Essays on Multinationals and International Spillovers

by

Nitya Pandalai Nayar

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2016

Doctoral Committee:

Associate Professor Andrei A. Levchenko, Chair
Professor Alan Deardorff
Assistant Professor Kyle Handley
Professor Linda L. Tesar

To my mother.

ACKNOWLEDGEMENTS

I am extremely grateful to my advisers, Andrei Levchenko, Alan Deardorff, Kyle Handley and Linda Tesar for their continuous guidance, support, wisdom and encouragement. This thesis is a reflection of what I learned from them. I am also thankful to have had excellent co-authors, Christoph Boehm, Aaron Flaaen and Andrei Levchenko, during this process. They motivated me to work harder, and were the source of many inspiring conversations.

The work in these papers has benefited from discussions and comments from several people, including Dan Ackerberg, Regis Barnichon, Javier Cravino, Jordi Galí, Stefania Garetto, Rob Johnson, Sebnem Kalemli-Ozcan, Lutz Kilian, Joan Monras, Eduardo Morales, Eric Sims, Isaac Sorkin, Sebastian Sotelo, Jaume Ventura and Kei-Mu Yi as well as from seminar participants at several institutions and conferences. I also thank the Michigan Institute for Teaching and Research in Economics (MITRE) for financial support, and the Michigan Census RDC for support with the data used in this thesis.

Finally, I am deeply indebted to my parents, sister and grandmother for their patient and continued support throughout this process. Their belief in me made this possible. Chris has helped me through these years in more ways than I can count.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

CHAPTER I

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

with Christoph E. Boehm and Aaron Flaaen

1.1 Introduction

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

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

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

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

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

The increase in imports is largest when the plant is shut down.

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

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

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

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

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

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

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

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

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

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

1.2 Data and Stylized Facts

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

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

1.2.1 Data

Much of this paper relies on a number of restricted-use Census datasets that we have augmented with indicators of multinational affiliate and ownership status. Studying the manufacturing sector over a period of time that spans two distinct industrial classification systems is a challenging task. To create a consistent definition of manufacturing for the period 1993-2011, we apply a new concordance between the SIC/NAICS classification changes that is described in *Fort and Klimek* (2015). We supplement this concordance with our own set of fixes to account for known data issues, and apply it to the Longitudinal Business Database (LBD), a longitudinally-consistent dataset comprising the universe of all business establishments in the U.S. See Appendix A.1.1 for more details on the construction of the consistent manufacturing sample.

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

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

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

intended for further manufacture (intermediate goods) and those destined for consumption (final goods). See *Boehm et al.* (2014a) or Appendix A.1.4 for more details on this procedure.

1.2.2 Facts on Foreign Sourcing and Employment Decline

An aggregate picture of the decline in manufacturing emerges from basic statistics pertaining to our sample. The number of establishments we classify as manufacturing falls from nearly 355,000 in 1993 to under 259,000 in 2011. Table 1.1 shows that the annual rates of decline have been highest in U.S. multinationals and purely domestic, non-trading establishments. The only group to have experienced an increase in net establishments during this period is foreign multinational firms. This group serves as a reminder that supply chain restructuring could also stimulate U.S. employment.[2]

The employment counts in Table 1.2 show a similar picture of aggregate decline. Total manufacturing employment in our sample decreases from nearly 16 million workers in 1993 to 10.26 million in 2011. U.S. multinational establishments constituted 33.3% of the 1993 manufacturing employment but contributed 41% of the subsequent overall decline. While employment at other exporting and importing establishments grew in the first decade of the sample, U.S. multinationals have experienced a steady secular decline throughout our sample.

Concurrent with this employment decline has been a large increase in the participation of trade by U.S. firms. We document the fraction of firms participating in intermediate input sourcing, separately based on whether it occurs at arms length or intra-firm, in Table 1.4. We split the firms into U.S. multinationals and other firms (this group includes the few foreign multinationals in our sample). The fraction of U.S. multinationals participating in arms-length input sourcing from developing countries has increased by nearly 30 percentage

[2]Table 1.3 shows that the decline in multinational firms has not been as severe as the decline in multinational-owned establishments. In the next section, we will show that the extensive margin of establishment shutdown plays an important role in understanding the decline of employment in U.S. multinationals.

points, and the fraction sourcing related party inputs from these countries has doubled. In contrast, the share of firms sourcing from developed countries has only increased about 10 percentage points during our sample period. This fact motivates our analysis in later sections, which will look at sourcing patterns for developing and developed country groups separately.

1.2.2.1 Overall Employment Growth Differential of Multinationals

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

$$e_{it} = \alpha + \beta M_{it} + \mathbf{\Gamma X}_{it} + u_{it} \tag{1.1}$$

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

[3]See *Haltiwanger et al.* (2013) for a recent example. In Appendix A.2.1 we decompose the within-group employment patterns into job creation/destruction rates, separated by intensive and extensive margins.

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

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

Table 1.5 presents the results from this specification, pooled across all years of our sample. The inclusion of records of zero employment before births and after deaths determines whether the measured effect captures the establishment level entry and exit margin. When pooling across years (1994-2011), and focusing only on the intensive margin, we find that multinational establishments have a slightly positive growth rate differential of 1.9 percentage points relative to non-multinational establishments. Once the extensive margin is accounted for, however, this differential changes sign and becomes significantly negative. This is consistent with the strong negative net job-destruction rates at the extensive margin in the analysis in Appendix A.2.1, and points to establishment closure as key to understanding employment declines in multinationals.

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

Table A.6 displays results from this specification with different subsamples and additional controls for robustness. We conclude this set of stylized facts by examining employment and trade of establishments which become part of a multinational firm around the transition time.

[6]A simple aggregation exercise based on our employment weighted regression results tells us the number of jobs lost in U.S. multinational firms relative to the control group. The growth rates implied by the employment weighted specification can be directly applied to multinational employment in the sample year by year, to arrive at this number. Our estimates imply 2.02 million jobs were lost in these firms relative to a narrowly defined control group. Further details are provided in Appendix A.2.

1.2.2.2 Evidence Using an Event-Study Framework

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

We first divide establishments into four mutually exclusive groups: purely domestic and non-exporting, exporting, owned by a U.S. multinational or owned by a foreign multinational. An establishment's state is then defined by the group it belongs to. We next explore whether changes in establishment state are an important feature of our data. To calculate the average transition rates of establishments, we divide the number of establishments transitioning from one state to another in year $t+1$ (including those that retain state) by the total establishments of that type in year t. Table 1.6 reports the results.

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

[7] Table 1.6 shows that multinationals have relatively high exit rates, and there are low transition rates into and out of multinational status overall. However, the large number of establishments per year in our sample provides sufficiently many transitions for our analysis.

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

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

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

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

$$e_{ik}^{y} = \mathbf{\Gamma_{ik}^{y}} \mathbf{X_{i}^{y}} + \sum_{k=-5, k \neq 0}^{10} \delta_k T_{ik}^{y} + u_{ik}^{y}, \tag{1.2}$$

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

An establishment can appear multiple times in this specification. If the establishment exists for several years as a non-multinational until it transitions into multinational status, the establishment would show up in (potentially) several different event years: First as part of a control group for some other transitioning establishment, and then, once, as part of a

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

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

"treated" group of plants in the year of its own transition. This fact has implications for the way that standard errors are calculated. It implies the need to cluster in cells that include the event year — in order to account for potentially correlated errors across the event — in addition to clustering by plant. We utilize the methodology for two-way clustering described in *Cameron et al.* (2011), which also allows for high-dimensional fixed effects.[12]

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

Figure 1.2 shows the estimates of δ_k. Establishments that transition into multinational status experience a relative increase in their employment growth rates in the first two years. This behavior is consistent with the notion that an expansion of international activity is positively correlated with business outcomes for that firm. Subsequent years, however, show a persistently negative effect in employment growth, on the order of roughly 3-6 percentage points relative to the control group. This slight increase followed by a persistently lower employment growth rates could be explained by an initial domestic growth that coincides with (and serves to support) multinational expansion. Such growth could include time spent by the firm learning to replicate processes within the establishment abroad. Following a successful expansion, the firm may then choose to shut down or downsize duplicated firm activities.[14] In future work, we will attempt to disentangle these competing explanations.

Our results point to the importance of studying a long horizon to understand the conse-

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

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

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

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

To examine the role of import substitution in this decline, we estimate equation (1.2) after replacing the left hand side with firm-level intermediate imports (split by related party and arms-length). Figure 1.3 shows estimates of δ_k pertaining to imports. The figure demonstrates that transitions are associated with sizeable increases in both related-party and arms-length intermediate imports. This evidence suggests significant substitution between foreign imports and domestic employment.[15][16]

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

1.2.2.3 Why multinationals?

The stylized facts above demonstrated that multinationals have a consistent negative employment growth rate differential, and that establishments transitioning into multinational

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

[16] We demonstrate the robustness of this result to alternative specifications of firm-level trade in Appendix A.2.

status reduce their employment while their parent firms increase their imports. But is it the ownership (partial or total) of establishments abroad that leads to such employment declines? Or is it simply supply chain restructuring through foreign sourcing of inputs?[17] In other words, why multinationals?

To assess the role of supply chain restructuring overall relative to that occurring within multinationals, we re-run our analysis in Section 1.2.2.2, but consider non-multinational importer transitions instead of the multinational transitions. Intuitively, if the presence of *any* arms length imports are a sufficient indicator of significant supply chain restructuring, employment in these transitioning establishments should display a similar time-path to the multinational transitions in Figure 1.2. The results for employment growth differentials relative to a control group – consisting of non-multinational domestic firms based on the narrowly-defined cells of establishment characteristics as before – are shown in Figure 1.4. Clearly, these establishments do not display such persistent relative employment differentials. This evidence rules out the hypothesis that the presence of some arms-length imports is sufficient to predict relative employment declines. Further, we note that multinationals import the vast majority (over 90% on average) of intermediate inputs in our sample, as shown in Table 1.11 and discussed in the next section.[18] The importing transitions here therefore assess employment outcomes at firms that are primarily importing final goods, or importing small quantities of intermediates. This could account for the lack of employment differentials, as the degree of supply chain restructuring is minimal.[19] In fact, the tight overlap between foreign sourcing of intermediates and multinationals does not permit us to separately identify a multinational employment effect from the employment effects of foreign sourcing.

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

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

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

Why is foreign sourcing of intermediates concentrated in multinationals? Our data permit a closer look at whether there is a relationship between inter and intra-firm imports which lead to a greater degree of overall global production sharing in multinationals. While the share of related-party imports of multinationals is not significantly different to that of arms-length (roughly 53 vs 47 percent on average in our sample), perhaps there exist complementarities between intra- and inter-firm imports. We explore this hypothesis by estimating the following regression for the sample period 1993-2011:

$$\log IMP_{ijkt}^{AL} = \alpha_{ijt} + \gamma_{kt} + \beta \log IMP_{ijkt}^{RP} + \epsilon_{ijkt}. \tag{1.3}$$

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

The results from this regression confirm that sourcing inputs within the firm in a particular foreign location induces more arms-length sourcing as well — even in narrowly defined product categories. This complementarity helps explain the concentration of imports within multinationals in our sample (see Table 1.7), and is presumably the reason their supply chain restructuring is large enough to show large employment effects. Underlying explanations for this finding could include network effects that enable firm sourcing closely related products from suppliers in the same countries both at arms-length or intra-firm, or lower fixed costs of joint arms-length/related-party imports than of each approach separately. We incorporate the last dimension in our structural model in the following section.

1.2.2.4 Linking firm-level employment to imports

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

1.3 A Framework of Offshoring

We next build a structural model featuring firms' choices of supply chain structure to explore whether foreign sourcing can explain the observed changes in employment. Firms can select a sourcing *location* for their intermediates, as well as a sourcing *mode*: whether to produce intra-firm or to source from outside the firm. Intra-firm production abroad is the

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

defining characteristic of a multinational in this model, reflecting the vertical supply chain structure of U.S. multinationals in our data. Much of the literature assumes perfect technology transfer within a firm or to its suppliers, but empirical evidence for this assumption is lacking. We therefore adopt a more general specification that allows for imperfect technology transfer across sourcing locations and modes.

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

1.3.1 Demand for Manufacturing Goods

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

$$X = \left(\int_{\omega \in \Omega} [s(\omega)]^{\frac{1}{\sigma}} [x(\omega)]^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}} \tag{1.4}$$

subject to the constraint

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

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

$$x(\omega) = s(\omega) E P_X^{\sigma-1} p(\omega)^{-\sigma} \tag{1.5}$$

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

$$P_X = \left(\int_{\omega\in\Omega} s(\omega) p(\omega)^{1-\sigma} d\omega \right)^{\frac{1}{1-\sigma}}. \tag{1.6}$$

1.3.2 Firms

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

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

$$x(\varphi, \mathbf{f}, s) = \left(\int_0^1 x(\nu, \varphi, \mathbf{f}, s)^{\frac{\rho-1}{\rho}} d\nu \right)^{\frac{\rho}{\rho-1}}. \tag{1.7}$$

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

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

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

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

$$c(\varphi, \mathbf{f}, s) = \left(\int_0^1 (p(\nu, \varphi, \mathbf{f}, s))^{1-\rho} \, d\nu \right)^{\frac{1}{1-\rho}}. \tag{1.8}$$

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

1.3.2.1 Supply chains

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

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

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

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

$$x_j(\nu, \varphi) = \frac{h_j(\varphi)}{a_j(\nu)} l_j(\nu, \varphi). \tag{1.9}$$

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

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

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

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

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

$$p_j(\nu,\varphi)\frac{x_j(\nu,\varphi)}{\tau_j} - l_j(\nu,\varphi)\,w_j \tag{1.10}$$

subject to the production function (1.9). Here, w_j and τ_j denote wages and iceberg transport costs. If the quantity demanded is positive and finite, optimality requires that the potential producer sets price equal to marginal cost

$$p_j(\nu,\varphi) = \frac{\tau_j a_j(\nu)\,w_j}{h_j(\varphi)}. \tag{1.11}$$

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

1.3.2.2 Basic model implications

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

$$\chi_j(\varphi,\mathbf{f}) = \frac{T_j h_j(\varphi)^{\theta}(\tau_j w_j)^{-\theta}}{\sum_{k\in J(\varphi,f)} T_k h_k(\varphi)^{\theta}(\tau_k w_k)^{-\theta}}. \tag{1.12}$$

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

Optimal input sourcing also implies that the unit cost function (1.8) becomes

$$c(\varphi, \mathbf{f}) = (\gamma)^{\frac{1}{\theta}} \left[\Phi(\varphi, \mathbf{f})\right]^{-\frac{1}{\theta}} \tag{1.13}$$

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

$$\Phi(\varphi, \mathbf{f}) = \sum_{j \in J(\varphi, \mathbf{f})} T_j h_j(\varphi)^{\theta} (\tau_j w_j)^{-\theta}. \tag{1.14}$$

Equation 1.14 summarizes the firm's efficiency at producing its unique variety. We refer to this term as the firm's (overall) production efficiency. As is intuitive, firms of higher types and firms with more sourcing locations/modes have greater values of Φ and lower unit costs. Notice that neither the cost shares (1.12) nor the unit costs depend on the quantity the firm produces.

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

$$\tilde{\pi}(\varphi, \mathbf{f}) = p(\varphi, \mathbf{f}) x(\varphi, \mathbf{f}) - c(\varphi, \mathbf{f}) x(\varphi, \mathbf{f}) \tag{1.15}$$

subject to the demand function (2.8). The firm optimally sets its price to a constant markup over marginal cost, $p(\varphi, \mathbf{f}) = \frac{\sigma}{\sigma-1} c(\varphi, \mathbf{f})$. It is then possible to express revenues as

$$R(\varphi, \mathbf{f}) = s \Sigma P_X^{\sigma-1} \left[\Phi(\varphi, \mathbf{f})\right]^{\frac{\sigma-1}{\theta}}, \tag{1.16}$$

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

1.3.2.3 The choice of the firm's sourcing strategy

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

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

$$\mathbb{E}\left[s\right]\frac{\Sigma}{\sigma}P_X^{\sigma-1}\left[\Phi\left(\varphi,\mathbf{f}\right)\right]^{-\frac{1}{\theta}} - w_H f_J. \tag{1.17}$$

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

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

1.3.3 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

$(\varphi, \mathbf{f}, s)$ with sourcing strategy $J(\varphi, f)$ is

$$l_{HI}(\varphi, \mathbf{f}, s) = \Theta P_X^{\sigma-1} \frac{sE}{w_H} \underbrace{\frac{T_{HI} h_{HI}(\varphi)^{\theta} (w_H)^{-\theta}}{\Phi(\varphi, \mathbf{f})}}_{\chi_{HI},\ \text{Reallocation effect}} \underbrace{\Phi(\varphi, \mathbf{f})^{\frac{\sigma-1}{\theta}}}_{\text{Size effect}}, \tag{1.18}$$

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

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

First, increased foreign competitiveness shifts a greater fraction of intermediate production towards that location — a reallocation effect. This decreases χ_{HI} and reduces labor demand. On the other hand, greater foreign competitiveness increases the firm's optimal size through an increase in production efficiency Φ. This has a positive effect on labor demand. While the elasticity of χ_{HI} with respect to production efficiency Φ is -1, the elasticity of firm size with respect to Φ is $\frac{\sigma-1}{\theta}$, as is also evident in the expression for revenues (equation 1.16). The net effect on employment therefore depends on the sign of $\frac{\sigma-1}{\theta} - 1$. If it is negative, the model implies that the reallocation effect dominates and employment declines.

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

$\sigma - 1 - \theta < 0$.

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

1.3.4 Structural Estimation

Combining equations (1.12) and (1.18), the firm's labor demand at home (scaled by w_{HI}/χ_{HI} and logged) can be expressed as

$$\ln \frac{w_{HI} l_{HI} (\varphi, \mathbf{f}, s)}{\chi_{HI} (\varphi, \mathbf{f})} = \Psi_j - \frac{\sigma - 1}{\theta} \ln \chi_j (\varphi, \mathbf{f}) + (\sigma - 1) \ln h_j (\varphi) + \ln s, \; j \in J \subset \mathcal{J} \quad (1.19)$$

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

The intuition behind the estimating equation (1.19) is closely related to the scale and reallocation effects discussed above. Since the model predicts that the reallocation effect is independent of parameters (recall that the elasticity of l_{HI} with respect to χ_{HI} is one), it is sufficient to estimate the scale effect. We can do so by focusing on $\frac{w_{HI} l_{HI}}{\chi_{HI}}$ rather than labor demand directly. It is easily verified that this ratio is proportional to firm revenues. Note that the intuition behind this estimating equation is closely related to the key insight of *Blaum et al.* (2015), who show that knowledge of firm domestic expenditure shares and revenues is sufficient to measure decreases in unit costs due to imported inputs in a large class of models of importing firms.[25] Here our equation implies that knowledge of the cost shares

[25] Our model falls in this class, and we extend this insight further to all cost shares of a particular firm.In contrast to *Blaum et al.* (2015), our estimation strategy uses all the cost shares of the firm rather than only

of the firm and firm revenues (or expenditure on domestic labor) is sufficient to estimate the scale effect.

Suppose for the moment that $h_j(\varphi)$ was observed. Then, under the assumptions made on s, (in particular, that it is independent of type φ, and fixed costs $\mathbf{f}$, and that it is revealed to the firm only after sourcing decisions are made,) the parameters in equation (1.19) could be consistently estimated by ordinary least squares. Controlling for the remaining variables, $-\frac{\sigma-1}{\theta}$ captures the scale effect. Intuitively, a smaller share χ_j reflects greater production efficiency resulting in greater firm scale. In contrast, large shares imply a smaller scale.

Unfortunately, $h_j(\varphi)$ is not observed and the estimation of (1.19) when $h_j(\varphi)$ is subsumed into the error term yields a biased estimate of $-\frac{\sigma-1}{\theta}$. If $\chi_j(\varphi, \mathbf{f})$ is positively correlated with $h_j(\varphi)$, then the estimate of $-\frac{\sigma-1}{\theta}$ is biased upward. Conversely, if $\chi_j(\varphi, \mathbf{f})$ is negatively correlated with $h_j(\varphi)$, the estimate of $-\frac{\sigma-1}{\theta}$ is biased downward.

The model implies that, conditional on a particular sourcing strategy J, the terms $\chi_j(\varphi, \mathbf{f})$ and $h_j(\varphi)$, $j \in J$ cannot all be positively or all be negatively correlated. The reason is that the sum of a firm's shares over all locations/modes in its sourcing strategy must be one. Therefore, if there exists a sourcing location/mode $j \in J(\varphi, \mathbf{f})$ for which the share $\chi_j(\varphi, \mathbf{f})$ is increasing in firm type φ, some other share, say $\chi_k(\varphi, \mathbf{f})$, $k \in J$, $k \neq j$, must be decreasing in φ. The estimation of (1.19) by OLS therefore does deliver useful information about $\frac{\sigma-1}{\theta}$. If we condition on a particular sourcing strategy J and estimate (1.19) for all $j \in J$, the true value must (asymptotically) lie between the highest and the lowest estimate.

This bounding procedure can be refined further, and provides us with a range for the structural constant. We discuss the relevant details in Appendix A.3. Our first approach to learn about $\frac{\sigma-1}{\theta}$ is to compute the tightest possible bounds, which are reported in the next section.

the domestic cost share, which helps us bound $\frac{\sigma-1}{\theta}$ in the absence of known firm productivity.

While the bounds we obtain are useful, a point estimate of $\frac{\sigma-1}{\theta}$ is naturally preferable. Indeed, under certain conditions it is possible to express φ in terms of observables in equation (1.19) and to estimate $\frac{\sigma-1}{\theta}$ directly. By dividing two sourcing shares from m and k, $m \neq k$, as given in equation (1.12) by one another and rewriting the result, we obtain

$$\frac{h_m(\varphi)}{h_k(\varphi)} = \left(\frac{\chi_m(\varphi, \mathbf{f})}{\chi_k(\varphi, \mathbf{f})} \frac{T_k}{T_m}\right)^{\frac{1}{\theta}} \frac{\tau_m w_m}{\tau_k w_k}. \tag{1.20}$$

We next define $\eta_{m,k}(\varphi) = h_m(\varphi) / h_k(\varphi)$. If $\eta_{m,k}$ is invertible, a point we return to below, then it is possible to rewrite (1.19) for all j and $m \neq k$ as

$$\ln \frac{w_{HI} l_{HI}(\varphi, \mathbf{f})}{\chi_{HI}(\varphi, \mathbf{f})} = \Psi_j - \frac{\sigma-1}{\theta} \ln \chi_j(\varphi, \mathbf{f}) + (\sigma-1) \ln h_j \left(\eta_{m,k}^{-1}\left(\left(\frac{\chi_m(\varphi, \mathbf{f})}{\chi_k(\varphi, \mathbf{f})} \frac{T_k}{T_m}\right)^{\frac{1}{\theta}} \frac{\tau_m w_m}{\tau_k w_k}\right)\right) + \ln s. \tag{1.21}$$

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

A number of semiparametric methods have been developed to consistently estimate $\frac{\sigma-1}{\theta}$ in equation (1.21).[26] Below we report the results for two approaches. First, we approximate the unknown function of $\chi_m(\varphi) / \chi_k(\varphi)$ by a truncated series expansion (see, e.g. Andrews, 1991), using polynomials as basis functions. Second, we approximate the unknown function of χ_m/χ_k with a step function.[27]

A necessary condition in the derivation of equation (1.21) is that the function $\eta_{m,k}(\varphi) = h_m(\varphi) / h_k(\varphi)$ is invertible. If this were not the case, the ratio of shares χ_m/χ_k would not provide useful information about the firm's type. It turns out that although h_m and h_k are

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

[27] More precisely, we partition the range of χ_m/χ_k into fifty percentiles and define fifty indicator variables taking the value one if χ_m/χ_k falls between two consecutive percentiles. We then replace the unknown function of χ_m/χ_k in equation (1.21) by a step function based on these fifty indicator variables.

unknown, our prior analysis allows us to tell whether $\eta_{m,k}$ is invertible. To see this, consider the following procedure. First, fix a particular sourcing strategy J. Next, estimate equation (1.19) separately for all $j \in J$. Denote by m the sourcing location/mode for which the largest (most upward-biased) estimate of $-\frac{\sigma-1}{\theta}$ is obtained and by k the sourcing location/mode for which the smallest (most downward-biased) estimate of $-\frac{\sigma-1}{\theta}$ is obtained. It then must be that the ratio χ_m/χ_k is strictly increasing in φ and that $\eta_{m,k}(\varphi)$ is invertible (equation 1.20). Notice that if the highest and lowest estimates of $-\frac{\sigma-1}{\theta}$ from the procedure above are very close together, then $\eta_{m,k}$ is not invertible, but the bias of $-\frac{\sigma-1}{\theta}$ is negligibly small.

Finally, we discuss two practical issues regarding the estimation of $-\frac{\sigma-1}{\theta}$. To allay concerns about measurement error in the shares χ_j, we estimate several specifications using the firm's shares from the previous year as instruments. For robustness, we also estimate equations (1.19) and (1.21) after replacing the left hand side variable with the log of firm revenues (recall that the model predicts that revenues are proportional to $\frac{w_{HI} l_{HI}}{\chi_{HI}}$).

Linking the Model and the Data

The structural estimation requires data on firm revenues and cost shares from the various sourcing methods $j \in J(\varphi)$. Revenues and cost share information are constructed from the Census of Manufacturers (CMF) merged with import information from the LFTTD. For revenues, we use the total value of shipments of the firm's manufacturing establishments. Total costs are constructed from information on the cost of materials inputs, firm interplant transfers and total machinery expenditures of the firm. We identify intermediate input imports of the firms using a product-level classification method based on the firm's industry. This method is discussed in detail in appendix A.1.4 and in *Boehm et al.* (2014a).

We use lagged values of the cost shares as instruments in robustness exercises. As the Census is quinquennial and only available in 1997, 2002 and 2007, the lagged values have to be constructed using data from the Annual Survey of Manufacturers (ASM) in 1996, 2001 and 2006. The ASM includes information on all large manufacturing establishments, but is not a complete sample of the smaller plants. Therefore we first construct a firm level total cost using establishments sampled in the ASM, and then scale up this variable using the information on total employment captured within the ASM relative to the total employment in our baseline sample. As our baseline sample is built from the LBD, it contains information on all manufacturing establishments of a firm in a year. The assumption implicit in this procedure is that the firm's cost function is the same across the establishments not captured by the ASM as it is in the surveyed portion of the firm.[28] Table 1.10 contains mean cost shares for multinationals sourcing from all locations for the three Census years in our sample.

1.3.4.1 Results and Discussion

We estimate the model in three separate cross-sections in 1997, 2002 and 2007. We find that our estimates of $\frac{\sigma-1}{\theta}$ are remarkably robust both to the method used and to the time period. Table 1.8 presents the bounds on $\frac{\sigma-1}{\theta}$ by year. As discussed in detail in Appendix A.3 our procedure implies a large number of bounds. To reduce the likelihood of statistical outliers, we report the 80th percentile lower and upper bounds.[29] The widest interval for $\frac{\sigma-1}{\theta}$ is $(0, 0.86]$ in 2002, implying the true parameter value is likely in the range where foreign sourcing is a substitute for domestic employment in a firm.

Prior to estimating $\frac{\sigma-1}{\theta}$ using a semi-parametric regression, we must first show that there indeed exists a function $\eta_{m,k}$ that is invertible. In Appendix A.3, we show that χ_{HO} is strictly increasing and χ_{HI} is strictly decreasing in φ. This implies that $\eta_{HO,HI}$ is invertible. When

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

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

estimating equation (1.21), we therefore control for the unknown term using polynomials or step functions in $\frac{\chi_{HO}}{\chi_{HI}}$.

Table 1.9 presents the baseline results for 1997, 2002 and 2007 using (a) polynomials as basis functions and (b) fifty dummy variables representing size bins. The lower panel of the table contains results for the same specifications with the cost shares instrumented by lagged values.[30] The point estimates obtained for each year and specification lie within the bounds in Table 1.8. The estimates range from 0.08 to 0.23, confirming that foreign sourcing is a strong substitute for domestic employment for the firms we study. The results are robust to using revenues instead of scaled payroll as a dependent variable (see Table A.3.3). We further estimate $\frac{\sigma-1}{\theta}$ by industry, to allow for sector-level differences in the scale effect. A kernel density of the estimates is shown in Figure 1.6. While the industry level estimates vary, for our sample of manufacturing industries they are never larger than one, implying that foreign sourcing is a substitute for domestic employment for all manufacturing industries.

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

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

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

timated both parameters separately in various contexts, and the range of estimates is wide. Our estimates, as well as those in *Antràs et al.* (2014) are consistent with earlier findings.[32]

1.3.5 Aggregation

We briefly explore the aggregate implications of our empirical model (1.21) in partial equilibrium. To do so, we consider the population of U.S. manufacturing firms in 1997 and predict the aggregate employment decline implied by the difference in their sourcing shares between 1997 and 2007 and our estimates of $\frac{\sigma-1}{\theta}$. We use an estimate of 0.2, which is at the upper end of our range of estimates. In this exercise we predict employment changes within firms sourcing from abroad, and first-order effects on their U.S. arms-length suppliers (sourcing from HO).

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

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

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

1.4 General equilibrium

1.4.1 A general equilibrium extension

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

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

Home country

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

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

Firms in the X sector Firms in the X sector set up supply chains and produce as described in Section 2.4. In this general equilibrium extension we assume that the number of firms is endogenous and determined by the following entry problem which has three stages. In the first stage, there is an unbounded mass of potential entrants who can pay fixed costs f_E to learn their type φ. In equilibrium, the number of entrants M is determined by a zero expected profit condition. Second, after learning their types, entrants must pay an additional fixed cost f_H to set up production in the Home country. Only firms with sufficiently high types φ find it profitable to do so. The lowest type that enters is φ_{LB}. Finally, those firms that produce in the Home country face the problem discussed in Section 2.4.

Market clearing Labor market clearing in the Home country requires that

$$
\begin{aligned}
L_H = L_H^Z \quad &+ \quad M \Bigg[\int_0^{\infty} \int_{\varphi_{LB}}^{\infty} \left[l_{HI}(\varphi, s) + l_{HO}(\varphi, s) \right] dG_{\varphi}(\varphi)\, dG_s(s) \\
&+ \quad f_E + f_H \left(1 - G_{\varphi}(\varphi_{LB})\right) + \int_{\varphi_{LB}}^{\infty} f_{J(\varphi)} dG_{\varphi}(\varphi) \Bigg] . \qquad (1.22)
\end{aligned}
$$

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

North

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

$$L_N = L_N^Z + M \int_0^\infty \int_{\varphi_{LB}}^\infty \left[l_{NI}(\varphi, s) + l_{NO}(\varphi, s) \right] dG_\varphi(\varphi)\, dG_s(s)\,. \tag{1.23}$$

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

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

1.4.2 Calibration

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

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

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

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

Table 1.11 shows the fraction of firms in the data that source according to each of these strategies. Despite this restriction we capture almost all firms (92.6% in 1997) and almost all trade (95.3% in 1997).[34]

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

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

We next assume that firm types have a Pareto distribution with a lower bound of unity and curvature parameter α_{φ}. The demand elasticity σ is set to 2.3 and the dispersion parameter θ to 6. These values imply that $(\sigma - 1)/\theta$ is 0.217, roughly consistent with the upper end of our point estimates. We also set $\tau_H = 1$ and $\tau_N = \tau_S = 1.15$. Although these parameters are not important for any of the model's predictions, we note that ρ is set to 1.5 and $\mathbb{E}[s]$ to 1.0025. Finally, we must assume a functional form for $h_j(\varphi)$. We choose a simple exponential, $h_j(\varphi) = \varphi^{\kappa_j}$. We note that this choice provides a reasonable fit to the mid-range of the firm size distribution (see Appendix A.3 for a discussion). We set κ_{HI} to one, and choose values for $\kappa_j, j \neq HI$ that are close to κ_{HI}.[35] Table 1.12 summarizes the

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

[35]The exponential assumption is only an approximation to the true functional form of h_j. We therefore have minimal guidance on calibrating these parameters. Choosing values close to 1 ensures these parameters

values of the preset parameters.

The remaining parameters of the model are chosen to match key features of our data in 1997. These parameters are T_j, $j \in \mathcal{J}$, $j \neq \text{HI}$, the fixed cost parameters f_J, $J \subset \tilde{\mathcal{J}}$, f_E, and f_H, the Pareto curvature parameter α_φ as well as the expenditure share β. The targets and the fit of the model in equilibrium are summarized in Table 1.13.

1.4.3 Quantitative exercises

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

The first panel of Table 1.14 reports the percent change of manufacturing employment and multinational employment when the technology parameters $T_j, j \in \{NO, NI, SO, SI\}$ are changed by one percent, one at a time. In response to changes in each of these parameters, aggregate manufacturing employment falls.

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

will not significantly influence the quantitative analysis. For robustness, Appendix A.4, presents the results from the quantitative exercises with alternate choices for κ_j.

the mass of firms M is ambiguous. Importantly, for our calibration the general equilibrium effects do not overturn the partial equilibrium result that foreign sourcing substitutes for domestic employment.

Turning to multinational employment, two offsetting effects are in play. As shown in the partial equilibrium model in Section 2.4, production of intermediates is reallocated towards the location/mode with the technological improvement. However, this effect can be offset by the movement of firms into or out of strategies with that sourcing location/mode. The net effect is therefore ambiguous. For our calibration the second effect implies that multinational employment rises as T_{NI} increases.

The second panel of Table 1.14 shows the employment change in response to lowered fixed costs for each sourcing strategy f_J. With the exception of f_{NI}, the sign of the response of employment is the same as in the case of a technological improvement. Here, the result is driven by the extensive margin: firms enter the sourcing strategy that has lower fixed costs, reducing domestic employment in response. As above, in most cases general equilibrium effects do not change the predictions in partial equilibrium.[36] The changes in multinational employment are governed by the transition of firms between sourcing strategies. In the case of a decrease in f_{NI} non-multinational firms enter multinational status, leading to an increase in multinational employment. In contrast, decreases in f_{SO} and f_{SI} largely induce firms to switch between sourcing strategies while maintaining their multinational status.

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

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

simply by matching observed import patterns. Table 1.13 illustrates the fit of our model to our calibration targets for 1997 and 2007.[37]

To match the observed trade patterns in 2007 the technology parameters $T_j, j \in \{HO, HI, NO, NI, SO,$ uniformly increase. This is shown in Table 1.15. T_{SO} and T_{SI} increase the most, reflecting the fact that imports from the South grew rapidly over the period 1997 - 2007. Although the fraction of multinationals increased over this time period, fixed costs of sourcing strategies increase between our two calibrations. In this model, firms respond to better technology abroad by entering the sourcing strategies that include foreign sourcing. To match the data – where the fraction of firms in these sourcing strategies has only shown small increases – the model has one counterbalancing force. Fixed costs increase to prevent the fraction of multinationals from rising beyond what is observed in the data. While this might appear counterintuitive, we note that this rise in fixed costs might reflect more complex production structures, which are harder to initially offshore. Further, in this model an increase in T_j is not separable from a decrease in τ_j or w_j. Therefore, the calibrated technology increases reflect a composite change in foreign wages and the variable costs of offshoring.

Targeting 2007 trade patterns results in an employment loss within multinational firms by 28%, slightly larger than that observed in the data (see Table 1.16). Total manufacturing employment falls by 13% which accounts for roughly half of the observed decline between 1997 and 2007.[38] In addition to the direct employment loss within multinationals, increased foreign sourcing reduces the demand for intermediates from domestic suppliers. Confronted with less demand for their products these suppliers scale down production and thereby contribute to the employment decline as well.

We advice some caution should be taken in the application of these general equilibrium

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

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

results. This exercise quantifies the decline in manufacturing employment due to increased foreign sourcing alone, and does not account for other factors – such as increased foreign demand for U.S. goods, or foreign multinationals – that could serve to mitigate the overall negative employment results. Indeed as we have shown in 1.2.2, both the number and employment of foreign multinationals in the U.S. increased during our sample. Finally, our analysis has focused on the effects on manufacturing, and it is important to note that one might suspect U.S. multinationals have increased their non-manufacturing employment. We hope to explore these effects in future research. In Appendix A.4, we discuss some counterfactual exercises where we only allow techonology parameters or fixed costs to change.

1.5 Conclusion

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

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

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

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

Table 1.1: Summary Statistics: Establishment Counts by Type: 1993-2011

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

Source: LBD-LFTTD-DCA-UBP as explained in text.
This table reports the establishment counts pertaining to the "constant" manufacturing sample used in section 3.3.

Table 1.2: Summary Statistics: Employment Counts by Type: 1993-2011

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

Source: LBD-LFTTD-DCA-UBP as explained in text.
This table reports the employment counts pertaining to the "constant" manufacturing sample used in section 3.3.

Table 1.3: Summary Statistics: Firm Counts by Type: 1993-2011

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

Source: LBD-LFTTD-DCA-UBP as explained in text.
This table reports the firm counts pertaining to the “constant” manufacturing sample used in section 3.3.

Table 1.4: Percentage of Firms Participating in Foreign Input Sourcing: 1993-2011

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

Source: LBD-LFTTD-DCA-UBP as explained in text.
This table reports the fraction of U.S. multinationals and non U.S. multinationals that sourced inputs from foreign countries. These are non-exclusive shares of the total number of firms in 1.3. Non U.S. multinationals includes foreign multinationals and other trading firms.

Table 1.5: Pooled Regression Results

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

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

Source: LBD-LFTTD-DCA-UBP as explained in text.
This table reports the pooled regression results from 1.1 at the establishment and firm level.

Table 1.6: Average Establishment-Level Transition Probabilities: 1993-2011

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

Source: LBD, DCA, and UBP
This table reports average probability of transition from state i in t to j in $t+1$ where $\{i, j \in D, X, MH, MF, Entry, Exit\}$. The average number of establishments corresponding to each type is in Table 1.1.

Table 1.7: Inter-Firm and Intra-Firm Sourcing

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

Source: LFTTD
This table reports the results from equation 1.3. The dependent variable is the log of a firm's inter-firm imports from a particular country or industry within a country.

Table 1.8: Estimation Results: Bounding

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

Source: LBD,LFTTD, CMF and ASM
This table reports bounds on $\frac{\sigma-1}{\theta}$ implied by the bounding procedure in section 1.3.4. The upper bound is the 80th percentile of all lower bounds calculated by applying the procedure to different sourcing strategies, as discussed in appendix A.3.

Table 1.9: Estimation Results: Semiparametric Regressions

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

Source: LBD,LFTTD, CMF and ASM
This table reports point estimates for $\frac{\sigma-1}{\theta}$ from the polynomial approximation and size bin approaches discussed in 1.3.4.We use fifty size bins for the approximation. The lower panel displays results where the cost shares are instrumented with lagged values.

Table 1.10: Cost Shares for Firms with $J = \{HO, HI, NO, NI, SO, SI\}$

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

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

Table 1.11: Firm Sourcing Patterns

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

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

Table 1.12: Calibration Stage 1

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

This table summarizes the first stage of the baseline calibration.

Table 1.13: Quantitative Exercises: Model Fit

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

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

Table 1.14: Quantitative Exercises: Local Effects

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

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

Table 1.15: Quantitative Exercises: Parameter changes

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

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

Table 1.16: Quantitative Exercises: Manufacturing Decline

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

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

Figure 1.1: Share of Trade and Firm Participation in Trade, by Type

Value of Trade by Firm Type

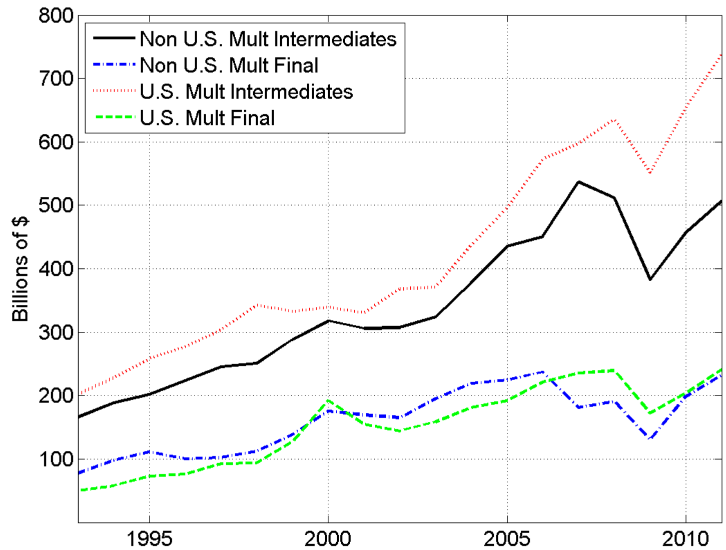

Source: LFTTD-DCA-UBP as explained in text.
These figures report the value of intermediate and final goods trade by firm type, as well as the share of intermediate inputs imported from low income countries by U.S. multinationals.

Figure 1.2: Employment Growth Differential of Multinational Transitions

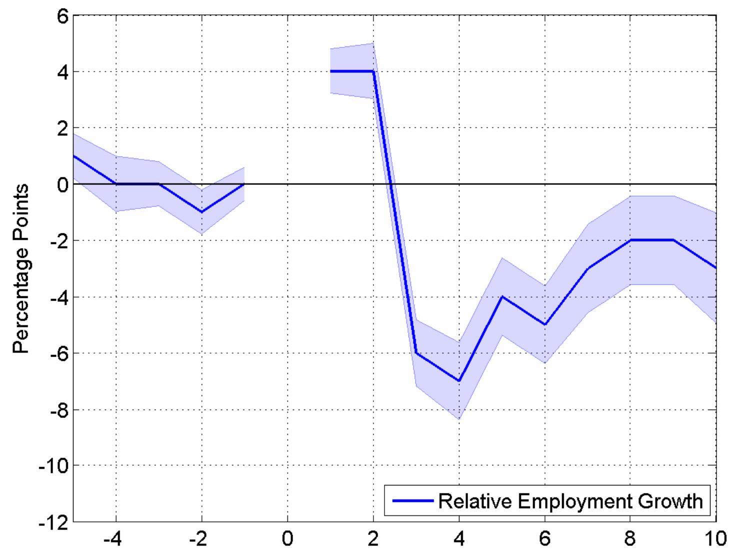

Source: LFTTD-DCA-UBP as explained in text.
This figure plots the pre and post annual deviations in the employment growth rate of establishments that transition into part of a multinational firm in year ($t = 0$), relative to a control group based on interacted effects of firm age, establishment size, and industry (in year $t = -1$). The control group consists of establishments that are not part of a multinational firm in year $t = 0$. See equation 1.2. The shaded area corresponds to a 95 percent confidence interval.

Figure 1.3: Importing Differentials of Multinational Transitions

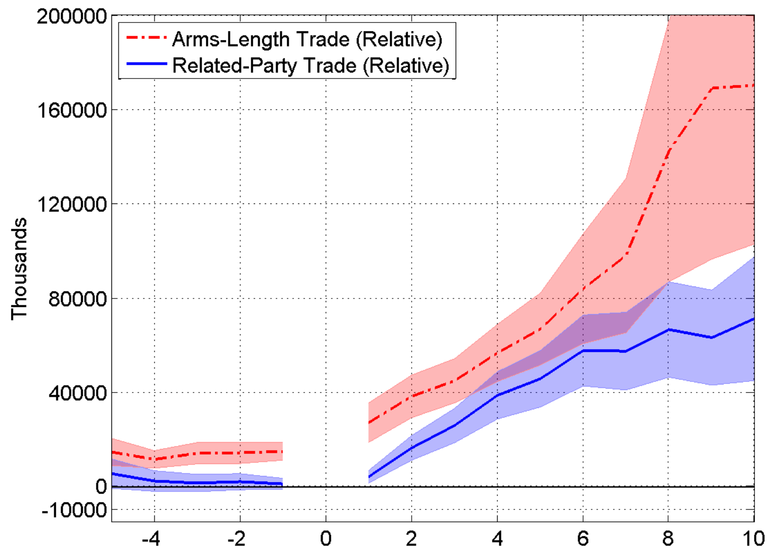

Source: LFTTD-DCA-UBP as explained in text.
This figure reports the related-party and arms-length intermediate input imports of the parent firm of an establishment that transitions into part of a multinational firm in year ($t = 0$), relative to a control group based on interacted effects of firm age, establishment size, and industry (in year $t = -1$). See equation 1.2, modified to reflect firm-level imports as dependent variables. The shaded area corresponds to a 95 percent confidence interval.

