}$, f_E , and f_H . To calibrate the taste parameters in the non-manufacturing aggregators $b_{HN}, b_{NH}, b_{SH}, b_{SN}, b_{HS}, b_{NS}$, we compute services import shares in final use for 1997 from the Johnson and Noguera (2016) Trade in Value Added database. The targets and the fit of the model in equilibrium are summarized in Table D2.

D.3 Quantitative exercises

We consider two counterfactuals. In the first exercise, we compute the employment and welfare changes relative to 1997 when we allow only the fixed cost parameters f_J , $J \subset \{(H,N), (H,N,S)\}$ to change to match 2007 trade flows. We stress that our framework cannot identify which shocks

⁵⁰In fact, similar to Antràs, Fort, and Tintelnot (2017), there are patterns 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. Second, if a firm sources from the South, it almost always also sources from the North. Appendix Table B7 shows the fraction of firms in the data that source according to each of these strategies.

⁵¹Adding further sourcing strategies, e.g. (H,S) would break the complete ordering and complicate the solution somewhat. Empirically, this sourcing strategy is unimportant. Arkolakis and Eckert (2017) demonstrate how to solve such models.

⁵²We choose 1997 as the baseline year in much of the paper as all of the data required for calibration and estimation is available in this year.

Table D1: Calibration Stage 1

Parameter	Value	Note
σ	4	Demand elasticity
θ	5.97	Frechet shape parameter
b_ϕ	1	Lower bound of the Pareto distribution
α_ϕ	4.5	Pareto shape parameter
β	0.75	Non-manufacturing expenditure share
A_H, A_N	14.32	Skill-adjusted wages in Home, from the ILO
A_S	1.02	Average skill-adjusted wages in South from the ILO
L_H, L_N	0.301	Skill-adjusted labor force in Home, from the ILO
L_S	2.35	Total skill-adjusted labor force in South from the ILO
τ	1.15	Transport costs of manufacturing intermediates
ρ	1.5	Elasticity of substitution of tasks
ϵ	4	Elasticity of substitution of manufactured bundles from H and N
a_{HN}, a_{NH}	0.0331	Weights in manufactured bundle (Calibrated to match share of imports in total manufacturing output)
a_{HH}, a_{NN}	1	Weights on own good in manufactured bundle (normalization)
a_{SH}, a_{SN}	1	South' weights on Home and North goods (normalization)
b_{HH}, b_{NN}, b_{SS}	1	Weights on own good in non-manufactured bundle (normalization)

Table D2: Quantitative Exercises: Model Fit and Parameter Changes

Moments	1997		2007	
	Model	Targets	CF 1 (only f_j changes)	Targets
N imports/Manufacturing sector sales	0.0593	0.0550	0.0760	0.076
S imports/Manufacturing sector sales	0.0313	0.0320	0.0740	0.074
Manufacturing employment share	0.1856	0.1686	0.1766	
Fraction of firms with $J = \{H, N\}$	0.0149	0.1250	0.0426	
Fraction of firms with $J = \{H, N, S\}$	0.0000	0.0570	0.0025	
Non-manufacturing shares:				
H import share from N	0.0099	0.0079	0.0100	
H import share from S	0.0051	0.0036	0.0021	
S import share from H, N	0.0722	0.0732	0.1452	
Not Targeted:				
Fraction of manufacturing employment in offshoring firms	0.2726	0.3070	0.3713	
Parameters				
% change in f_{HN}			-52.20%	
% change in f_{NH}			-52.20%	
% change in f_{HNS}			-99.57%	

Notes: This table summarizes the fit of the model to calibration targets in 1997 and 2007. The second panel displays the changes in parameters in the counterfactual.