Figure 1.4: Employment Growth Differential of Importer Transitions

Source: LFTTD-DCA-UBP as explained in text.
This figure plots the pre and post annual deviations in the employment growth rate of establishments that begin importing from abroad year ($t = 0$), relative to a control group based on interacted effects of firm age, establishment size, and industry (in year $t = -1$). The control group consists of establishments that are not part of a multinational firm in year $t = 0$, nor have recorded positive imports in the period ($t - 3, t = 0$). See equation 1.2. The shaded area corresponds to a 95 percent confidence interval.

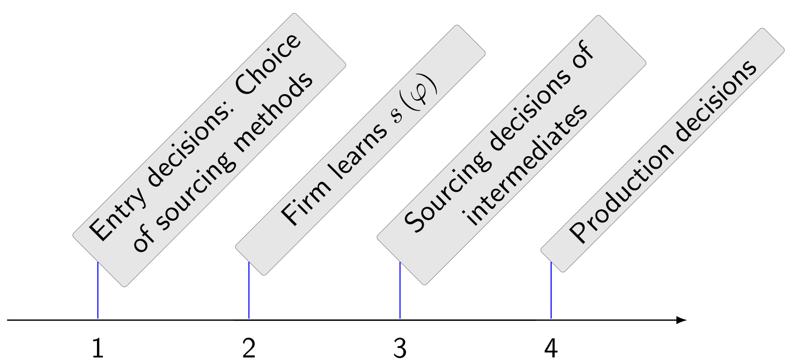

Figure 1.5: Stages of the Firm's Problem

Figure 1.6: Estimation results by industry

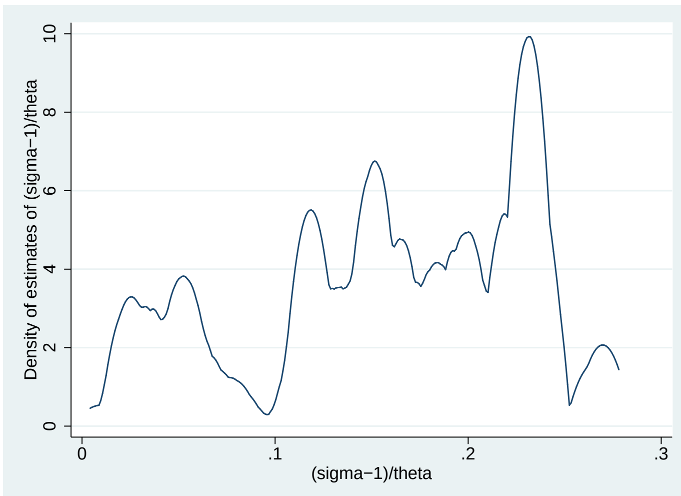

Source: CMF-LFTTD as explained in text.
This figure plots the kernel density of the results of the estimation of $\frac{\sigma-1}{\theta}$ using equation 1.21 by industry in 1997. The results are similar for other estimation years 2002 and 2007.

CHAPTER II

Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake

with Christoph E. Boehm and Aaron Flaaen

2.1 Introduction

The spillover effects of trade and financial linkages has been a preeminant topic in international economics in recent decades. The large expansions in trade and foreign direct investment (FDI) in the past twenty years have generated much discussion on whether they increase volatility (*di Giovanni and Levchenko*, 2012), increase comovement (*Frankel and Rose*, 1998; *Burstein et al.*, 2008) or lead to less diversified production and specialization (*Imbs*, 2004). Identifying the micro-foundations underlying the role of these linkages in the increased interdependence of national economies is challenging. Advanced economies are highly connected, and most variables influenced by any candidate mechanism are often correlated with other developments in the source and destination countries. There is often little in the way of exogenous variation to isolate any particular mechanism from a host of confounding factors. Moreover, the requisite data to examine these issues at the necessary detail and disaggregation have been, until recently, unavailable.

This paper provides empirical evidence for the cross-country transmission of shocks via the rigid production linkages of multinational firms. The principal mechanism at work is not new; the idea of input-output linkages as a key channel through which shocks propagate through the economy dates back to at least *Leontief* (1936) or *Hirschmann* (1958). Two advances in this paper permit a new quantitative evaluation of the nature and magnitude of these linkages. First, we utilize a novel dataset that, for the first time, links restricted U.S. Census Bureau microdata to firms' international ownership structure. This information permits a forensic focus on particular firms and their underlying behavior. Second, we utilize the March 2011 Tōhoku earthquake and tsunami as a natural experiment of a large and exogenous shock disrupting the production linkages originating from Japan.

We study the role of imported intermediate inputs in the transmission of this shock to the United States economy. Because disruptions to imports of final goods would be unlikely to affect U.S. production, we develop a new methodology for isolating firm-level imports of intermediate inputs. We show that the U.S. affiliates of Japanese multinationals are the most natural source of this transmission, due to their high exposure to imported intermediates from Japan. The scope for shocks to these imported inputs to pass through and affect the firm's U.S. production depends on how substitutable they are with inputs from alternative sources. In other words, the role of imported inputs in the transmission of shocks is governed by the elasticity of substitution with respect to domestic factors of production.

We estimate this elasticity using the relative magnitudes of high frequency input and output shipments in the months following the Tōhoku earthquake/tsunami. This proceeds in two steps. First, reduced form estimates corresponding to Japanese multinational affiliates on average show that output falls, without a lag, by a comparable magnitude to the drop in imports. These results suggest a near-zero elasticity of imported inputs. Second, we structurally estimate a firm-level production function that allows for substitution across different types of inputs. The structural estimation procedure we use is uniquely tailored to the experiment. In an initial period prior to the Tōhoku disruption, we infer information on the firm's productivity and optimal input mix. Then, applying this production function to the period of the disruption, we estimate the elasticity parameters based on how changes in the firm's input mix translate into changes in output.

This estimation strategy has a number of attractive features. Most importantly, it relies on very few assumptions. Direct estimation of the production function circumvents the many difficulties associated with specifying a firm's optimization problem in the period after the shock. Second, it yields transparent parameter identification. This is an advantage over traditional estimation strategies as it does not suffer from omitted variables and endogeneity concerns arising from correlated shocks. Third, it allows for the estimation across different subgroups of firms.

The structural estimates are broadly in agreement with the results from our reduced form exercise. For Japanese multinationals, the elasticity of substitution across material inputs is 0.2 and the elasticity between material inputs and a capital/labor aggregate is 0.03. For non-Japanese firms using inputs from Japan, the estimates of the elasticity of substitution across material inputs are somewhat higher at 0.42 to 0.62. While the high cost share and particularly low elasticity for Japanese affiliates explains their predominant contribution to the direct transmission of this shock to the U.S., the elasticity estimates for non-Japanese firms are still substantially lower than typical estimates used in the literature. We argue that the substantial share of intra-firm intermediate trade implies greater complementarities in aggregate trade than is currently recognized.

There are a number of important implications for such low values of the elasticity of substitution. This parameter appears in various forms in a wide span of models involving the exchange of goods across countries. As discussed by *Backus et al.* (1994) and *Heathcote* *and Perri* (2002) among others, this parameter is critically important for the behavior of these models and their ability to match key patterns of the data. Prior estimates of this parameter were based on highly aggregated data that naturally suffered from concerns about endogeneity and issues of product composition.[1] Reflecting the uncertainty of available estimates for the elasticity of substitution, it is a common practice to evaluate the behavior of these models along a wide range of parameter values.

It is well known that a low value for this parameter (interpreted as either substitution between imported and domestic goods in final consumption or as intermediates in production) improves the fit of standard IRBC models along several important dimensions. In particular, the elasticity of substitution plays a role in two highly robust failings of these models: i) a terms of trade that is not nearly as variable as the data, and ii) a consumption comovement that is significantly higher than that of output, whereas the data show the opposite relative ranking.[2]

To understand the relationship between the elasticity and comovement, it is helpful to recall that these models generate output comovement by inducing synchronization in factor supplies, a mechanism that by itself generally fails to produce the degree of comovement seen in the data. Complementarities among inputs together with heterogeneous input shocks will generate direct comovement in production, augmenting the output synchronization based on factor movements. *Burstein et al.* (2008) show that a low production elasticity of substitution between imported and domestic inputs reduces substitution following relative price movements, and thereby increases business cycle synchronization.[3] It is also relatively straightforward to see how a lower elasticity increases volatility in the terms of trade. When

[1]For a very useful compendium of this research from this era, see *Stern et al.* (1976). More recently, work by *Halpern et al.* (2011) and *Goldberg et al.* (2010) demonstrate that materials inputs from foreign countries are imperfectly substitutable with domestic inputs for Hungary and India respectively.

[2]Due to the robust nature of these shortcomings, *Backus et al.* (1995) refer to them as the "price anomaly" and "quantity anomaly" respectively.

[3]Although they do not estimate this parameter, the value they advocate (0.05) is indeed close to our estimates.

two inputs are highly complementary, deviations from the steady state mix are associated with large changes in their relative prices. In the words of *Heathcote and Perri* (2002, page 621): "greater complementarity is associated with a larger return to relative scarcity."

The estimates in this paper have implications for the role of trade in firm-level and aggregate volatility. Other research has argued that firms can diversify risk arising from country specific shocks by importing (*Caselli et al.* (2014)) or that firms with complex production processes of several inputs are less volatile as each input matters less for production (*Koren and Tenreyro* (2013)). On the other hand, there is a well-established fact that complementarities and multi-stage processing can lead to the amplification of shocks as in *Jones* (2011) and *Kremer* (1993). We discuss the potential for measured amplification in our context in Section 2.5.

This paper is also a contribution to the empirical evidence on the role of individual firms in aggregate fluctuations, emanating from the work of *Gabaix* (2011). Other related evidence comes from *di Giovanni et al.* (2014), who use French micro-data to demonstrate that firm-level shocks contribute as much to aggregate volatility as sectoral and macroeconomic shocks combined. The so-called granularity of the economy is very much evident in our exercise; though the number of Japanese multinationals is small, they comprise a very large share of total imports from Japan, and are arguably responsible for a measurable drop in U.S. industrial production following the Tōhoku earthquake (see Figure 2.3).

The strong complementarity across material inputs implies that non-Japanese input use falls nearly proportionately, thereby propagating the shock to other upstream (and downstream) firms in both the U.S. economy and abroad. Many suppliers were thus indirectly exposed to the shock via linkages with Japanese affiliates that had i) high exposure to Japanese inputs and ii) a rigid production function with respect to other inputs. Network effects such as these can dramatically magnify the overall transmission of the shock (both across countries and within). And while such effects are commonly understood to exist, this paper provides unique empirical evidence of the central mechanisms at work.

As is the case with most research based on an event-study, some care should be taken in generalizing the results to other settings. Although we have already highlighted the aggregate implications of the effects we estimate, one might worry that the composition of Japanese trade or firms engaged in such trade is not representative of trade linkages more broadly. We believe the results we obtain are informative beyond the context of this particular episode for two reasons. First, the features of Japanese multinationals that are underlying the transmission of this shock are common to all foreign multinational affiliates in the United States.[4] Second, estimates corresponding to all firms in our sample also exhibit substantial complementarities, and as a whole these firms account for over 70 percent of total U.S. manufacturing imports.

The next section describes the empirical strategy and data sources used in this paper, section 2.3 presents reduced form evidence in support of a low production elasticity of imported inputs for Japanese multinational affiliates. In Section 2.4, we expand the scope of parameters we identify with a structural model of cross-country production linkages. We estimate the parameters of this model across several different subgroups. Section 2.5 discusses the implications of these estimates, and details a number of checks and robustness exercises. The final section offers concluding thoughts.

2.2 Empirical Strategy and Specification

This section outlines the empirical approach of using an event-study framework surrounding the 2011 Tōhoku event to estimate the production elasticity of imported inputs. We discuss the relevant details of this shock, document the aggregate effects, and then outline the empirical specification for the firm-level analysis.

[4]Intra-firm trade accounts for a large majority of the trade of Japanese affiliates. More generally, the intra-firm share of imported intermediates for all foreign affiliates in the U.S. is 71 percent.

2.2.1 Background

The Tōhoku earthquake and tsunami took place off the coast of Northeast Japan on March 11, 2011. It had a devastating impact on Japan, with estimates of almost twenty thousand dead or missing (*Schnell and Weinstein* (2012)) and substantial destruction of physical capital. The magnitude of the earthquake was recorded at 9.0 on the moment magnitude scale (M_w), making it the fourth largest earthquake event recorded in the modern era.[5] Most of the damage and casualties were a result of the subsequent tsunami that inundated entire towns and coastal fishing villages. The effects of the tsunami were especially devastating in the Iwate, Miyagi, and Fukushima prefectures. The Japanese Meteorological Agency published estimates of wave heights as high as 7-9m (23-29ft), while the Port and Airport Research Institute (PARI) cite estimates of the maximum landfall height of between 7.9m and 13.3m (26-44ft).

Figure 2.1 shows the considerable impact of the Tōhoku event on the Japanese economy. Japanese manufacturing production fell by roughly 15 percentage points between February and March 2011, and did not return to trend levels until July. Much of the decline in economic activity resulted from significant power outages that persisted for months following damage to several power plants – most notably the Fukushima nuclear reactor.[6] Further, at least six Japanese ports (among them the Hachinohe, Sendai, Ishinomaki and Onahama) sustained significant damage and were out of operation for more than a month, delaying shipments to both foreign and domestic locations. It should be noted, however, that the largest Japanese ports (Yokohama, Tokyo, Kobe) which account for the considerable majority of Japanese

[5]Since 1900, the three earthquakes of greater recorded magnitude are: the 1960 Great Chilean earthquake (magnitude 9.5), the 1964 Good Friday earthquake in Prince William Sound, Alaska (magnitude 9.2); and the 2004 Sumatra-Andaman earthquake (magnitude 9.2).

[6]For precautionary reasons, all nuclear power plants were immediately shut down following the earthquake, and remained largely offline until 2014 or later. Because the electricity infrastructure exists on two separate grids (a 60Hz to the south and west, and 50Hz to the north and east), the reduction in power supply in Northeast Japan was not easily remedied, and power outages persisted for months.

trade, re-opened only days after the event.

As expected, the economic impact of the event was reflected in international trade statistics, including exports to the United States. Figure 2.2 plots U.S. imports from Japan around the period of the Tōhoku event, with imports from the rest of the world for comparison. The large fall in imports occurs during the month of April 2011, reflecting the several weeks of transit time for container vessels to cross the Pacific Ocean. The magnitude of this drop in imports is roughly similar to that of Japanese manufacturing production: a 20 percentage point drop from March to April, with a full recovery by July 2011.

More striking is the response of U.S. industrial production in the months following the event. Figure 2.3 demonstrates that there is indeed a drop in U.S. manufacturing production in the months following the Japanese earthquake. Although the magnitudes are obviously much smaller — roughly a one percentage point drop in total manufacturing and almost two percentage points in durable goods — the existence of a measurable effect is clear.[7]

Though tragic, the Tōhoku event provides a glimpse into the cross-country spillovers following an exogenous supply shock. This natural experiment features many characteristics that are advantageous for this type of study. It was large and hence measurable, unexpected, and directly affected only one country. The shock was also short-lived, which rules out immediate supplier restructuring and allows for an estimate of the elasticity for a given supply chain.[8] On the other hand, the short duration of the shock presents a challenge for measurement as it limits the available datasets with information at the required frequency. We utilize a novel firm-level dataset to uncover the mechanisms at work behind the transmission of this shock.

[7] At the level of total U.S. GDP, both Deutsche Bank and Goldman Sachs revised 2nd quarter U.S. estimates down by 50 basis points explicitly due to the events in Japan.

[8] It also rules out large balance sheet effects that would make differential credit conditions an operative feature.

2.2.2 Data

Several restricted-use Census Bureau datasets form the core of our firm-level analysis. The Longitudinal Business Database (LBD) collects the employment, payroll, and major industry of all establishments operating in the United States, and is maintained and updated as described by *Jarmin and Miranda* (2002). Longitudinal linkages allow the researcher to follow the establishment over time, and the annual Company Organization Survey (COS) provides a mapping from establishments to firms. All of the analysis in this paper will be at the firm-level.

The Longitudinal Foreign Trade Transactions Database (LFTTD), which links individual trade transactions to firms operating in the United States. Assembled by a collaboration between the U.S. Census Bureau and the U.S. Customs Bureau, the LFTTD contains information on the destination (or source) country, quantity and value shipped, the transport mode, and other details from point-of-trade administrative documents. Importantly for this study, the LFTTD includes import and export trade transactions at a *daily* frequency, which is easily aggregated to monthly-level trade flows. A number of important papers have utilized this resource, such as *Bernard et al.* (2007) and *Bernard et al.* (2006).

We utilize two novel extensions to this set of Census data products. First, a new link between a set of international corporate directories and the Business Register (BR) of the Census Bureau provides information on the international affiliates of firms operating in the United States. These directories provide information, for the first time, to identify those U.S. affiliates part of a foreign parent company, as well as those U.S. firms with affiliate operations abroad. This information is an important resource for identifying the characteristics of U.S. firms affected by the Tōhoku event. For information on these directories and the linking procedure used, please see Appendix B.2.1.

The second novel data resource is a system to classify firm-level import transactions as intermediate or final goods. Although intermediate input trade represents as much as two-thirds of total trade (see *Johnson and Noguera* (2012)), the LFTTD does not classify a trade transaction based on its intended use. To overcome this limitation, we use information on the products produced by U.S. establishments in a given industry to identify the set of products intended for final sale *for that industry.*[9] The remaining products are presumably used by establishments in that industry either as intermediate inputs or as capital investment. Details on this classification procedure are available in Appendix B.2.2. In the aggregate, this firm-level classification procedure yields estimates of the intermediate share of trade that are consistent with prior estimates: 64 percent of manufacturing imports are classified as "intermediates" in 2007.

Finally, we utilize geographic information on the severity of the earthquake/tsunami that is compiled by the U.S. Geological Survey (USGS). By geocoding the Japanese addresses of firms with U.S. operations, we construct an earthquake intensity measure for each Japanese affiliate location. We then apply such information to the U.S. operations as a way to further measure the sample of firms plausibly affected by the shock. Please see Appendix B.2.3.2 for details. Figure 2.4 shows the geographic distribution of one such USGS measure — the modified mercalli index (MMI) — along with the geocoded affiliate locations.

The ideal dataset to evaluate the transmission of the Tōhoku event on U.S. firms would consist of high frequency information on production, material inputs, and trade, separated out by geographic and ownership criteria. Unfortunately, Census data on production and material inputs at the firm-level is somewhat limited. The Annual Survey of Manufacturers (ASM) contains such information, but at an annual frequency and only for a subset of manufacturing firms. On the other hand, firm-level trade information is available at a nearly daily frequency, and covers the universe of firms engaged in exporting/importing. For the

[9]Note that products intended for final sale for a given industry may still be used as intermediates for other firms in a different industry. Alternatively, such "final goods" can be sold directly to consumers for ultimate consumption.

purposes of characterizing the shock to firm-level imports of intermediate goods, the LFTTD (and supplements identified above) is ideal. There remain significant gaps in information on a firm's domestic input usage, a limitation we discuss in subsequent sections.

Because of the challenges of high-frequency information on firms' U.S. production, we utilize a proxy based on the LFTTD — namely the firm's exports of goods to North America (Canada and Mexico). The underlying assumption of this proxy is that all firms export a fixed fraction of their U.S. output to neighboring countries in each period. The advantage of this approach is the ability to capture the flow of goods at a specific point in time. There are few barriers to North American trade, and transport time is relatively short. Moreover, exporting is a common feature of these firms, of which exports to North America is by far the largest component. The obvious disadvantage of this approach is that it conditions on a positive trading relationship between firms in the U.S. and Canada/Mexico. We will assess the quality of this measure as a proxy for output in section 2.5.3.1.[10]

2.2.3 Basic Theory

Before moving to our firm-level analysis, it is useful to describe the basic theoretical structure of the features of firm-level production that we estimate. The transmission of shocks within a firm's production chain is governed by the flexibility of production with respect to input sourcing. Rather than model these complex networks directly, the literature typically summarizes this feature with the well-known elasticity of substitution within a C.E.S. production function. Our identification of this elasticity will rely on the relative impacts on output and imported inputs following the shock. To be concrete, consider the

[10]Another consideration with the use of this proxy is whether it more accurately reflects production or sales, as the two are distinct in the presence of output inventories. In our case, this depends on whether the inventories are held in the U.S. or Canada/Mexico. Without further evidence, we interpret the proxy to be capturing some mix between production and sales. The structural estimation in section 2.4 will allow for such a mix.

C.E.S. production function

$$x = \left[(1-\mu)^{\frac{1}{\psi}} \left[F_D \right]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} \left[IM \right]^{\frac{\psi-1}{\psi}} \right]^{\frac{\psi}{\psi-1}} \tag{2.1}$$

where output consists of combining a domestic bundle of factors F_D (e.g. capital and labor) with a foreign imported input IM. The parameter μ reflects the relative weight on the input IM in production, conditional on prices and a given elasticity value. Suppose the firm purchases its inputs in competitive markets with prices p_D and p_M, respectively, and sells its good at price p_x. Our approach in this section will be to estimate the parameter ψ governing the degree of substitution between these inputs, using information on the output elasticity with respect to imported inputs, $\frac{\partial \ln p_x x}{\partial \ln p_M M}$, in the months following the shock.

The first order conditions imply that

$$\frac{F_D^*}{IM^*} = \frac{1-\mu}{\mu} \left(\frac{p_M}{p_D} \right)^{\psi}, \tag{2.2}$$

where F_D^* and IM^* denote the optimal quantities of inputs. We would like to show the theoretical foundations underlying the intuitive result that a one-for-one drop in output with the fall in imported inputs implies an elasticity of zero. To do this, we make the following assumptions, all of which we will relax to some degree in the estimation framework in Section 2.4:

1. Imported inputs shipments are disrupted, such that the firm receives a suboptimally low quantity of IM: $IM < IM^*$;

2. The firm is unable to adjust domestic inputs F_D^* or its price p_x after learning that it receives IM;

3. The firm does not shut down.

Given these assumptions, the following result holds:

Result 1. *Under assumptions 1) to 3):*

$$\frac{\partial \ln p_x x}{\partial \ln p_M IM} = \frac{1}{1 + \left(\frac{IM^*}{IM}\right)^{\frac{\psi-1}{\psi}} \left(\frac{1-\mu}{\mu}\right) \left(\frac{p_M}{p_D}\right)^{\psi-1}} \in (0, 1) \tag{2.3}$$

for any $\psi \in (0, \infty)$.

Proof. See Appendix B.1.1 for details. □

An immediate implication of this result is that the output elasticity is unity only when ψ approaches zero.[11] In this case $\left(\frac{IM^*}{IM}\right)^{\frac{\psi-1}{\psi}} \rightarrow 0$ (recall that $IM < IM^*$) and hence $\lim_{\psi \rightarrow 0} \frac{\partial \ln p_x x}{\partial \ln p_M IM} = 1$. Hence, observing a one-for-one drop in the value of output with the value of imported intermediates, we infer that ψ is close to zero. It is also straightforward to show that conditional on a value for $\psi \in (0, \infty)$, the output elasticity in (2.3) is increasing in the parameter μ. That is, conditional on a given drop in the imported input, a larger weight on this input leads to a larger percent response in output.

Our use of the natural experiment is critical for observing the effects of suboptimal input combinations (F_D^*, IM). To see this, suppose the firm could freely adjust F_D after learning it will receive $IM < IM^*$. Then, it would choose F_D such that $\frac{F_D}{IM} = \frac{F_D^*}{IM^*}$ and the firm would contract one-for-one with the drop in imports. It is a well-known fact that constant returns to scale production functions in competitive environments lead to indeterminate firm size. This has the implication that:

$$\frac{\partial \ln (p_x x)}{\partial \ln (p_M IM)} = \frac{\partial \ln (p_x x)}{\partial \ln (p_D F_D)} = \frac{\partial \ln (p_D F_D)}{\partial \ln (p_M IM)} = 1. \tag{2.4}$$

[11] There is a second case which we do not examine, where $\psi \rightarrow \infty$ and $p_M < p_D$ and thus the firm only uses IM. We discard this scenario because such a firm would not show up in our data (i.e. this case implies zero U.S. employment).

In this case it is not possible to learn anything about ψ from the joint behavior of output and the value of intermediate inputs. We provide evidence below that firms did not significantly adjust their domestic labor force following the disruption, so that a constant F_D is indeed a reasonable assumption in this simple framework. To be sure, there are a number of alternative frameworks where such behavior would not hold. We discuss some of these in Appendix A, and show that the mapping $\lim_{\psi \to 0} \frac{\partial \ln p_x x}{\partial \ln p_M IM} = 1$ is more general.

2.3 Reduced Form Evidence

This section will provide intuitive reduced-form evidence on the elasticity of substitution corresponding to the U.S. affiliates of Japanese multinationals. We discuss our strategy for understanding this elasticity via firm-behavior in the months following the Tōhoku event, and then report the results.

2.3.1 Framework

Our analysis of the production function (2.1) above demonstrates that a natural measure to evaluate the potential conduits of the Tōhoku shock to the United States would be the degree of reliance on Japanese imported inputs. This is best expressed as the cost share of inputs from Japan, and can be constructed in a Census year by taking a firm's Japanese imported inputs and dividing by all other inputs (which includes production worker wages and salaries, the cost of materials, and the cost of new machinery expenditures). Exposure to Japanese imported inputs is heavily concentrated among Japanese affiliates. In the year 2007, which is the closest available Census year, this cost share was nearly 22% on average for Japanese affiliates (see Table 2.1), compared to just 1% for other firms. For more detail on the heterogeneity across and within these firm groups, we construct a density estimate of such an exposure measure for the Japanese affiliates and non-Japanese multinationals. The results, shown in Figure 2.5, show little overlap between these distributions: there are few Japanese affiliates with low exposure to Japanese inputs, and few non-Japanese firms have substantial exposure.[12]

We now estimate the relative impacts on imported inputs and output for the Japanese affiliates as a group. To do this, we implement a dynamic treatment effects specification in which a firm is defined as being treated if it is owned by a Japanese parent company.[13] The effect on these firms can be inferred from the differential impact of the variable of interest relative to a control group, which soaks up common seasonal patterns and other demand-driven factors in the U.S. market. While there are a number of competing methodologies for this type of estimation, we use normalized propensity score re-weighting due to the relatively favorable finite-sample properties as discussed in *Busso et al.* (2014), as well as for its transparent intuition. Consistent estimation of the average treatment effect on the treated requires the assumption of conditional independence: the treatment/control allocation is independent of potential outcomes conditional on a set of variables. As the average Japanese firm differs considerably from other firms in the data, we use other multinational firms – both US and non-Japanese foreign- as our baseline control group prior to reweighting. To compute the propensity scores for reweighting, we control for size and industry, which ensures the control group has a similar industrial composition and size distribution as our treated sample.[14] Table 2.2 reports summary values for the sample, including statistics on the balancing procedure using the normalized propensity score.

[12]The exposure measure used in Figure 2.5 is from 2010 and does not include the cost of domestic material usage.

[13]We could have also used a threshold of Japanese input usage for the classification of treatment status. Doing so yields estimates that are very similar, which is due to the patterns evident in Figure 2.5. We have also tried conditioning on our geographic information (i.e. the firm-level Japanese MMI index) in defining a Japanese firm as being treated. The results are largely unchanged from those we report here, and for the sake of clarity we report results pertaining to the full sample.

[14]Using the predicted values (p) from the first stage regression, the inverse probability weights are $\frac{1}{1-p}$ for the control group and $\frac{1}{p}$ for the treated group. To normalize the weights such that the treated firms have weights equal to one, we then multiply each set of weights by p.

The magnitude of the shock for a representative Japanese multinational is captured by the effect on total imported intermediate products at a monthly frequency.[15] Including non-Japanese imported intermediates is important for applying the control group as a counterfactual, and the shares by source-country gives the necessary variation for identification: as shown in Table 2.2 the share of imported inputs from Japan is 70% of the total for Japanese firms and only 3.5% for non-Japanese multinationals. Let $V_{i,t}^M$ be the value of intermediate imports of firm i in month t, after removing a firm-specific linear trend through March 2011. We fit the following regression:

$$V_{i,t}^M = \alpha_i + \sum_{p=-4}^{9} \gamma_p E_p + \sum_{p=-4}^{9} \beta_p E_p \mathrm{JPN}_{i,p} + u_{i,t} \tag{2.5}$$

where α_i are firm fixed-effects, γ_p are monthly fixed effects (with the indicator variables E_p corresponding to the calendar-months surrounding the event), and $u_{i,t}$ is an error term. The baseline sample will consist of January 2009 to December 2011. We denote March 2011 as t=0.

The β_p coefficients are of primary interest. The $\mathrm{JPN}_{i,t}$ is an indicator variable equal to one if the firm is owned by a Japanese parent company. Interacting these indicator variables with each month of the panel allows for a time-varying effect of Japanese ownership on a firm's overall intermediate input imports, particularly during and after the Tōhoku event. The β_p coefficients will estimate the differential effect of the Tōhoku event on Japanese multinational affiliates in the U.S., compared to the control group of non-Japanese firms. A useful interpretation of the $\{E_p \mathrm{JPN}_{i,p}\}$ variables is as a set of instruments that captures the exogeneity of imports during these months, reflecting the source-country share of imports from Japan as evident in Table 2.2. To evaluate the differential impact on production for Japanese firms, we simply replace the dependent variable in equation (2.5) with the firm's

[15]We consider Japanese and non-Japanese intermediate imports separately in section 2.4.

North American exports, denoted $V_{i,t}^{NA}$.

It is important to highlight that equation (2.5) is in levels. There are several reasons for doing so, as opposed to using log differences or growth rates. First, allowing for the presence of zeros is important when the data are at a monthly frequency, particularly given the magnitude of the shock to imports for Japanese firms. The second reason is more conceptual. Because we are interested in calculating the average effect of these firms that represents (and can scale up to) the aggregate impact on the U.S. economy, it is appropriate to weight the firms based on their relative size. The levels specification does exactly this: the absolute deviations from trend will be greater for the bigger firms and hence will contribute disproportionately to the coefficient estimates.[16] In section 2.4, we evaluate this framework with the results one would obtain when estimating the effect on a firm-by-firm basis.

In addition to the Conditional Independence Assumption highlighted earlier, the β_p coefficients are valid estimates of the mean effect for Japanese affiliates only in so far as the control group is not itself impacted by the shock. This Stable Unit Treatment Value Assumption (SUTVA) implies that general equilibrium effects or peer effects (e.g. strategic interaction) do not meaningfully effect the estimates. The share of imported inputs from Japan is low for the control group, and thus the shock is unlikely to have a measurable effect on imported inputs as a whole. We discuss strategic interaction in section 2.5.3.4.

2.3.2 Results: Total Manufacturing Sector

The top panel of Figure 2.6 plots the β_p coefficients from equation (2.5) for the months surrounding the Tōhoku event. Relative to the control group, there is a large drop in total intermediate input imports by Japanese firms in the months following the earthquake. The drop in intermediate inputs bottoms out at 4 million USD in $t = 3$ (June 2011) and the

[16] See Appendix B.3.1 for more discussion, as well as results obtained using other specifications. Importantly, in a reduced sample abstracting from zeros, a weighted regression using percentage changes directly yields estimates that are very close to those presented here.

point estimates do not return back to the pre-shock trend until month $t = 7$ (October 2011).

More interesting are the results from panel B of Figure 2.6, which looks for evidence of the production/sales impact of this shock on Japanese firms via their North American exports. The differential time-path of N.A. exports also exhibits a substantial drop following the Tōhoku event, hitting a trough of 2 million USD below baseline in $t = 2$ (May 2011). The standard errors, which are clustered at the firm level, are themselves interesting. As made clear via the 95-percent confidence bands on the point estimates of Figure 2.6, the standard errors increase dramatically in the months following the shock, a feature we interpret to reflect heterogeneous incidence and timing of the shocks (as well as the recoveries) for the Japanese multinationals.

To gain a sense of the average percentage drops of these two data series for Japanese multinationals as a group, we take the two plots of the differential dollar amounts from Figure 2.6 and divide by the average pre-shock level for these firms (see Table 2.2). The results, plotted jointly in Figure 2.7, show the fraction below pre-shock trend levels for these firms, on average. There is a remarkable correlation between these two series – whereby there is essentially a one-for-one drop in output for a given drop in intermediate imports. Using the mapping from Result 1, these reduced form results suggest a production function that is essentially Leontief in the imported input.

One potential concern with the interpretation of these results is separating out the intermediate input channel with other channels, such as a direct "productivity shock" affecting the U.S. operations of Japanese affiliates. Separating an ownership channel from an imported input channel is difficult due to lack of substantial overlap we identified above: few Japanese firms have low input exposure and few non-Japanese firms have high input exposure. In appendix B.3.2 we present results using a binary response model to disentangle the defining features of the import and output disruptions during this time.

2.4 Structural Estimation of Cross Country Input Linkages

The relative movements of imported inputs and output of Japanese multinational firms point to little substitutability of Japanese intermediate inputs. In this section we expand our analysis by structurally estimating the production function of firms affected by the Tōhoku shock. Unlike in the previous section, which used a set of instruments related to the differential import share of intermediates coming from Japan, this estimation relies on leveraging the high degree of exogenous variation in Japanese inputs coming from the Tōhoku event, while also fully specifying the production function under study. This estimation serves multiple purposes. First, it is reassuring to find elasticities that are consistent with the heuristic evidence implied by our reduced-form results, when imposing a conventional production function framework. Second, by imposing additional structure, we are able to distinguish two elasticities: one between Japanese material inputs and other material inputs, and another between an aggregate bundle of material inputs and domestic capital and labor. Finally, by using an estimation procedure not relying on a control group we can obtain separate estimates for Japanese and non-Japanese firms. The results corroborate the claim that the supply chains of Japanese and non-Japanese exhibit different degrees of rigidity.

The estimation procedure will utilize information from two distinct periods: the six months preceding and the six months following the March 11 event. The pre-period, which we denote by $\tau - 1$, yields information on the production function of the firm under profit-maximizing conditions. In the post-period, denoted τ, we do not impose that the firm is optimizing over its input use, due to the fact that shipments from Japan are to some extent beyond the control of the firm.

2.4.1 Framework

We assume that the firm's technology in any period t is given by the nested CES aggregate

$$x_{i,t} = \phi_i \left[\mu_i^{\frac{1}{\zeta}} \left(K_{i,t}^{\alpha} L_{i,t}^{1-\alpha} \right)^{\frac{\zeta-1}{\zeta}} + (1-\mu_i)^{\frac{1}{\zeta}} M_{i,t}^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}}, \tag{2.6}$$

where

$$M_{i,t} = \left(\nu_i^{\frac{1}{\omega}} \left(m_{i,t}^{-J} \right)^{\frac{\omega-1}{\omega}} + (1-\nu_i)^{\frac{1}{\omega}} \left(m_{i,t}^{J} \right)^{\frac{\omega-1}{\omega}} \right)^{\frac{\omega}{\omega-1}}. \tag{2.7}$$

In this production function $x_{i,t}$, $K_{i,t}$, and $L_{i,t}$ denote the output, capital, and labor of firm i. The variable $M_{i,t}$ denotes an aggregate of intermediate inputs consisting of materials sourced from Japan ($m_{i,t}^{J}$) and materials sourced from all places other than Japan ($m_{i,t}^{-J}$), including domestic materials. We are interested in estimating ω and ζ, which parameterize the substitutability between Japanese and non-Japanese materials and that between the capital-labor aggregate and the aggregate of intermediate inputs. The parameters μ_i and ν_i are firm-specific weights and ϕ_i parameterizes the firm's productivity, all of which we assume are constant over the short time horizon we consider. Further, we assume that the firm is monopolistically competitive and faces a CES demand function

$$p_{i,t}^{x} = \left(\frac{Y_{i,t}}{x_{i,t}} \right)^{\frac{1}{\varepsilon}}. \tag{2.8}$$

As usual, $Y_{i,t}$ is the bundle used or consumed downstream and serves as a demand shifter beyond the control of the firm.

2.4.1.1 Pre-Tsunami period

Period τ corresponds to the period April-September 2011, and $\tau-1$ the period September 2010 - February 2011. We exclude the month of March 2011. In period $\tau-1$ the firm operates in a standard environment, choosing capital, labor, and materials so as to maximize

$$p^x_{i,\tau-1} x_{i,\tau-1} - w_{\tau-1} L_{i,\tau-1} - R_{\tau-1} K_{i,\tau-1} - p^{-J}_{i,\tau-1} m^{-J}_{i,\tau-1} - p^{J}_{i,\tau-1} m^{J}_{i,\tau-1}$$

subject to (2.6), (2.7), and (2.8). The firm takes all factor prices as given. Material prices $p^J_{i,\tau-1}$ and $p^{-J}_{i,\tau-1}$ are firm-specific to indicate that different firms use different materials.

It is straightforward to show that this optimization problem implies

$$K_{i,\tau-1} = \frac{\alpha}{1-\alpha} \frac{w_{\tau-1} L_{i,\tau-1}}{R_{\tau-1}}, \tag{2.9}$$

$$\nu_i = \frac{\left(p^{-J}_{i,\tau-1}\right)^{\omega} m^{-J}_{i,\tau-1}}{\left(p^{J}_{i,\tau-1}\right)^{\omega} m^{J}_{i,\tau-1} + \left(p^{-J}_{i,\tau-1}\right)^{\omega} m^{-J}_{i,\tau-1}}, \tag{2.10}$$

$$\mu_i = \frac{\left(\left(\frac{R_{\tau-1}}{\alpha}\right)^{\alpha} \left(\frac{w_{\tau-1}}{1-\alpha}\right)^{1-\alpha}\right)^{\zeta} K^{\alpha}_{i,\tau-1} L^{1-\alpha}_{i,\tau-1}}{\left(P^M_{i,\tau-1}\right)^{\zeta} M_{i,\tau-1} + \left(\left(\frac{R_{\tau-1}}{\alpha}\right)^{\alpha} \left(\frac{w_{\tau-1}}{1-\alpha}\right)^{1-\alpha}\right)^{\zeta} K^{\alpha}_{i,\tau-1} L^{1-\alpha}_{i,\tau-1}}, \tag{2.11}$$

where

$$P^M_{i,\tau-1} = \left[\nu_i \left(p^{-J}_{i,\tau-1}\right)^{1-\omega} + (1-\nu_i)\left(p^{J}_{i,\tau-1}\right)^{1-\omega}\right]^{\frac{1}{1-\omega}}.$$

We will use these relationships in the structural estimation that follows below.

2.4.1.2 Post-Tsunami period

At the beginning of period τ many firms' production processes in Japan are disrupted. Obtaining the desired amount of shipments of materials from Japan may either be prohibitively expensive or simply impossible. Modeling firm behavior in this environment therefore requires modifications to the previous setup. One possibility is to assume that the quantity of materials that firms obtain from Japan is exogenous and that firms freely choose non-Japanese materials, capital and labor. This option is unattractive for two reasons. First, due to existing contracts it is unlikely that a firm is able to adjust the quantities of non-Japanese materials, capital, and labor without costs in such a short time frame. One remedy would be to add adjustment costs to the model. Although straightforward, this approach would require us to estimate additional parameters. Second, and more importantly, the materials sourced from Japan ($m^J_{i,t}$) may not be exogenous for *every* firm. Some suppliers in Japan may have been unaffected by the earthquake and tsunami such that materials could be shipped as desired. Hence, using this approach would require us to distinguish between firms whose supply chains are disrupted and those whose are not. That is, we would have to classify firms based on an endogenous outcome.

For these reasons we prefer an alternative approach, namely to estimate the production function without specifying the full optimization problem. We only assume that in period τ, firms operate the same technologies given by (2.6) and (2.7), and that no firm adjusts its capital stock such that $K_{i,\tau} = K_{i,\tau-1}$. Conditional on knowing the time-invariant features of the production function (ϕ_i, μ_i, ν_i), we next describe an estimation procedure that allows us to find the elasticity parameters most consistent with the observed input choices and output evident in the data.

2.4.2 Estimation

Recall that we use North American exports as a proxy for a firm's output $p^x_{i,t}x_{i,t}$, with the underlying assumption that the former is proportional to the latter. We continue here in the same spirit, though we now make this assumption explicit. Let $V^{NA}_{i,t}$ be the value of North American exports at time t and define

$$\kappa_i = \frac{V^{NA}_{i,\tau-1}}{p^x_{i,\tau-1}x_{i,\tau-1}}. \tag{2.12}$$

In words, κ_i is the fraction of firm i's shipments exported to Canada and Mexico in the six months preceding the tsunami. We next make two assumptions that allow us to construct an estimation equation. First, we assume that a relationship analogous to (2.12) continues to hold in period τ, except for a log-additive error $u_{i,\tau}$. That is,

$$\ln V_{i,\tau}^{NA} = \ln \kappa_i p_{i,\tau}^x x_{i,\tau} + u_{i,\tau}. \tag{2.13}$$

The second assumption is that $\mathbb{E}\left[u_{i,\tau}|X_i\right] = 0$ where X_i is a vector of all right-hand-side variables. Setting the conditional mean of $u_{i,\tau}$ to zero is a standard exogeneity assumption requiring that, loosely speaking, the error is uncorrelated with all right-hand-side variables. It rules out, for example, that in response to a fall in Japanese intermediate imports firms export a fraction of their shipments to Canada and Mexico that systematically differs from κ_i. We provide evidence in section 2.5.3.1 that demonstrates that this is a reasonable assumption.

Using equation (2.6) we can rewrite (2.13) as

$$\ln\left(V_{i,\tau}^{NA}\right) = \ln\left(\kappa_i \phi_i\right) + \ln\left(p_{i,\tau}^x \left[\mu_i^{\frac{1}{\zeta}}\left(K_{i,\tau}^{\alpha} L_{i,\tau}^{1-\alpha}\right)^{\frac{\zeta-1}{\zeta}} + \left(1-\mu_i\right)^{\frac{1}{\zeta}}\left(M_{i,\tau}\right)^{\frac{\zeta-1}{\zeta}}\right]^{\frac{\zeta}{\zeta-1}}\right) + u_{i,\tau}. \tag{2.14}$$

Values for ν_i and μ_i are obtained from equations (2.10) and (2.11).[17] Using (2.12), the intercept can be constructed from the previous period

$$\kappa_i \phi_i = \frac{V_{i,\tau-1}^{NA}}{p_{i,\tau-1}^x \left[\mu_i^{\frac{1}{\zeta}}\left(K_{i,\tau-1}^{\alpha} L_{i,\tau-1}^{1-\alpha}\right)^{\frac{\zeta-1}{\zeta}} + \left(1-\mu_i\right)^{\frac{1}{\zeta}}\left(M_{i,\tau-1}\right)^{\frac{\zeta-1}{\zeta}}\right]^{\frac{\zeta}{\zeta-1}}}.$$

Notice that κ_i and ϕ_i are not separately identified. Under standard assumptions, we can consistently estimate equation (2.14) using, e.g., nonlinear least squares. The only parameters to calibrate are the rental rate of capital R_τ and the capital share in the capital/labor aggregate α. We estimate the two elasticities, ζ and ω. Notice that ω appears in the intermediate aggregate $M_{i,\tau}$ as shown in equation (2.7). The estimates $(\hat{\zeta}, \hat{\omega})$ solve $\min\limits_{\{\zeta,\omega\}} \sum_{i=1}^{N} (u_{i,\tau})^2$.

[17] After constructing μ_i according to equation (2.11) we average by industry to reduce the level of noise.

Why do we restrict the sample to the year surrounding the Tōhoku event? To understand this, recall that the principal difficulty of estimating production functions lies in unobserved inputs and productivity. Since both are unobserved by the econometrician, they are absorbed into the error term. However, because they are known to the firm, other input choices depend on them. Hence, right-hand-side variables and the error term will generally be correlated, rendering estimates inconsistent.[18]

By restricting the sample period to a single year, the assumption of constant firm productivity seems appropriate. If productivity is constant, it cannot be correlated with the error term, thereby ruling out one of the concerns.[19] The fact that the Tōhoku event was an unexpected shock negates much of the concern about endogeneity arising from unobserved inputs. To see why, consider the case when the firm anticipates a supply chain disruption in a future period. Firm adjustment of unobserved inputs in expectation of this shock will impact input choices – leading to an endogeneity problem where inputs are correlated with the shock. Put simply, the unexpected nature of the Tōhoku event works towards equalizing the information sets between the econometrician and the firm because factor choices are not affected prior to the shock being realized.[20]

Before turning to the data we briefly discuss the intuition of parameter identification. Unlike other approaches to estimating elasticities of substitution (e.g. *Feenstra et al.* (2014)), our method does not rely on the response of relative values to a change in relative prices.[21] In fact, in an econometric sense, our approach treats all inputs as *independent* variables.

A simple example illustrates how the parameters are identified. Consider the production

[18]This problem is discussed in greater detail in, for example, *Ackerberg et al.* (2006).

[19]Of course, the size and exogeneity of the shock also helps with this concern: any idiosyncratic productivity movements during this time are surely subsumed by the earthquake/tsunami.

[20]An unobserved input that could remain operative in our case is that of factor utilization. Since the scope for substantial adjustment along this dimension seems quite limited, we remain confident that our estimates would be robust to the inclusion of this missing ingredient.

[21]Given that we observe little systematic variation in prices (see section 2.5.3.4), we believe that our approach is more appropriate in this setting.

function (2.6) and suppose that, for a particular firm, the initial period yields a value of $(1-\mu) = 0.4$. The elasticity ζ determines how deviations from this measure of the optimal input mix between the intermediate aggregate $M_{i,\tau}$ and the capital labor aggregate translate into measured output. Thus, if we observe comparatively fewer intermediates $M_{i,\tau}$, reflecting a different mix of inputs than that given by 0.4, we obtain an elasticity estimate for ζ that best matches the response in output. Because the estimates for μ, ν, and $\kappa_i\phi_i$ are themselves functions of the elasticities, this procedure must iterate across the parameter space to find the estimate most consistent with the data. Similar reasoning applies for the identification of the ω elasticity based on relative movements in Japanese materials, non-Japanese materials, and output. The estimates we obtain are the best fit across the firms in each sample.

2.4.3 Connecting Model and Data

Estimation of the model requires data on employment, Japanese and non-Japanese material inputs, as well as on exports to North America and output prices for periods $\tau - 1$ and τ. Since data on firm-specific capital stocks are hard to obtain and likely noisy, we use equation (2.9) to construct it from firm payroll and a semi-annual rental rate of 7 percent for period $\tau - 1$.[22] Recall that the capital stock is not adjusted over this time horizon so that $K_{i,\tau} = K_{i,\tau-1}$. The parameter α is calibrated to 1/3.[23]

Quarterly employment information comes from the Business Register, which we adjust to reflect the average value over the 6 month periods we study, as they do not align with the quarters defined within a calendar year.[24] As discussed in earlier sections, the LFTTD contains firm-level data of Japanese imports and North American exports. For non-Japanese

[22] This comes from assuming a real interest rate of 4 percent, combined with an annual depreciation rate of 10 percent, and then adjusting for a semi-annual frequency. The estimates are insensitive to alternative values of the rental rate.

[23] In principle it is possible to construct a firm-specific value for α, using value-added information available in a census year. We are currently exploring the feasibility of this option.

[24] Specifically: $L_{\tau-1} = \frac{1}{6}Emp_{2010Q3} + \frac{1}{2}Emp_{2010Q4} + \frac{1}{3}Emp_{2011Q1}$ and $L_{\tau} = \frac{1}{2}Emp_{2011Q2} + \frac{1}{2}Emp_{2011Q3}$.

material inputs, we would ideally combine the non-Japanese imported materials with information on domestic material usage for these firms. As information on domestic material inputs is not available in Census data at this frequency, we utilize information on the *total* material expenditures from the Census of Manufacturers (CM) to construct a firm-level scaling factor to gross up non-Japanese intermediate imports. Put differently, we impute non-Japanese material inputs from non-Japanese input imports. For each firm, we construct the scaling factor as

$$\frac{P_i^M M_i - p_i^J m_i^J}{p_i^{-J} m_i^{-J}} \tag{2.15}$$

from the latest CM year. Because the closest available CM year is 2007 in our data, there is some concern about missing or outdated information for this factor. We mitigate this by using industry-specific means for missing values, and winsorizing large outliers at the 90th/10th percentiles.

Regarding information on prices, the LFTTD records the value and quantity of each trade transaction (at the HS10 level), and thus it is possible to construct the associated price, or "unit-value" of each shipment directly.[25] Aggregating up these shipments into a firm-month observation is complicated, of course, by the differing quantity units. Lacking any better alternative, we simply average the transaction prices using the dollar value of each transaction as weights.

Finally, we restrict the sample of firms to those that have regular imports from Japan and non-Japan over the periods we study, as well as regular North American exports.[26] While this substantially limits the number of firms in each sample, the shares of trade represented by these firms in each category remains very high (see Table 2.3).

We obtain standard errors using bootstrap methods, which also allow us to account for the

[25]Those transactions with missing or imputed quantity information are dropped. Future efforts will evaluate whether it is possible to recover the quantity values from prior transaction details.

[26]Specifically, we drop any firm that has more than 3 months of zeros for any of these values, over the period $\tau - 1$ or the period τ.

uncertainty implied by the imputation of non-Japanese material inputs. We draw randomly with replacement from our set of firms to construct 5000 different bootstrap samples. For each of these samples, the non-Japanese materials share is imputed as described above before the estimation proceeds.

2.4.4 Summary of Results

The results of the estimation are shown in Table 2.3. The elasticity between material inputs for Japanese affiliates is 0.2, while the elasticity between the aggregate material input and capital/labor is 0.03. Together, these estimates are indeed consistent with the reduced-form evidence for the (ψ) elasticity from section 2.3.2. The relative magnitudes are also intuitive: while Japanese imported inputs are strong complements with other material inputs — consistent with the high share of intra-firm transactions comprising this trade — there is even less scope for substitution between material inputs and domestic capital/labor.

The estimation procedure also allows us to estimate these elasticities for two samples of non-Japanese firms: non-Japanese multinationals and non-multinational firms. While the estimates for the ζ elasticity are indeed very close for these other samples, the elasticity estimates corresponding to material inputs are higher, at 0.6 and 0.4 respectively. The lower share of intra-firm imports from Japan for the non-Japanese multinationals aligns with the argument that this type of trade is the key source of non-substitutability in the short-run. On the other hand, the low estimates for non-multinational firms, which have essentially zero intra-firm imports, may point to other mechanisms at work beyond the role of intra-firm trade. More generally, however, the estimates for these parameters are all significantly lower than what is commonly assumed (typically unity or higher) in the literature.

Although the number of firms included in this estimation is small (550 firms in total across the three subgroups) , they account for a large share of economic activity in the United States. Looking at their combined share of total trade, these firms account for over

80% of Japanese intermediate imports, 68% of non-Japanese intermediate imports, and well over 50% of North American exports. Such high concentration of trade among relatively few firms is consistent with other studies using this data (see *Bernard et al.* (2007)).

2.5 Discussion

The structural estimates of the model are broadly in agreement with the evidence in section 2.3.2: imported inputs are strong complements with other inputs in the production function. The rigidity of the production function for multinational firms in particular is likely due to i) the high degree of intra-firm trade in what is presumably highly specialized inputs, and ii) .[27] Our results have a number of important implications for how economists should think about multinational firms in general, as well as aggregate topics such as volatility and business cycle co-movement.

2.5.1 Aggregation

Before relating our estimates to macroeconomic topics, it is important to discuss aggregation. Indeed, in any study utilizing micro-level estimates to inform macro-level objects of interest, the details of aggregation and heterogeneity are of critical importance. Work by *Imbs and Méjean* (2011) argues that imposing homogeneity across sectors when estimating consumption elasticities can be overly restrictive, creating a heterogeneity bias which can be quantitatively large. In our case one could discuss aggregation along various dimensions: across products, industries, firms, and so on. We examine the effects of product-level aggregation in section 2.5.3.3 below.

A primary concern is how to translate the results from the firm-level subsamples into estimates that would pertain to macro-oriented models. As a first step, the final column

[27] The vertical integration of production across countries, within the firm, has shown to be a key driver of the decline in joint ventures (see *Desai et al.* (2004)).

in Table 2.3 shows the elasticity estimates when aggregating across all firms in the sample. The results are consistent with the estimates by subgroup, suggesting substantial complementarities across inputs. All estimates in Table 2.3, however, correspond to the average across firms in each group, and do not take into account heterogeneity in firm size within the groups. It is relatively straightforward to modify our estimation procedure to weight firms according to their relative size.[28] We report the results from this modified estimation in Panel B of Table 2.4. When comparing the results to those in Table 2.3, it is evident that the weighted estimates are not substantially different than the unweighted estimates. Although the samples of firms comprising these estimates do not amount to the total manufacturing sector of the United States, they do account for the considerable majority of U.S. trade.

2.5.2 Implications

The rigid production networks of foreign-owned multinationals will have direct consequences on the destination (host) economy. Previous literature has hypothesized that input linkages could generate business-cycle comovement, but supportive empirical evidence has been difficult to find. This paper can be seen as a first step in establishing empirical evidence for a causal relationship between trade, multinational firms, and business cycle comovement. In a companion paper (*Boehm et al.* (2014b)), we evaluate the quantitative importance of such complementarities of imported inputs by multinational affiliates. When separately accounting for intermediate input trade by multinationals and traditional trade in final goods, the model distinguishes between the production elasticity of imported inputs and the traditional "Armington" elasticity used to bundle together international goods for consumption. The complementarities in import linkages by multinationals increases value-added comove-

[28]Since the appropriate measure of size in our context is output, we follow our convention and use the relative amounts of North American exports in the period before the shock as the weights.

ment in the model by 11 percentage points relative to a benchmark without such firms.

This model shares similarities with several other existing models, particularly *Burstein et al.* (2008). A key advantage of *Boehm et al.* (2014b), however, is a tight link to Census data for matching other features of multinationals and trade. *Johnson* (2014) also looks at the role of vertical linkages on comovement, but applies greater input-output structure on the model. Such features will generate increases in value-added comovement in his model, the magnitude of which becomes significant only when the elasticity of substitution among inputs is sufficiently low. Other work also identifies multinationals as a key source of the transmission of shocks: *Cravino and Levchenko* (2014) demonstrates that foreign multinational affiliates can account for about 10 percent of aggregate productivity shocks.[29]

The low value for ω indicates the presence of spillovers beyond the immediate effect from Japan. That is, imports from non-Japanese locations are lower as a result of the shock in Japan[30], and we would presume this applies to suppliers within the United States as well. Specifically, upstream suppliers (in countries other than Japan as well as within the U.S.) were affected indirectly due to the exposure of Japanese affiliates to the shock combined with the rigidity of their production with respect to those inputs. Downstream suppliers that rely on the inputs from the disrupted firms would likewise be adversely affected. The presence of such spillovers combined with the large network of input linkages can indeed magnify the total effect of the transmission of the shock to the U.S. market. Such effects are also evident in a related paper, *Carvalho et al.* (2014), which finds large spillovers in both upstream and downstream firms in Japan following the 2011 earthquake.

Another branch of literature on the diversification of risk has studied whether firms using complex production structures with several intermediates could be less volatile (*Koren and*

[29]Of course, shocks can be passed through to affiliates through other means as well. See *Peek and Rosengren* (1997) and *Peek and Rosengren* (2000) for the case of U.S. affiliates from Japan.

[30]To confirm this, see Figure **??**, which replicates the results in Figures 2.6 (Panel A) and 2.7, but only for non-Japanese imports.

Tenreyro (2013)). *Kurz and Senses* (2013) establish that firms with substantial imports and exports have lower employment volatility than domestic firms in the medium to long term, which they attribute partly to the diversification of risk.[31] The key result in this paper points to a possibly overlooked fact: the extent of the benefits from diversification depends heavily on the substitutability of inputs. Conditional on a given number of inputs used in production, a firm will likely experience greater volatility if each input is key to the production process and inputs are subject to heterogeneous shocks.[32] Conceptually, an increase in the use of imported inputs should not be viewed necessarily as diversification. A fragmentation of production can lead to an increased supply chain risk that is an important counterweight to whatever efficiencies such complex input sourcing might afford, particularly when the production elasticities are low.

The rigid production networks of multinational firms also influences our understanding of why firms segment production across country borders. In a related paper, **?** shows that despite the presence of substantial and complex import linkages with the source country (consistent with a vertical framework of FDI), the motive for multinational production appears to be to serve the domestic market (consistent with the horizontal framework of FDI). The result could be called "horizontal FDI with production sharing."[33] The evidence for strong complementarities in this production sharing, however, presents a puzzle. Why does the firm replicate only select portions of the supply chain, considering the penalties for disruptions and mismatched inputs are so great? It is perhaps the case that the segments of the production chain that remain in the source country have a location-specific component

[31] An interesting result from *Kurz and Senses* (2013) is that firms that only import are actually more volatile than the domestic-only benchmark.

[32] *Pravin and Levchenko* (2014) outline theoretical results showing that for a given elasticity value (in their case, Leontief), volatility in output per worker should be actually decreasing in the number of inputs used.

[33] *Ramondo et al.* (2014) is another recent example arguing for a more nuanced interpretation of multinational production.

that is not easily transferable when the firm moves production abroad.[34] Understanding the dynamics behind these sourcing decisions is an area in need of further research.[35]

2.5.3 Robustness and Extensions

2.5.3.1 Mis-measurement of Firm Production

A natural concern with our analysis is the use of N.A. exports as a proxy for firm-level production. Perhaps it is the case that shipments abroad fall disproportionately more than domestic shipments following a shock to production. If this were the case, the N.A. exports would indeed be a poor proxy for production, and its usefulness in evaluating a production elasticity substantially compromised.

To evaluate this concern, we narrow our study to the automotive sector, which has data on production, sales, and inventory at a monthly frequency. Using the Ward's electronic databank, which reproduces the published series in the annual Automotive Yearbook, we obtain plant-level information on production, and model-line information on inventory, sales, and incentives.[36] The baseline specification is the same as in equation (2.5), where the dependent variable is now Q_{jit}: production of plant j of firm i in month t. The Japanese multinational firms are, in this case, those automakers with plants located within North America but whose parent company is headquartered in Japan.[37]

Figure 2.8 shows the results, where we once again divide by pre-shock levels to gain a sense of the percentage effects of these changes. Relative to their U.S. counterparts, Japanese automakers in the United States experienced large drops in production following the Tōhoku

[34] The model of knowledge sharing in *Keller and Yeaple* (2013) is one attempt to analyze the dynamics between such transfers being accomplished in embodied (intra-firm trade) or disembodied (direct communication) form. Alternatively, domestic content requirements may provide incentives to produce specified inputs in one location over another.

[35] For a recent example of how such investment and sourcing decisions can alter a country's comparative advantage over time, see *Alviarez* (2014).

[36] Appendix B.3.7 details further features of this data and explains how the sample was constructed.

[37] These firms are Honda, Mitsubishi, Nissan, Toyota, and Subaru.

event. Production bottomed out in May of 2011 — two months after the event — at almost 60 percent below trend. [38] The point estimates return to a level near zero in September of 2011, implying that the shock affected production for nearly 6 months.[39] We interpret these results to be largely supportive of the results obtained using the exports-based proxy for production. The percentage drops in the two series are remarkably similar: a trough of 59% at $t = 2$ in the automotive data vs 53% at $t = 2$ using the proxy. We conclude that, at least for this exercise, the proxy appears to be providing valuable information on a firm's U.S. production behavior.[40]

2.5.3.2 Intermediate Input Inventories

Inventories are another obvious feature that should influence the relationship between input shipments, production, and the elasticity of substitution. In particular, inventories of intermediate inputs allow the firm to absorb unforeseen shocks to input deliveries without an impact on the production process.[41] As it relates to the production elasticity, however, the presence of these inventories should serve to diminish or delay the production impact, thereby *increasing* the elasticity relative to what it would be without such inventories.

In fact, it is striking the extent to which we do not see any evidence for the role of

[38]The average monthly plant-level production at these firms during December 2010 through February 2011 was about 12,200 units a month. The magnitude of the drop in May was -7200 units.

[39]We describe additional results on the behavior of inventories, sales, incentives, and production in Japan in an appendix.

[40]In addition, one might be concerned that the N.A. exports series may be contaminated with Japanese imports whose country of ultimate destination is Canada/Mexico (a.k.a "in-transit shipments" – imports to Canada/Mexico via U.S.). These shipments should not be picked up in the reporting systems underlying the LFTTD. According to section 30.2(d)(1) of the U.S. Code of Federal Regulations, "In-transit shipments of goods from one foreign country to another where such goods do not enter the consumption channels of the United States are excluded from filing the Electronic Export Information (EEI)." Additionally, the Army Corps of Engineers has suspended the requirement to file the Form 7513, Shippers Export Declaration (SED) for In-transit Goods leaving the United States via vessel. Finally, the corroborating results from section 2.5.3.1 should also serve to allay such concerns.

[41]The existence of final good inventories, on the other hand, makes a distinction between the production and sales of a particular product. Here, the presence of final good inventories implies that the firm can continue to sell from existing inventory stocks even while production is temporarily affected.

intermediate input inventories in the production impacts of Figure 2.6 (Panel B) or Figure 2.8. The effect on production appears to be almost immediate, indicating that the stock of inventories of imported intermediates is low (less than one month's supply) for these firms.

We obtain a rough sense of the degree of inventory holdings from the Census of Manufacturers micro-data. Combining information on the beginning period stock of materials inventories with the annual usage of materials, we calculate the average monthly supply of inventories for each firm.[42] Panel A of Table 2.1 calculates the production-weighted averages over a select set of firm groups.[43] We see that on average, Japanese multinationals hold a little over 3-weeks supply of intermediate inputs as inventory. This is slightly less than non-multinational firms, a fact that aligns with the oft-cited "lean" production processes made famous by Japanese firms in previous decades. Though these data are for the year 2007, there is little reason to believe these relative magnitudes have changed substantially over a period of a few years. For completeness, Panel A of Table 2.1 also reports the corresponding estimates for output inventories.[44]

Low inventory holdings combined with an inelastic production function suggests that firms are willing to tolerate some degree of expected volatility in their production. Either the costs of holding inventories or diversifying sources of supply are sufficiently high, or firms believe the probability of disruption is low. In either case, these lean production strategies carry a greater potential for the propagation of shocks across countries, perhaps affecting firms with limited knowledge of their indirect exposure through complicated production chains.

[42] Unfortunately, the CM data does not report imported materials inventory separately.

[43] These numbers are broadly similar, though somewhat lower than other estimates in the literature. See *Ramey* (1989) for one example.

[44] At first glance, the average monthly supply of these output inventories looks surprisingly low. On the other hand, it is probably the case that inventories are held jointly by the manufacturer and wholesale/retail establishments Thus, considering the inventories of manufacturers alone could potentially under-represent the "true" level of output inventories available for smoothing out production disturbances.

2.5.3.3 Multi-Products and Sub-Optimal Mix

In the frameworks used in sections 2.3.1 and 2.4, we consider the aggregate bundles of imported intermediates, abstracting away from product-level detail. In reality, the firms in our dataset often import many distinct intermediate inputs from Japan. The structure of a CES production function implies that if each of these within-country inputs was non-substitutable with one another (a further, nested Leontief structure), the production impact of a disruption in the supply of just one input could be amplified relative to the value of that input.[45] We evaluate this possibility below.

This is particularly true given the heterogeneous impact of the Tōhoku event across Japan (see Figure 2.4). This could translate into considerable dispersion in the impact on the *products* imported by a particular U.S. firm or Japanese affiliate. With product-level shocks, considering the effect on the aggregate import bundle amounts to assuming either 1) perfect substitutability among products, or 2) that the firm maintains an optimal within-country product mix at all times.

To be concrete, it may be more accurate to view the M_t in equation (2.1) as a further C.E.S. function of multiple products. Thus, we can define the proper measurement of this variable as

$$V_{i,t}^{M} = P_{i,t}^{M} \left(\sum_{n=1}^{N} \eta_n^{\frac{1}{\chi}} (m_{n,i,t}^{J})^{\frac{\chi-1}{\chi}} \right)^{\frac{\chi}{\chi-1}}, \tag{2.16}$$

where now $V_{i,t}^{M}$ is the value based on a combination of N distinct products, with weights η_n and elasticity χ.

Product-level heterogeneity in the production impact of the shock combined with imperfect coordination among input suppliers implies that the aggregate (measured) import bundle for a particular firm may turn out to be suboptimal. In this case, we are measuring $\widehat{V^M}_{i,t} = \sum_{n=1}^{N}(p_{n,t}^{m} m_{n,t}) \geq V_{i,t}^{M}$. And the lower the elasticity of substitution among products,

[45] This point has been made in somewhat differing contexts, by *Kremer* (1993) and *Jones* (2011).

the more severe the disconnect between the measured imports and the "effective" imports — that which is actually useful in downstream production.

A suboptimal product mix indicates that measured imports ($\widehat{V^M}_{i,t}$) are greater than the effective imports ($V^M_{i,t}$). As a result the measured output response to the import shock will be larger than otherwise, resulting in a downward "bias" in the elasticity estimates from section 2.3.1 and 2.4.[46] Such an effect is decreasing in the product-level elasticity parameter χ, as complementarity itself is the driving force between differences in $\widehat{V^M}_{i,t}$ and $V^M_{i,t}$. In addition, the effect is also increasing in the degree of deviation from the optimal product mix.

Does this exert a quantitatively large effect on our point estimates? Given the emphasis on low inventories and lean production processes in downstream operations, one might expect that across-product adjustment would take place before sending the inputs abroad. To analyze this empirically, we analyze whether there are significant deviations in the product composition of Japanese imports during the months following the Tōhoku event. To do this, we construct a measure of the distance of a firm's import bundle from a benchmark, which we will interpret to be the optimal bundle. Let $t = s^*$ be such a benchmark date. Then, using the product-level information in the LFTTD data we construct for each firm i, the share of total imports from Japan for a given product code n. Defining this share to be $s_{n,i,t}$, we then construct the average product-level distance from optimum $DO_{i,t}$ as:

$$DO_{i,t} = \frac{1}{N^i} \sum_{n=1}^{N^i} \left(|s_{n,i,t} - s_{n,i,s^*}| \right) \tag{2.17}$$

where N^i is the total number of products imported by firm i. We define the period s^* to be the months of April-June of 2010, and then evaluate DO_i at a monthly frequency, with particular interest in the months following the Tōhoku event. While there may be

[46]Because the source of this downward pressure on the estimate for ψ (or ω) is itself a very low product-level elasticity, it is unclear whether this should be considered a bias in the traditional sense.

natural movements in the bundle of products imported from Japan, evidence for substantial coordination failure in product composition or heterogeneity in product-level shocks would come from any abnormal jumps in this index in the months of the disruption. One can calculate this at various levels of product aggregation (i.e. HS4, HS6, HS8, HS10), though we report results using the HS6 level.[47]

The results of this exercise are shown in Figure 2.9. We plot the average DO_i across Japanese firms for each month (the figure shows a 3-month moving average) during the period 2009-2011. Mechanically, this measure should be relatively close to zero in the months consisting of the benchmark (April-June 2010). While there is a secular rise in this measure on either side of this benchmark period, there do not appear to be any large jumps in the months directly following the Tōhoku event. More interesting, perhaps, are the considerably larger values for this measure during early 2009, which might reflect the effects of the trade collapse associated with the Great Recession. We interpret Figure 2.9 as evidence that the potential for suboptimal mix across products from Japan does not pose a serious problem to our measurement in previous sections.

2.5.3.4 Other Considerations

Strategic Behavior: Another possibility that could affect the interpretation of the results from Figure 2.7 might be strategic behavior, particularly on the part of the competitors of Japanese firms in the United States. These firms could raise production or prices following the negative supply shock affecting their competitors, which would serve to bias downward the β_p coefficients from the equation with $X_{i,t}^{NA}$ as the dependent variable.[48] To evaluate this

[47]The level of aggregation we use attempts to balance concerns along two dimensions. With less product aggregation (i.e. HS10 level), one might be concerned with the inherent lumpiness of product-level firm imports. More product aggregation, on the other hand, could mask important product differences within a particular product grouping.

[48]Specifically, in equation (2.5) the γ_p's would be higher than would be expected without the shock, and hence the β_p's artificially low.

possibility, we turn once again to the automotive data. Here, we can look directly at the production of non-Japanese automakers in the months directly following the Tōhoku event. Appendix Figure A1 plots the relative production of these firms, using time-series variation only. There appears to be no quantitatively meaningful responses in the months following March 2011. This should not come as a surprise given capacity constraints and utilization adjustment costs, particularly given the short time horizon. We provide evidence on the role of prices next.

Prices: Traditionally, estimating the elasticity of substitution is accomplished via price and quantity data for products over extended periods of time. For the short horizon we consider in this paper, there are several reasons why prices may not have the scope to adjust. Many supplier relationships negotiate prices for longer periods of time than one or two months. Second, and perhaps more importantly, Table 2.2 demonstrates that the large majority of imported intermediate inputs are intra-firm. The observed prices of these transactions are transfer-prices (within firm) and not likely to change reflecting any short-term disturbance. However, because the LFTTD contains both quantity and price information, we can confirm whether or not prices remained relatively stable during this period. The results in Appendix Table B.6 confirm that there are few significant price movements on import or export transactions for either Japanese or non-Japanese multinationals surrounding the Tōhoku event.[49]

Domestic Inputs: It is also possible to evaluate the response of domestic inputs directly, using the limited information we have on quarterly firm-level employment and payroll information, taken from the Census Bureau's Business Register (BR).[50] We consider the evidence in Appendix B.3.5 and find no significant effects on either employment or payroll

[49]Further details on the construction of the data underlying the analysis of unit values is available in Appendix B.3.6.

[50]The BR itself receives quarterly payroll and employment information for business and organizational employers from the IRS: Form 941, the Employer's Quarterly Federal Tax Return. For more information on the BR (formerly the SSEL), see *Walker* (1997).

for Japanese firms in the quarter(s) following the shock (see Table B.5). Of course, there are a number of reasons — principally labor adjustment costs — why one would expect little, if any, impacts on employment following this short-lived shock. Press releases dispatched by the Japanese automakers during this time indicated that no layoffs would occur. Rather, the firms indicated that they would use the production stoppages for employee skill and safety training.

2.5.4 External Validity

Finally, we discuss the external validity of this result. The exogenous variation we use to identify this elasticity is tied to a particular event in time, making generalization subject to some caveats. On the other hand, there are few, if any, estimates of this parameter in the existing literature. The critical question is whether the mechanisms underlying the elasticity estimates are operative beyond the circumstances surrounding this event study.

The pattern of strong intermediate input linkages with the source country is not restricted to Japanese affiliates only. As shown in *Flaaen* (2013a), over 45 percent of the imports for all foreign multinational affiliates are sourced from the country of the parent firm. The cost share of imported intermediates from the source country is 0.12 for all foreign affiliates, which is lower than the 0.22 for Japanese affiliates but still much larger than the representative importing firm in the United States. The cost share of all imported inputs is actually quite close: 35 percent for Japanese affiliates vs 32 percent for all foreign affiliates.

A related concern is whether the estimates for Japanese affiliates are driven solely by the automotive sector. The ideal check would be to run industry-by-industry subgroup estimates for the elasticities, thereby generating heterogeneity that could be assessed relative to expectations. Unfortunately, the small number of firms applicable for this analysis, combined with disclosure requirements associated with the Census Bureau data usage, prevents this degree of detail. Instead, we address this concern by splitting the sample into a motor vehicle and non-motor vehicle subsample. We do this for the Japanese multinationals as well as the total sample of all firms. The results for these four subsamples are reported in Panel C of Table 2.4. Using the published data from the B.E.A., the automotive sector is a large but not overwhelming percentage of total Japanese manufacturing affiliates in the U.S. The entire motor vehicle sector as a whole comprises significantly less than half of value-added (roughly 40 percent) for the Japanese manufacturing affiliates.

When viewed in light of the substantial fraction of intra-firm imports comprising multinational affiliate trade, the low elasticity of substitution should not come as a surprise. One would not expect close substitutes for the sort of specialized products reflecting firm-specific knowledge that likely comprises this trade. Moreover, such a low estimate for an elasticity of this nature is not without precedent. Using different methodologies, recent work by *Atalay* (2014) highlights strong complementarities between intermediate inputs, using industry-level data for the United States.[51]

Any elasticity estimate is tied to the time-horizon to which it corresponds. *Ruhl* (2004) emphasizes the difference between elasticities implied by responses to temporary vs permanent shocks. Larger values are calculated for an elasticity following a permanent shock, owing in part to firm responses along the extensive margin. In our context, we estimate the elasticity subject to a short-lived shock where the structure of the supply chain is plausibly fixed and extensive margin movements of supplier relationships would not apply. For this reason the elasticity parameters (ω, ζ) should likely generalize to other contexts of this horizon and for shocks of this general duration. Even for a long-lived shock, the estimated elasticities would remain relevant while the firm makes changes to its network of suppliers. Evaluating whether there is evidence for long-term supply-chain reorganization following the Tōhoku event is an area of ongoing work.

[51]The point estimate for the elasticity of substitution among intermediate inputs from *Atalay* (2014) is 0.03.

2.6 Conclusions

Using a novel firm-level dataset to analyze firm behavior surrounding a large exogenous shock, this paper reveals the mechanisms underlying cross-country spillovers. We find complementarities in the international production networks of Japanese affiliates, such that the U.S. output of these firms declined dramatically following the Tōhoku earthquake, roughly in line with an equally large decline in imported inputs. The elasticity of substitution between imported and domestic inputs that would best match this behavior is very low – nearly that implied by a Leontief production function. The reliance on intra-firm imports by multinational affiliates from their source country is the most plausible explanation for such strong complementarities in production. Structural estimates of disaggregated elasticities are similarly low, and imply spillovers to upstream and downstream firms in the U.S. and abroad (non-Japan). The large impacts to Japanese affiliates together with the propagation to other U.S. firms explains the large transmission of the shock to the U.S. economy in the aggregate.

These elasticities play a critical role in the way international trade impacts both source and destination economies. Such complementarities between domestic and foreign goods have been shown to improve the ability of leading theoretical models to fit key moments of the data. We emphasize here the distinction between substitutability between domestic and foreign final goods (a "consumption" elasticity of substitution, or the so-called Armington elasticity) and substitutability between domestic and foreign intermediate goods (a "production" elasticity of substitution). In a companion paper (*Boehm et al.* (2014b)), we document the behavior of a model with such complementarities in imported intermediates, and discuss how these elasticity parameters interact. Calibrating this model to the share of multinational affiliate trade in intermediates yields an increase in value-added comovement of 11 p.p.

Such rigid production networks will also play a substantial role in aggregate volatility, productivity growth and dispersion, and the international ownership structure of production. The novel datasets described in this paper may help to shed light on these and other areas of research in the future.

Table 2.1: Summary Statistics: Imported Inputs and Inventories by Firm Type

	Japanese Multinationals	Non Multinationals
Panel A: Avg. Monthly Supply of Inventories		
Inputs	0.83	1.08
Output	0.31	0.45
Panel B: Cost Share Of Imported Inputs		
from Japan	21.8	1.0
from all countries	35.0	17.5

Source: CM, LFTTD, DCA, and UBP as explained in the text. The data are for year 2007. This table reports the average monthly supply of inventories [(usage/12)/beginning period inventory stock] for materials and output, as well as the cost share of imported products.

Table 2.2: Summary Statistics

Panel A: Cost Share Of Imported Inputs

	Japanese Firms	Non Multinationals
from Japan	21.8	1.0
from all countries	35.0	17.5

Panel B: Treatment Effects Sample Details

	Japanese Firms	Other Multinationals	Balancing Tests t	Balancing Tests $p > \|t\|$	% Reduct |bias|
N.A. Exports	3,504,894	3,413,058	0.38	0.706	79.1
share intra-firm	72.0	52.2			
Intermediate Input Imports	8,075,893	7,596,761	0.87	0.384	88
share from Japan	70.0	3.5			
share intra-firm	86.0	21.7			
Industry (Avg)	–	–	0.009	0.965	97.8

Source: LFTTD, DCA, and UBP as explained in the text.
Panel A data are for year 2007. Panel B reports the baseline average values of N.A. exports and intermediate input imports, as well as the characteristics of that trade, for the two groups of firms: Japanese affiliates and other multinational firms. The statistics are calculated in the three months prior to the Tōhoku earthquake: Dec. 2010, Jan. 2011, and Feb 2011. The control group of other multinational firms has been re-weighted using the normalized propensity score, from a specification including the level of N.A. exports, int imports, and industry dummies. The final three columns report balancing tests of the equality of the means between the treated and control group.

Table 2.3: Firm-Level Estimation: Results and Sample Details

Panel A: Calibration

Parameter	Value
R_t	0.07
α	1/3

Panel B: Estimation Results

	Japanese Multinationals	Non-Japanese Multinationals	Non-Multinationals	All Firms
ω	0.201	0.624	0.423	0.552
	[0.02 0.43]	[0.16 0.69]	[0.26 0.58]	[0.21 0.62]
ζ	0.032	0.038	0.032	0.037
	[0.030 0.673]	[0.035 0.508]	[0.029 1.68]	[0.034 0.038]
	Sample Details			
Weight on K/L Aggregate ($\bar{\mu}$)	0.223	0.514	0.278	0.409
Weight on JPN Materials ($1 - \bar{\nu}$)	0.173	0.044	0.147	0.096
Number of Firms	105	304	141	550
Share of Total Trade				
JPN int imports	0.60	0.23	0.03	0.86
Non-JPN int imports	0.02	0.66	0.01	0.69
N.A. exports	0.08	0.47	0.01	0.56

Source: CM, LFTTD, DCA, and UBP as explained in the text.
This table reports the results from the firm-level estimation detailed in section 2.4. Panel A outlines the parameters that are calibrated prior to estimation. The top two rows of Panel B reports the point estimates of the elasticities, and the corresponding 95 percent confidence intervals using a bootstrapping procedure. (See Appendix B.3.3 for more details on the measurement of dispersion for these estimates.) Rows 3 and 4 report other estimates related to the calculated production functions. The final rows of Panel B describe features of the estimation samples.

Table 2.4: Firm-Level Estimation: Other Results

Panel A: Calibration

Parameter	Value
R_t	0.07
α	1/3

Panel B: Estimation Results (Weighted)

	Japanese Multinationals	Non-Japanese Multinationals	Non-Multinationals	All Firms
ω	0.157	0.611	0.543	0.606
	[0.02 0.40]	[0.30 1.23]	[0.305 0.57]	[0.28 0.70]
ζ	0.241	0.038	0.032	0.037
	[0.03 0.884]	[0.034 0.51]	[0.029 0.55]	[0.034 0.038]
Number of Firms	105	304	141	550

Panel C: Estimation Results: MV Sector

	Japanese Mult.		All Firms	
	Motor Vehicles	Non-Motor Vehicles	Motor Vehicles	Non-Motor Vehicles
ω	0.311	0.094	0.414	0.574
	[0.019 0.398]	[0.016 0.59]	[0.27 0.60]	[0.16 0.66]
ζ	0.032	0.071	0.037	0.037
	[0.030 0.48]	[0.028 1.27]	[0.031 0.64]	[0.033 0.037]
Number of Firms	35	70	100	450

Source: CM, LFTTD, DCA, and UBP as explained in the text.
This table reports additional estimation results. Panel B recalculates the results from Table 2.3 using a vector of weights to assign larger firms a greater share in the estimation. Panel C divides the samples based on the motor vehicle industry.

Figure 2.1: Index of Japanese Industrial Production: Manufacturing Jul.2010 - Jan.2012

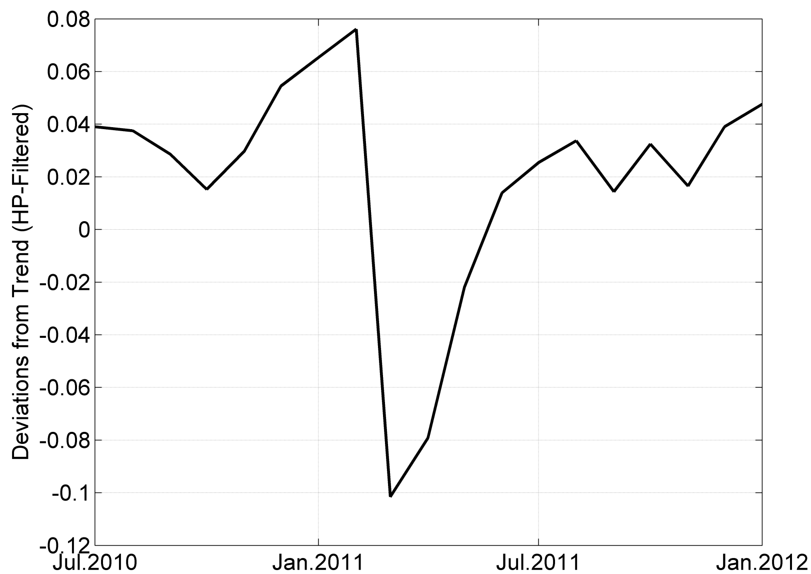

Source: Japanese Ministry of Economy, Trade, and Industry (METI). The series are logged, HP-Filtered, after seasonally adjusting.

Figure 2.2: U.S. Imports from Japan and Rest of World, Jul.2010 - Jan.2012

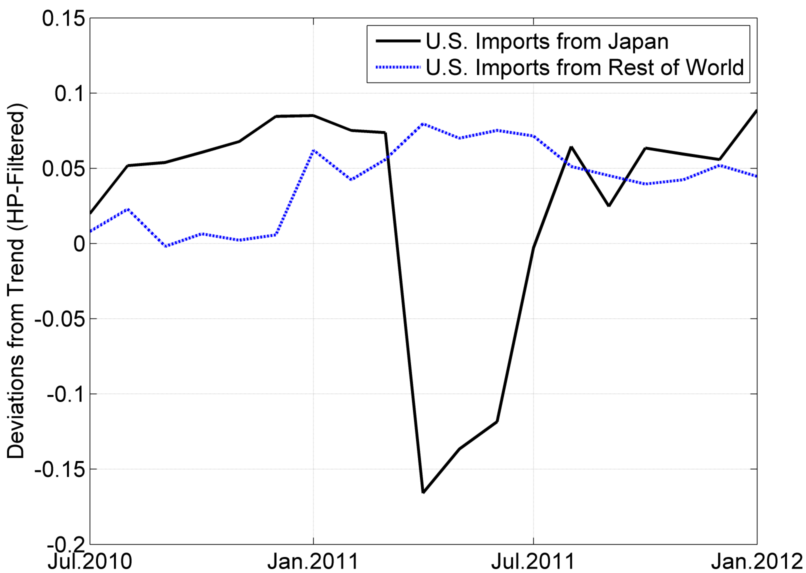

Source: U.S. Census Bureau (FT900: U.S. International Trade in Goods and Services). The series are logged, HP-Filtered, after seasonally adjusting.

Figure 2.3: U.S. Industrial Production: Manufacturing and Durable Goods

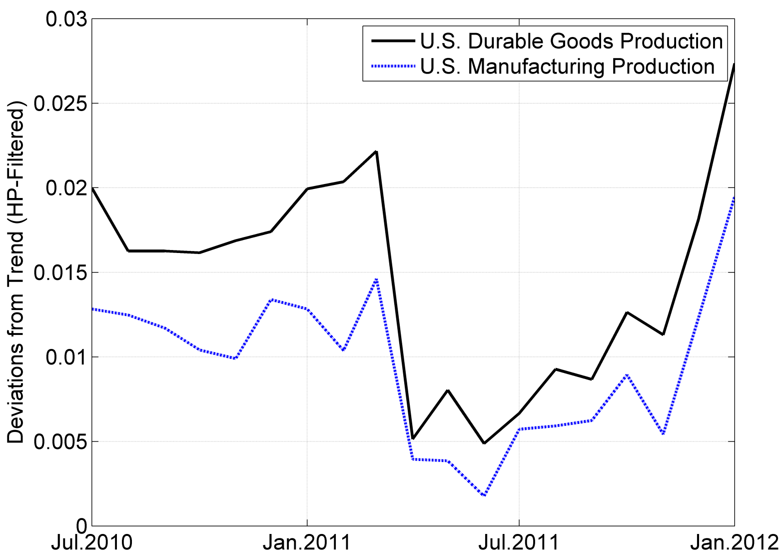

Source: Federal Reserve Board, Industrial Production and Capacity Utilization - G.17 Series . Series is Seasonally Adjusted.

Figure 2.4: Geographic Distribution of Earthquake Intensity and Affiliate Locations

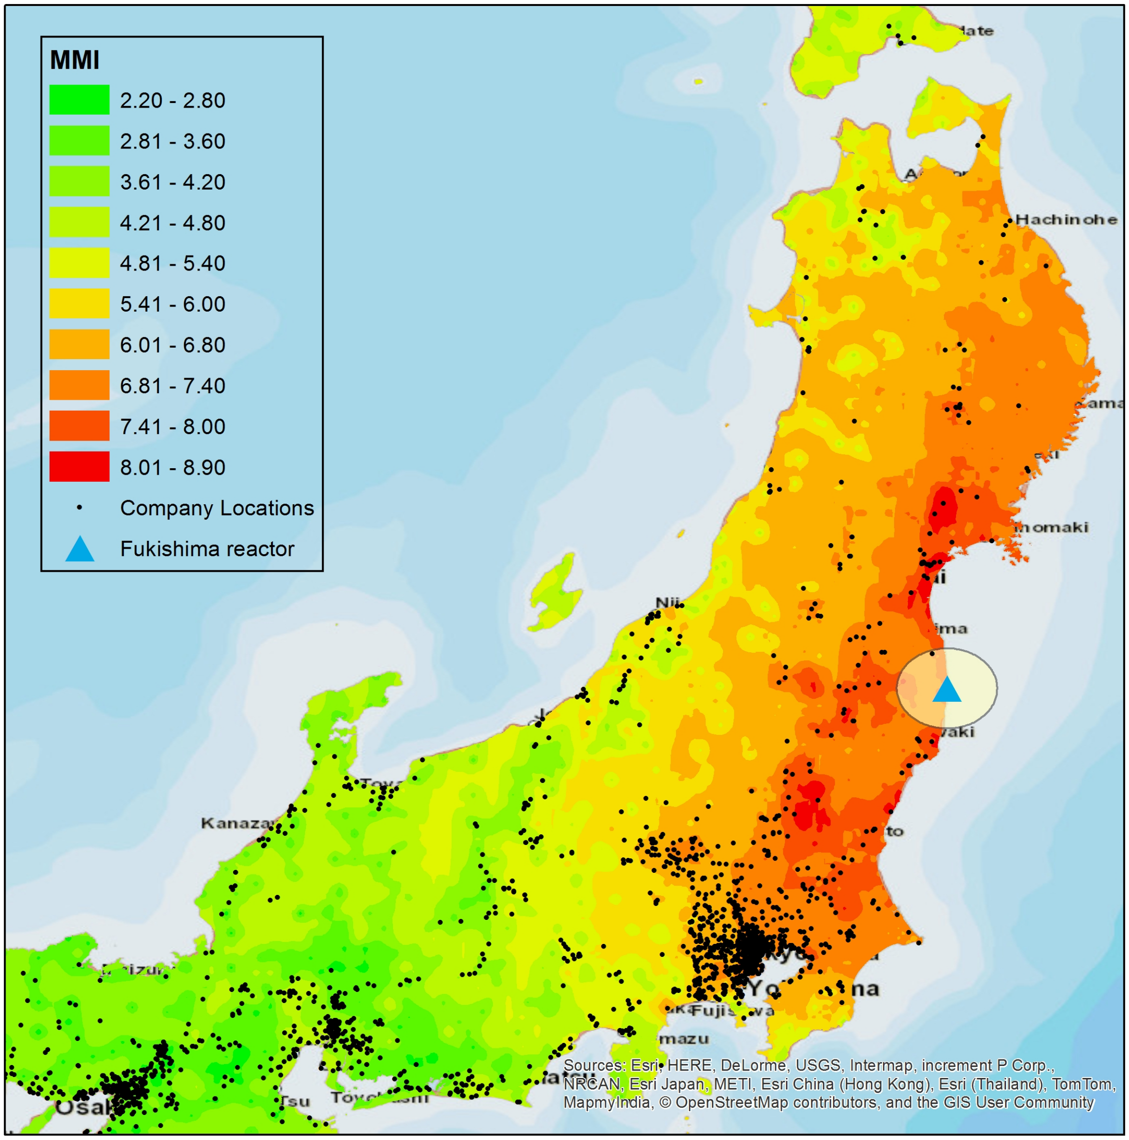

Source: USGS and DCA/Uniworld Directories
This figure plots the geographic distribution of the Tōhoku earthquake, based on recorded measurements taken directly after the event. The "Modified Mercalli Intensity" (MMI) scale is constructed based on a relation of survey response and measured peak acceleration and velocity amplitudes from prior major seismic events. Each dot corresponds to a geocoded Japanese affiliate location corresponding to a firm with U.S. operations. For more details, see Appendix B.2.3.2.