(technology, fixed costs, variable costs) occurred between 1997 and 2007 to generate observed changes in foreign sourcing – all these shocks have qualitatively the same implications for firm sourcing shares in our model.⁵³ The purpose of this exercise is to illustrate the role the observed changes in the foreign sourcing of U.S. multinationals played in the aggregate manufacturing decline, when general equilibrium effects are present. We simply choose fixed cost declines as one method of generating the observed aggregate changes in foreign sourcing. The benefit of this approach is that we do not need to recalibrate the entire model to match the data in 2007; rather we can be explicit about relating the changes in fixed costs to changes in employment and welfare.

Table D2 illustrates the fit of our model to our calibration targets for 1997 and 2007, and summarizes the changes in fixed costs necessary to match aggregate input imports in both periods. To match the observed trade patterns in 2007 the fixed cost parameters $f_j, j \in \mathcal{J}$ uniformly decrease. f_{HNS} decreases the most, reflecting the fact that imports from the South grew rapidly over the period 1997-2007.

The results of this exercise are shown in Table D3. Targeting 2007 trade patterns results in Home manufacturing employment falling by 4.9%, which accounts for around one fifth of the observed decline between 1997 and 2007.⁵⁴ In this exercise Home manufacturing employment declines as more firms choose to source intermediates from abroad. However, as these firms lower their unit costs, the Home and North manufacturing price indexes fall and demand for the manufactured goods increase. This raises manufacturing employment. On net the effect is similar to the previous “naive” counterfactual in Section 4.1. Home welfare increases modestly.

Our second counterfactual exercise implements an asymmetric policy where Home firms are not permitted to offshore production, while North firms continue to do so. In practice, we maintain all parameters at their 1997 levels and assume that fixed costs faced by Home firms increase to prevent Home offshoring as an equilibrium outcome. The results of this exercise are in Table D3. Home firms face higher unit costs relative to the baseline model. This reduces their scale and increases the price of the Home manufactured bundle relative to the North bundle. Despite higher prices consumers in all countries demand the Home manufactured good, and the existing manufacturing firms now reallocate all their intermediate good production to the Home country. The net effect is a small increase in Home manufacturing employment of 2.6%. Due to the increase in the manufacturing price index, Home welfare falls slightly. Fewer firms enter manufacturing in Home.

The final row of Table D3 considers the same policy change when the elasticity of substitution between the Home and North good equals 6. This results in a smaller increase in the Home manufacturing price index, as the consumer substitutes towards the relatively cheaper North good. The increased substitutability also leads to an even smaller welfare loss. Manufacturing employment faces several offsetting effects – Home firms are less productive and have less demand worldwide, but these firms reallocate all input production to Home. Further, North multinationals that source inputs from Home face increased demand and increase their employment worldwide. The net effect is a slightly smaller increase in Home manufacturing employment than with $\epsilon = 4$.

We advise 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 structural change or

⁵³In fact in this model an increase in T_j is not separable from a decrease in τ_{jk} or w_j . Therefore, any calibrated technology increases would reflect a composite change in foreign wages and the variable costs of offshoring.

⁵⁴We 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.

frictions in labor mobility between sectors – that could alter the overall negative employment results we find. Further, 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. Addressing such questions in future research would require more data than currently available to properly calibrate a large multi-country model of multinational sourcing patterns.

Table D3: Quantitative Exercises: Manufacturing Decline

	L_H^X	M_H	P_Y^X	H Welfare
Data (1993 - 2011)	-36%			
Data (1997 - 2007)	-25%			
Counterfactual 1: Only fixed costs change	-4.9%	1.1%	-26.7%	0.3%
Counterfactual 2: No Home offshoring	2.6%	-1.7%	8.6%	-0.3%
Counterfactual 3: No Home offshoring, $\epsilon = 6$	2.1%	-1.9%	2.5%	-0.2%

Notes: This table summarizes the changes in aggregate manufacturing employment and other aggregate variables within the counterfactuals. All changes are reported as percent changes relative to the initial equilibrium. L_H^X is Home manufacturing employment, M_H is the mass of firms in Home and P_Y^X is the Home manufacturing price index. 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.