Figure 2.5: Density of Firm-Level Exposure to Japanese Imported Inputs: By Firm Type

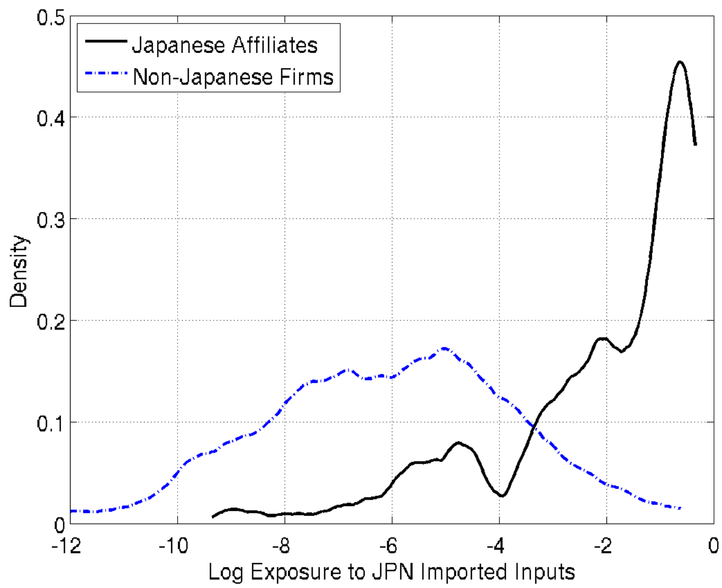

Source: LFTTD-DCA-UBP as explained in text. The estimates correspond to year 2010. This figure displays density estimates of the log exposure measure to Japanese imported inputs, separately for Japanese affiliates and non-Japanese multinational firms. The measure is defined as the ratio of Japanese imported inputs to total imported inputs plus U.S. salaries and wages. Estimates at either tail are suppressed for confidentiality purposes.

Figure 2.6: Dynamic Treatment Effects: Japanese Firms

A. Relative Intermediate Input Imports of Japanese Firms

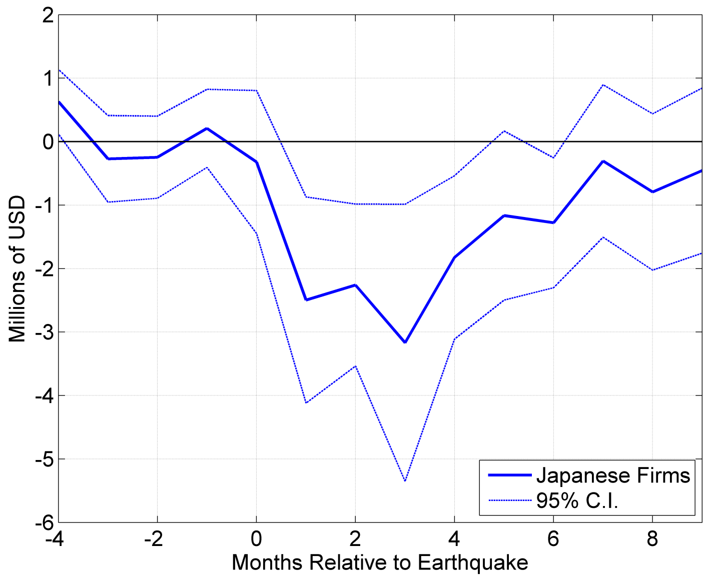

B. Relative North American Exports of Japanese Firms

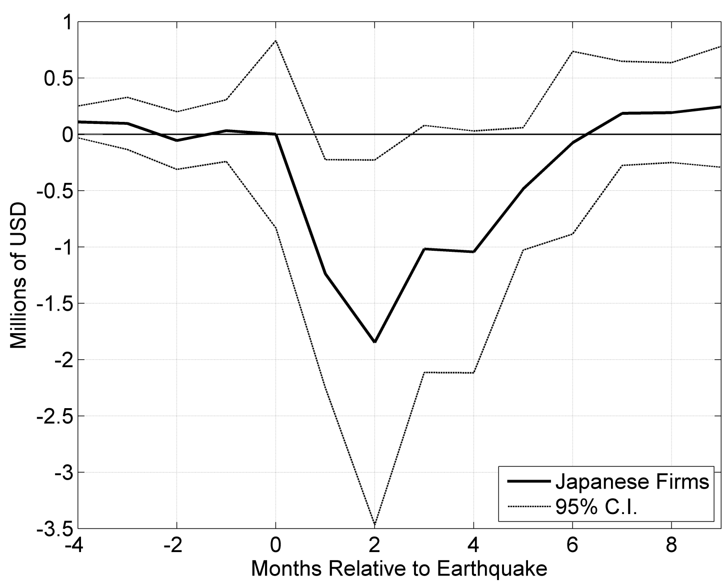

Source: LFTTD-DCA-UBP as explained in text.
These figures report the intermediate imports and North American exports of the U.S. affiliates of Japanese firms relative to a control group of other multinational firms. The values are coefficient estimates taken from an interaction of a Japanese-firm dummy with a monthly dummy – additional baseline monthly dummies remove seasonal effects. See equation 2.5 in the text. Standard errors are clustered at the firm level.

Figure 2.7: Relative Imported Inputs and Output (Proxy) of Japanese Firms: Fraction of Pre-Shock Level

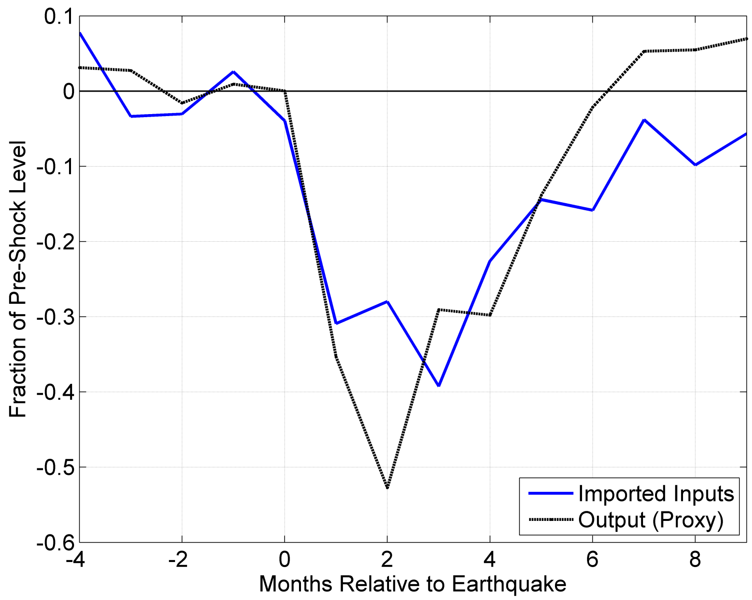

Source: LFTTD-DCA-UBP as explained in text.
This figure reports the intermediate imports and output proxy (North American exports) of the U.S. affiliates of Japanese firms relative to a control group of other multinational firms. The values are percent changes from the pre-shock level of each series, defined as the average of the months December 2010, January 2011, and February 2011.

Figure 2.8: Assessing the Output Proxy Using Monthly Automotive Production

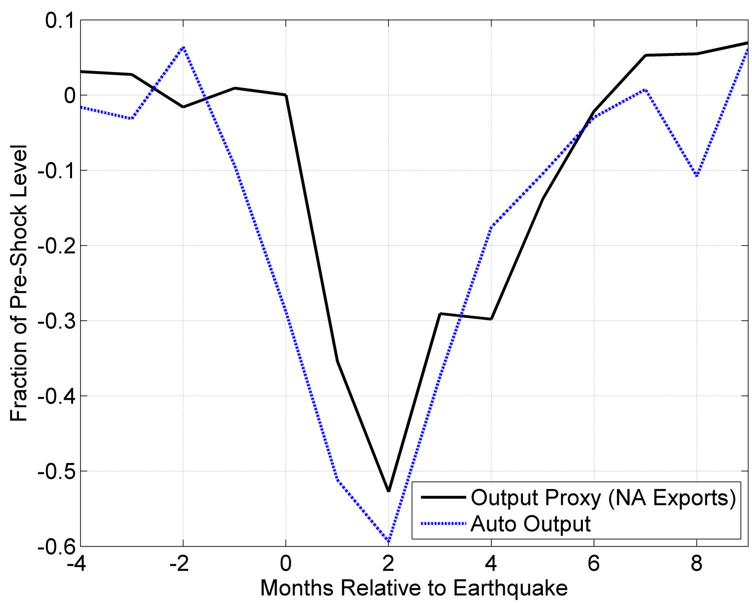

Source: Ward's Automotive Database
This figure reports the production levels of Japanese auto plants relative to a control group of non-Japanese auto plants. The values are percent changes from a pre-shock level, defined as the average of the months December 2010, January 2011, and February 2011. See equation B.14 in the text. For purposes of comparison, we also include the equivalent measure corresponding to total manufacturing of Japanese affiliates using the output proxy from Census data (from Figure 2.7). The Japanese automakers are Honda, Mazda, Mitsubishi, Nissan, Toyota, and Subaru. For the sake of clarity, we suppress the standard errors for the automotive series, though there are 4 months with below zero production based on a 95 percent confidence interval. See Appendix B.3.7 for more details.

Figure 2.9: Japanese Products: Average Distance from Benchmark Cost Shares: JPN Multinationals

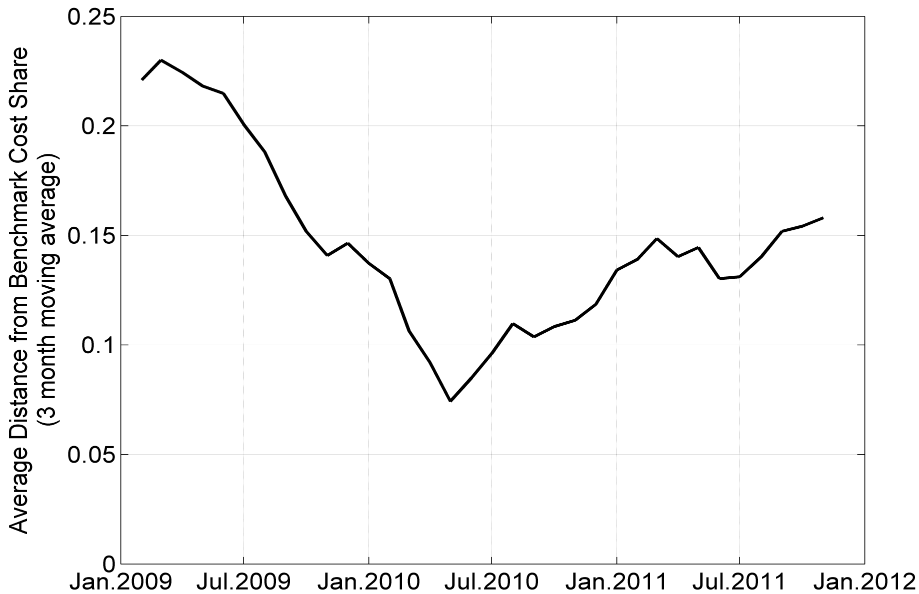

Source: LFTTD-DCA-UBP as explain in the text
Underlying this figure is the calculation of the average total (absolute) deviations from a benchmark measure of a firm's cost shares across input products from Japan. See equation 2.17 in the text. The figure reports the mean across the Japanese multinationals used in the section 2.4.

CHAPTER III

TFP, News, and "Sentiments:" The International Transmission of Business Cycles

with Andrei A. Levchenko

3.1 Introduction

Business cycles in advanced economies exhibit strong positive comovement. A complete empirical and theoretical account of positive cross-border comovement remains elusive. The International Real Business Cycle (IRBC) literature, going back to *Backus et al.* (1992) develops quantitative models in which fluctuations are driven by surprise TFP shocks, and assesses their performance in generating comovement. However, a series of empirical contributions in the closed-economy literature have argued that the bulk of (short-run) business cycle fluctuations is actually accounted for by non-technology shocks, customarily referred to as "demand" shocks. (For a number of different approaches that reach this conclusion, see *Blanchard and Quah*, 1989; *Galí*, 1999; *Canova and de Nicoló*, 2003; *Basu et al.*, 2006). It is thus a natural conjecture that international business cycle comovement can be driven by transmission of non-technology as well as technology shocks across borders. Indeed, international business cycle models are more successful at matching basic moments in the data when augmented with demand shocks (*Stockman and Tesar*, 1995; *Wen*, 2007).[1]

This paper investigates empirically the relative importance of the cross-border transmission of both technology and non-technology shocks. It uses US and Canada as a laboratory to study these issues. These two economies are closely integrated, and very asymmetric in size. The latter feature implies that identified US shocks are unlikely to be "contaminated" by endogenous US responses to Canadian shocks.[2]

We begin by identifying three types of US shocks in a structural vector auto-regression (VAR) setting. The first is a shock to contemporaneous TFP. This shock is identified as the reduced-form TFP shock, assuming that the TFP series is ordered first. New to the study of the international business cycle, the TFP series we use is adjusted for unobserved input utilization. *Basu et al.* (2006) show that the utilization adjustment has a large impact on both the properties of the TFP series itself, and on the impulse responses of US macroeconomic aggregates to the TFP shock. The second shock is a news shock about future TFP (*Beaudry and Portier*, 2006), identified following *Barsky and Sims* (2011) as the shock that has no contemporaneous TFP impact and explains the maximum of the forecast error variance of the utilization-adjusted TFP series.

Most importantly, we propose a new identification strategy for a non-technology business cycle shock. The VAR includes an expectation variable, alternatively a GDP forecast from the Philadelphia Fed's Survey of Professional Forecasters or the Michigan/Reuters Consumer Confidence variable. The non-technology shock is identified as the shock orthogonal to both the surprise-TFP and the news-TFP shocks that explains the maximum of the residual

[1]An obvious alternative is that international comovement is generated by transmission of policy or credit shocks. Available evidence suggests that the importance of these shocks in fluctuations is limited. *Kim* (2001) and *Maćkowiak* (2007) show that shocks to the US monetary policy explain only very limited share of forecast error variance of other countries' output, while *Ilzetzki and Jin* (2013) show that even the sign of the impact is not stable over time. In a similar vein, *Helbling et al.* (2011), *Kollmann* (2013), and *Eickmeier and Ng* (2015) show that the share of variance of other countries' GDP accounted for by the US credit shocks and bank shocks is small as well.

[2]This approach has been adopted by *Cushman and Zha* (1997), *Schmitt-Grohé* (1998), *Justiniano and Preston* (2010), and *Miyamoto and Nguyen* (2014), among others.

forecast error variance of this expectational variable. Because the shock is identified explicitly from data on expectations after controlling for shocks to current and future TFP, we label this shock "sentiment" as an homage to the recent literature on non-technology driven business cycles (e.g., *Angeletos and La'O*, 2013; *Benhabib et al.*, 2015b; *Huo and Takayama*, 2015). It is important to underscore that this is only a shorthand, as we do not identify the precise mechanisms that produce fluctuations in these theoretical contributions. Our shock can be driven by anything that makes agents expect better/worse times, conditional on available information about current and future productivity.

We identify the three shocks in the US data. We then estimate the impact of these shocks on Canadian macro aggregates, included as non-core variables in the baseline VAR. The results can be summarized as follows. The sentiment shock generates a US business cycle and accounts for an important share of forecast error variance in the US macro aggregates. GDP, consumption, and hours (as well as expectations) all increase on impact, peak within a year, and revert back to pre-shock values in the medium run. These dynamics are consistent with the sentiment shock being a transitory "demand" shock. The sentiment shock drives the bulk of short-run fluctuations in the US. It accounts for 65-75% of the forecast error variance in GDP at short frequencies (one year or less). At short frequencies, it also accounts for about one-fifth of the forecast error variance of consumption and about 70% of the forecast error variance in total hours. The finding that a non-technology shock is responsible for a large share of short-run fluctuations is of course consistent with results from other ways of identifying "demand" shocks (see, e.g., *Galí*, 1999; *Canova and de Nicoló*, 2003, among others).

Our main results concern the cross-border transmission of shocks to Canada. The first important finding that sets the stage for the rest of the results is that Canadian utilization-adjusted TFP does not react to any of the three identified US shocks. This makes us confident that the business cycle impact of US shocks on Canada is not contaminated by an underlying correlation between US shocks and Canadian TFP. On the rest of the Canadian variables, the three identified shocks have very different impacts.

The common theme is that in the short run, Canadian aggregates react much more strongly to the sentiment shocks than to the surprise and news TFP shocks. Following a sentiment shock, Canadian GDP rises instantaneously and peaks within one year. By contrast, the response of Canadian GDP to the US surprise TFP or news shocks is positive but takes place with a lag of 2-3 quarters. Canadian consumption and hours follow the same pattern. Canadian exports to the US and US exports to Canada both rise instantaneously following a sentiment shock, peak at 1 or 2 quarters, and then fall back to steady state. By contrast, there is not much of a trade response to surprise TFP shocks. US news do not generate a positive trade response for over 1 year following the shock. There is suggestive evidence that Canadian imports from and exports to the US actually fall on impact in response to a US news shock.

Among the three identified US shocks, the sentiment shock is by far the most important in accounting for the forecast error variance of the Canadian variables. At short frequencies, it accounts for 20-40% of the forecast error variance of Canadian GDP, 8-12% of Canadian consumption, 20-35% of Canadian hours, and 25-44% of Canada-US trade flows. By contrast, the surprise TFP shock accounts for less than 6% of the forecast error variance of Canadian GDP and hours across all frequencies between 1 quarter and 5 years, and for less than 10% of Canadian consumption. The (lack of) importance of the US news shock is similar at short frequencies, though the news shock does become more important for Canadian output and consumption at frequencies longer than 2 years.

Finally, we examine the role of the three US shocks in business cycle comovement between US and Canada by means of computing conditional correlations between the variables due to each shock following the approach in *Galí* (1999). The correlation of the US and Canadian GDP conditional on surprise TFP shocks is 0.47. The surprise TFP shock actually generates a negative correlation in consumption (-0.13) and hours (-0.47) between US and Canada. Conditional GDP correlations due to news (0.99) and sentiment (0.99) shocks are much higher. These shocks generate positive instead of negative correlations of consumption and hours as well. The sentiment shock generates a conditional correlation in consumption of 0.80 and in hours of 0.98, which is substantially higher than that generated by the news shock.

The bottom line is that at short frequencies, the non-technology shocks generate a much stronger cross-border impact of US shocks and account for a higher share of Canadian fluctuations. The sentiment shocks also generate much higher conditional correlations between US and Canadian aggregates than surprise TFP shocks. At the same time, news shocks are also important for international comovement at medium frequencies. An empirical account of observed international comovement therefore requires knowledge of the impact of both types of shocks, coupled with the understanding that the surprise TFP innovation central to most IRBC models is actually a shock that does not generate substantial comovement.

The results are robust to wide range of additional controls and alternative empirical models. Augmenting the VAR with the federal funds rate, we identify the sentiment shock alongside a monetary policy shock. Accounting for monetary policy shocks does not change either the properties of the sentiment shock nor its explanatory power. We also show that the sentiment shock is not an oil shock or an uncertainty shock. We add to the VAR a number of variables to increase the information set used to identify shocks: stock prices, consumer prices, the real exchange rate, as well as an estimated factor variable. Adding these variables does not alter the features of the sentiment shock or diminish noticeably its importance.

Our analysis is most closely related to empirical assessments of cross-border transmission of shocks, in particular non-technology shocks. *Canova* (2005) examines the impact of US supply and demand shocks on Latin America, while *Corsetti et al.* (2014) assess the reaction of externally-oriented variables – such as real exchange rates and foreign assets – to US supply and demand shocks. Both of these papers identify supply and demand shocks using sign restrictions. Our paper contributes a novel identification strategy for supply and demand shocks, based on expectational variables (for demand) and utilization-adjusted TFP (for supply). Importantly, we separate news about future TFP – which can look like a demand shock in the short run – from sentiment shocks unrelated to TFP.

Our paper draws heavily on the recent closed-economy empirical and theoretical literature on "demand"-driven fluctuations (see, among others, *Galí*, 1999; *Beaudry and Portier*, 2006; *Lorenzoni*, 2009; *Barsky and Sims*, 2011; *Angeletos and La'O*, 2013; *Blanchard et al.*, 2013; *Benhabib et al.*, 2015b,a; *Huo and Takayama*, 2015). Two recent papers in particular identify shocks that are interpreted as sentiments. *Angeletos et al.* (2014) extract a shock that explains the most of the forecast error variance of key macroeconomic aggregates, and show that it has the properties consistent with being a confidence shock. *Angeletos et al.* (2014) and *Milani* (2014) structurally estimate fully-specified DSGE models that incorporate sentiment shocks, and show that the sentiment shocks identified within the structure of those models can explain a large fraction of the US business cycle fluctuations. Our empirical strategy complements both of these approaches. In contrast to both of these alternatives, we explicitly separate a strictly non-technology sentiment shock from the TFP news shock. Relative to the data-driven exercise in *Angeletos et al.* (2014), our identification strategy is based on explaining the variation only in an explicit expectational variable. Our strategy thus "ties our hands behind our back" to a much greater extent, as we are not extracting a shock that by construction explains the bulk of fluctuations in the key macro aggregates. We complement the fully structural DSGE estimation approach by performing a more data-driven exercise. It is reassuring that our findings regarding the importance of "sentiments" in the US business cycle are consistent with these alternative approaches. Substantively, of course, our focus is on the international dimension of shock transmission.

The rest of the paper is organized as follows. Section 3.2 discusses the empirical strategy and estimation methods. Section 3.3 describes the data. Section 3.4 presents the main results, while Section 3.5 discusses interpretation and relates our analysis to the literature. Section 3.6 concludes.

3.2 Empirical Strategy

3.2.1 Identification of Shocks

Our identification strategy builds on *Uhlig* (2003, 2004) and *Barsky and Sims* (2011). As an illustration of why it is important to separate non-technology shocks from news TFP shocks, suppose that the TFP process in the US is affected by only two innovations: an unanticipated 'surprise' TFP shock and a 'news' shock. An example of a process that would satisfy these conditions is:

$$TFP_t = \lambda_1 \epsilon_t^{sur} + \lambda_2 \epsilon_{t-s}^{news}, \tag{3.1}$$

where ϵ^{sur} and ϵ^{news} are the surprise and anticipated innovations in TFP and the agents learn about the news shock $s > 0$ periods in advance.[3]

Further, assume that expectations of future economic activity are influenced not only by the surprise innovation in TFP and the anticipated future improvement in TFP, but also by 'sentiments,' as the agents rationally expect a positive sentiment shock to lead to a temporary boom in the economy and increase output. Forward-looking agents also respond to other changes in the economy that could stimulate GDP, but we assume that the bulk of the variation in expectations of future activity is due to these three shocks. A simple process for expectations F_t that satisfies this assumption is:

$$F_t = \lambda_1^F \epsilon_t^{sur} + \lambda_2^F \epsilon_{t-s}^{news} + \lambda_3^F \epsilon_t^{sent} + \zeta_t, \tag{3.2}$$

[3]This TFP process can clearly be modified to include a persistent component.

with ϵ^{sent} the sentiment shock.

Expectations of better future economic conditions, controlling for current fundamentals, can be due to either news of high future TFP, or to positive 'sentiment.' Clearly, in order to extract a non-technology shock from data on expectations, we must control for news of future productivity. It would not be possible to identify the three shocks of interest from movements in TFP and expectations alone. We therefore consider the processes for these variables together with other forward-looking macroeconomic aggregates in a VAR. Let Y_t denote the $k \times 1$ vector of observables in levels. For much of our analysis, this will be US TFP, real GDP, consumption, hours, and forecasts of GDP. The moving average representation of this k-variable VAR is:

$$Y_t = B(L)\,\mathbf{u}_t,$$

where $\mathbf{u}_t$ is the vector of reduced-form disturbances, L denotes the lag operator and $B(L)$ is the matrix of lag order polynomials.

To identify the structural shocks, we assume that there exists a linear relationship $\mathbf{u}_t = A\boldsymbol{\epsilon}_t$ where $\boldsymbol{\epsilon}_t$ is the vector of structural shocks and A is the impact matrix. This implies that the structural representation of the VAR is

$$Y_t = A(L)\,\boldsymbol{\epsilon}_t,$$

where $A(L) = B(L)\,A$. Clearly, assuming that the structural shocks each have unit variance, $AA' = \Sigma$, where Σ is the covariance matrix of u. It is well known that the Choleski decomposition $\tilde{A}$ of Σ provides one candidate for A, but this is just one among many. For any orthonormal $k \times k$ matrix D such that $DD' = I$, $\tilde{A}D$ will provide an identification of the structural shocks.

The forecast error h steps ahead is defined as

$$Y_{t+h} - E_{t-1}Y_{t+h} = \sum_{\tau=0}^{h} B_\tau \tilde{A} D \boldsymbol{\epsilon}_{t+h-\tau},$$

where B_τ is the reduced-form matrix of lag-τ moving average coefficients. Since the elements of $\boldsymbol{\epsilon_t}$ are independent, this equation illustrates that the forecast error variance of a particular variable i at horizon h is the sum of the contributions of the k structural shocks. Let $\Omega_{i,j}(h)$ denote the contribution of shock j to the forecast error variance of variable i at horizon h. The assumption that only two shocks (surprise and news) affect true TFP then implies:

$$\Omega_{1,sur}(h) + \Omega_{1,news}(h) = 1 \quad \forall h. \tag{3.3}$$

The unexpected TFP innovation ϵ_t^{sur} in (3.1) is identified as the reduced-form innovation in a VAR with TFP ordered first. By identifying the reduced-form innovation in TFP as the first structural shock, we effectively fix $\Omega_{1,1}(h)$ at all horizons. The news shock ϵ_{t-s}^{news} is true news about future changes in TFP s periods ahead. Without loss of generality, assume the second structural shock is the news shock, and thus the second column of $\tilde{A}D$ is its impact vector. The news shock is identified as the linear combination of the remaining VAR innovations that maximizes the residual forecast error variance of TFP, $1 - \Omega_{1,1}(h)$, over a finite horizon H^{news}.[4] Of course, in practice (3.3) is unlikely to hold as an identity for all $h \leq H^{news}$. Thus, given the Choleski decomposition $\tilde{A}$, *Barsky and Sims* (2011) choose the vector γ^{news} (the second column of D), such that this second shock maximizes the residual forecast error variance of the TFP process over horizon H^{news}.

[4] In the empirical implementation we select $H^{news} = 40$, or a ten-year horizon.

Formally, we select γ^{news} as the solution to the problem:

$$\gamma^{news} = \text{argmax} \sum_{h=0}^{H^{news}} \Omega_{1,2}(h) = \text{argmax} \sum_{h=0}^{H^{news}} \left(\frac{\sum_{\tau=0}^{h} B_{1,\tau} \tilde{A} \gamma^{news} \gamma^{news\prime} \tilde{A}' B'_{1,\tau}}{\sum_{\tau=0}^{h} B_{1,\tau} \Sigma B'_{1,\tau}} \right)$$

subject to

$$D(1,i) = 0 \quad \forall i \neq 1 \tag{3.4}$$

$$D \text{ is orthonormal,} \tag{3.5}$$

where the lower-triangular matrix $\tilde{A}$ is the Choleski decomposition (so $\tilde{A}(1,m) = 0 \; \forall m > 1$).

We next proceed to the identification of the sentiment shock. As this shock cannot be inferred from movements to TFP, our identification will rely on its impact on expectational variables. These will be alternately forecasts of GDP by professional forecasters or consumer confidence. Further, we impose that this shock does not affect true TFP. The procedure outlined above naturally builds in this assumption: by allowing only the first two shocks to affect TFP, we minimize the TFP impact of the remaining $k-2$ structural shocks, which includes the sentiment shock.

Let the expectational variable F_t be ordered 5th in the VAR, and without loss of generality assume that the sentiment shock is the 3rd shock. Note that by equating the first reduced-form shock to the surprise innovation to TFP and then identifying the news shock as in *Barsky and Sims* (2011), we have in effect fixed $\Omega_{5,1}(h)$ and $\Omega_{5,2}(h)$ at all horizons. We therefore select the sentiment shock as the linear combination of the remaining $k-2$ reduced-form innovations that maximizes the forecast error variance of F_t, where k is the total number of core and non-core variables in the VAR.[5] Because the sentiment shock is short-run, we

[5]Note, we do not allow the reduced form shock to the Canadian variable, ordered k, to affect the identification of the sentiment shock. Hence, the sentiment shock is identified from $k-2$ reduced form shocks (as the surprise TFP innovation also does not affect it).

select it to maximize the forecast error variance for a 2-quarter horizon ($H^{sent} = 2$). Formally:

$$\gamma^{sent} = \text{argmax} \sum_{h=0}^{H^{sent}} \Omega_{5,3}(h) = \text{argmax} \sum_{h=0}^{H^{sent}} \left(\frac{\sum_{\tau=0}^{h} B_{5,\tau} \tilde{A} \gamma^{sent} \gamma^{sent\prime} \tilde{A}' B_{5,\tau}'}{\sum_{\tau=0}^{h} B_{5,\tau} \Sigma B_{5,\tau}'} \right)$$

subject to

$$D(1,i) = 0 \quad \forall i \neq 1 \tag{3.6}$$

$$D \text{ is orthonormal} \tag{3.7}$$

$$D(:,2) = \gamma^{news}. \tag{3.8}$$

Both the news and sentiment identification steps are conditional on an arbitrary orthogonalization, the Choleski decomposition $\tilde{A}$. The first restriction – (3.4) and (3.6) – common to both problems specifies that none of the $k-1$ structural shocks has a contemporaneous impact on TFP. The second restriction, (3.5) and (3.7), states that the matrix D remains orthonormal throughout, and thus the identified shocks are orthogonal to each other. Restriction (3.8) ensures that identification of the sentiment shock holds identification of the news shock constant. We expect the surprise TFP and the news shocks, as informative about true fundamentals, to explain the movements in the forecast of GDP. The sentiment shock identified in this manner simply captures patterns in the residual variance of the forecast of GDP, once supply-side determinants are accounted for. The identification strategy for both shocks is robust to the reordering of the remaining $k-1$ variables in the VAR other than TFP.[6]

Our strategy relies on 'medium-run' identification. It might appear that the natural

[6] In a recent paper, *Angeletos et al.* (2014) adopt a closely related identification strategy to extract a factor that explains most of the business cycle variation in hours and investment at frequencies of 6-32 quarters. In contrast to our approach, that paper obtains an expression for the share of the variance of a variable due to a shock at this frequency through a spectral decomposition, and then chooses a linear combination of shocks that maximizes the variance of the selected variables. TFP is not included in their VAR. In short, they sum across variables, while we maximize the residual forecast error variance of a single, expectational, variable – either GDP forecast or confidence – over several horizons.

identification of the sentiment shock would make use of a 'long-run' restriction, namely that it has no long-run impact on output or forecasts. We prefer the method here as several papers have emphasized that long-run restrictions are problematic in VARs of finite order, where the coefficient estimates are biased (*Faust and Leeper*, 1997). Medium-run identification has shown better behavior in finite samples (*Francis et al.*, 2014).[7]

3.2.2 Estimating International Transmission

We estimate the impact of the US shocks on various Canadian aggregates in turn, treating them as 'non-core' variables in the VAR. The Canadian variables are included one at a time and are ordered last in a six-variable VAR with 5 US series. The matrices of coefficients are restricted to allow no current or lagged impact of the Canadian variable on the five US variables. We believe this assumption is reasonable given the small size of the Canadian economy relative to the US (Canadian GDP is about one-tenth that of the US). Section 3.4.1 shows that the results are robust to allowing lagged Canadian variables to affect US variables.

The impulse responses of Canadian variables to the identified US shocks are interpreted as evidence of cross-border transmission of those shocks to Canada, rather than a correlation of underlying Canadian shocks with the US shocks. A useful check presented below is to construct the impulse responses of Canadian TFP to these identified shocks, and ascertain that Canadian TFP does not comove with the identified US shocks. Section 3.5 also checks for the possibility of correlated sentiment shocks, which would not be visible in TFP movements, and finds little evidence that the impulse responses of Canadian aggregates to US shocks are due to a correlated Canadian shock.

[7] An alternative approach to long-run identification in VARs uses the spectral factorization of the variance matrix at frequency zero. This does not circumvent the issues related to long-run restrictions in general, however. We do not pursue the spectral approach in this paper as we are not aware of methods by which we would be able to identify the three shocks, while maintaining the medium-run identification structure.

We estimate the reduced-form VAR with estimated generalized least squares (EGLS) using a method adapted from *Lütkepohl* (2005). The VAR in p lags is:

$$Y_t = C_0 + C_1 L Y_t + ... + C_p L^p Y_t + u_t$$

where C_j are $k \times k$. If the Canadian variable is ordered last, the restriction that Canadian variables are have no impact on US variables amounts to $C_j\,(1 : k-1, k) = 0\ \forall C_j$. Rewrite the VAR in compact form as $Y = CZ + U$, where $Y = [Y_1, ..., Y_T]$, $Z_t = [1, Y_t, .., Y_{t-p+1}]$, $Z = [Z_0, ...Z_{T-1}]$, $C = [C_0, ...C_p]$, and $U = [u_1, ...u_T]$.

Let the constraints on the coefficients of the six-variable VAR be written as $\boldsymbol{\beta} = vec\,(C) = R\mathbf{b} + \mathbf{r}$, where R is a known matrix of rank M, $\mathbf{r}$ is a vector of constants, and $\mathbf{b}$ is the $(M \times 1)$ vector of unknown parameters to be estimated. Appropriately pick R (size $k\,(kp+1) \times M$) and $\mathbf{r}$ such that the desired constraints on C_j hold. Clearly, linear restrictions of the type we are interested in can easily be expressed in this form.

The EGLS estimate of $\mathbf{b}$ is then:

$$\mathbf{b} = \left[R' \left(ZZ' \otimes \Sigma_u^{-1}\right) R\right]^{-1} R \left(Z \otimes \Sigma_u^{-1}\right) z \tag{3.9}$$

where

$$z = vec\,(Y) - (Z \otimes I_K)\,\mathbf{r}$$

and Σ_u is any consistent estimator of the unknown covariance matrix of $vec\,(U)$. We initialize Σ_u as

$$\hat{\Sigma}_u = \frac{1}{T - kp - 1} \hat{U}_{ols} \hat{U}'_{ols}$$

where $\hat{U}_{ols}$ are the residuals from an unconstrained ordinary least squares estimation of the six-variable VAR(p). We use an iterative procedure, in which we compute a new covariance matrix from the first stage EGLS residuals to replace Σ_u in the computation of the next value of $\mathbf{b}$ and iterate to convergence. This procedure is asymptotically more efficient than standard multivariate least squares, and under the assumption of Gaussian errors the estimator for $\mathbf{b}$ in (3.9) is the same as the maximum likelihood estimator. Using estimates of $\mathbf{b}$ it is then straightforward to compute the impulse response functions of each Canadian macro aggregate to the three shocks of interest. Note that the identification of the shocks is unaffected by this procedure.

Following the recommendation of *Hamilton* (1994), the model is specified in levels, since parameter estimates in levels are still consistent even in the presence of cointegration, while the vector error correction model might be misspecified when the cointegration is of unknown form. The baseline implementation uses $p = 4$ lags, the optimal lag length according to the Akaike Information Criterion. All standard errors are constructed from 2000 bias-corrected bootstraps as in *Kilian* (1998).

3.3 Data

The time period covered by our data is 1968:Q4 to 2010:Q3. All variables are logged. For a measure of US productivity, we use the quarterly, utilization-adjusted TFP series from *Fernald* (2014). The series is the quarterly version of the annual series developed by *Basu et al.* (2006). That paper constructs a modified Solow residual from industry-level data, allowing for both non-constant returns to scale and varying unobserved capital and labor utilization. The identification of the three structural shocks in our VAR relies on an accurate measure of US technology. Clearly, accounting for measurement issues arising from changes is utilization is crucial. *Basu et al.* (2006) find that the detrended utilization-adjusted TFP is both less correlated with output, and less volatile than the standard Solow residual. Unfortunately the industry-level data required for controlling for non-constant returns to scale are not available quarterly, so the *Fernald* (2014) series corrects only for variable capital and labor utilization.

US population and hours data are from the BLS. For population, we use the civilian non-institutionalized population age 16 and over. Aggregate hours are the total hours of wage and salary workers on non-farm payrolls. For consumption and output, we use the National Income and Product Accounts (NIPA) tables from the BEA. Output is measured as quarterly real GDP, chain-weighted, from NIPA table 1.1.6. As a chain-weighted series for non-durables and services consumption is not available, we construct a series using the Tornqvist approximation (see *Whelan*, 2000, for details on chain-weighting in the BEA data). For this procedure, we use the nominal shares of spending on non-durables and services from NIPA table 1.1.5. Chain-weighting reduces the dependence of a series on the choice of base year, and is the current standard for macroeconomic series constructed by all major statistical agencies. All variables are converted into per capita terms.

The data on the forecasts of US GDP come from the Survey of Professional Forecasters (SPF), provided by the Federal Reserve Bank of Philadelphia. For NIPA variables, the survey contains quarterly forecasts at several horizons as well as longer-term forecasts. We use the one quarter ahead growth rate forecast. The perturbation to US sentiment that we are interested in identifying is not related to true technological progress, and we would expect the effects of this shock to be very short-lived. The survey provides mean and median levels forecasts as well as growth rates. The base year for the levels forecasts changes periodically throughout the survey. To avoid issues related to rebasing the forecasts ex-post, we construct an index of implied GDP levels forecasts from the mean forecast of the one quarter ahead growth rate. We check the sensitivity of our results to using a two- or three-quarter ahead growth rate forecast, as well as different horizons $H^{sent} = 4, 8, 16$ over which we expect the sentiment shock to contribute to the forecast error variance of the GDP forecast variable, and find no significant differences in the shape of the responses.

In addition, we re-do the analysis using an index of consumer confidence from the University of Michigan Survey of Consumers instead of the SPF GDP forecast. We use the consumer confidence series E12Y, constructed from the responses to the question *'And how about a year from now, do you expect that in the country as a whole, business conditions will be better, or worse than they are at present, or just about the same?'*

A consistent measure of quarterly hours for the length of our sample is not easily available for most countries. For Canada, we use a new dataset assembled by *Ohanian and Raffo* (2012), constructed from the OECD's Main Economic Indicators database and other sources. Our Canadian hours measure is the total hours worked in Canada divided by the Canadian population. The population data are taken from CANSIM (the Statistics Canada database), and is the quarterly estimate of total population in all provinces and territories of Canada. Canadian real GDP and consumption are taken from the OECD Economic Outlook and are also converted into per capita terms. For the bilateral exports and imports series, we use data from the IMF's Direction of Trade Statistics (DOTS) database. The series are deflated with a US GDP deflator and deseasonalized using the X-12 ARIMA program developed by the US Census Bureau.

3.3.1 Utilization-Adjusted TFP for Canada

The last critical variable for the analysis is a measure of Canadian TFP. Ideally, we would use a utilization-adjusted series with further adjustments for non-constant returns to scale, similar to the *Basu et al.* (2006) series for the US. Unfortunately, such a series to our knowledge is not available for any other country. The data required to construct such a series are also not available at the quarterly frequency for Canada. Therefore we build our own utilization-adjusted TFP series for Canada, following the approach in *Imbs* (1999). This method uses a similar insight, namely that with a constant returns to scale production function the first-order conditions for capital and labor are informative about the choices of capital utilization and the workweek of labor. As data on the capital stock are also not available at the quarterly frequency, we use the perpetual inventory method to construct an initial capital stock series, given data on investment from the OECD and a constant depreciation rate. This produces a starting utilization series. We then use an iterative procedure to construct a time-varying depreciation rate, capital stock, and implied utilization series consistent with the observed investment in the data. We construct labor utilization from information on hours worked, wages, and consumption in Canada. The wage data is from the OECD Main Economic Indicators (MEI). The utilization-adjusted Solow residual is then $\log TFP = \log Y_t^{Can} - (1-\alpha)(\log K_t + \log u_t) - \alpha(\log N_t + \log e_t)$, where e_t is labor utilization,u_t is capital utilization, Y_t^{Can} is output, K_t is capital and N_t is hours worked. Details of the procedure are in Appendix C.1.

We present the impulse response functions for both the utilization-adjusted TFP series and the implied capital utilization series.[8]

3.4 Results

Our baseline specification identifies the news shock at a horizon of ten years, the sentiment shock at a horizon of two quarters, and uses the forecast of GDP one quarter ahead as the fifth variable in the VAR. We begin by discussing the responses to the surprise TFP, news, and sentiment shocks on the US economy (Figures 3.1, 3.2, and 3.3), followed by an analysis of the transmission to Canada. Section 3.5 places the results in the context of standard business cycle models and some variants proposed in the literature.

The surprise TFP innovation signals a deviation in TFP from trend of about 0.8%. The effects of the shock die out slowly, with TFP decreasing but staying significantly above trend

[8]We check the responses of the unmodified Solow residual as well, and find it does not move in response to the shocks. However we think it is still important to correct for utilization, as it is a channel through which the Canadian economy could respond.

for 12 quarters. The responses of other domestic variables to this shock are consistent with other empirical investigations (*Basu et al.*, 2006; *Barsky and Sims*, 2011). Output increases temporarily before falling below trend after two years. Consumption stays constant on impact, and declines with output.

Our identified news shock signals a slowly building increase in utilization-adjusted TFP, beginning in quarter 2. Consumption increases slightly on impact and continues for two years, after which it exhibits a very slight decline. There is an impact decrease in hours, qualitatively consistent with the results in *Barsky and Sims* (2011). The response of hours turns positive one year after the shock, peaking at about Q9. There is no significant impact effect on output. Rather, the response of output builds slowly, similar to technology (but stronger). The peak increase is later than for surprise TFP, two years after the shock. Reassuringly, the forecasts of GDP track the responses of actual GDP quite well, with the response of the forecast variable peaking about one quarter before GDP.

Overall, these responses are in line with *Barsky and Sims* (2011). As in that paper, the impact decrease in hours is consistent with a strong wealth effect, and indicates that the news shock does not solve the impact comovement problem of hours, consumption, and output.[9] It therefore cannot explain the unconditional positive comovement of these variables in the data. As *Barsky and Sims* (2011) point out, however, the responses to the news shock shown here are consistent with the predictions of a simple neoclassical growth model augmented with news shocks. As the response of hours is eventually positive, our news shock does generate comovement a few periods after impact, indicating that it is an important component of business cycle fluctuations in the medium term. On the other hand,

[9]This problem has been commonly observed in response to estimated TFP shocks (*Galí*, 1999), and news shocks were originally discussed as a possible solution. For instance *Beaudry and Portier* (2006) identify news shocks as the innovation in stock prices orthogonal to current TFP and find that the identified shock does generate positive comovement on impact. The news shocks identified in that paper capture a much longer-term improvement in technology, and therefore dissimilar to those in *Barsky and Sims* (2011) and our paper. Furthermore, the *Beaudry and Portier* (2006) identification scheme has been shown to deliver non-unique dynamic paths when extended to several variables (*Kurmann and Mertens*, 2014).

Barsky et al. (2014) argue that it is unclear whether the comovement in the dynamic paths of all variables is due to the news shock itself or the realized productivity growth.

The impulse responses to the sentiment shock look noticeably different. There is an impact increase in output, consumption, hours, and the forecast variable. There is a very small and insignificant decrease in measured TFP, which might be due to the quarterly series not perfectly correcting for utilization as discussed in Section 3.3. The business cycle generated by the shock lasts approximately three years. A substantial empirical literature beginning with *Galí* (1999) has previously argued that demand shocks are promising for explaining business cycles. Ours is (to our knowledge) the first paper to directly measure these shocks based on forecast or confidence data while ensuring they are uncorrelated with both current and future technological change. We discuss the relationship of our identified shock to the literature on demand shocks in Section 3.5.

The top panels of Tables 3.1-3.3 report the share of the forecast error variances of the US macro aggregates accounted for by the TFP, news, and sentiment shocks respectively. At short frequencies, the sentiment shock appears most important. It accounts for 65-75% of the variation in GDP, 18-22% in consumption, and 62-71% in hours at horizons 1 year or less. By contrast, at these frequencies surprise TFP shocks explain less than 8-12% of the variation in GDP, 2% in consumption, and 2-8% in hours. The news shock does a little bit better for consumption (36-48%), but is about equally unimportant for GDP and hours. Not surprisingly, at longer frequencies the news shock increases in importance. *Barsky and Sims* (2011) reach a qualitatively similar conclusion about the news and surprise innovations, and point out that unexplained shocks were responsible for most of the variation at business cycle frequencies in domestic aggregates. Our analysis has now identified one such shock.

International Transmission. Figure 3.4 sets the stage for the remainder of the results. It shows the impulse responses of Canadian utilization-adjusted TFP to the three identified US shocks. None of the three identified shocks have a perceptible impact on Canadian technology (note also the different scale of the y-axis compared to the other figures). The news shock actually leads to a barely visible, though persistent and significant increase in Canadian TFP beginning about five quarters ahead. This might indicate the presence of technology spillovers, but the magnitude is quantitatively tiny. Thus, whatever impact of US shocks on Canada that we find below is not accompanied by a change in Canadian productivity.

Figure 3.5 shows that the three shocks lead to very different reactions of Canadian GDP. Neither shock to true TFP leads to an impact increase in GDP. The surprise TFP innovation in the US generates the smallest visible spillovers, with a slight increase in output three quarters after impact. The increase is short-lived, peaking at four quarters, after which Canadian output quickly returns to trend. In contrast, the news shock leads to more persistent Canadian output growth. GDP starts to increase two quarters after impact, lagging one quarter behind its US counterpart. The effects of the shock are more long-lived, with GDP peaking a little over two years after impact. At five years, output is still significantly above steady state.

The most striking is the response to the sentiment shock. Canadian GDP jumps on impact, in sync with US output. It increases further for two quarters, before gradually returning to steady state. The effects of the shock are significant for two and a half years, demonstrating that the sentiment shock has the potential to generate output comovement at high frequencies.

As it is clear that Canadian TFP is not affected, we propose one channel, consistent with our results, through which US sentiment shocks could generate spillovers. As Figures 3.6 and 3.7 show, Canadian exports to the US and imports from the US show the strongest responses to the sentiment shock. Both series jump on impact, a two percent deviation from trend. They demonstrate a strong hump-shaped pattern: the increase in Canadian exports peaks at one quarter. However they stay significantly above trend for two years. Since the US is Canada's largest trade partner and the sentiment shock generates increased demand in the US, this response is unsurprising.

The increased exports do not come through lowered Canadian consumption. Rather, as Figures 3.8 and 3.9 show, the factors of production are used more intensively following a US sentiment shock: Canadian hours increase, as does capital utilization. The increased production for export increases GDP, and generates an income effect which leads to higher consumption on impact (Figure 3.10). Demand for imports increases as well as a result of the higher consumption, and US exports to Canada rise.[10] The empirical evidence clearly suggests that the sentiment shock has the potential to not just generate a domestic business cycle, but explain both international synchronization as well as the positive correlation between exports and imports (*Engel and Wang*, 2011).

The news shock also generates comovement between Canadian exports and imports, but the impact effect is actually negative. The impact of higher future demand in the US contains both a substitution effect and an income effect. Holding TFP and production constant, the news shock would increase the price of future Canadian output and lead to a substitution effect towards consumption today. That said, cheaper future imports lower the price of future output and induce a negative substitution effect. However, the income effect from the future prolonged period of high export demand should unambiguously increase consumption and decrease hours. Each of these effects cannot be isolated in our framework, but the net effect is a slight decrease in Canadian hours after a US news shock, an insignificant decrease in GDP and a decrease in exports on impact. Consumption does not jump, so the wealth effect is not dominant, but it also begins to increase at about Q3.[11] After one year,

[10]Of course this is only one plausible channel. *Schmitt-Grohé* (1998) finds that exports are not a strong enough channel for the transmission of a generic shock to US output to Canada. That paper does not distinguish between the types of shocks that affect the US, however.

[11]Other explanations such as habit formation could also play a role here.

there is positive comovement among the key US and Canadian aggregates following a news shock. This implies that news shocks could also be an important component of comovement at medium- to long-term frequencies. It is unclear why US exports fall on impact. One possible explanation is weak demand in Canada coupled with the decreased production in the US.

The discussion above points to the different and complementary roles of the news and sentiment shocks in generating business cycle spillovers. Forecast error variance decompositions provide additional support for the importance of sentiment shocks at shorter frequencies, and of the news shocks at longer frequencies, internationally as well as domestically. The bottom panels of Tables 3.1-3.3 report the shares of forecast error variances of the Canadian macro aggregates accounted for by the three identified US shocks. At short frequencies, the sentiment shock is by a large margin the most important of the three. The sentiment shock contributes substantially to the forecast error variance of US-Canada trade, explaining up to 44% of the variance of Canadian exports and 41% of imports at the one year horizon. It also explains a large fraction of the forecast error variance of Canadian output (41% at one year), hours and utilization (over one third), and consumption (8-12%). The impact of this short run US 'demand' shock on a smaller trading partner is persistent, as it still accounts for 36% of the variance of output at 10 years. The small share of Canadian TFP variation attributed to the sentiment shock at 10 years is likely due "leakage" in the utilization adjustment, as our procedure for Canada is even coarser than the *Fernald* (2014) method on the US data.

In contrast the news shock is only responsible for very long run variation in TFP, output, and consumption, and does not contribute much to explaining the forecast error variance of other Canadian variables. The surprise TFP shock contributes very little to the forecast error variance of the Canadian aggregates at any frequency.

As further evidence on the importance of both sentiment and news shocks for international comovement, we construct correlations of key variables conditional on only one type of shock. As in *Galí* (1999), these correlations can be inferred directly from the structural impulse response coefficients. Formally, the correlation of variables j and k conditional on shock i, ρ^i_{jk}, is

$$\rho^i_{jk} = \frac{\sum_{h=0}^{\infty} A^h_{ji} A^h_{ki}}{\sqrt{\sum_{h=0}^{\infty} \left(A^h_{ji}\right)^2} \sqrt{\sum_{h=0}^{\infty} \left(A^h_{ki}\right)^2}},$$

where A^h_{ji} is the lag-h, (j, i)-th element of the matrix $A(L)$ of lag order polynomials of the structural moving average representation of the VAR, that captures the impulse response of variable j to shock i at lag h. In practice, we compute these correlations for a finite but large maximum horizon of 10000 periods.

The results of this exercise are in Table 3.4. The sentiment and news shocks both generate high correlations (both 0.99) of US and Canadian output, while the surprise TFP innovation delivers a much lower correlation than observable in the data. The surprise TFP shock actually generates a slightly negative US-Canada correlation of consumption (-0.13) and a strongly negative (-0.47) US-Canada correlation of hours. While both news and sentiment shocks deliver strongly positive consumption correlations, the correlation of hours due to the news shock is too low at only 0.46, but due to the sentiment shock it is too high at 0.98. The sentiment shock comes the closest to explaining the unconditional cross-correlations of exports from Canada with US output.[12]

In summary, the impulse response functions, variance decompositions, and conditional correlations show that surprise TFP innovations, which are usually assumed to be the key driver of IRBC models, play a negligible role in the international transmission of shocks. Sentiment shocks are important for transmission at higher frequencies, while news shocks

[12]Interestingly, the surprise TFP innovation does a reasonable job of reproducing the cross-correlations of US output, consumption, and hours, despite the impact impulse responses being inconsistent with closed economy RBC models. As *King and Rebelo* (1999) point out, data generated by feeding utilization-adjusted TFP into a model with sufficient internal propagation mechanisms does a reasonable job of matching historical US time series. The news and sentiment shocks also match the correlation of output and consumption well, but both undershoot the hours and output correlation.

play a stronger role at medium/long frequencies.

We conclude this section by discussing the role of all three shocks in recent business cycles. Figures 3.11– 3.14 display the historical decompositions of the key US and Canadian macro aggregates into the components due to the three identified US shocks. While the TFP, sentiment, and news components all contributed to the great recession in the US, the fall in output in Canada appears driven entirely by the sentiment shock. Similar patterns are visible for Canadian consumption and hours, as well as exports and imports, indicating the sentiment shock played a key role in the transmission of the recent recession. The sentiment shock does not appear to contribute equally to all recessions however, with the dips in output and consumption in the 1981-82 recession driven primarily by news.

3.4.1 Robustness

We check the responses of all variables to variations in the horizons of identification of the sentiment shock, and find no significant qualitative difference for $H^{sent} = 4$, 8, or 16. We also vary the forecast variable used in identification, using forecasts of GDP two quarters ahead and three quarters ahead. The qualitative shape of the dynamic responses remains the same. To conserve space the results are not reported here, but are available on request.

Consumer confidence. To check robustness of the results to the choice of expectational variable, we replace the GDP forecasts the VAR with the E12Y variable from the Michigan Survey of Consumers, constructed from the responses to the question *'And how about a year from now, do you expect that in the country as a whole, business conditions will be better, or worse than they are at present, or just about the same?'* The results from this exercise are in Appendix Figures C.1-C.10. Reassuringly, the patterns described above are robust to the expectational variable used to identify the sentiment shock.

Monetary policy shocks. Our empirical strategy permits the separate identification of a monetary policy shock. We extend the core VAR to include the federal funds rate, ordered fifth. The identification of the monetary policy shock is then standard (*Christiano et al.*, 1999): we assume that the shock has no contemporaneous impact on the variables ordered above the federal funds rate (TFP, Consumption, GDP, and Hours), but does have an impact on the variables below (the expectational variable). This simply requires certain zero restrictions on the impact matrix $\tilde{A}D$.[13] The identification of the news and sentiment shocks then follows as in the baseline with the added monetary policy shock restriction.

Figure 3.15 plots the impulse responses of the main US macro aggregates and Canadian GDP to the monetary policy shock and the sentiment shock side by side. Controlling explicitly for monetary policy shocks does not change the properties of the sentiment shock. As in the baseline, consumption, output, and hours all increase on impact in response to the sentiment shock. Interestingly, the federal funds rate increases following a sentiment shock, pointing to policy tightening in response to increased demand generated by positive "sentiments." The federal funds rate rises slightly on impact, and then increases sharply further in the second quarter. It then stays flat for about two years, and the subsequent decline is slower than following a monetary policy shock.

Figure 3.15 also makes clear that the identified sentiment shock is not a traditional monetary policy shock. In the figure, the monetary policy shock is an unexpected decrease in the federal funds rate, and is thus expansionary. The responses of the macroeconomic aggregates are as expected following a monetary policy loosening (*Christiano et al.*, 1999). Output and consumption are flat initially and then rise. Hours decrease for about 5 quarters, and then increase. Turning to comovement, the path of Canadian GDP tracks US GDP in both cases, so both shocks generate spillovers. However, the dynamic responses are different:

[13]If the monetary policy variable is ordered in position j, the restrictions simply imply that the jth column of $\widetilde{AD}$ must have zeros in rows 1 through $j-1$.

a negative shock to the Federal Funds rate generates a gradual expansion, while the sentiment shock generates a short-lived expansion on impact.

Table 3.5 reports the variance decompositions for the two shocks. Augmenting the analysis with the monetary policy shocks does not affect the share of the forecast error variance attributed to the sentiments. The sentiment shock is still the dominant shock for the variability of output and hours at time horizons of less than two years. The monetary policy shock explains most of the variance of the federal funds rate (88% at one quarter) but does not contribute much to variance of the other variables: it explains 5% of the variance of US output and 9% of the variance of Canadian output at a horizon of five years.

Additional controls: Stock Prices, Oil Prices, FAVAR. *Beaudry and Portier* (2006, 2014) identify news shocks with a long-run restriction in a VAR with TFP and an index of stock prices. Our identification of news shocks is medium-run and based on the information content in forward-looking real variables. Stock prices are also forward-looking, so we test the robustness of our identified shock to adding stock prices as an additional control. We use the index of stock prices from *Beaudry and Portier* (2014), ordered second, but we maintain the medium-run identification strategy.[14] Figures 3.16 and 3.17 display the impulse responses of US and Canadian GDP and hours, respectively, to a sentiment shock while augmenting the VAR with the stock price variable. The results are very similar to the baseline specification. In particular, the impact effects are almost identical. The addition of stock prices as a control leads a slightly different dynamic path, but the difference is slight.

Another potential concern is whether our sentiment shock might be picking up oil-price shocks. That is, perhaps not including a measure of oil prices would lead to omitted variable bias in our specification. We test this by augmenting the core VAR with the oil price index

[14]We can increase H^{news} to an arbitrarily large number to approximate the long-run restriction in *Beaudry and Portier* (2006), and we do find that the responses of key variables to the news shock approach their findings (results available on request). However, as long-run restrictions can be problematic (*Faust and Leeper*, 1997), we favor our medium-run approach.

available from FRED.[15] The IRFs of US and Canadian GDP and hours are reported in Figures 3.16 and 3.17. The responses of the core variables to the sentiment shock are almost unchanged relative to the baseline. We also test the responses to adding the US-Canada real exchange rate or US CPI as additional controls. We construct the bilateral real exchange variable using the nominal Canadian-US dollar exchange rate and the US and Canadian consumer price indices from the International Financial Statistics. The units are Canadian basket/US basket, so an increase in the variable is a US appreciation. Figures 3.16 and 3.17 report the responses of US and Canadian GDP and hours when including the real exchange rate, and show that the main results are unaffected.

We augment the core VAR with the first factor identified in *Forni et al.* (2014) to increase the information available about the macroeconomy in identifying news and sentiment shocks. Including this factor further mitigates the possible omitted variables issues in the VAR. Figures 3.16 and 3.17 present the impulse responses of GDP and hours to the sentiment shock in the FAVAR. Reassuringly, we find very similar responses of all core variables with the FAVAR, though the point estimates of the dynamic responses are smaller for longer frequencies.

Table 3.6 reports the share of the forecast error variance of selected core variables attributed to the sentiment shock, while including each of the additional controls. The importance of the sentiment shock for accounting for the forecast error variance of US GDP and hours and Canadian GDP does not differ appreciably from the baseline in each case.

Response of prices. Figure 3.18 displays the impulse responses to the three identified shocks of the key price series: US CPI, US stock prices, oil prices, and the US-Canada real exchange rate.[16] The response of the price variables to the three shocks is consistent

[15]We use a seasonally adjusted consumer price index for all urban consumers, fuel oil and other fuels, series ID CUSR0000SEHE.

[16]The response of US and Canadian GDP and hours in the VAR including US CPI are not substantially different from the other robustness exercises reported in Figures 3.16 and 3.17 so they are omitted to conserve

with theories of news and demand shocks (particularly, demand shocks not driven by shocks to the price variable itself). The US consumer price index increases slightly following the sentiment shock, and the increase is persistent. This response supports the notion that the demand shock embodied in the identified sentiment shock is inflationary. By contrast, there is no response of US CPI to the surprise TFP shock, and prices fall following a news TFP shock. This difference is further illustration that the sentiment shock affects the economy differently from disturbances to technology.

There is no impact response of oil prices to the sentiment shock, ruling out the possibility that the sentiment shock is an oil price shock. Two quarters following the sentiment shock, oil prices if anything rise modestly, indicating that times of positive sentiment do not systematically coincide with low oil prices. In response to the news shock, oil prices fall and stay low, consistent with the decline in inflation documented in *Barsky and Sims* (2012).

The information content of stock prices is evident in the response to the news shock. On impact, there is a large jump in the stock price index, with a further increase for about five quarters followed by slow reversion. However, at the maximum horizon plotted (20 quarters) the index is still substantially above trend. Stock prices also display an impact increase in response to the sentiment shock, but the increase is more muted. This suggests that the sentiment shock is indeed a shock to higher-order beliefs about the economy, which are rational though not based on expected changes to TFP.

The bilateral real exchange rate displays an impact increase only in response to the news shock (this is followed by a gradual decline that approximately coincides with the actual increase in TFP). With the sentiment shock, there is no response for two quarters, and then a slight but persistent US appreciation. The surprise TFP shock leads to a gradual depreciation of the real exchange rate. This is similar to the results in *Nam and Wang* (2015), who estimate the response of real exchange rates to news and surprise TFP shocks.

space.

Spillovers or correlated shocks. Our three US shocks are identified using only US data. Therefore, the possibility that the observed impulse responses of Canadian variables to US shocks are due to exogenously correlated shocks affecting both countries simultaneously cannot be ruled out. For the two technology shocks, this is unlikely to be a problem: Figure 3.4 shows that Canadian TFP does not respond to US surprise and news TFP shocks. However, it may still be that there are exogenous common shocks to US and Canadian sentiments. We evaluate this hypothesis by identifying a surprise TFP innovation, a news shock, and a sentiment shock in Canadian data alone. We then check the correlation of these identified shocks with their US counterparts. Note that if there are indeed spillover effects from US to Canada, this is not a clean exercise: Canadian expectations of future Canadian economic activity will rise following a US sentiment shock, not because optimism increased exogenously in Canada, but because Canadian agents know that a positive sentiment shock in the US will increase Canadian GDP via cross-border transmission. In this sense, identifying a Canadian sentiment shock as if it were a closed economy stacks the deck against us, as those shocks might embody Canadian agents' endogenous revisions of expectations following an increase in US sentiment.

The Canadian expectational variable is an index constructed from the responses to the question "Do you expect overall economic conditions in Canada six months from now to be: Better/Same/Worse", and comes from the Conference Board of Canada. As in *Barsky and Sims* (2012), we construct the composite index by subtracting the percentage of responses answering 'worse' from those answering 'better' and adding 100. This series corresponds most closely to the US confidence series from the Michigan Survey of Consumers. Therefore we compare the Canadian shocks identified with these data with those identified from the five variable core US VAR using the consumer confidence series. Table 3.7 presents the correlations between the US and Canadian shocks identified this way. The correlation of the surprise TFP innovations is 0.16. The US and Canadian sentiment shocks are actually slightly positively correlated, while the news shocks are negatively correlated. However, these correlations are low (0.18 for the sentiment shock and -0.17 for the news shock), indicating that the spillovers observed in the estimated impulse responses are unlikely to be driven primarily by exogenous common shocks.

Our baseline identification strategy does not allow feedback effects from the Canadian variables to the US variables in the VAR coefficients. These restrictions are testable. For each Canadian variable, we perform likelihood ratio tests comparing restricted (baseline) and unrestricted VARs. When the Canadian variables are TFP, output, hours, exports, or imports, we fail to reject the null that the restricted VAR is the true model. For Canadian consumption, however, the null is rejected. Therefore Figure 3.19 presents the responses to the identified shocks when Canadian consumption is the sixth variable and there are no restrictions on the lagged coefficients. Substantively, this does not change the baseline results for any shock, indicating the addition of Canadian consumption as a core variable is not extremely informative for the news or sentiment shocks.

We also attempted to test an alternative model where the Canadian variables were the core entries in the VAR and the US variables were treated as non-core. However, this model does not converge for any US variable. As we cannot estimate the restricted version of this model, we could not evaluate this alternative setup. This is supportive of our assumption that while shocks to the US matter for Canada, the converse is not true.

Additional exercises. The robustness checks above show that the sentiment shock is not a monetary policy shock or an oil price shock. The sentiment shock has characteristics suggesting it is similar to a rational "optimism" shock, where the optimism is not related to a change in productivity. While a proof that this is a pure shock to agents' higher-order expectations is not feasible in this empirical context, we also examine whether this shock is related to uncertainty (second moment) shocks using the correlation of the identified shock series with the VIX index, which measures market expectations of near-term volatility. We obtain a quarterly VIX index using the average aggregation method from FRED, beginning in 1990:Q1. The sentiment shocks are moderately negatively correlated with the growth rates of the VIX index, with a correlation of -0.24. This indicates positive sentiment shocks are more likely to occur when uncertainty growth is low, but the link is not very strong. The sentiment shock could therefore partially capture an uncertainty shock, but it is highly unlikely that fluctuations in uncertainty are entirely responsible for the impact of our sentiment shocks.

Finally, we also examine the properties of the forecast variables used in estimation. If the forecast is a very accurate predictor of future GDP, then the sentiment shock by construction would maximize a substantial fraction of the forecast error variance of GDP as well. Table 3.8 presents the correlations of the forecast and consumer confidence variables, in growth rates, with GDP growth. The forecast of GDP growth is highly correlated with contemporaneous GDP growth (correlation of 0.93), but substantially less correlated with realized future GDP growth one quarter ahead (correlation of only 0.38). The growth of consumer confidence does display the highest correlation with one quarter ahead GDP growth, but this correlation is still very low at 0.25. This is evidence that selecting the sentiment shock as the shock that maximizes its contribution to the forecast error variance of the forecast/consumer confidence variables is not equivalent to simply selecting a shock that by construction explains the bulk of the variation in GDP at business cycle frequencies.

3.5 Discussion

News and noise shocks. Signals about future TFP are likely to be riddled with noise. *Blanchard et al.* (2013) point out that news shocks cannot be separated from noise shocks (unfounded signals about future TFP) in a structural VAR setting, since if the econometrician can extract different paths of variables in response to a noise shock, so can the consumer.

It is clear that these noise shocks are not related to our sentiment shock, which is identified as a fully rational change in forecasts or sentiments orthogonal to surprise and news TFP. That is, in our economy neither the forecaster/consumer nor the econometrician will believe the sentiment shock to be either news or a noisy signal of future TFP. Further, *Barsky and Sims* (2012) assess the importance of noise or 'animal spirits' shocks and find that they do not account for a substantial portion of the relationship between confidence and output. This supports the notion that the responses to the news TFP shocks in Figure 3.2 are informative of the impact of true news.

Demand shocks. Our identified shock could be a combination of several shocks traditionally considered 'demand' shocks. The strategy simply relies on the forecasters rationally expecting an increase in GDP that is not due to current TFP innovations or news/noise about future TFP. To the extent that an increase in 'demand' leads agents to forecast an increase in economic activity, it would be identified as a sentiment shock in our framework. However, our shock is not consistent with all of the results associated with demand shocks previously identified in the literature. For instance, the demand shock identified by *Galí* (1999) as the shock orthogonal to changes in long run labor productivity leads to a temporary increase in labor productivity. Monetary shocks are also commonly proposed as demand shocks. As we demonstrated in Figure 3.15, controlling for monetary policy shocks does not significantly change the shape of the impulse responses to the sentiment shock. Further, monetary policy and sentiment shocks generate very different impulse responses both in the US and Canada.[17] The response of US consumption to the sentiment shock would be consistent with models with taste shocks. For instance, *Stockman and Tesar* (1995) show that the addition of a preference shock increases the volatility of consumption.

[17] *Angeletos et al.* (2014) find that monetary policy shocks deliver closed-economy business cycle moments very similar to their sentiment shock, but require the assumption that they affect the economy in an implausibly large way.

Investment-specific technology shocks also exhibit properties that would appear similar to demand shocks, despite being shocks to technology. The dynamics of output and hours following our sentiment shock do bear a resemblance to those in response to the investment-specific technology shocks as identified in *Fisher* (2006) post 1982:Q3. The magnitudes are very different however. More importantly, the investment-specific shock generates increases in labor productivity. From the responses of output and hours in Figure 3.3, this does not appear to be true for our sentiment shock, increasing our confidence that it is not the investment-specific technology shock.

Relationship to standard neoclassical and New Keynesian models. A large body of work on closed-economy business cycles has established that (i) RBC models driven by a technology shock do well at matching the key moments of US data; and (ii) estimated technology shocks do not deliver impulse responses similar to those in RBC models, calling into question the mechanisms driving the model's success (see also *Galí and Rabanal*, 2005). While surprise TFP shocks have proven of questionable value in explaining US business cycles, they have been even worse at explaining international transmission. The seminal work of *Backus et al.* (1992) showed that even with correlated shocks, the gap between theory and data was large (see also *Justiniano and Preston*, 2010, for a recent statement of this result). Many variants of the original model have been proposed to improve on these results, with limited success. Part of the reason for this failure is that with uncorrelated technology shocks investment increases in the country receiving a positive TFP surprise. As a result, consumption is more highly correlated across countries than output, contrary to the data.

News shocks identified from structural VARs have proven a better fit to the predictions of the neoclassical model. However, due to the wealth effect on labor supply they do not generate the desired impact comovement in the closed economy. A similar outcome appears in an open economy setting: *Kosaka* (2013) finds that news shocks do help generate an international business cycle, but only when the model is parameterized with a low elasticity of substitution between domestic and foreign goods and no wealth effect on labor supply. *Beaudry et al.* (2011) also rely on a low elasticity between the goods produced by different countries in order to generate positive comovement in response to a news shock.

Nominal rigidities have also been proposed as an explanation for the drop in hours in response to a technology shock. The intuition is simple: with sticky prices, only a fraction of firms adjust their prices downward in response to a productivity shock. Therefore, aggregate demand (price level) will rise (fall) less than proportionately to the shock, and hours will fall as a consequence (*Galí*, 1999; *Galí and Rabanal*, 2005). *Galí and Rabanal* (2005) also find in their estimated New Keynesian model that a pure preference shock accounts for the bulk of the variation in the output growth, hours, and the nominal interest rate. However, *Angeletos et al.* (2014) estimate a medium-scale DSGE model allowing for nominal rigidities and find a shock that generates impulse responses for the US economy alone similar to our sentiment shock. Allowing for nominal rigidities helps, but does not seem crucial in their results. At this stage, it is unclear whether sticky prices will be indispensable in a model of the international transmission driven by news and sentiment shocks.

VAR invertibility. For a structural VAR to uncover the "true" shocks in an economic model, *Fernández-Villaverde et al.* (2007) demonstrate that the matrices associated with the state space representation of the model must satisfy a certain invertibility condition. In particular, this is necessary for the VAR to have a moving average representation. This condition can be violated if the structural shocks are not fundamental, i.e. they cannot be recovered from the current and lagged values of the variables. This is a particular concern for news shocks, which are shocks that contain information about future TFP innovations.[18]

[18] This issue would be of less concern if the true VAR only contained surprise TFP innovations and sentiment shocks, as neither are signals of future shocks. However, it would not be possible to identify

This problem is mitigated if the VAR has sufficient information (*Forni and Gambetti*, 2014). *Beaudry and Portier* (2014) discuss this issue in detail, and argue that while VARs with news shocks can be non-fundamental, they don't have to be. Furthermore, *Sims* (2012) shows that even if the VAR is non-fundamental, the resulting impulse response functions can be good approximations of the true impulse responses. In our context, the baseline VAR includes several forward-looking variables, and the results are robust to including additional variables such as stock prices or a factor, that would further mitigate this concern.

3.6 Conclusion

We introduce a novel identification scheme to uncover the effects of surprise TFP innovations, news shocks, and "sentiment" shocks. These shocks have very different implications for international comovement in US and Canadian data. The bulk of high-frequency business cycle comovement can be attributed to the sentiment shocks, while the news shocks are important for medium- to long-term synchronization. Surprise TFP innovations, which are the most common driver of IRBC models, are found to be nearly irrelevant for international business cycle synchronization. Future work will include estimating a dynamic two-country model to quantify the effects of the different shocks.

sentiment shocks without first controlling for news shocks.

Figure 3.1: The Impulse Responses to the US Surprise TFP Shock

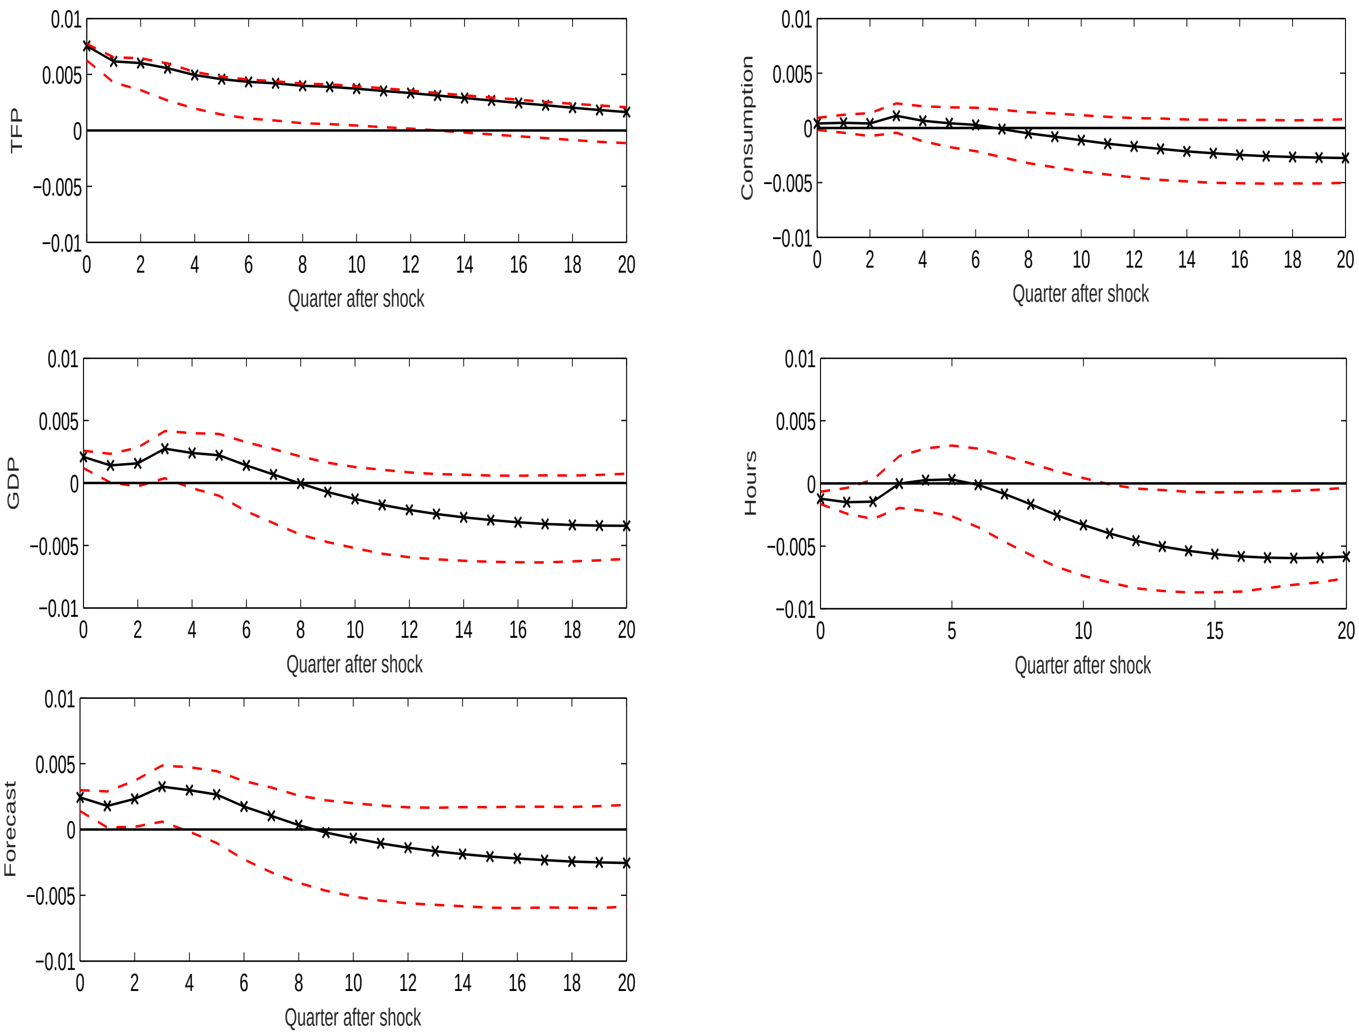

Notes: This figure plots the impulse responses of US TFP, GDP, consumption, hours, and the forecast of US GDP in response to the surprise TFP shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3.2: The Impulse Responses to the US News TFP Shock

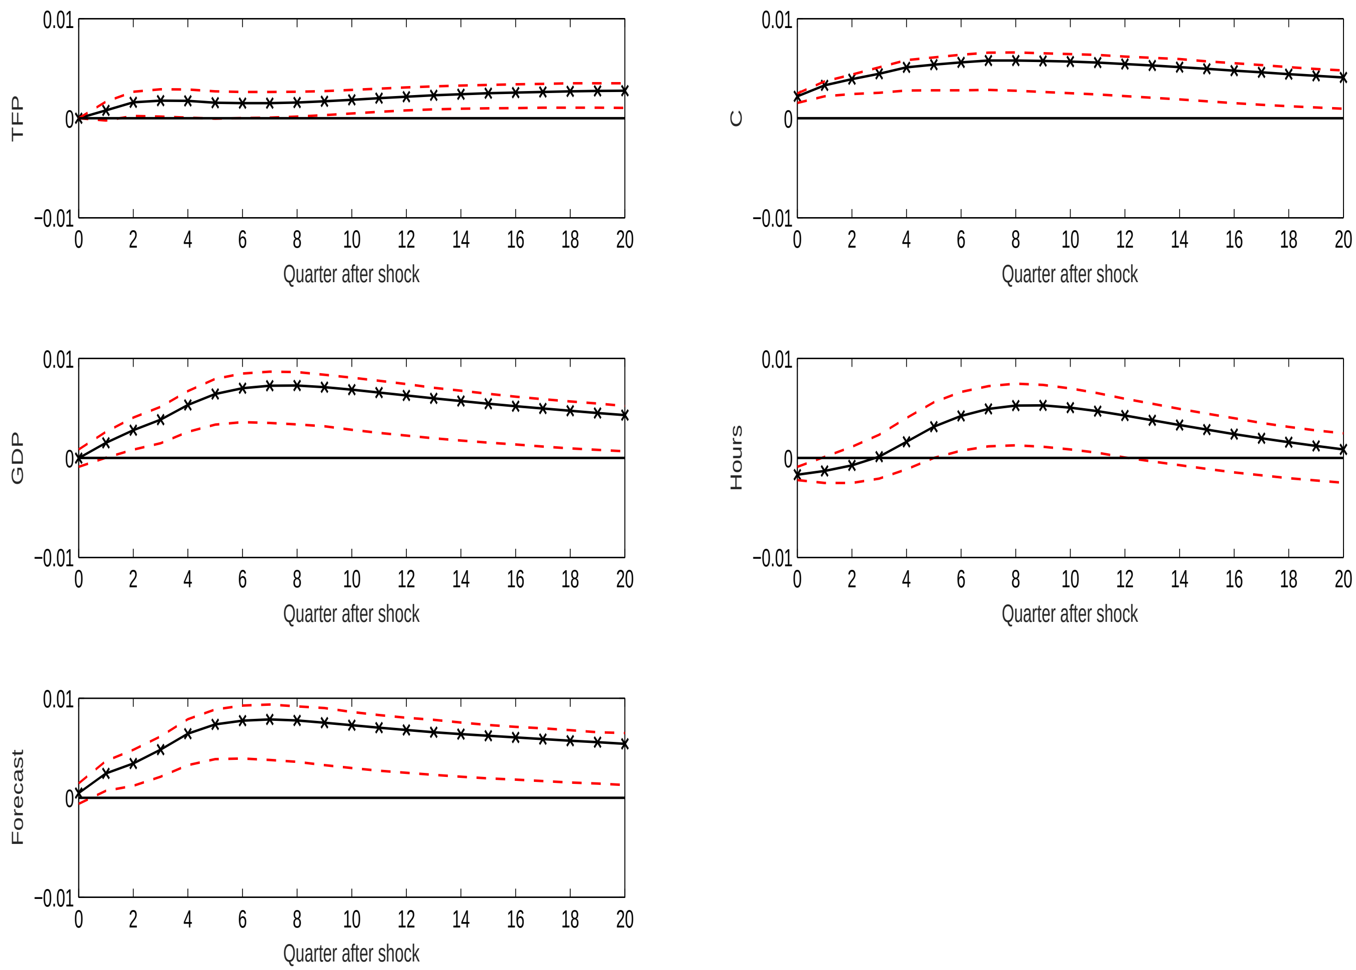

Notes: This figure plots the impulse responses of US TFP, GDP, consumption, hours, and the forecast of US GDP in response to the news shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3.3: The Impulse Responses to the US Sentiment Shock

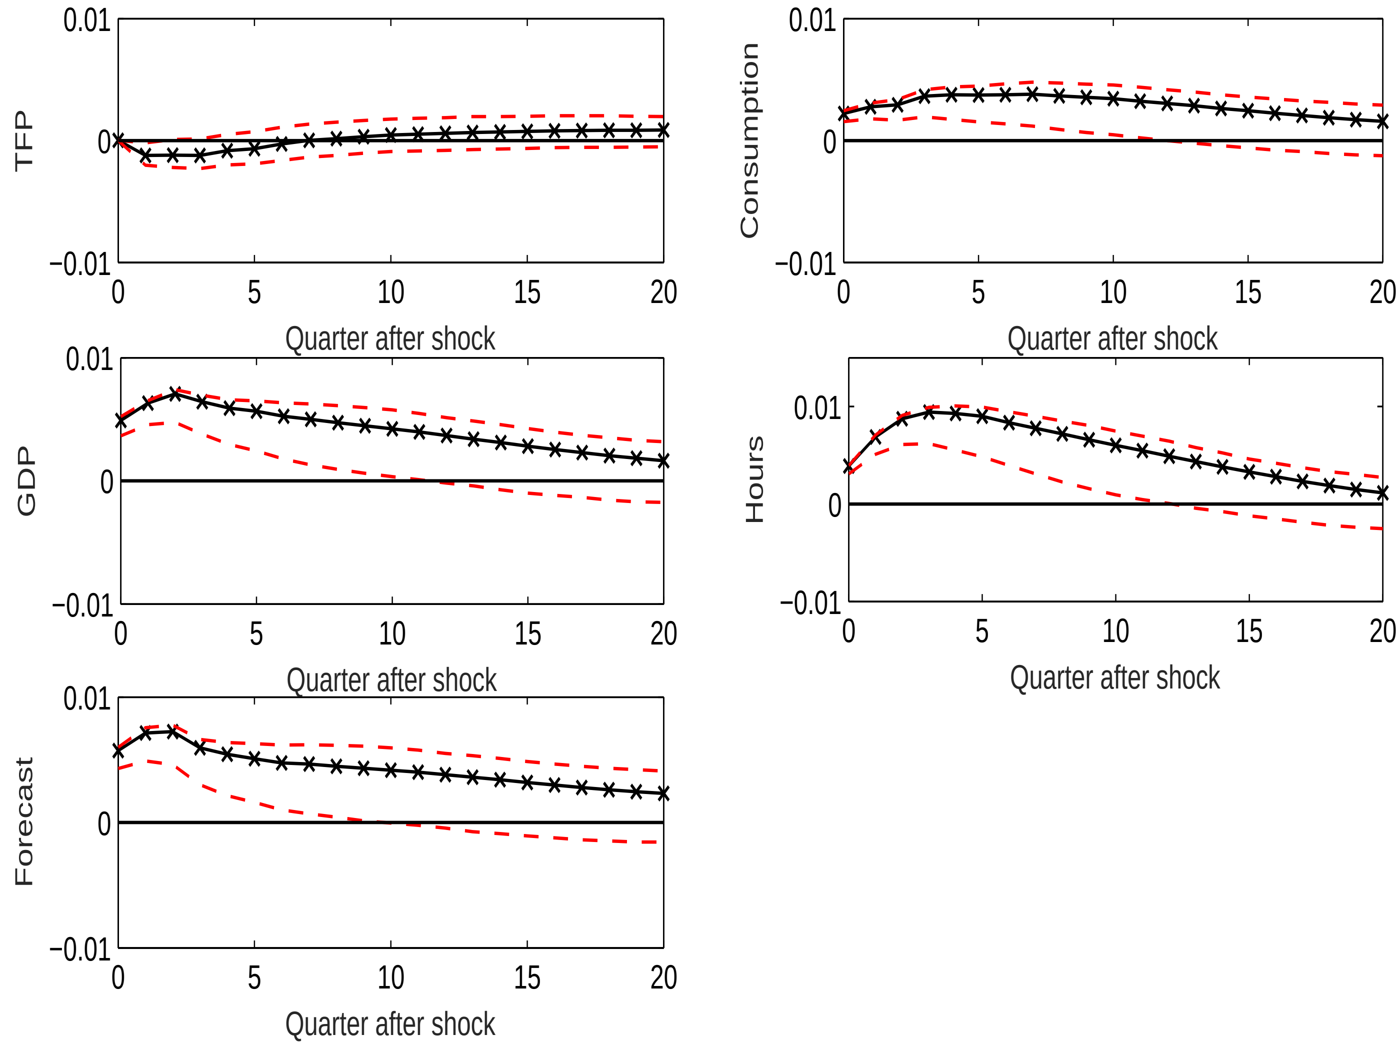

Notes: This figure plots the impulse responses of US TFP, GDP, consumption, hours, and the forecast of US GDP in response to the sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3.4: The Impulse Responses of Canadian TFP to the Three US Shocks

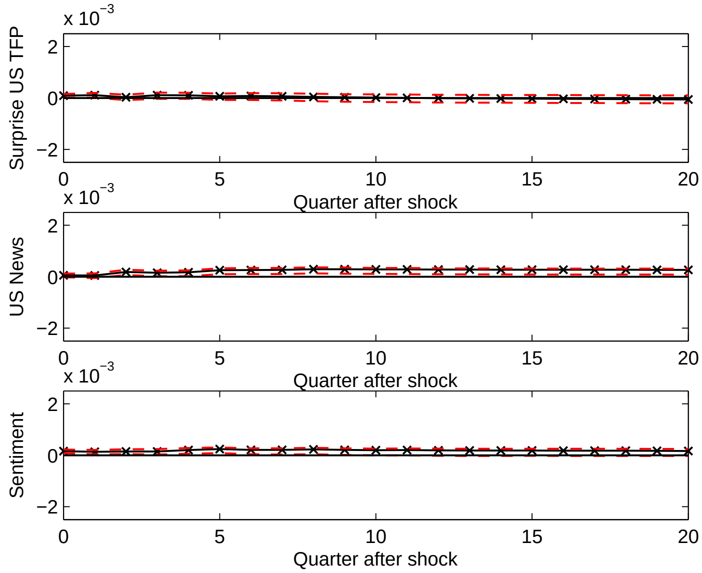

Notes: This figure plots the impulse responses of Canadian utilization-adjusted TFP to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3.5: The Impulse Responses of Canadian GDP to the Three US Shocks

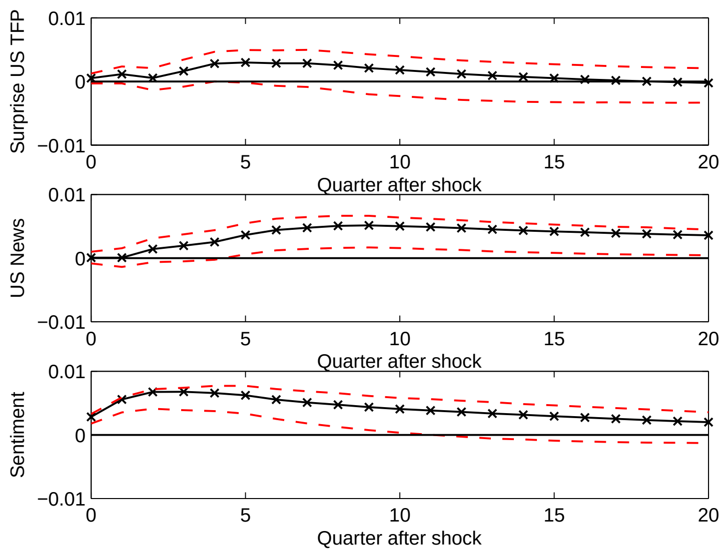

Notes: This figure plots the impulse responses of Canadian GDP to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3.6: The Impulse Responses of Canadian Exports to the US to the Three US Shocks

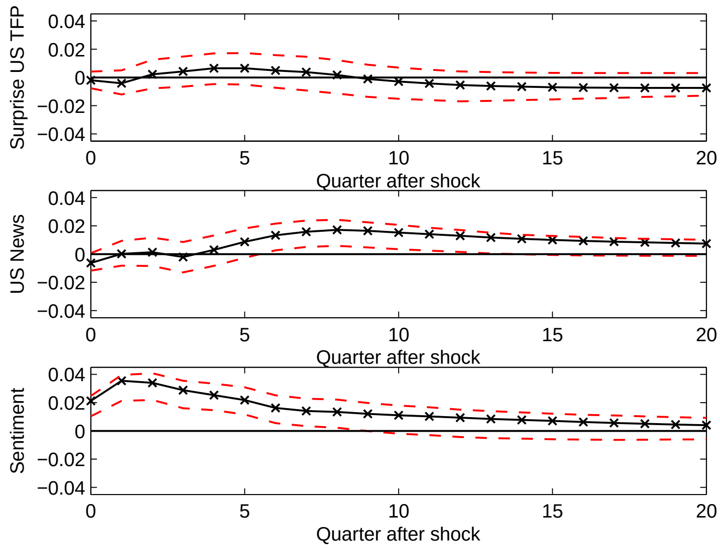

Notes: This figure plots the impulse responses of Canadian exports to the US to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3.7: The Impulse Responses of Canadian Imports from the US to the Three US Shocks

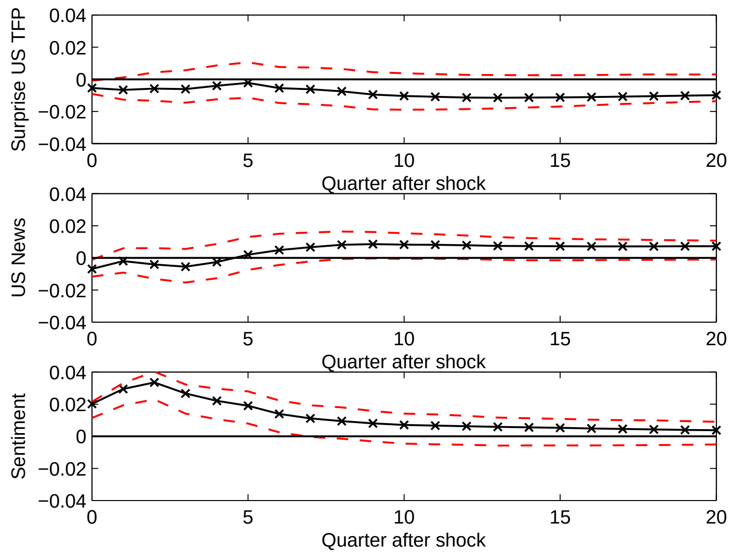

Notes: This figure plots the impulse responses of US exports to Canada to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3.8: The Impulse Responses of Canadian Hours to the Three US Shocks

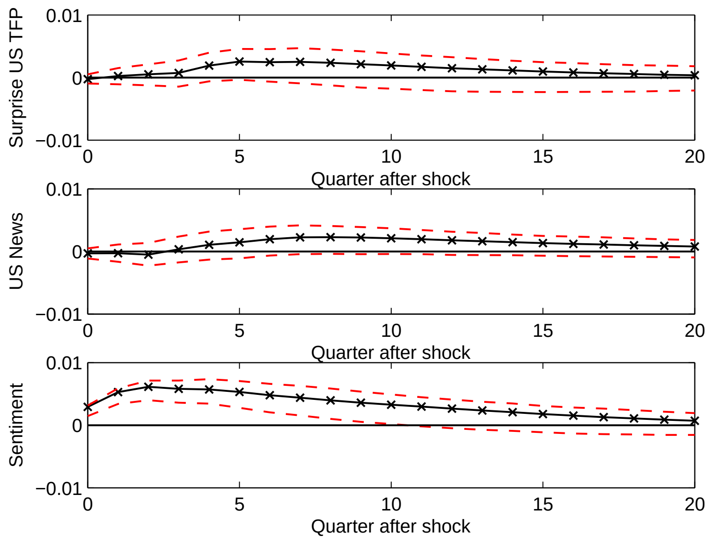

Notes: This figure plots the impulse responses of Canadian total hours to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3.9: The Impulse Responses of Canadian Utilization to the Three US Shocks

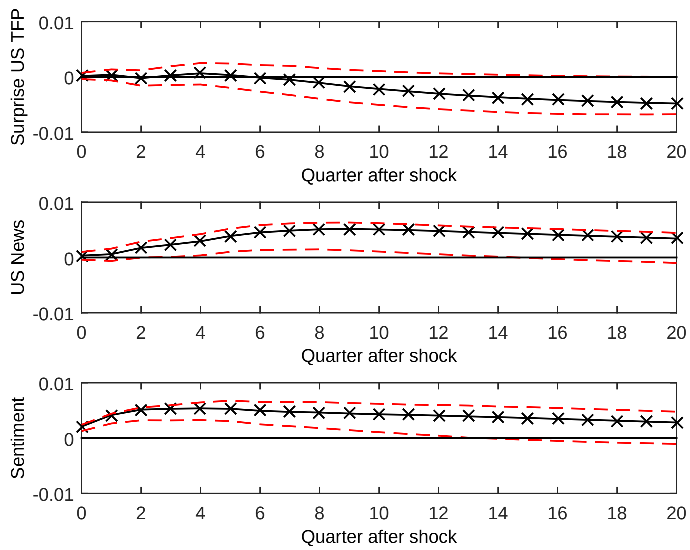

Notes: This figure plots the impulse responses of Canadian utilization to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3.10: The Impulse Responses of Canadian Consumption to the Three US Shocks

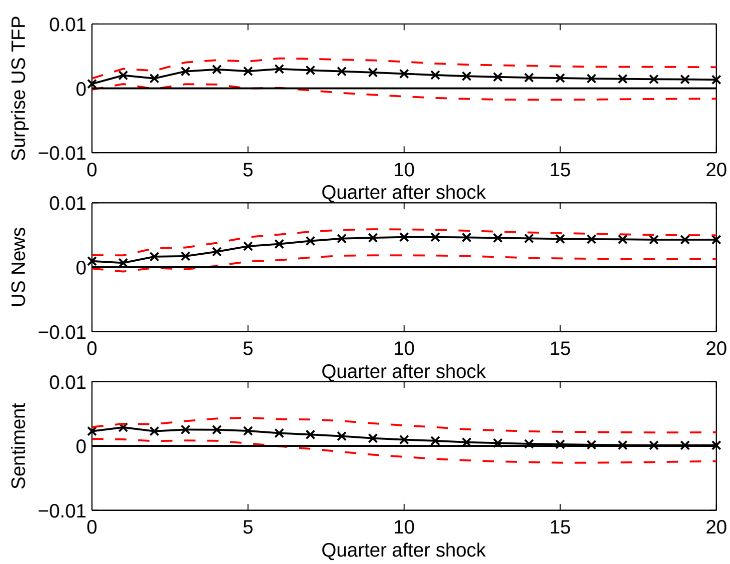

Notes: This figure plots the impulse responses of Canadian consumption to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3.11: Historical Decompositions, Part 1

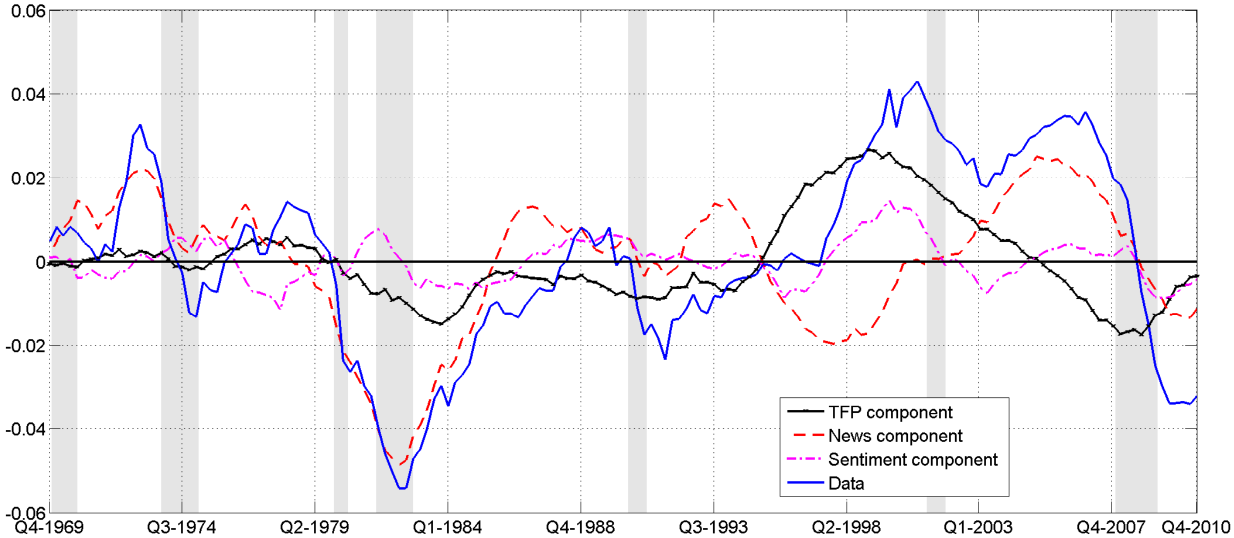

(a) US Output

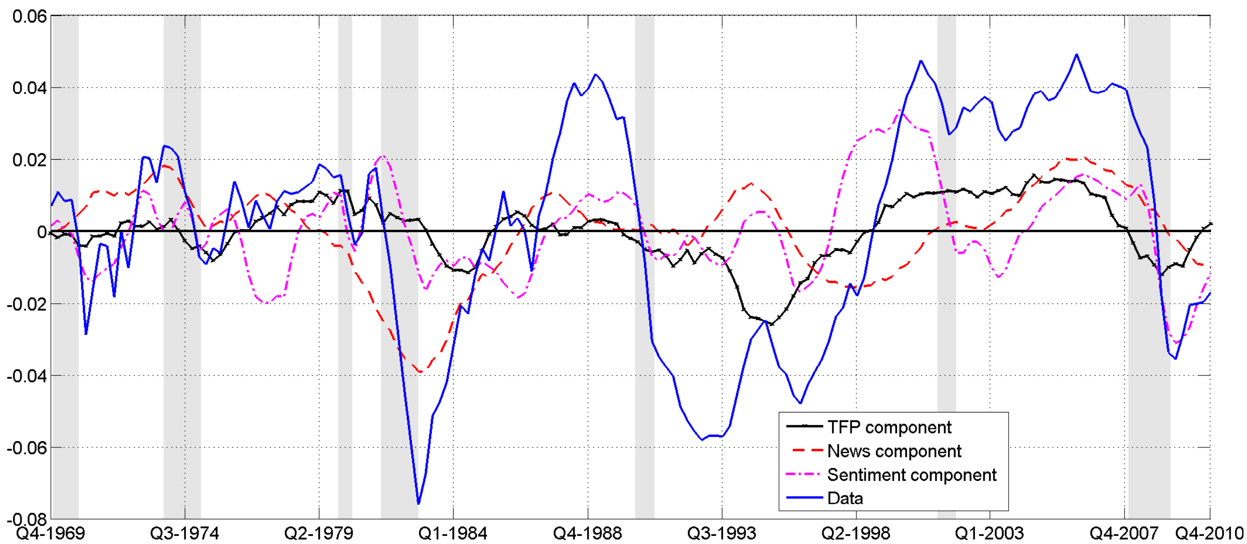

(b) Canadian Output

Notes: These figures show the decomposition of historical data into components due to the three identified shocks. The shaded areas are US NBER recession dates.

Figure 3.12: Historical Decompositions, Part 2

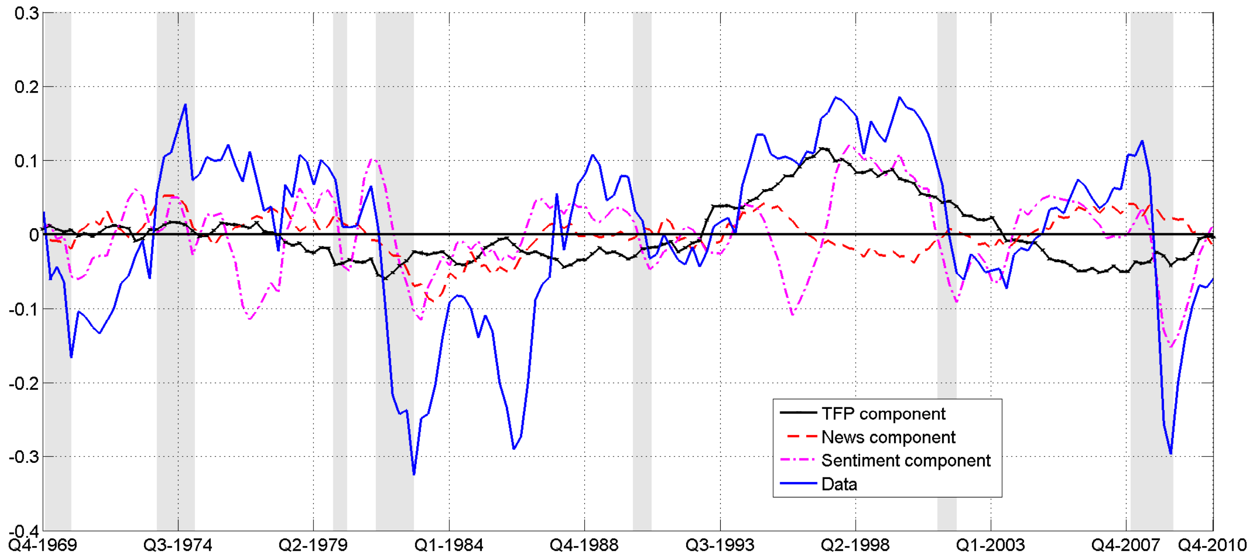

(a) Canadian Imports from the US

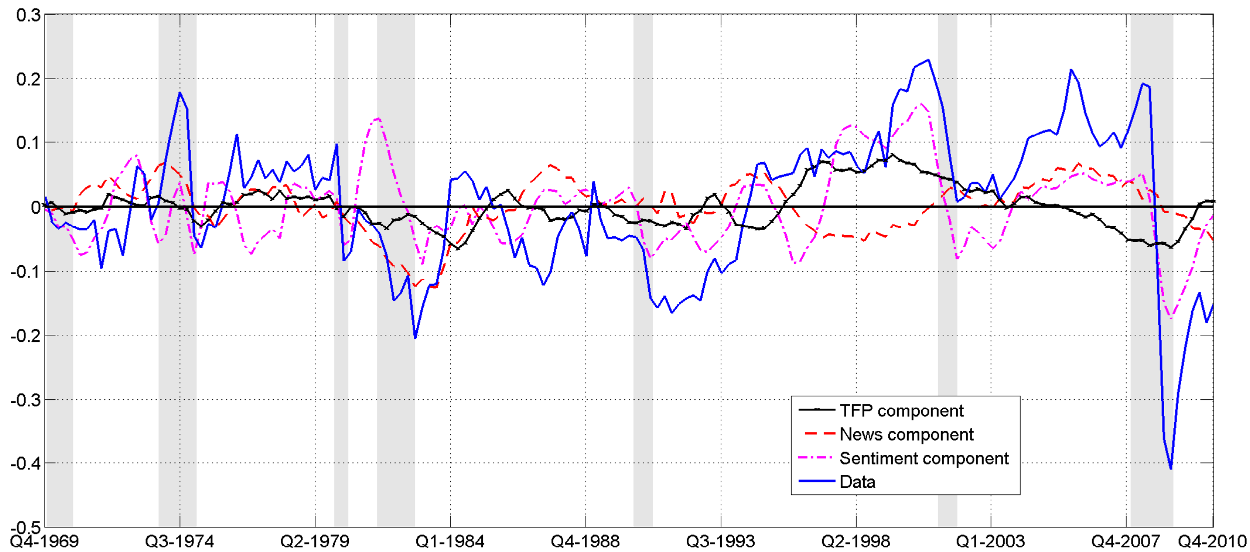

(b) Canadian Exports to the US

Notes: These figures show the decomposition of historical data into components due to the three identified shocks. The shaded areas are US NBER recession dates

Figure 3.13: Historical Decompositions, Part 3

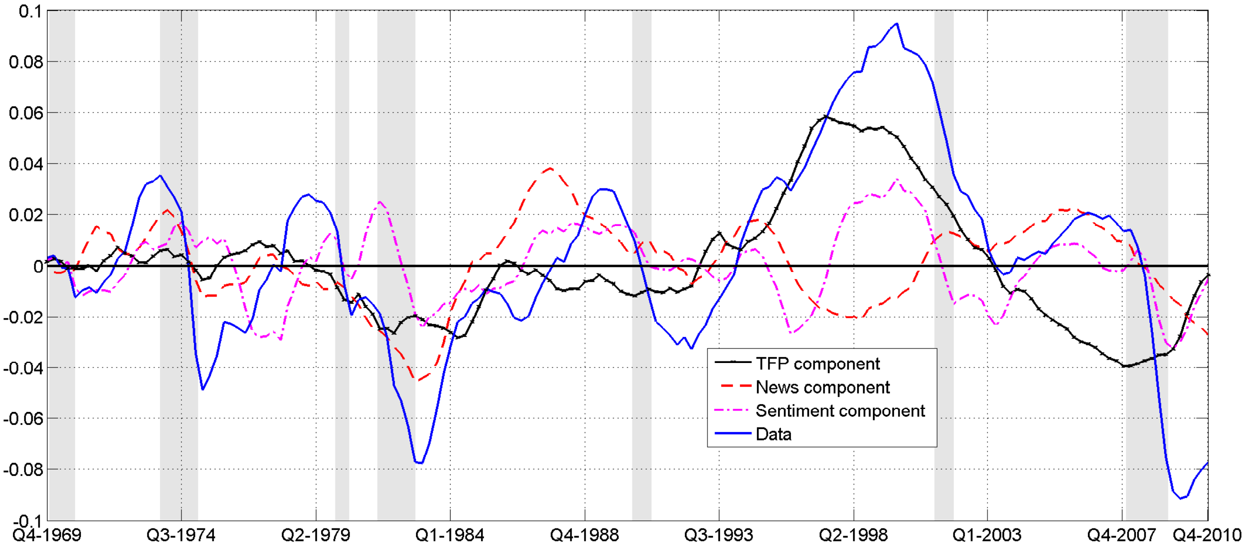

(a) US Hours

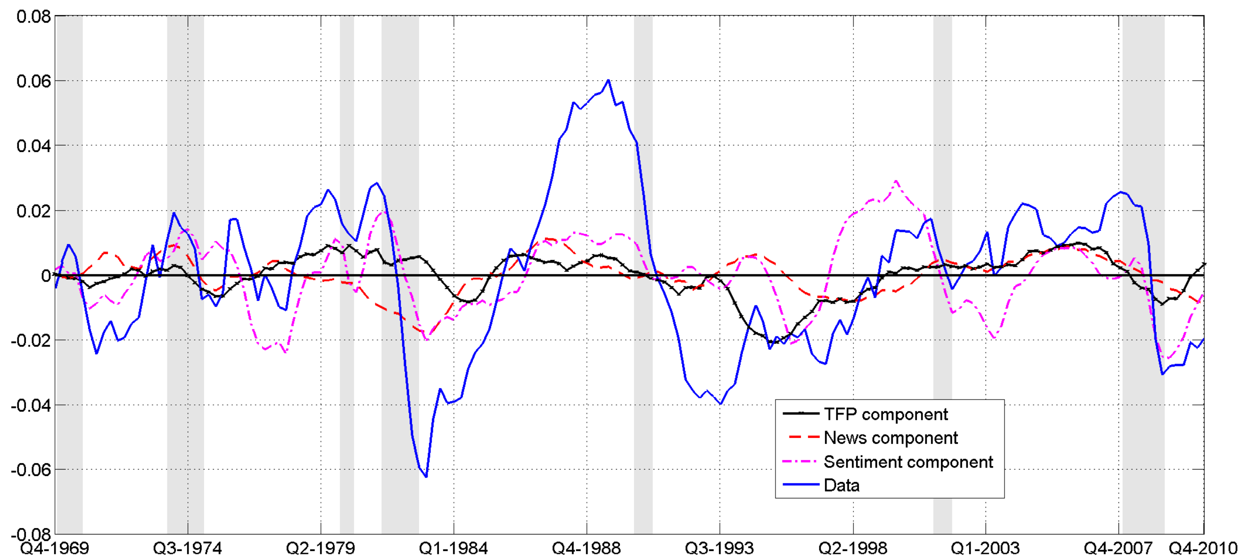

(b) Canadian Hours

Notes: These figures show the decomposition of historical data into components due to the three identified shocks. The shaded areas are US NBER recession dates

Figure 3.14: Historical Decompositions, Part 4

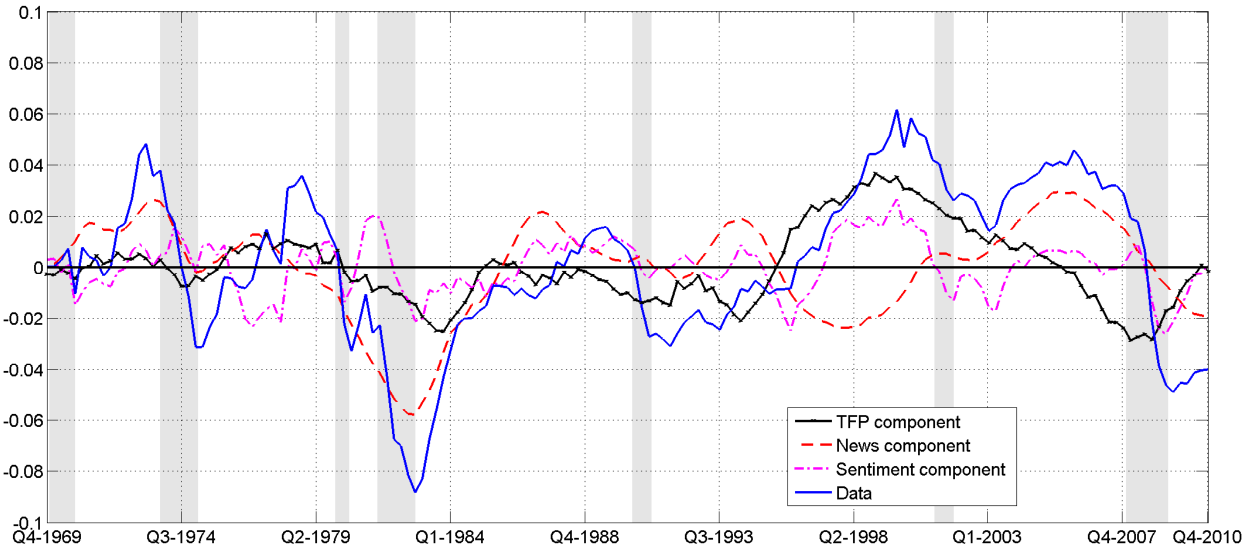

(a) US Consumption

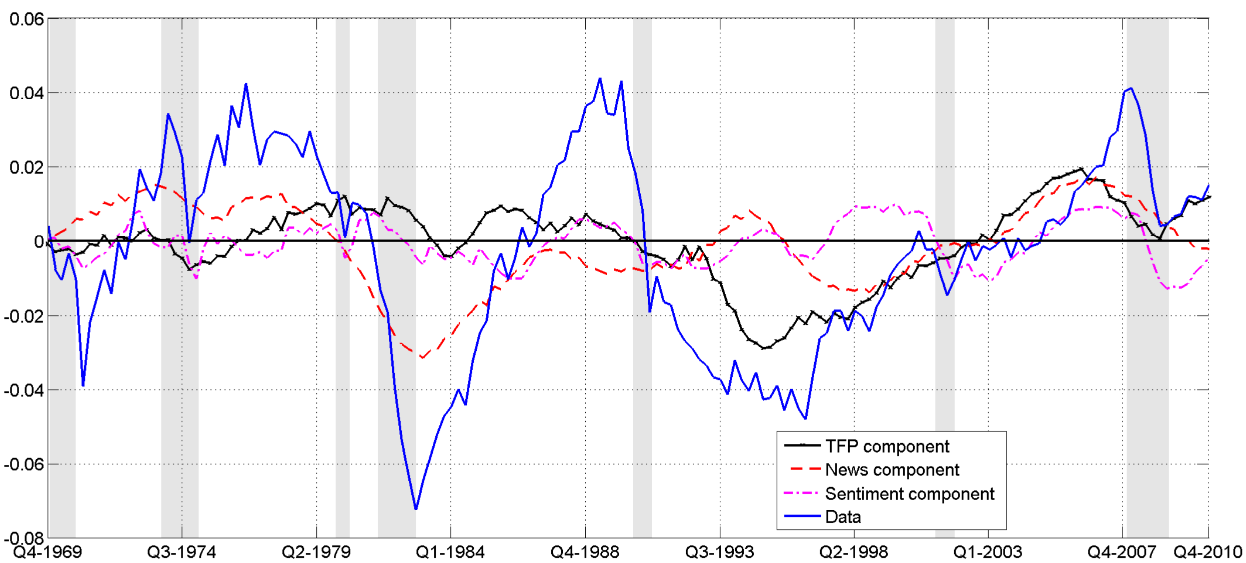

(b) Canadian Consumption

Notes: These figures show the decomposition of historical data into components due to the three identified shocks. The shaded areas are US NBER recession dates.

Figure 3.15: The Impulse Responses of the Core US Variables and Canadian GDP to a Sentiment and Monetary Policy Shock

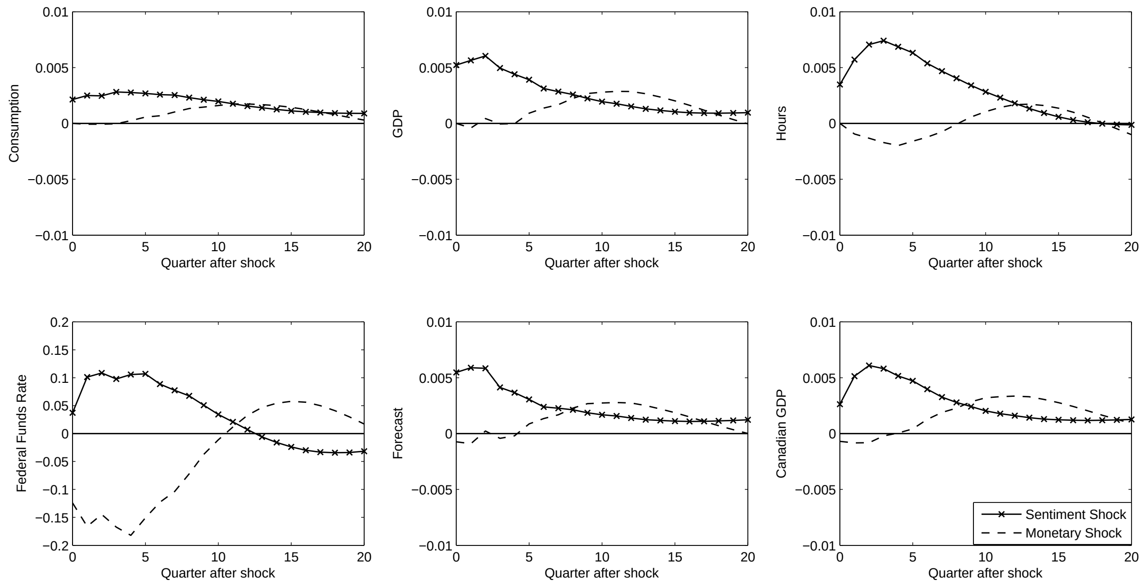

Notes: This figure plots the impulse responses of US Consumption, GDP, Hours, the Federal Funds rate, the Forecast variable and Canadian GDP to a sentiment shock and a monetary policy shock, identified as discussed in Section 3.4.1. Dashed lines plot the responses to the monetary policy shock and lines marked by -x- plot the responses to the sentiment shock.

Figure 3.16: The Impulse Responses US and Canadian GDP to the Sentiment Shock in a VAR with Additional Controls

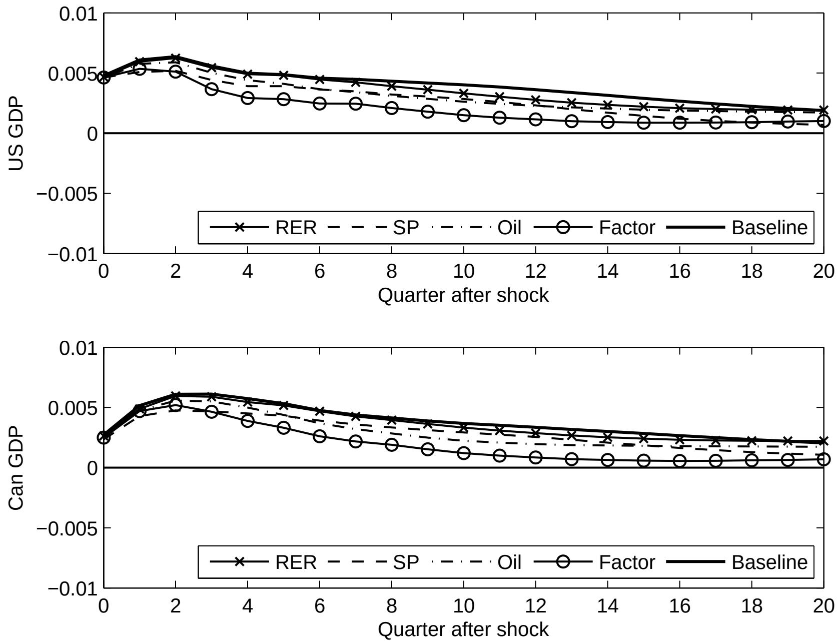

Notes: This figure plots the impulse responses of US and Canadian GDP to the sentiment shock, in a VAR with additional controls identified as discussed in 3.4.1. The additional controls are a measure of stock prices (labeled SP), an oil price index (Oil), the real exchange rate (RER) and a US factor (Factor).

Figure 3.17: The Impulse Responses of US and Canadian Hours to the Sentiment Shock in a VAR with Additional Controls

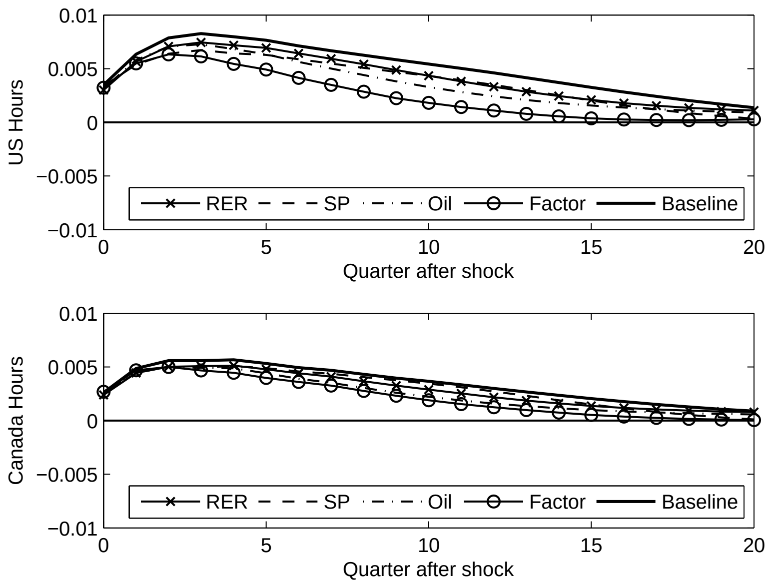

Notes: This figure plots the impulse responses of US and Canadian GDP to the sentiment shock, in a VAR with additional controls identified as discussed in 3.4.1. The additional controls are a measure of stock prices (labeled SP), an oil price index (Oil), the real exchange rate (RER) and a US factor (Factor).

Figure 3.18: The Responses of Price Variables to the Three Shocks

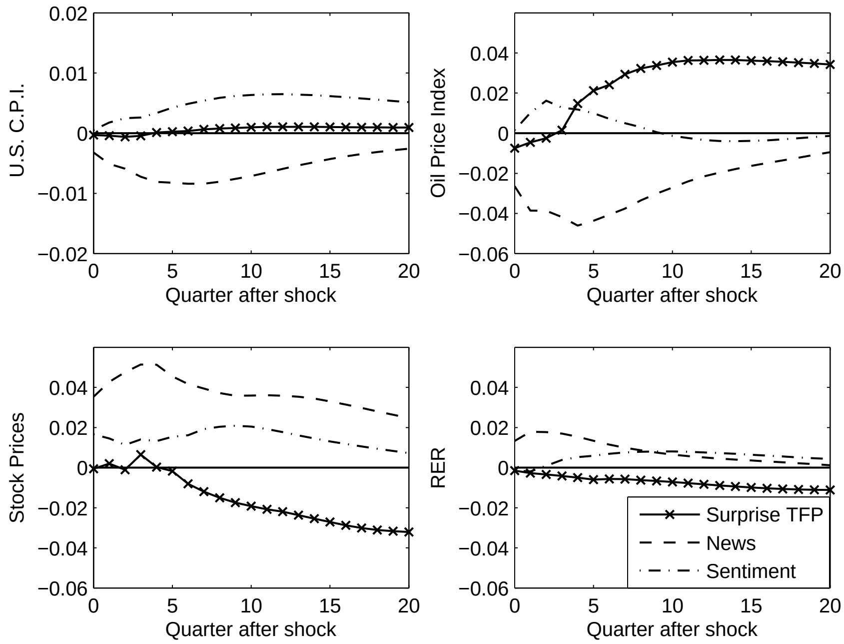

Notes: This figure plots the responses of the US CPI, the oil price index, the stock price index, and the US-Canada real exchange rate to the three shocks.

Figure 3.19: Responses with Canadian Consumption as a Core Variable

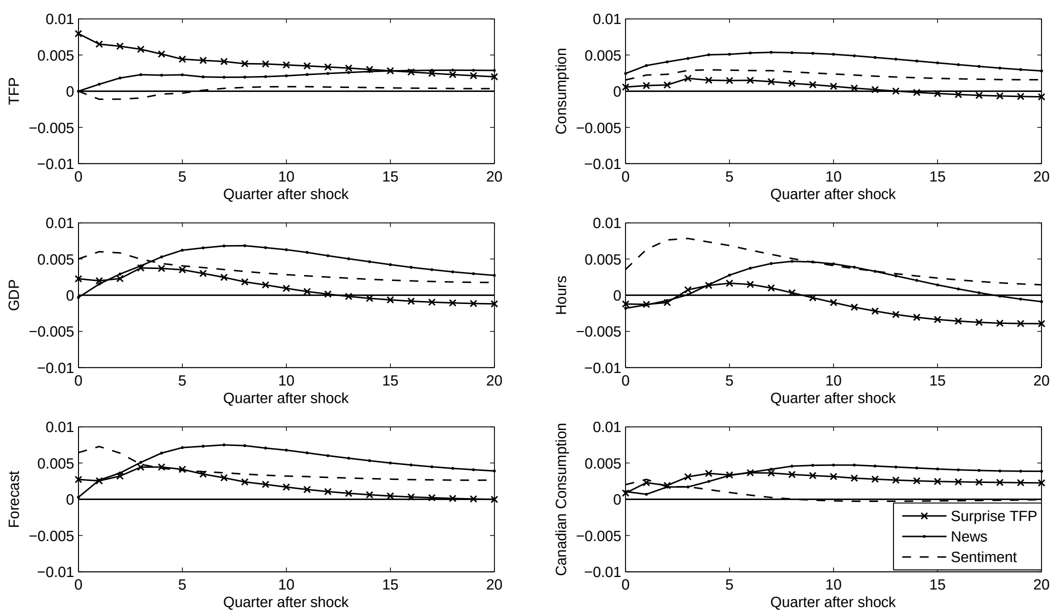

Notes: This figure plots the impulse responses of all key variables to the three shocks in a six variable VAR where Canadian Consumption is treated as a core variable (see section 3.4.1 for details).

Table 3.1: Surprise TFP Shock: Variance Decomposition

Panel A: US

Horizon	TFP	GDP	Consumption	Hours	Forecast
1Q	1.00	0.12	0.01	0.08	0.13
	(0.00)	(0.05)	(0.03)	(0.04)	(0.05)
2Q	0.98	0.08	0.01	0.05	0.08
	(0.02)	(0.04)	(0.03)	(0.04)	(0.04)
1Y	0.93	0.08	0.02	0.02	0.09
	(0.05)	(0.05)	(0.03)	(0.03)	(0.06)
2Y	0.91	0.05	0.01	0.01	0.06
	(0.07)	(0.05)	(0.04)	(0.03)	(0.05)
5Y	0.79	0.08	0.06	0.17	0.05
	(0.12)	(0.10)	(0.11)	(0.13)	(0.08)
10Y	0.59	0.13	0.12	0.28	0.07
	(0.13)	(0.12)	(0.13)	(0.14)	(0.10)

Panel B: Canada

Horizon	Output	Consumption	Hours	Exports	Imports	TFP	Utilization
1Q	0.01	0.01	0.00	0.00	0.02	0.02	0.00
	(0.02)	(0.02)	(0.01)	(0.01)	(0.02)	(0.03)	(0.01)
2Q	0.01	0.05	0.00	0.00	0.02	0.04	0.00
	(0.03)	(0.05)	(0.01)	(0.02)	(0.03)	(0.04)	(0.02)
1Y	0.02	0.07	0.00	0.01	0.02	0.04	0.00
	(0.03)	(0.07)	(0.02)	(0.02)	(0.03)	(0.04)	(0.02)
2Y	0.06	0.10	0.04	0.01	0.02	0.04	0.00
	(0.07)	(0.10)	(0.06)	(0.04)	(0.05)	(0.06)	(0.04)
5Y	0.05	0.10	0.06	0.04	0.11	0.02	0.11
	(0.08)	(0.11)	(0.09)	(0.07)	(0.10)	(0.07)	(0.12)
10Y	0.05	0.08	0.06	0.07	0.16	0.05	0.29
	(0.09)	(0.11)	(0.09)	(0.08)	(0.11)	(0.09)	(0.16)

Notes: This table shows the contribution of the surprise TFP innovation to the forecast error variance of all variables at different horizons. Standard errors are from 2000 bootstrap repetitions.

Table 3.2: News Shock: Variance Decomposition

Panel A: US					
Horizon	TFP	GDP	Consumption	Hours	Forecast
1Q	0.00	0.00	0.36	0.14	0.00
	(0.00)	(0.01)	(0.07)	(0.06)	(0.02)
2Q	0.01	0.03	0.45	0.06	0.05
	(0.01)	(0.03)	(0.08)	(0.04)	(0.04)
1Y	0.04	0.11	0.48	0.02	0.15
	(0.04)	(0.07)	(0.09)	(0.03)	(0.07)
2Y	0.06	0.32	0.52	0.08	0.35
	(0.06)	(0.11)	(0.11)	(0.06)	(0.11)
5Y	0.18	0.45	0.56	0.13	0.46
	(0.11)	(0.12)	(0.13)	(0.08)	(0.12)
10Y	0.36	0.45	0.54	0.12	0.47
	(0.13)	(0.13)	(0.15)	(0.08)	(0.15)

Panel B: Canada							
Horizon	Output	Consumption	Hours	Exports	Imports	TFP	Utilization
1Q	0.00	0.02	0.00	0.02	0.03	0.01	0.01
	(0.01)	(0.03)	(0.01)	(0.03)	(0.03)	(0.02)	(0.02)
2Q	0.00	0.01	0.00	0.01	0.02	0.01	0.01
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
1Y	0.02	0.03	0.00	0.01	0.01	0.07	0.05
	(0.04)	(0.04)	(0.02)	(0.02)	(0.03)	(0.05)	(0.05)
2Y	0.11	0.12	0.02	0.05	0.02	0.18	0.16
	(0.08)	(0.08)	(0.05)	(0.04)	(0.03)	(0.09)	(0.09)
5Y	0.26	0.32	0.06	0.15	0.06	0.39	0.24
	(0.11)	(0.13)	(0.07)	(0.07)	(0.06)	(0.12)	(0.11)
10Y	0.33	0.51	0.06	0.17	0.10	0.47	0.20
	(0.13)	(0.14)	(0.07)	(0.08)	(0.08)	(0.14)	(0.11)

Notes: This table shows the contribution of the news shock to the forecast error variance of all variables at different horizons. Standard errors are from 2000 bootstrap repetitions.

Table 3.3: Sentiment Shock: Variance Decomposition

Panel A: US

Horizon	TFP	Output	Consumption	Hours	Forecast
1Q	0.00	0.65	0.18	0.62	0.85
	(0.00)	(0.06)	(0.05)	(0.07)	(0.05)
2Q	0.02	0.75	0.21	0.71	0.81
	(0.02)	(0.06)	(0.06)	(0.07)	(0.06)
1Y	0.03	0.61	0.22	0.69	0.62
	(0.03)	(0.09)	(0.07)	(0.08)	(0.09)
2Y	0.02	0.36	0.21	0.52	0.35
	(0.04)	(0.10)	(0.09)	(0.12)	(0.10)
5Y	0.02	0.25	0.18	0.35	0.26
	(0.05)	(0.10)	(0.10)	(0.13)	(0.10)
10Y	0.03	0.21	0.15	0.29	0.22
	(0.05)	(0.10)	(0.10)	(0.12)	(0.10)

Panel B: Canada

Horizon	Output	Consumption	Hours	Exports	Imports	TFP	Utilization
1Q	0.19	0.08	0.19	0.25	0.25	0.01	0.18
	(0.06)	(0.05)	(0.06)	(0.07)	(0.07)	(0.01)	(0.06)
2Q	0.32	0.12	0.29	0.38	0.36	0.01	0.28
	(0.08)	(0.06)	(0.08)	(0.08)	(0.08)	(0.02)	(0.07)
1Y	0.41	0.09	0.35	0.44	0.41	0.04	0.33
	(0.09)	(0.06)	(0.10)	(0.09)	(0.09)	(0.04)	(0.09)
2Y	0.34	0.05	0.36	0.39	0.36	0.03	0.25
	(0.10)	(0.05)	(0.11)	(0.10)	(0.10)	(0.04)	(0.10)
5Y	0.29	0.03	0.37	0.32	0.29	0.06	0.17
	(0.10)	(0.05)	(0.12)	(0.10)	(0.10)	(0.06)	(0.10)
10Y	0.26	0.02	0.37	0.30	0.26	0.08	0.12
	(0.10)	(0.05)	(0.11)	(0.10)	(0.10)	(0.07)	(0.09)

Notes: This table shows the contribution of the sentiment shock to the forecast error variance of all variables at different horizons. Standard errors are from 2000 bootstrap repetitions.

Table 3.4: Conditional Correlations

	Data	*Correlation Conditional on:* TFP	News	Sentiment
US, Canada Output	0.73	0.47	0.99	0.99
		(0.28)	(0.07)	(0.05)
US, Canada Consumption	0.51	-0.13	0.94	0.80
		(0.47)	(0.13)	(0.23)
US, Canada Hours	0.68	-0.47	0.46	0.98
		(0.56)	(0.45)	(0.30)
Exports from Canada, US Output	0.66	0.95	0.97	0.89
		(0.19)	(0.16)	(0.09)
Canadian Imports, US Output	0.58	0.79	0.91	0.91
		(0.31)	(0.28)	(0.12)
US Output, US Consumption	0.84	0.97	0.99	0.96
		(0.05)	(0.03)	(0.04)
US Output, US Hours	0.87	0.81	0.20	0.73
		(0.39)	(0.44)	(0.31)

Notes: This table shows conditional correlations of various macroeconomic aggregates in response to the three shocks. Standard errors are from 2000 bootstrap replications. The data column refers to the unconditional correlations from HP-filtered data with a smoothing parameter of 1600.

Table 3.5: Sentiment vs. Monetary Policy Shocks: Variance Decomposition

Panel A: Sentiment Shock							
Horizon	TFP	C	GDP	H	FF	Forecast	Can GDP
1Q	0.00	0.36	0.83	0.64	0.08	0.70	0.18
2Q	0.02	0.32	0.83	0.64	0.20	0.62	0.33
1Y	0.02	0.25	0.64	0.61	0.23	0.47	0.42
2Y	0.02	0.18	0.34	0.42	0.22	0.23	0.33
5Y	0.02	0.10	0.17	0.23	0.20	0.12	0.21
10Y	0.03	0.08	0.12	0.17	0.18	0.08	0.17
Panel B: Monetary Policy Shock							
Horizon	TFP	C	GDP	H	FF	Forecast	Can GDP
1Q	0.00	0.00	0.00	0.00	0.88	0.01	0.01
2Q	0.01	0.00	0.00	0.01	0.75	0.01	0.01
1Y	0.04	0.00	0.00	0.02	0.65	0.01	0.01
2Y	0.10	0.01	0.01	0.02	0.55	0.01	0.01
5Y	0.14	0.03	0.05	0.02	0.47	0.04	0.09
10Y	0.12	0.03	0.04	0.07	0.45	0.03	0.07

Notes: This table reports the forecast error variance contributions of the sentiment shock and the monetary policy shock in a VAR that includes the federal funds rate.

Table 3.6: Other Controls: Sentiment Shock Variance Decomposition

Panel A: US GDP				
Horizon	Factor	RER	Stock Prices	Oil
1Q	0.65	0.63	0.66	0.64
2Q	0.72	0.76	0.72	0.76
1Y	0.54	0.66	0.49	0.63
2Y	0.23	0.45	0.26	0.34
5Y	0.11	0.27	0.15	0.18
10Y	0.09	0.22	0.12	0.15
Panel B: US Hours				
Horizon	Factor	RER	Stock Prices	Oil
1Q	0.59	0.47	0.66	0.53
2Q	0.70	0.58	0.70	0.65
1Y	0.66	0.58	0.57	0.66
2Y	0.37	0.48	0.36	0.47
5Y	0.14	0.30	0.23	0.24
10Y	0.11	0.24	0.19	0.18
Panel C: Canadian GDP				
Horizon	Factor	RER	Stock Prices	Oil
1Q	0.17	0.19	0.15	0.18
2Q	0.28	0.31	0.25	0.29
1Y	0.33	0.43	0.28	0.39
2Y	0.21	0.39	0.21	0.30
5Y	0.12	0.23	0.16	0.20
10Y	0.10	0.20	0.14	0.18

Notes: This table shows the contribution of the sentiment shock to the forecast error variance of the key variables when the core VAR is extended with various controls, discussed in Section 3.4.1. Factor refers to the FAVAR, RER refers to the real exchange rate, and Stock Prices and Oil refer to VARs that include an index of stock prices and an oil price index respectively

Table 3.7: Correlations of Shocks: US and Canada

Contemporaneous Correlations	US TFP	US News	US Sentiment
Canada TFP	0.16	0.01	0.17
Canada News	-0.04	-0.02	-0.17
Canada Sentiment	-0.02	0.11	0.18

Notes: This table shows the contemporaneous correlation of the three identified shocks between the US and Canada. They are identified in separate VARs with only the core variables corresponding to each country.

Table 3.8: Cross-Correlations of Forecast and Confidence with GDP

	Lags										
Variable	-5	-4	-3	-2	-1	0	1	2	3	4	5
GDP Forecast	-0.07	0.04	0.01	0.18	0.20	0.93	0.38	0.24	0.14	0.11	-0.03
Consumer Confidence	-0.17	-0.12	-0.21	-0.13	-0.15	0.21	0.25	0.18	0.12	0.25	0.03

Notes: This table shows the cross-correlation of the GDP forecast and the Consumer Confidence variable with GDP at leads and lags. All variables are in growth rates.

APPENDICES

APPENDIX A

Chapter 1 Appendices

A.1 Chapter1: Data Appendix

A.1.1 Creating Constant Manufacturing Sample

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

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

Figure A.1 illustrates how our constant manufacturing sample compares to manufacturing employment from two other sources: published totals from the Current Employment Survey and *Pierce and Schott* (2013).

A.1.2 Identifying Plants Owned by Multinationals

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

A.1.2.1 External Sources of Information

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

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

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

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

A.1.2.2 The Matching Procedure: An Overview

The matching procedure uses a set of record linking utilities described in *Wasi and Flaaen* (2014). This program uses a bigram string comparator algorithm on multiple variables with differing user-specified weights.[1] The primary variables for matching include the establishment name along with geographic indicators of street, city ,zip code, and state.

Recognizing the potential for false-positive matches, w use a relatively conservative criteria for identifying matches between the directories and the Census Bureau data. In practice, the procedure generally requires a match score exceeding 95 percent, except in those cases

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

where ancillary evidence provides increased confidence in the match.[2] This matching proceeds in an iterative fashion, in which a series of matching procedures are applied with decreasingly restrictive sets of matching requirements. In other words, the initial matching attempt uses the most stringent standards possible, after which the non-matching records proceed to a further matching iteration, often with less stringent standards. In each iteration, the matching records are assigned a flag that indicates the standard associated with the match.

See Table B.1 for a summary of the establishment-level match rate statistics by year and type of firm. Table B.2 lists the corresponding information for the Uniworld data.

A.1.3 Creating Panel of Multinational Plants

The external directories allow for relatively easy categorization of the multinational status of U.S. plants. If the parent firm contains addresses outside of the United States, but is headquartered within the U.S., we designate this establishment as part of a U.S. multinational firm. If the parent firm is headquartered outside of the United States, we designate 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.[3] 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.

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

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

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.1.3.1 Within-Year Rules

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

A.1.3.2 Checks and Rules for Across Years

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

A.1.4 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.[4] 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.[5] Such product-level production data exists as part of the "Products" trailer file of the Census of Manufacturers. As detailed in *Pierce and Schott* (2012) (see page 11), combining import, export, and production information at a product-level is useful for just such a purpose.

It is important to acknowledge that the Census data on trade exists at the firm level, while the other information used in this paper is , principally, at the plant level. Utilizing the establishment industry information, however, will allow us to parse a firm's trade based on the intermediate/final distinction for a given establishment, thereby generating some heterogeneity in firm trade across establishments.[6]

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

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

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

A.1.4.1 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 of product p in industry j can be written as:

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

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

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

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

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

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

Acknowledging this concern, we set an exogenous threshold level W such that any p in a given j with $S_{pj} > W$ is classified as a final good product for that industry. The upper portion of Table B.3 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.[7] 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.1.4.2 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.

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

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

$$\left.\begin{aligned} M_{ij}^{\text{int}} &= \sum_{p \notin P_j} M_{ipj} \\ \\ M_{ij}^{\text{fin}} &= \sum_{p \in P_j} M_{ipj} \end{aligned}\right\} \qquad \text{where} \qquad P_j = \{p \mid S_{pj} \geq W\}\,. \tag{A.1}$$

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

A.1.5 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.1.3). As the Corporate Directories are matched at the establishment level, when aggregating up to the firm, there are occasional conflicts in the definition of a firm between the Census and the Directories. We rely on the Census definition of a firm. Conflicts are resolved as follows:

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

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

Note, firm identifiers in the Census are sometimes problematic longitudinally. An example is that the firm identifier changes when the firm goes from being a single unit to a multi-unit establishment. Further, mergers and acquisitions can lead in some cases to the birth of a new firm identifier, and in others to the continuation of one of the merged identifiers. As such, results pertaining to the extensive margin that use the firm identifier as the basis of analysis will be overstated. This is a problem faced by all longitudinal firm-level analysis using Census Bureau data. We do not use longitudinal information in classifying U.S. and foreign multinationals, or non multinational firms. However, some of our analysis in 3.3 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 A.1: DCA Establishments and Match Rates, by Firm Type

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

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

Figure A.1: Comparison of Constant Manufacturing Employment Samples: 1993-2011

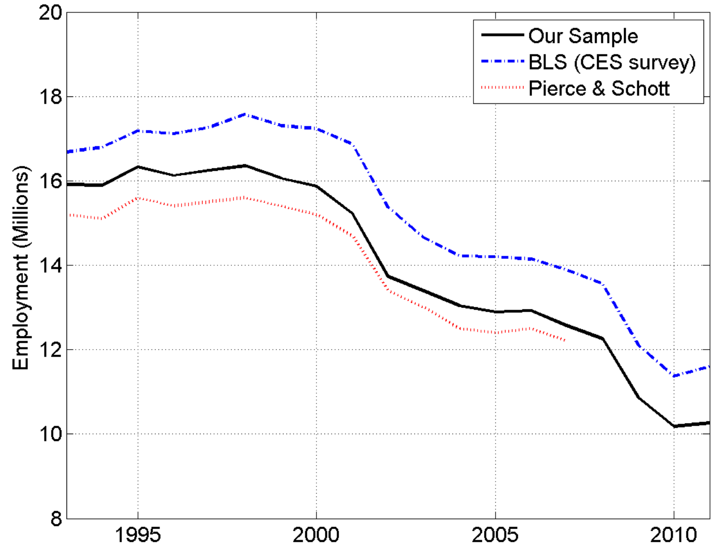

Source: BLS, *Pierce and Schott* (2013) and the LBD.

A.2 Chapter 1: Additional Results

A.2.1 Within-Group Decompositions

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

Employment growth is affected jointly by the rate of job creation and that of job destruction, and further by the extent to which this pertains to establishment births and deaths rather than employment changes at continuing establishments. Following the common practice exemplified by *Davis and Haltiwanger* (2001) we decompose the changes in within-group

Table A.2: Uniworld Match Statistics: 2006-2011

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

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

Table A.3: Appendix Table Comparing the Results from Threshold Values W

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

This table is applicable to the year 2007.

employment into job creation/destruction rates, separated by intensive and extension margins. Formally, let employment at establishments in group $S \in \{D, X, MH, MF\}$ in time t be denoted as $E_{S,t}$. Defining S^{+}_{t-1} and S^{-}_{t-1} as the set of establishments in S that increase (decrease) employment between t-1 and t, we can then define the job creation ($JC_{S,t}$) and destruction ($JD_{S,t}$) rates as:

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

$$\text{Job Destruction Rate:} \qquad JD_{S,t} = \frac{\sum_{i \in S^{-}_{t-1}} |\Delta e_{i,S,t}|}{\left(E_{S,t} + E_{S,t-1}\right)/2} \tag{A.3}$$

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

Figures A.2 report the intensive job creation/destruction rates of the three relevant groups we study. In Panel A, the job creation rates show both cyclicality and a secular decline for both domestic and exporting establishments.[8] The job creation rate for multinational firms is lower and slightly less cyclical than the other groups. The high cyclicality of job destruction rates is very much evident in Panel B of Figure A.2. Taking into account that both JC and JD rates are known to decrease with both firm size and firm age (see *Davis and Haltiwanger* (1992) and *Haltiwanger et al.* (2013)), and that multinationals are 3 times (20 times) larger than exporting (domestic) establishments, it is striking how similar the job destruction rates for multinationals are to the other two groups. With this in mind, it appears that job destruction plays a more important role for multinationals relative to non-multinational establishments, and has been an important driver of the observed aggregate decline in employment in this group.

Figure A.3 translates the job creation and destruction rates into a net measure of employment gains by type of establishment. Panel A shows that multinational establishments have had lower net growth rates than the domestic/exporting groups in nearly every year of our sample. While domestic/exporting firms were on net adding jobs following the 2001 recession in the U.S., the multinational establishments continued to shed jobs through the 2008/2009 financial crisis. In this way multinationals are shown to be a contributor to the "jobless recovery" of the 2003-2007 expansion.

A.2.2 Other Results on Transitions

A.2.2.1 Assumptions of Firm-Level trade Following an Establishment Death

There are at least two distinct approaches to account for the role of establishment death on the import activity at the *firm*-level. The estimates in Figure 1.3 fill in the post-death values for a given establishment with the actual imports of the firm associated with that

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

establishment. [9] This approach better captures the import substitution that may occur if a plant is closed in response to offshore activities. If this was the case, we would see a larger import differential relative to the benchmark calculation. On the other hand, if establishment deaths are associated with broad firm decline, then this differential import measure would be smaller relative to the benchmark.

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

Below we demonstrate the effects of these assumptions on our estimates of import behavior surrounding the event study . In our baseline sample underlying Figure 1.3, 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 A.4 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.

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

A.2.2.2 Other Trade Effects Following Multinational Transitions

We estimate equation (1.2) using various types of firm-level trade corresponding to establishments that transition into part of a multinational firm. The results pertaining to related-party and arms-length intermediate imports are shown in Figure 1.3. 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 A.5. Perhaps more surprisingly, we also find strong growth in export volumes in the years following a multinational transition. The increase in exports (shown in Figure A.6), together with the broad increase in importing activity, demonstrates the overall modifications of the production structure of these firms that accompany expansions abroad.

What do our results imply in the aggregate? To convert the estimates from Figure 1.2 into a measure of total job gain/loss from new multinational activities.

Further details are available in Appendix A.2.3.1.

A.2.3 Quantifying Job Loss: Back-of-the-Envelope Calculations

A.2.3.1 Job Loss from Multinational Transitions

This section describes how we convert the estimates on relative employment growth rates of new multinational plants into a measure of the aggregate net gains of employment. The coefficients from Figure 1.2 represent relative employment effects, expressed in percentage points, of a transitioning plant. These effects represent averages that span the entire period (1993-2011) for which plants may be transitioning into a multinational firm. To translate these percentage points into jobs, one challenge is to identify the appropriate base on which to apply the relative percentage differentials. Unfortunately, the average size of transitioning plants is not currently available. However, using the productivity/size ordering of firms implied by models such as *Helpman et al.* (2004), and confirmed using similar data sources in *Flaaen* (2013a), 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 (1.2) run out. (Essentially, we assume that the growth rate differentials in all years $t > 10$ are zero.) Of course, extrapolating the estimates beyond year 10 would magnify the job losses – adding an additional percentage point or two in accounting for the total job loss – resulting from multinational transitions.

Formally, we compute the job loss as

$$\sum_{t=1994}^{2010} T_t E_t \sum_{i=1}^{\min\{10,2010-t\}} \delta_i \prod_{j=1}^{i-1} (1+\delta_j) \tag{A.4}$$

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 (1.2). Table A.4 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.

A.2.3.2 Job Loss from all Multinational Activity: Total

A similar exercise can be done using the coefficient estimates from Table 1.5. This calculation is somewhat easier in that we simply apply the employment growth rate differential to the average establishment size of multinationals, and then multiply by the total number of multinational establishments in each year. Table A.5 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.

A.2.4 Regression Evidence: Robustness

Table A.6 presents results from running the specification in equation 1.1 for various subsamples of our data. The results are also robust to including lagged establishment or firm employment growth rates as controls (available upon request).

Figure A.2: Job Creation and Destruction Rates by Group: Intensive Margin

A. Job Creation Rates

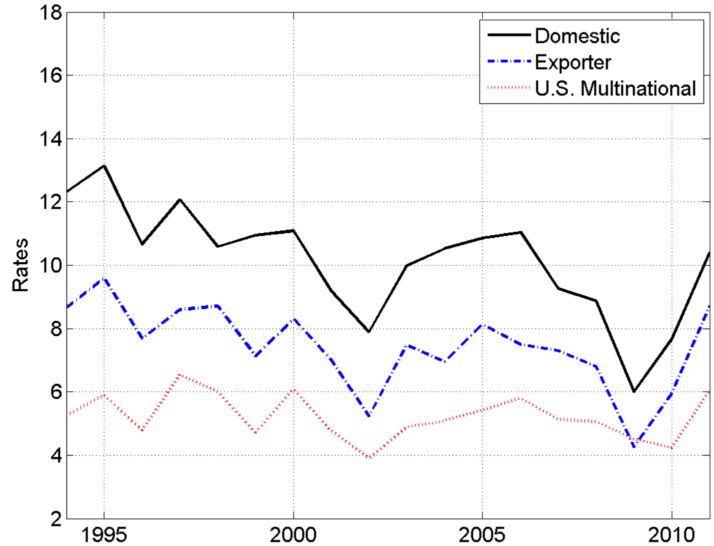

B. Job Destruction Rates

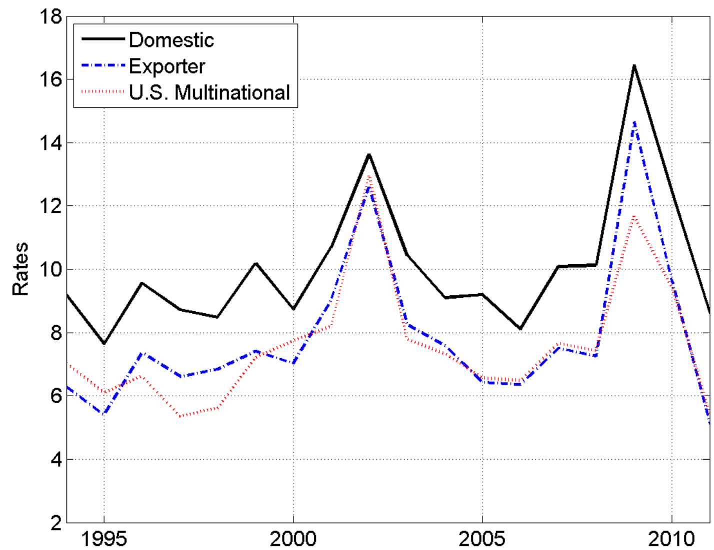

Source: LFTTD-DCA-UBP as explained in text.
These figures report the decomposition of within-group growth rates of employment at the intensive margin. See equation A.2 in the text.

Figure A.3: Net Growth Rates by Group:

A. Intensive Margin

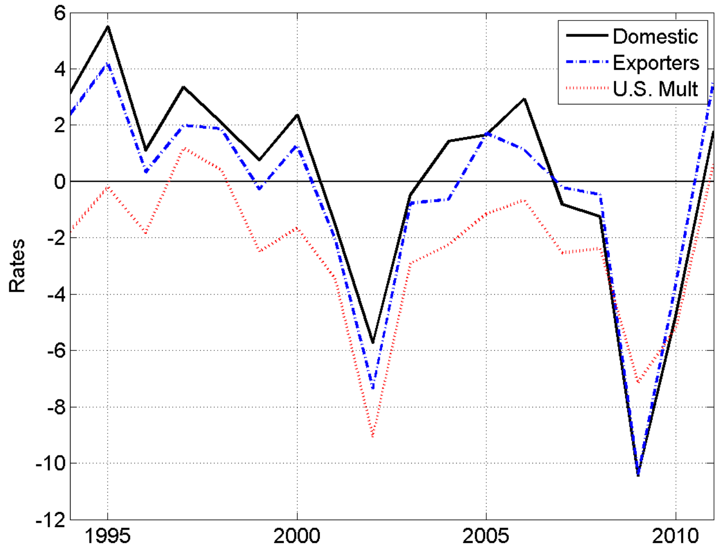

B. Extensive Margin

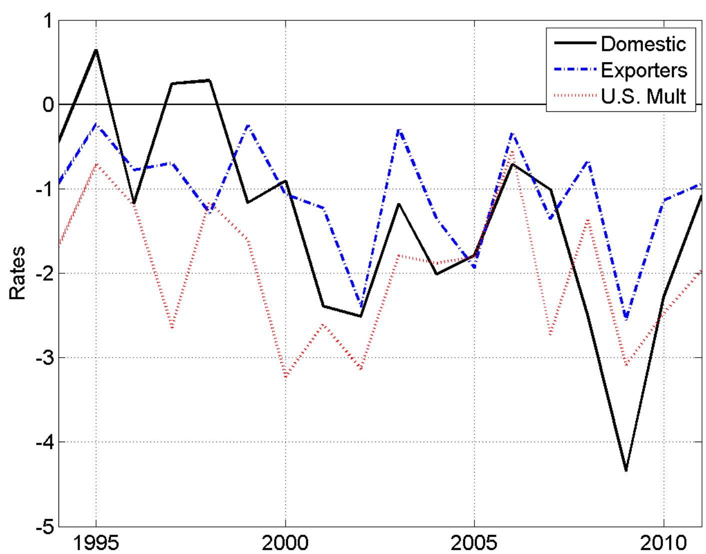

Source: LFTTD-DCA-UBP as explained in text.
These figures report the decomposition of within-group growth rates of employment at the intensive margin. See equation A.2 in the text.

Figure A.4: Importing Differentials of Multinational Transitions, Balanced Panel

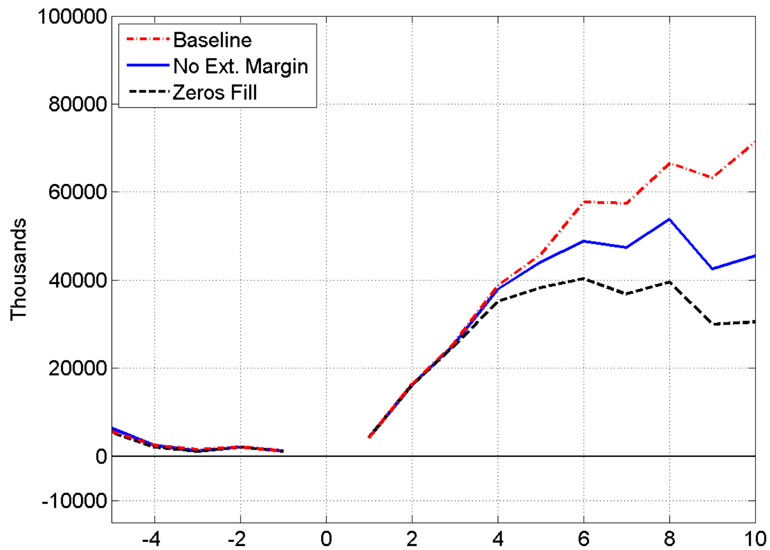

Source: LFTTD-DCA-UBP as explained in text.
This figure reports the related-party intermediate input imports of the parent firm of the transitioning establishment relative to a control group, as outlined in equation 1.2. *Zero Fill* refers to a balanced panel with zeros for trade after an establishment death. *No Ext. Margin* refers to the sample with no extensive margin effects following the establishment death.

Figure A.5: Final Goods Importing Differentials of Multinational Transitions

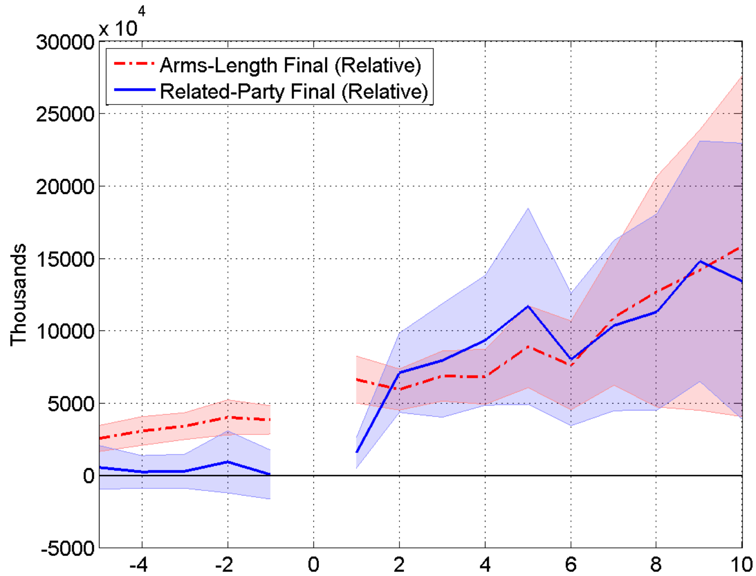

Source: LFTTD-DCA-UBP as explained in text.
This figure reports the related-party and arms-length final goods imports of the parent firm of an establishment that transitions into part of a multinational firm in year ($t = 0$), relative to a control group based on interacted effects of firm age, establishment size, and industry (in year $t = -1$). See equation 1.2, modified to reflect firm-level imports as dependent variables. The shaded area corresponds to a 95 percent confidence interval.

A.3 Chapter 1: Structural estimation appendix

A.3.1 Estimation

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

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

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

Figure A.6: Exporting Differentials of Multinational Transitions

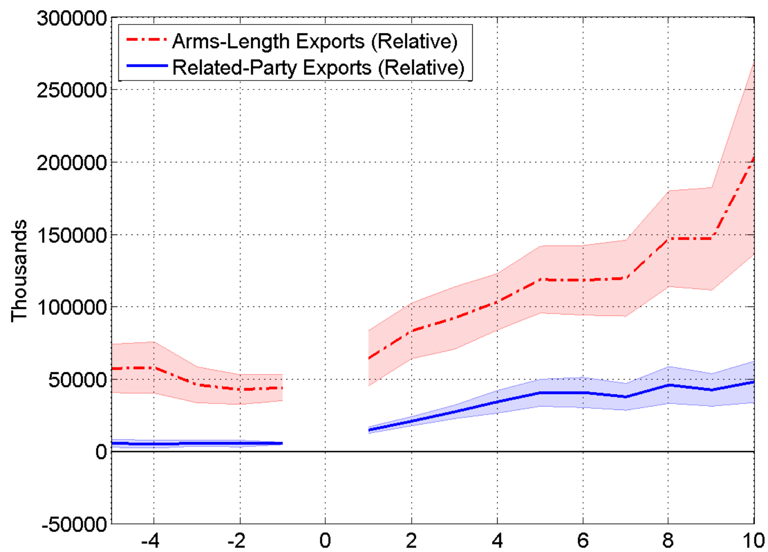

Source: LFTTD-DCA-UBP as explained in text.
This figure reports the related-party and arms-length exports of the parent firm of an establishment that transitions into part of a multinational firm in year ($t = 0$), relative to a control group based on interacted effects of firm age, establishment size, and industry (in year $t = -1$). See equation 1.2, modified to reflect firm-level imports as dependent variables. The shaded area corresponds to a 95 percent confidence interval.

Revenues of a firm of type φ are given by

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

and the sourcing share from location/mode j is

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

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

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

Table A.4: Appendix Table Detailing Aggregate Job Loss from New Multinational Plants

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

Source: Estimates based on Table 1.1, Table 1.2, and Figure 1.2.

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

$$\ln R(\varphi) = \Psi_I - \frac{\sigma - 1}{\theta} \ln \sum_{j \in I} \chi_j(\varphi) + \frac{\sigma - 1}{\theta} \ln \left(\sum_{j \in I} T_j \left[h_j(\varphi) \right]^{\theta} (\tau_j w_j)^{-\theta} \right) + \ln s(\varphi) \quad \text{(A.5)}$$

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

Table A.5: Appendix Table Detailing Aggregate Job Loss from All Multinational Plants

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

Source: Estimates based on Table 1.1, Table 1.2, and Table 1.5.
[1]This column applies the coefficient estimates from the intensive/extensive and weighted estimates from Table 1.5. [2]This column applies the coefficient estimates from the intensive/extensive and unweighted estimates from Table 1.5.

Next, we fix a particular sourcing strategy J and partition it into the strict subsets $I_1, ..., I_S$. We then estimate equation (A.5) for all $I_1, ..., I_S$ and obtain S estimates of $-\frac{\sigma-1}{\theta}$. Now the same logic as described in the text applies. As the sample size tends to infinity, the true value of $-\frac{\sigma-1}{\theta}$ must lie between the smallest and the largest estimates we obtain. Of course, in practice these bounds are estimated with error.

Table A.6: Regression Results: Subsamples

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

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

A.3.2 Invertibility of $\eta_{m,k}$

To show that $\eta_{m,k}$ is invertible for some m and k, we estimate specification 1.19 separately for all $j \in \mathcal{J}$. Note, the technology transfer function is not conditional on a sourcing strategy, but only on a location/mode j. Unlike the bounding procedure, therefore we do not condition on a particular sourcing strategy $J \subset \mathcal{J}$, but pool all observations that source from a given location/mode. The results of the estimation are shown in Table A.7.

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

We next estimate 1.21 for $j = HO$ and plot the semi-parametric component $(\sigma - 1)$.. as a function of $\frac{\chi_{HO}}{\chi_{HI}}$. The result is shown in Figure ??. As expected, the technology transfer function is increasing in the share ratio.

Table A.7: Bias in single share estimates of $\frac{\sigma-1}{\theta}$

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

Source: LBD, LFTTD and CMF
This table reports the results from estimating 1.19 for all firms in 1997. The single shares are instrumented with lagged shares. F statistics for the first stage are significant at conventional levels. The results for years 2002 and 2007 (not shown) are similar.

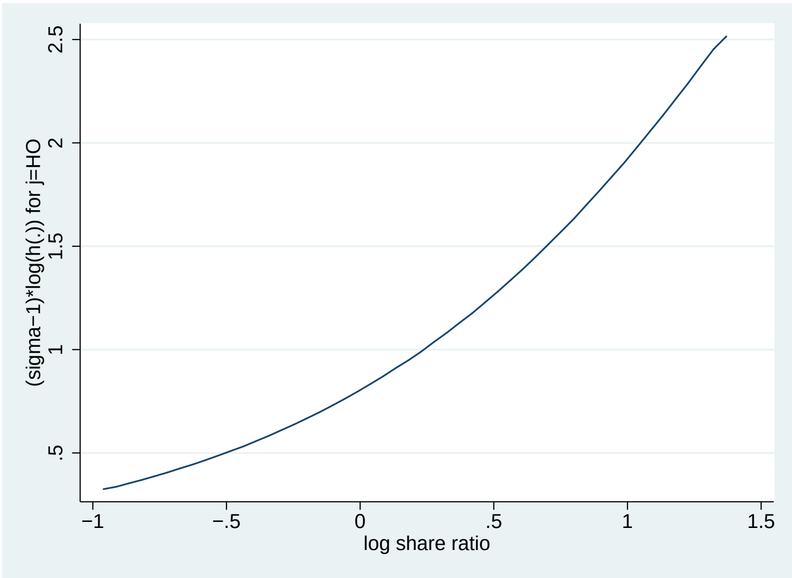

Source: LBD, LFTTD and CMF
This figure displays the results of plotting $h_j(\varphi)$ on $\frac{\chi_{HO}}{\chi_{HI}}$ as discussed in the text. The size distribution of $\frac{\chi_{HO}}{\chi_{HI}}$ is truncated at the 15th and 85th percentiles.

A.3.3 Estimation results: robustness

Table A.8: Estimation Results: Semiparametric Regressions (Robustness)

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

Source: LBD,LFTTD, CMF and ASM
This table reports point estimates for $\frac{\sigma-1}{\theta}$ from the polynomial approximation and size bin approaches discussed in 1.3.4, where the dependent variable is firm revenues. The lower panel displays results where the cost shares are instrumented with lagged values.

Table A.9: Robustness to κ_j

	Baseline	All $\kappa_j = 1, j \neq HO$	$\kappa_j = 1, \in \{NI, SI\}$
Manufacturing Employment	-0.13	-0.14	-0.13
Multinational Employment	-0.28	-0.27	-0.27
Non-MN Employment Employment	-0.07	-0.08	-0.07

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

A.4 Chapter 1: Quantitative Exercises Appendix

A.4.1 Robustness to choices of κ_j

This section presents results of fitting the model in Section 1.4 to the calibration targets in Table 1.13 with alternative choices for $\kappa_j, j \in \{HO, NO, NI, SO, SI\}$. We present the declines in employment between 1997 and 2007 implied by the model when (a) all technology transfer parameters with the exception of κ_{HO} are set to 1, and (b) all within firm technology transfer parameters are set to 1 ($\kappa_j = 1, \in \{NI, SI\}$).

Table A.10: Robustness: Quantitative Exercises

	Data (1997- 2007)	Baseline	Only T_j changes	Only fixed costs change
Manufacturing Employment	-0.25	-0.13	-0.06	-0.03
Multinational Employment	-0.27	-0.28	0.06	0.13
Non-MN Employment Employment	-0.24	-0.07	-0.12	-0.10

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

A.4.2 Counterfactual exercises

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

Table A.10 presents the results of these alternative calibrations. Notice that in both counterfactual exercises, we do not change any of the calibration targets, so we have more targets than parameters to fit the model. Manufacturing employment falls in aggregate in both cases, but by a smaller amount than in the baseline. Further, multinational employment actually increases, with the largest effect in the calibration where only fixed costs fall to match observed importing patterns. The large declines in fixed costs in this case result in entry into multinational activity, which dominates the within-firm effect of declining domestic employment due to import substitution. Similar reasoning applies to the case with only technological improvements, but the effect is smaller.[10]

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

APPENDIX B

Chapter 2 Appendices

B.1 Chapter 2: Basic Theory Appendix

B.1.1 Proof of Result 1

Suppose that the firm solves

$$\max p_x x - p_D F_D - p_M IM$$

subject to

$$x = \left[(1-\mu)^{\frac{1}{\psi}} [F_D]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} [IM]^{\frac{\psi-1}{\psi}}\right]^{\frac{\psi}{\psi-1}}$$

and

$$p_x = \left(\frac{Y}{x}\right)^{\frac{1}{\varepsilon}}$$

The first order conditions are

$$\left(1-\frac{1}{\varepsilon}\right)(Y)^{\frac{1}{\varepsilon}}(x)^{\frac{1}{\psi}-\frac{1}{\varepsilon}}(1-\mu)^{\frac{1}{\psi}}[F_D]^{-\frac{1}{\psi}} = p_D$$

$$\left(1-\frac{1}{\varepsilon}\right)(Y)^{\frac{1}{\varepsilon}}(x)^{\frac{1}{\psi}-\frac{1}{\varepsilon}}\mu^{\frac{1}{\psi}}[IM]^{-\frac{1}{\psi}}=p_M$$

Dividing one by the other gives

$$\frac{F_D^*}{IM^*}=\frac{1-\mu}{\mu}\left(\frac{p_M}{p_D}\right)^{\psi}.$$

The same equation can be obtained under perfect competition.

Now take the production function and multiply it by p_x

$$p_x x=p_x\left[(1-\mu)^{\frac{1}{\psi}}[F_D]^{\frac{\psi-1}{\psi}}+(p_M)^{-\frac{\psi-1}{\psi}}\mu^{\frac{1}{\psi}}[p_M IM]^{\frac{\psi-1}{\psi}}\right]^{\frac{\psi}{\psi-1}}$$

Taking logs gives

$$\ln(p_x x)=\frac{\psi}{\psi-1}\ln\left(p_x\left[(1-\mu)^{\frac{1}{\psi}}[F_D]^{\frac{\psi-1}{\psi}}+(p_M)^{-\frac{\psi-1}{\psi}}\mu^{\frac{1}{\psi}}[p_M IM]^{\frac{\psi-1}{\psi}}\right]\right) \tag{B.1}$$

$$=\frac{\psi}{\psi-1}\ln\left(p_x\left[(1-\mu)^{\frac{1}{\psi}}\exp\left(\frac{\psi-1}{\psi}\ln[F_D]\right)+(p_M)^{-\frac{\psi-1}{\psi}}\mu^{\frac{1}{\psi}}\exp\left(\frac{\psi-1}{\psi}\ln[p_M IM]\right)\right]\right) \tag{B.2}$$

Before differentiating, recall the assumption that the firm takes prices p_M as given and that it cannot change p_x after learning about the shock. Then

$$\frac{\partial\ln p_x x}{\partial\ln p_M M}=\frac{\psi}{\psi-1}\frac{p_x(P_M)^{-\frac{\psi-1}{\psi}}\mu^{\frac{1}{\psi}}\exp\left(\frac{\psi-1}{\psi}\ln[p_M IM]\right)\frac{\psi-1}{\psi}}{p_x\left[(1-\mu)^{\frac{1}{\psi}}\exp\left(\frac{\psi-1}{\psi}\ln[F_D]\right)+(p_M)^{-\frac{\psi-1}{\psi}}\mu^{\frac{1}{\psi}}\exp\left(\frac{\psi-1}{\psi}\ln[p_M IM]\right)\right]} \tag{B.3}$$

$$=\frac{1}{1+\left(\frac{1-\mu}{\mu}\right)^{\frac{1}{\psi}}\left[\frac{F_D}{IM}\right]^{\frac{\psi-1}{\psi}}} \tag{B.4}$$

We evaluate this elasticity at

$$\frac{F_D^*}{IM} = \frac{IM^*}{IM}\frac{1-\mu}{\mu}\left(\frac{p_M}{p_D}\right)^{\psi}$$

so that

$$\frac{\partial \ln p_x x}{\partial \ln p_M IM} = \frac{1}{1+\left(\frac{IM^*}{IM}\right)^{\frac{\psi-1}{\psi}}\frac{1-\mu}{\mu}\left(\frac{p_M}{p_D}\right)^{\psi-1}}$$

B.1.2 On Flexibility in Domestic Inputs

Under the assumption of perfect competition, the first order conditions are:

$$x\left(1-\mu\right) = \left(p_D\right)^{\psi} F_D$$

$$x\mu = \left(p_M\right)^{\psi} IM$$

If the firm takes prices p_x, p_M, and p_D as given, the following elasticities are immediate:

$$\frac{\partial \ln\left(p_x x\right)}{\partial \ln\left(p_D F_D\right)} = \frac{\partial \ln\left(p_x x\right)}{\partial \ln\left(p_M M\right)} = \frac{\partial \ln\left(p_D F_D\right)}{\partial \ln\left(p_M M\right)} = 1.$$

The above equations demonstrate that a constant returns to scale production function combined with these assumptions on market structure imply that the output elasticity will equal one for all values of the elasticity of substitution. For this reason, we require some assumptions limiting the flexibility of domestic inputs following the import disruption.

Below we show an alternative way of understanding the interaction of competitive factor markets, changes in domestic inputs, and the mapping of the output elasticity into parameter values for the elasticity of substitution. Consider the total derivative of $\ln(x)$:

$$d\ln x = \frac{\partial \ln x}{\partial IM} d\ln IM + \frac{\partial \ln x}{\partial F} d\ln F \tag{B.5}$$

$$d\ln x = \frac{\mu^{\frac{1}{\psi}}(IM)^{\frac{\psi-1}{\psi}}d\ln IM}{(1-\mu)^{\frac{1}{\psi}}[F_D]^{\frac{\psi-1}{\psi}}+\mu^{\frac{1}{\psi}}[IM]^{\frac{\psi-1}{\psi}}} + \frac{(1-\mu)^{\frac{1}{\psi}}(F_D)^{\frac{\psi-1}{\psi}}d\ln F_D}{(1-\mu)^{\frac{1}{\psi}}[F_D]^{\frac{\psi-1}{\psi}}+\mu^{\frac{1}{\psi}}[IM]^{\frac{\psi-1}{\psi}}} \tag{B.6}$$

Dividing by $d\ln IM$ yields:

$$\frac{d\ln x}{d\ln IM} = \frac{\mu^{\frac{1}{\psi}}(IM)^{\frac{\psi-1}{\psi}}}{(1-\mu)^{\frac{1}{\psi}}[F_D]^{\frac{\psi-1}{\psi}}+\mu^{\frac{1}{\psi}}[IM]^{\frac{\psi-1}{\psi}}} + \frac{(1-\mu)^{\frac{1}{\psi}}(F_D)^{\frac{\psi-1}{\psi}}}{(1-\mu)^{\frac{1}{\psi}}[F_D]^{\frac{\psi-1}{\psi}}+\mu^{\frac{1}{\psi}}[IM]^{\frac{\psi-1}{\psi}}}\frac{d\ln F_D}{d\ln IM}$$

Now, as before, combining the first order conditions from the profit maximization problem, we have:

$$\frac{F_D(\cdot)}{IM} = \frac{1-\mu}{\mu}\left(\frac{p_D}{p_M}\right)^{-\psi} \tag{B.7}$$

Log-differentiating this expression:

$$\begin{aligned}
d\ln\left(\frac{F_D}{IM}\right) &= -\psi d\ln\left(\frac{p_D}{p_M}\right) \\
d\ln F_D - d\ln IM &= -\psi d\ln\left(\frac{p_D}{p_M}\right) \\
\frac{d\ln F_D}{d\ln IM} &= 1-\psi\frac{d\ln\left(\frac{p_D}{p_M}\right)}{d\ln IM}
\end{aligned} \tag{B.8}$$

Finally, we have:

$$\frac{d\ln x}{d\ln IM} = \frac{\mu^{\frac{1}{\psi}}(IM)^{\frac{\psi-1}{\psi}}}{(1-\mu)^{\frac{1}{\psi}}[F_D]^{\frac{\psi-1}{\psi}}+\mu^{\frac{1}{\psi}}[IM]^{\frac{\psi-1}{\psi}}} + \frac{(1-\mu)^{\frac{1}{\psi}}(F_D)^{\frac{\psi-1}{\psi}}\left[1-\psi\frac{d\ln\left(\frac{p_D}{p_M}\right)}{d\ln IM}\right]}{(1-\mu)^{\frac{1}{\psi}}[F_D]^{\frac{\psi-1}{\psi}}+\mu^{\frac{1}{\psi}}[IM]^{\frac{\psi-1}{\psi}}} \tag{B.9}$$

Thus, if there is no change in the relative input price following the disruption in IM of the firm: $\frac{d\ln\left(\frac{p_M}{p_D}\right)}{d\ln IM} = 0$, then the output elasticity will be equal to one regardless of the value of ψ. On the other hand, any assumptions that yield a non-zero change in the relative input prices will then yield the result that $\frac{d\ln x}{d\ln IM} = 1$ provided $\psi \to 0$.

B.2 Chapter 2: Data Appendix

B.2.1 Matching Corporate Directories to the Business Register

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

B.2.1.1 Directories of International Corporate Structure

The LexisNexis Directory of Corporate Affiliations (DCA) is the primary source of information on the ownership and locations of U.S. and foreign affiliates. The DCA 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.

The second source used to identify multinational firms comes from Uniworld Business Publications (UBP). This company 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.

B.2.1.2 Background on Name and Address Matching

Matching two data records based on name and address information is necessarily an imperfect exercise. Issues such as abbreviations, misspellings, alternate spellings, and alternate name conventions rule out an exact merging procedure, leaving the researcher with probabilistic string matching algorithms that evaluate the "closeness" of match — given by a score or rank — between the two character strings in question. Due to the large computing requirements of these algorithms, it is common to use so-called "blocker" variables to restrict the search samples within each dataset. A "blocker" variable must match exactly, and as a result this implies the need for a high degree of conformity between these variables in the two datasets. In the context of name and address matching, the most common "blocker" variables are the state and city of the establishment.

The matching procedure uses a set of record linking utilities described in *Wasi and Flaaen* (2014). This program uses a bigram string comparator algorithm on multiple variables with differing user-specified weights.[1] This way the researcher can apply, for example, a larger weight on a near *name* match than on a perfect *zip code* match. Hence, the "match score" for this program can be interpreted as a weighted average of each variable's percentage of

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

bigram character matches.

B.2.1.3 The Unit of Matching

The primary unit of observation in the DCA, UBP, and BR datasets is the business establishment. Hence, the primary unit of matching is the establishment, and not the firm. However, there are a number of important challenges with an establishment-to-establishment link. First, the DCA (UBP) and BR may occasionally have differing definitions of the establishment. One dataset may separate out several operating groups within the same firm address (i.e. JP Morgan – Derivatives, and JP Morgan - Emerging Markets), while another may group these activities together by their common address. Second, the name associated with a particular establishment can at times reflect the subsidiary name, location, or activity (i.e. Alabama plant, processing division, etc), and at times reflect the parent company name. Recognizing these challenges, the primary goal of the matching will be to assign each DCA (UBP) establishment to the most appropriate business location of the parent firm identified in the BR. As such, the primary matching variables will be the establishment name, along with geographic indicators of street, city, zip code, and state.

B.2.1.4 The Matching Process: An Overview

The danger associated with probabilistic name and address procedures is the potential for false-positive matches. Thus, there is an inherent tension for the researcher between a broad search criteria that seeks to maximize the number of true matches and a narrow and exacting criteria that eliminates false-positive matches. The matching approach used here is conservative in the sense that the methodology will favor criteria that limit the potential for false positives at the potential expense of slightly higher match rates. As such, the procedure generally requires a match score exceeding 95 percent, except in those cases where ancillary evidence provides increased confidence in the match.[2]

This matching proceeds in an iterative fashion, in which a series of matching procedures are applied with decreasingly restrictive sets of matching requirements. In other words, the initial matching attempt uses the most stringent standards possible, after which the non-matching records proceed to a further matching iteration, often with less stringent standards. In each iteration, the matching records are assigned a flag that indicates the standard associated with the match.

See Table B.1 for a summary of the establishment-level match rate statistics by year and type of firm. Table B.2 lists the corresponding information for the Uniworld data.

B.2.1.5 Construction of Multinational Indicators

The DCA data allows for the construction of variables indicating the multinational status of the U.S.-based establishment. 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. We also retain the nationality of parent firm.[3]

There can be a number of issues when translating the DCA-based indicators through the DCA-BR bridge for use within the Census Bureau data architecture. First, there may be disagreements between the DCA and Census on what constitutes a firm, such that an establishment matches may report differing multinational indicators for the same Census-identified firm. Second, such an issue might also arise due to joint-ventures. Finally, incorrect matches may also affect the degree to which establishment matches agree when aggregated to a firm definition. To address these issues, we apply the following rules when using the

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

[3]The multinational status of firms from the UBP directories are more straightforward.

DCA-based multinational indicators and aggregating to the (Census-based) firm level. There are three potential cases:[4]

Potential 1: A Census-identified firm in which two or more establishments match to different foreign-country parent firms

1. Collapse the Census-identified firm employment based on the establishment-parent firm link by country of foreign ownership
2. Calculate the firm employment share of each establishment match
3. If one particular link of country of foreign ownership yields an employment share above 0.75, apply that link to all establishments within the firm.
4. If one particular link of country of foreign ownership yields an employment share above 0.5 and total firm employment is below 10,000, then apply that link to all establishments within the firm.
5. All other cases require manual review.

Potential 2: A Census-identified firm in which one establishment is matched to a foreign-country parent firm, and another establishment is matched to a U.S. multinational firm.

1. Collapse the Census-identified firm employment based on the establishment-parent firm link by type of DCA link (Foreign vs U.S. Multinational)
2. Calculate the firm employment share of each establishment match
3. If one particular type of link yields an employment share above 0.75, apply that link to all establishments within the firm.

[4] Some of these cases also apply to the UBP-BR bridge.

4. If one particular type of link yields an employment share above 0.5 and total firm employment is below 10,000, then apply that link to all establishments within the firm.

5. All other cases require manual review.

Potential 3: A Census-identified firm in which one establishment is matched to a non-multinational firm, and another establishment is matched to a foreign-country parent firm (or U.S. multinational firm).

Apply same steps as in Potential 2.

B.2.2 Classifying Firm-Level Trade

The firm-level data on imports available in the LFTTD does not contain information on the intended use of the goods.[5] 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 nature of FDI, and the role of imported goods in the transmission of shocks. Fortunately, the Census Bureau data contains other information that can be used to distinguish intermediate input imports from final goods imports. Creating lists of the principal products produced by firms in a given detailed industry in the United States 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.[6] 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.

[5]This is one advantage of the survey data on multinational firms available from the Bureau of Economic Analysis. There are, however, a number of critical disadvantages of this data source, as outlined in **?**.

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

B.2.2.1 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 of product p in industry j can be written as:

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

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

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

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

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

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

Acknowledging this concern, we set an exogenous threshold level W such that any p in a given j with $S_{pj} > W$ is classified as a final good product for that industry. The upper portion of Table B.3 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.[7] 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.

B.2.2.2 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. [8] While

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

[8] It is worth pointing out that the most obvious way that this would materialize is by vertical integration of the firm in its U.S. operations. Provided that the industry designation of the firm pertains to its most downstream operations, then this is would not serve to bias the firms' classification of imported goods, as the upstream products are not actually "final" goods for that firm.

recognizing the caution that should be used in this regard, we adopt the approach that is commonly used in such circumstances: the industry of the firm is defined as that industry encompassing the largest employment share.

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

$$\left.\begin{aligned} M_{ij}^{\text{int}} &= \sum_{p \notin P_j} M_{ipj} \\ \\ M_{ij}^{\text{fin}} &= \sum_{p \in P_j} M_{ipj} \end{aligned}\right\} \qquad \text{where} \qquad P_j = \{p \mid S_{pj} \geq W\}\,. \tag{B.10}$$

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

B.2.3 Sample Selection

B.2.3.1 Constructing the Baseline Dataset

This section will discuss the steps taken to construct the sample used in section 2.3.1.

Beginning with the raw files of the LFTTD export/import data, we drop any transactions with missing firm identifiers, and those pertaining to trade with U.S. territories. Next, we merge the LFTTD files with the HS-NAICS6 product concordance from *Pierce and Schott* (2012); if there is no corresponding NAICS6 code for a particular HS code, then we set NAICS6 equal to XXXXXX. We then aggregate up to the level of Firm-Country-Month-NAICS6, and then create extracts according to three sets of destinations/sources: Japan, Non-Japan, and North America (Canada and Mexico). Then, assigning each firm to an LBD-based industry (see below), we run the NAICS-based trade codes through the intermediate/final goods filter discussed in Appendix B.2.2. The firms' monthly trade can then be split into intermediate and final goods components. We repeat this step for years 2009, 2010, and 2011.

Using the Longitudinal Business Database, we drop inactive, ghost/deleted establishments, and establishments that are not in-scope for the Economic Census. To create the sample of manufacturing firms in the U.S., we first create a firm industry code defined as the industry encompassing the largest share of firm employment. We then drop non-manufacturing firms. Next, we merge the LBD for each year with the DCA-Bridge (see section B.2.1) containing multinational indicators. We then apply the rules specified above for clarifying disagreements with the DCA-based multinational indicators. After creating monthly copies of each firm, we merge by firm-month to the trade data. Missing information of trade data is altered to represent zeros. We repeat these steps for years 2009-2011, and then append the files together. Firms that do not exist in all three years are dropped from the sample.

B.2.3.2 GIS Mapping of Earthquake Intensity Measures to Affiliate Locations

As part of the Earthquake Hazards Program, the U.S. Geological Survey produces data and map products of the ground motion and shaking intensity following major earthquakes. The preferred measure to reflect the perceived shaking and damage distribution is the estimated "Modified Mercalli Intensity (MMI)" which is based on a relation of survey response and measured peak acceleration and velocity amplitudes. The USGS extends the raw data from geologic measurement stations and predicts values on a much finer grid using standard seismological inferences and interpolation methods. The result is a dense grid of MMI values covering the broad region affected by the seismic event. For more information on this methodology, see *Wald et al.* (2006).

To utilize this information, we take all Japanese addresses from the DCA/Uniworld directories that correspond to any U.S. operation via an ownership link. We geocode these addresses into latitude/longitude coordinates using the Google Geocoding API, and then compute the inverse distance-weighted mean of the relevant seismic intensity measure based on a 10km radius surrounding a given establishment. The firm identifiers within the corporate directories allow us to create firm-specific measures (average and maximum values, by manufacturing/non-manufacturing), which can then be brought into the baseline Census dataset via the bridges discussed in appendix B.2.1.

Table B.1: DCA Match Statistics: 2007-2011

	# of DCA Establishments	Matched to B.R.	Percent Matched
Total			
2007	112,346	81,656	0.73
2008	111,935	81,535	0.73
2009	111,953	81,112	0.72
2010	111,998	79,661	0.71
2011	113,334	79,516	0.70
U.S. Multinationals			
2007	22,500	16,396	0.73
2008	23,090	16,910	0.73
2009	22,076	16,085	0.73
2010	21,667	15,785	0.73
2011	21,721	15,557	0.72
Foreign Multinationals			
2007	10,331	7,555	0.73
2008	9,351	6,880	0.74
2009	11,142	8,193	0.74
2010	11,308	8,181	0.72
2011	11,619	8,357	0.72

Table B.2: Uniworld Match Statistics: 2006-2011

	# of Uniworld Establishments	Matched to B.R.	Percent Matched
Foreign Multinationals			
2006	3,495	2,590	0.74
2008	3,683	2,818	0.76
2011	6,188	4,017	0.65
U.S. Multinationals[1]			
2007	4,043	3,236	0.80
2009	4,293	3,422	0.80

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

Table B.3: Appendix Table Comparing the Results from Threshold Values W

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

B.3 Appendix: Other Results

B.3.1 Alternate Specifications for Treatment Effects Regressions

Our results from section 2.3.2 are based on a sample including all Japanese multinationals in manufacturing, and therefore uses a levels specification to allow for zeros in the firm-month observations. Because larger firms exhibit greater absolute deviations from trend, this roughly amounts to weighting firms based on size, such that the results correspond to a representative firm based on the aggregate effect of the group.

To see this, and to explore how the levels specification influences our interpretation, we repeat the analysis on a subset of the firms for which we can view the percentage changes directly. Specifically, we drop any firms with zeros in any month for intermediate imports or N.A. exports during the sample, and then take logs and HP-filter each series to obtain percentage deviations from trend for each firm.[9] The results of this exercise are shown in Panel A of Figure B.1. We suppress standard errors for the sake of clarity; the drops are significant at the 95% level for between 2-4 months following the shock. If we rerun these regressions while also weighting according to the pre-shock size of firms, we obtain a picture that looks much closer to Figure 2.7, see Panel B of Figure B.1.

These results indicate that the larger firms appear to be affected the most from this shock. This could be partly a result of our proxy being less effective for smaller firms that may not engage in consistent exports to North America.

B.3.2 Probit Model of Import/Output Disruptions

We specify a simple probit model to understand the relative importance of various firm-level characteristics in the import and output declines following the tsunami. The model is

[9] We re-weight the control group as described in section 2.3.1.

$$Pr(X_{ik}^D = 1) = \Phi\left[\beta_1 \text{JPN}_{ik} + \beta_2 Exposed_{ik} + \beta_3 MMI_{ik} + \beta_4 Port_{ik} + \gamma_k\right] \qquad \text{(B.11)}$$

where the dependent variable (X_{ik}^D) is an indicator equal to one if the N.A. exports of firm i in industry k are on average 20% below trend during the five months following the Tōhoku event. The independent variables are also indicators: JPN_{ik}, for affiliates of Japanese multinationals; $Exposed_{ik}$, for firms with an exposure to Japanese inputs above 0.05 of total material; MMI_{ik} for firms with an elevated MMI value pertaining to their average Japanese manufacturing locations; and $Port_{ik}$ for firms that typically rely on imports via ports damaged by the tsunami.[10] The γ_k term allows for industry-specific intercepts. To evaluate the determinants of an input disruption from Japan, we replace the dependent variable with J_{ik}^D, an indicator for a drop in Japanese imported inputs of 20% relative to trend.

Panel A of Table B.4 evaluates firm characteristics predicting a drop in U.S. output (X_{ik}^D), as measured by our proxy. The columns (1)-(4) show the results from different specifications with various combinations of the covariates in equation (B.11). Both Japanese ownership and high exposure to Japanese inputs significantly increase the probability of an output disruption, as expected. In columns (3) and (4), we demonstrate that Japanese ownership is substantially more indicative of an output decline than high input exposure alone. In Panel B, we replace the dependent variable with the binary measure of a drop in Japanese intermediate inputs (J_i^D). The results from these regressions indicate, unsurprisingly, that high exposure to Japanese imports are highly predictive of a subsequent disruption following the Tōhoku event. Apart from their exposure to imports from Japan, the Japanese affiliates

[10]Specifically, the $MMI_{ik} = 1$ if the average Japanese manufacturing establishment corresponding to a U.S. firm is above the median (roughly an MMI of 5.2) of all firms with Japanese manufacturing locations. The affected ports are: Onahama, Hitachi, Kashima, Haramachi, Shiogama, Sendai, Shimizu, Ishinomaki, Hashinohe, Miya Ko, Kamaishi, Ofunato, and Kessennuma.

are no more likely to suffer a disruption to these imports (see column 8).[11] While the results from Table B.4 are somewhat inconclusive, they nevertheless point to unique features of the production function of Japanese affiliates that yields direct pass-through of Japanese shocks to the U.S. economy. Our estimation procedure that follows should help to clarify this point further.

B.3.3 Bootstrapping Standard Errors

We use bootstrapping methods to compute measures of the dispersion of our point estimates. Using random sampling with replacement within each group of firms, we create 5000 new artificial samples and re-run the estimation procedure. The standard deviation of the point estimates across these bootstrap samples is shown in Table 2.3. To gain a more complete picture of the dispersion, we create density estimates for each sample of firms across the parameter space for the elasticities. These densities are shown in Figure B.3.

B.3.4 Effects on U.S. Exports to Japan

Another dimension of the transmission of the Tōhoku shock to the United States is U.S. exports back to Japan. To the extent that firms in the U.S. receive inputs from Japan for processing and re-shipment back to Japan, one might expect the U.S. exports to Japan may fall following the Tōhoku event. On the other hand, U.S. firms may have increased shipments to Japan following the shock in order to offsett what were large production and supply shortages within Japan. To evaluate this, we re-run the specification in equation (2.5) but replace $V_{i,t}^{M}$, the value of intermediate imports of firm i in month t, with $V_{i,t}^{JEXP}$, the value of Japanese exports of firm i in month t. The results are shown in Figure B.2. As is clear from the figure, we do not see strong evidence to support either hypothesis regarding this particular trade flow, at least as it pertains to Japanese multinationals in particular.

[11]The combined effect of the coefficients on Japan and JPN*Exp is -0.16, and not significant.

B.3.5 Effects on Employment and Payroll

The Standard Statistical Establishment List (SSEL) contains quarterly employment and payroll information for all employers (with some small exceptions) in the U.S. economy. This list is held separately as a single-unit(SSEL-SU) and multi-unit (SSEL-MU) file. The Report of Organization Survey (ROS) asks firms to list the establishments which report under a particular EIN, and this information is then recorded to the firm identifier on the Multi-Unit File. To build a quarterly employment series at the firm-level, we link the EIN variables on the SU file with the firm-identifier linked with each EIN on the MU file. In principle, the four quarters of payroll listed on the SSEL is combined by Census to create an annual payroll figure for each establishment, which is the value recorded in the LBD. Similarly, the employment variable corresponding to the 1st quarter (week of March 12) from the SSEL is that used by the LBD.

Once we merge the SSEL-based data with quarterly employment and payroll to the LBD for a particular year, we conduct a series of reviews to ensure that the annual payroll (and 1st quarter employment) roughly align. Any establishments with disagreements between the SSEL-based payroll and LBD-based payroll such that the ratio was greater than 2 or less than 0.6 were dropped.

After these modifications were made, the remainder of the data construction was similar to that in section B.2.3. We merge multinational indicators from the DCA, drop non-manufacturing firms, append the 2009, 2010, and 2011 files together, and keep only those firms that exist in each year. Using the same set of firms as a control group as specified in section 2.3.1, we run the following regression:

$$\Delta \text{emp}_{j,t} = \sum_{i=-3}^{3} \gamma_i E_i + \sum_{i=-3}^{3} \beta_i E_i D_{j,i} + u_{j,t} \tag{B.12}$$

where $\Delta\text{emp}_{j,t} \equiv ln(\text{emp}_{j,t}/\text{emp}_{j,t-4})$, where $\text{emp}_{j,t}$ indicates employment at firm j in quarter t. We also re-run the equation specified in equation B.12 using payroll $pay_{j,t}$ as the dependent variable (where $\Delta \text{pay}_{j,t} \equiv ln(\text{pay}_{j,t}/\text{pay}_{j,t-4})$. The qualitative results are shown in table B.5.

B.3.6 Effects on Unit Values (Prices) of Trade

The LFTTD contains information on quantities as well as values for each trade transaction, recorded at a highly disaggregated product definition (HS 10 digit). This allows for the construction of unit values (prices) for each firm-product-month observation, which allows for an analysis of price movements surrounding the Tōhoku event.

The majority of the data construction is identical to that in section B.2.3, however there are a number of modifications. First, we drop all transactions with missing or imputed quantities in the LFTTD, and then aggregate to the Firm-HS10-month frequency, separately for each type of trade transaction: 1) Related-Party imports from Japan; 2) Non Related-Party imports from Japan; 3) Related-Party exports to Canada/Mexico; and 4) Non Related-Party exports to Canada/Mexico. Next, we select only those firms identified as manufacturing in the LBD. We keep the related-party and arms-length transactions separate as one may expect these prices to behave differently following a shock. As above, we keep only manufacturing firms, append the annual files together, and then select only those firms identified as a multinational in either 2009, 2010, or 2011.

At the product level, there is little reason to suspect trends or seasonal variation over this short of a time period. Moreover, there is no concern here about accounting for zeros in the data. As such we take a firm j's imports (exports) of product p in month t, and run the following specification in logs ($\text{m}_{p,j,t} = log(\text{M}_{p,j,t})$:

$$\text{m}_{p,j,t} = \alpha_{pj} + \sum_{i=-19}^{9} \gamma_i E_i + \sum_{i=-19}^{9} \beta_i E_i D_{j,i} + u_{j,t} \tag{B.13}$$

where α_{pj} are firmXproduct fixed-effects, γ_i are monthly fixed effects (with the dummy variable $E_i' s$ corresponding to each calendar month), and $u_{j,t}$ are random effects. The variables $D_{j,t}$ are dummy variables equal to one if the firm is owned by a Japanese parent company.

A qualitative version of the results is shown in Table B.6.

B.3.7 Ward's Automotive Data

Ward's electronic databank offers a variety of data products for the global automotive industry at a monthly frequency. We obtain Japanese production (by model), North American production (by plant and model), U.S. inventory (by model), and North American sales (by model) all for the period January 2000 to December 2012. The inventory and sales data also contain the country of origin, so one can separate out these variables based on whether a particular model was imported vs domestically-produced. The series cover the universe of the assembly operations of finished cars and light trucks. Unfortunately, there is no information on input shipments.

For the plant-level analysis of production, the base sample consists of 167 plants active at some point during 2000-2012. We remove plants that were not continuously in operation during the period 2009-2012, and combine several plants that are recorded separately in the data, but are in effect the same plant. After these modifications, the sample reduces to 62 plants, 22 of which are owned by a Japanese parent. The average monthly production in the three months preceding the shock is 12,904 for Japanese plants, and 14,903 for Non-Japanese plants. The specification is identical to that in section 2.3.1:

$$Q_{i,t} = \alpha_0 + \alpha_i + \sum_{p=-14}^{9} \gamma_p E_p + \sum_{p=-14}^{9} \beta_p E_p \mathrm{JPN}_{i,p} + u_{i,t} \tag{B.14}$$

where here the variable $Q_{i,t}$ is auto production by plant i in month t, after removing a plant-specific trend though March 2011. Because these plants can be tracked with some confidence back in time, it is reasonable here to remove seasonality directly, rather than assume a shared seasonal component between the treated and control groups as in section 2.3.2. We use the X12-ARIMA model, provided by the National Bank of Belgium, and apply it to each series before correcting for trend. The results for the Japanese plants are mostly similar, as shown in table B.7.

Figure B.1: Relative Inputs and Output (Proxy) of Japanese Firms (Reduced Sample) Logged, HP-Filtered

A. No Size-Weighting

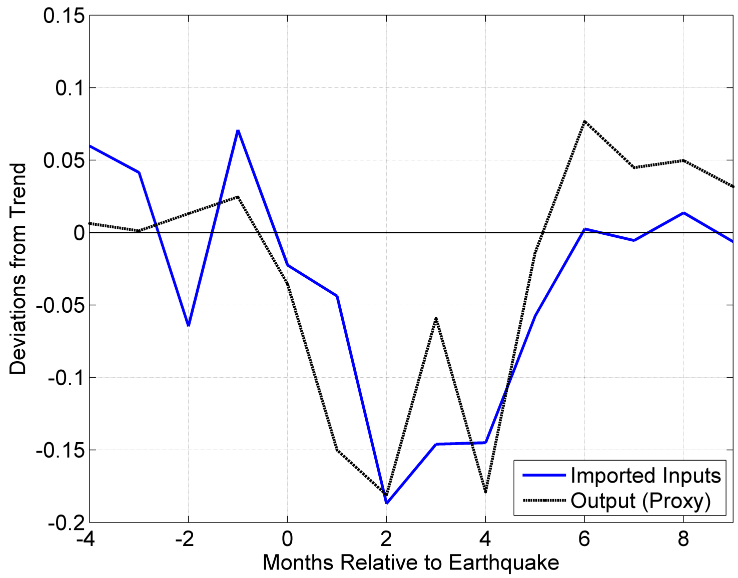

B. Size-Weighted

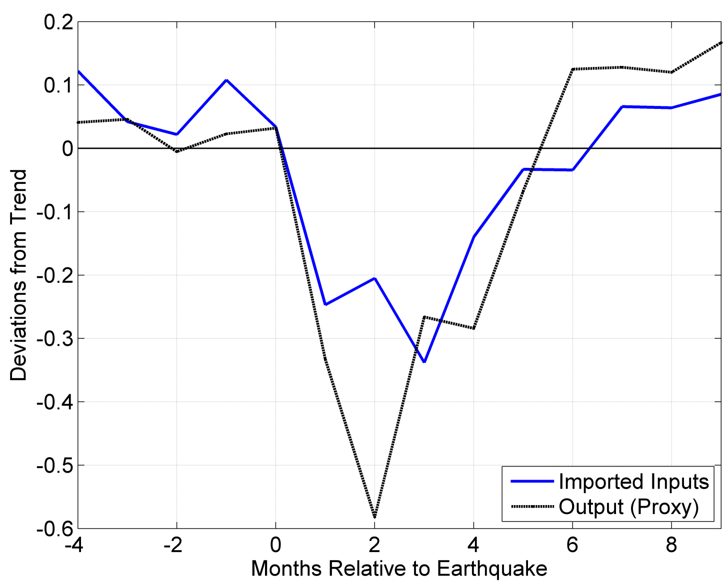

Source: LFTTD-DCA-UBP as explained in text.
These figures report the relative percentage deviations from trend of Japanese affiliates relative to a control group of other multinational firms. The values are coefficient estimates taken from an interaction of a Japanese-firm dummy with a monthly dummy – additional baseline monthly dummies remove seasonal effects. These results reflect a reduced sample with no firm-month zeros in imported inputs or N.A. exports. The data is logged, and HP-filtered using a monthly smoothing parameter.

Figure B.2: Dynamic Treatment Effects: Relative Japanese Exports of Japanese Firms

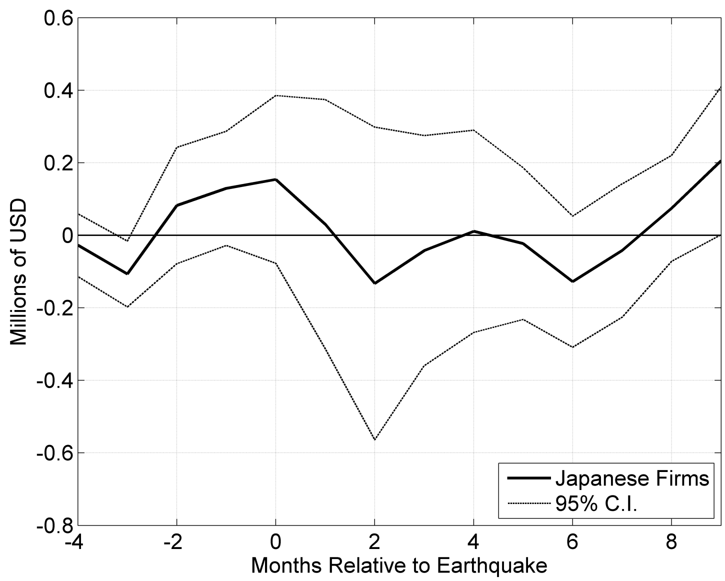

Source: LFTTD-DCA-UBP as explained in text.
These figures report the Japanese exports of the U.S. affiliates of Japanese firms relative to a control group of other multinational firms. The values are coefficient estimates taken from an interaction of a Japanese-firm dummy with a monthly dummy – additional baseline monthly dummies remove seasonal effects. Standard errors are clustered at the firm level.

Table B.4: Predicting Japanese Import and U.S. Output Disruption by Firm Characteristics

	Panel A: Disruption to U.S. Output (proxy) $X_i^D = 1$				Panel B: Disruption to Japanese Imports $J_i^D = 1$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Japan	0.443***		0.352***	0.347**	0.707***		0.310***	0.686***
	(0.0921)		(0.117)	(0.152)	(0.0917)		(0.115)	(0.150)
Exposed		0.351***	0.145	0.140		0.814***	0.636***	0.991***
		(0.0886)	(0.112)	(0.149)		(0.0880)	(0.110)	(0.144)
JPN*Exp				-0.00771				-0.848***
				(0.228)				(0.222)
MMI	-0.176***	-0.121*	-0.178***	-0.178***	0.346***	0.389***	0.341***	0.306***
	(0.0676)	(0.0646)	(0.0676)	(0.0683)	(0.0691)	(0.0667)	(0.0694)	(0.0704)
Ports	-0.174	-0.144	-0.197		0.248	0.217	0.168	0.174
	(0.224)	(0.225)	(0.226)		(0.211)	(0.212)	(0.213)	(0.213)
Constant	-0.674	-0.674	-0.674	-0.674	-4.672	-4.672	-4.672	-4.668
	(0.681)	(0.681)	(0.681)	(0.681)	(85.78)	(85.78)	(85.78)	(85.00)
Industry Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2451	2451	2451	2451	2451	2451	2451	2451

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Source: LFTTD, DCA, UBP, and USGS as explained in the text. This table reports the results of a probit model prediction of JPN import and N.A. exports (output) disruption based on firm characteristics. See section 2.3.1 for a definition of the variables.

Figure B.3: Density Estimates of Elasticities Across Bootstrap Samples

A. Japanese vs non-Japanese Multinationals: Materials Elasticity (ω)

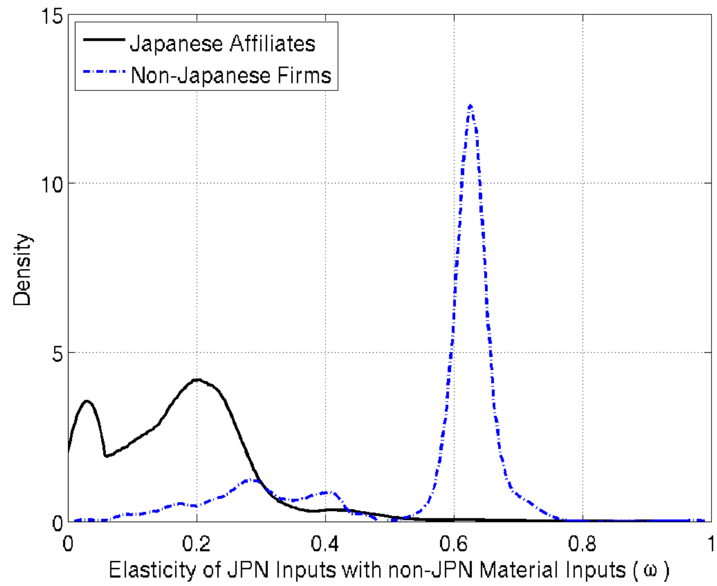

B. Japanese vs non-Japanese Multinationals: Materials-Capital/Labor Elasticity (ζ)

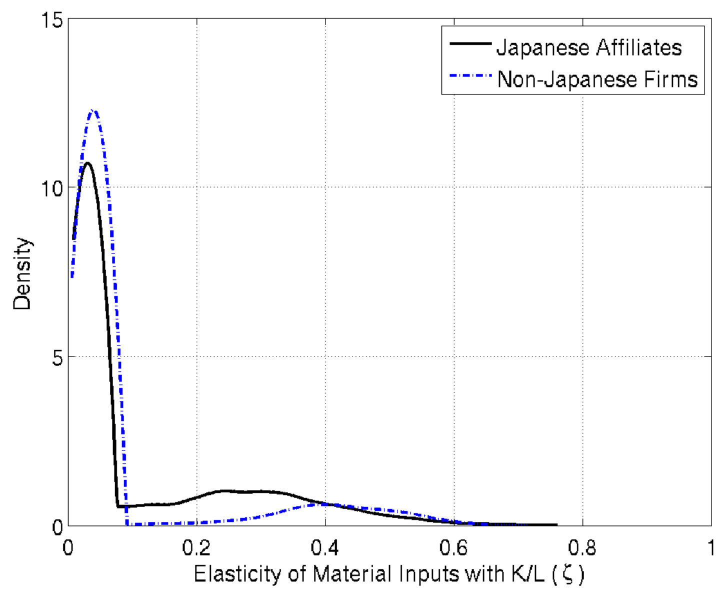

Figure B.3: Density Estimates of Elasticities Across Bootstrap Samples

C. Non-multinationals and All Firms: Materials Elasticity (ω)

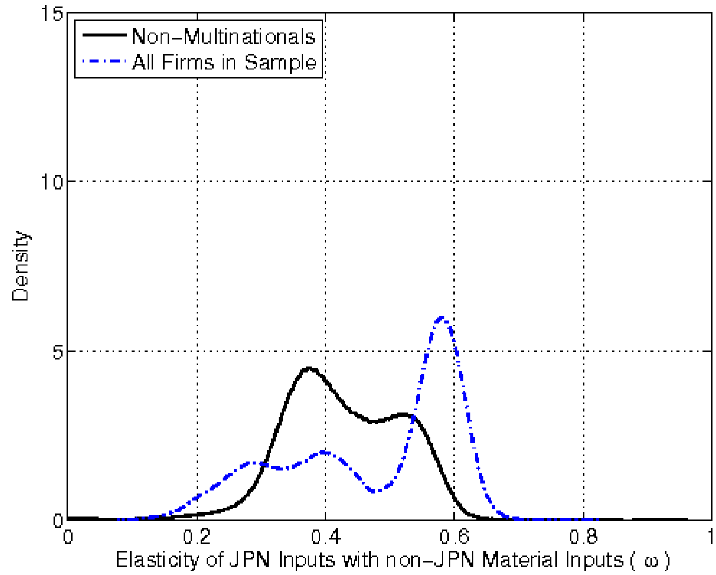

D. Non-multinationals and All Firms: Materials-Capital/Labor Elasticity (ζ)

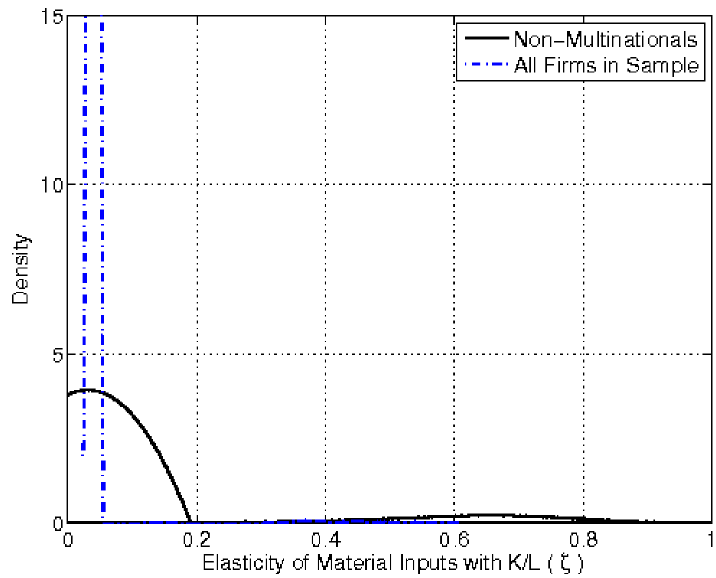

Source: LFTTD-DCA-UBP as explained in text.

Table B.5: Dynamic Treatment Effects: Quarterly Employment/Payroll Surrounding Tōhoku Event

	Log 4-Quarter Difference	
	Employment	Payroll
Independent Variables	(1)	(2)
Q2_2010 (t=-3)	pos***	pos***
Q3_2010 (t=-2)	pos***	pos***
Q4_2010 (t=-1)	pos***	pos***
Q1_2011 (t=0)	pos***	pos***
Q2_2011 (t=1)	pos***	pos***
Q3_2011 (t=2)	pos***	pos***
Q4_2011 (t=3)	pos***	pos***
JPNxQ2_2010 (t=-3)	neg	neg
JPNxQ3_2010 (t=-2)	neg	neg
JPNxQ4_2010 (t=-1)	neg	neg
JPNxQ1_2011 (t=0)	neg	neg
JPNxQ2_2011 (t=1)	neg	neg
JPNxQ3_2011 (t=2)	neg	neg
JPNxQ4_2011 (t=3)	neg	pos
constant	neg***	neg***
Firm Fixed Effects	Yes	Yes
Observations		
R-squared		

Source: SSEL and DCA as explained in the text.
Robust standard errors (clustered at the firmXProduct level) pertaining to each sign coefficient are indicated by: *** p<0.01, ** p<0.05, * p<0.1.
This table reports qualitative features of firm employment and firm payroll in the quarters surrounding the Tōhoku earthquake and tsunami. The first set of coefficients correspond to quarter dummies, whereas the second set (JPNx) correspond to the interaction of a Japanese firm dummy with quarter dummies. See equation B.12 in the text. The dependent variable is the four-quarter log difference of employment (payroll).

Table B.6: Dynamic Treatment Effects: Unit Values of Trade Surrounding Tōhoku Event

	Log Unit-Value of:			
	JPN Imports: Related Party	JPN Imports: Non-Related Party	N.A. Exports Related Party	N.A. Exports Non-Related Party
Independent Variables	(1)	(2)	(3)	(4)
Sep 2010 (t=-6)	neg**	pos	pos*	pos
Oct 2010 (t=-5)	pos	neg	pos**	pos
Nov 2010 (t=-4)	pos	pos	pos**	pos
Dec 2010 (t=-3)	pos	neg	pos	pos
Jan 2011 (t=-2)	neg	pos	neg	pos
Feb 2011 (t=-1)	pos	neg	pos**	pos
Mar 2011 (t=0)	neg	pos	pos	pos
Apr 2011 (t=1)	pos	pos	pos	pos
May 2011 (t=2)	neg	pos	neg	pos**
Jun 2011 (t=3)	pos**	neg	pos**	neg
Jul 2011 (t=4)	neg	neg	pos	neg
Aug 2011 (t=5)	pos	pos	neg	pos
Sep 2011 (t=6)	pos	pos	pos	pos**
Oct 2011 (t=7)	neg	neg	pos	pos
Nov 2011 (t=8)	pos	neg	pos	neg
Dec 2011 (t=9)	neg	pos	pos**	pos
JPNxSep 2010 (t=-6)	pos**	neg*	neg**	neg
JPNxOct 2010 (t=-5)	neg*	pos	pos	pos
JPNxNov 2010 (t=-4)	neg	pos	neg	neg
JPNxDec 2010 (t=-3)	neg	neg*	pos	pos
JPNxJan 2011 (t=-2)	pos	neg	neg	neg
JPNxFeb 2011 (t=-1)	neg	pos	pos	pos**
JPNxMar 2011 (t=0)	pos	pos	neg	neg
JPNxApr 2011 (t=1)	neg	pos	neg	neg
JPNxMay 2011 (t=2)	pos	neg	pos	neg
JPNxJun 2011 (t=3)	neg	pos*	neg	neg
JPNxJul 2011 (t=4)	pos	neg	pos	neg
JPNxAug 2011 (t=5)	neg*	neg*	neg	pos
JPNxSep 2011 (t=6)	neg	neg	neg	neg
JPNxOct 2011 (t=7)	pos	neg	neg	neg
JPNxNov 2011 (t=8)	neg	neg	neg	pos
JPNxDec 2011 (t=9)	neg	neg	pos	neg
constant	pos	neg	neg	neg
FirmXProduct Fixed Effect	Yes	Yes	Yes	Yes
Observations				
R-Squared				

Source: LFTTD, DCA, and UBP as explained in the text.
Robust standard errors (clustered at the firmXProduct level) pertaining to each sign coefficient are indicated by: *** p<0.01, ** p<0.05, * p<0.1 .
This table reports qualitative features of the unit values of trade surrounding the 2011 Tōhoku earthquake and tsunami. The first set of coefficients correspond to monthly dummies, whereas the second set (JPNx) correspond to the interaction of a Japanese firm dummy with monthly dummies. See equation B.13 in the text.

Figure B.4: Automotive Production, Inventory, Sales by Firm Type, Distributed Lag Model

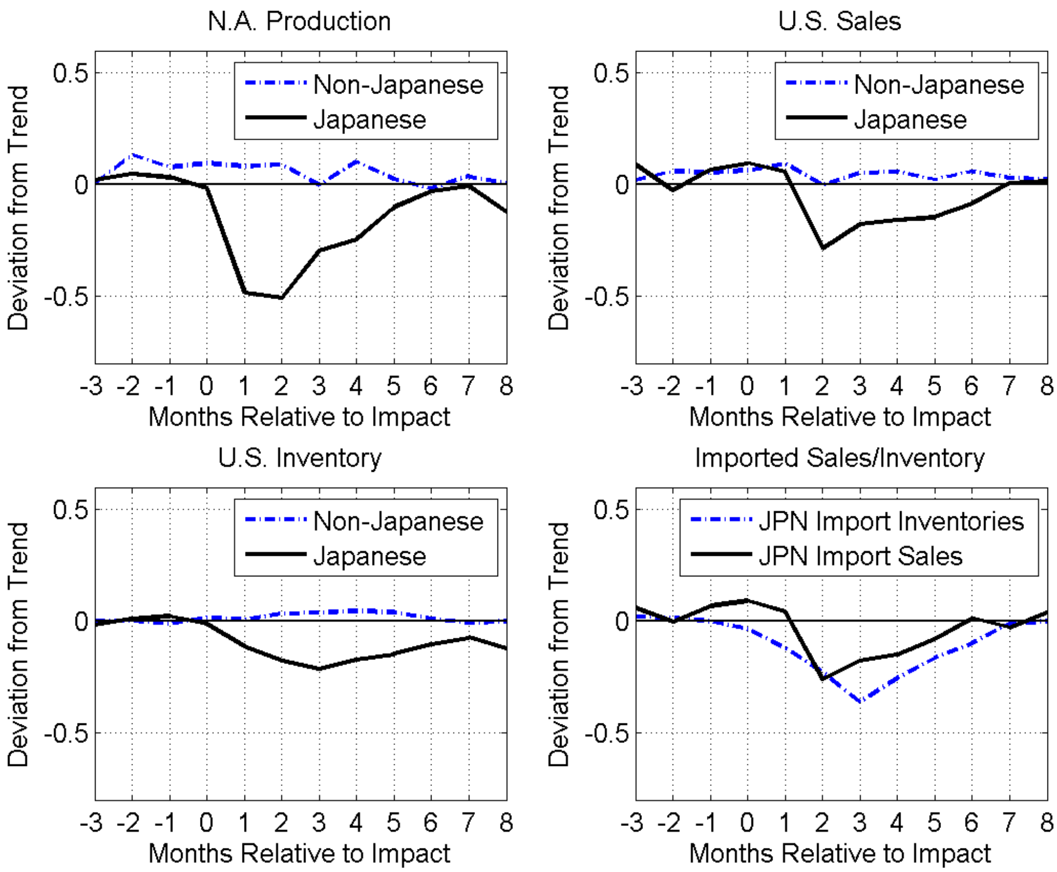

Source: Ward's Automotive Database
This figure reports North American production, and U.S. sales and inventory data according to firm type: Japanese and non-Japanese firms. The values are coefficient estimates taken from a distributed lag model, exploiting time-series variation only. The underlying series have been seasonally adjusted, logged, and HP-Filtered Standard errors are suppressed in the interests of clarity. The Japanese automakers are Honda, Mazda, Mitsubishi, Nissan, Toyota, and Subaru.

Table B.7: Dynamic Treatment Effects: N.A. Automotive Production

VARIABLES	(1) Prod	(2) Prod	VARIABLES (cont'd)	(1) Prod (cont'd)	(2) Prod (cont'd)
Nov 2010 (t=-4)	91.06	17.78	JPN x Nov 2010 (t=-4)	-195.8	-341.7
	(649.9)	(608.8)		(841.9)	(799.2)
Dec 2010 (t=-3)	-1,973***	310.3	JPN x Dec 2010 (t=-3)	-385.0	-408.3
	(467.5)	(497.5)		(736.5)	(706.4)
Jan 2011 (t=-2)	-611.5	1,083*	JPN x Jan 2011 (t=-2)	781.0	-1,092
	(637.3)	(618.7)		(792.1)	(804.6)
Feb 2011 (t=-1)	694.9*	756.3*	JPN x Feb 2011 (t=-1)	-1,142	-1,210*
	(401.9)	(394.7)		(696.2)	(666.8)
Mar 2011 (t=0)	4,356***	1,483***	JPN x Mar 2011 (t=0)	-3,515***	-2,592***
	(524.9)	(389.1)		(812.0)	(842.7)
Apr 2011 (t=1)	-216.2	305.5	JPN x Apr 2011 (t=1)	-6,239***	-6,099***
	(707.7)	(620.4)		(1,303)	(1,282)
May 2011 (t=2)	1,584***	799.1	JPN x May 2011 (t=2)	-7,244***	-6,625***
	(525.4)	(511.3)		(1,651)	(1,740)
Jun 2011 (t=3)	1,366**	-499.3	JPN x Jun 2011 (t=3)	-4,564***	-3,423**
	(623.6)	(594.9)		(1,248)	(1,320)
Jul 2011 (t=4)	-4,512***	123.3	JPN x Jul 2011 (t=4)	-2,143	-3,723***
	(878.4)	(606.2)		(1,430)	(1,045)
Aug 2011 (t=5)	685.6	-1,323**	JPN x Aug 2011 (t=5)	-1,275	-1,108
	(744.0)	(648.1)		(970.8)	(1,012)
Sep 2011 (t=6)	-836.5	-1,895***	JPN x Sep 2011 (t=6)	-359.4	40.37
	(663.7)	(641.5)		(930.7)	(959.8)
Oct 2011 (t=7)	-338.0	-1,434**	JPN x Oct 2011 (t=7)	93.27	-265.4
	(662.3)	(632.4)		(885.6)	(785.8)
Nov 2011 (t=8)	-1,393**	-1,443**	JPN x Nov 2011 (t=8)	-1,318	-2,059*
	(582.8)	(601.2)		(1,159)	(1,183)
Dec 2011 (t=9)	-4,511***	-1,619**	JPN x Dec 2011 (t=9)	759.1	24.95
	(774.4)	(655.5)		(1,105)	(803.9)
Constant				-1,535***	-1,683***
				(89.30)	(91.95)
Plant Fixed Effects				Yes	Yes
Remove Plant-Specific Pre-Shock Trend				Yes	Yes
Remove Seasonal Component				No	Yes
Observations				2,976	2,976
R-squared				0.260	0.272

Source: Ward's Automotive Yearbook
Robust standard errors (clustered at the plant level) in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

APPENDIX C

Chapter 3 Appendices

C.1 Chapter 3: Data Appendix

We describe the algorithm used to construct a utilization-adjusted TFP series for Canada. Our procedure is adapted from *Imbs* (1999), as the quarterly data necessary to construct a series with the *Fernald* (2014) methodology are not currently available for Canada for the requisite time period. The method in *Imbs* (1999) is in the spirit of *Basu et al.* (2006), in that it also relies on identifying movements in unobserved (aggregate) utilization from observed changes in inputs and output. Unlike *Basu et al.* (2006), this method does not control for sectoral differences or non-constant returns to scale. We briefly describe the steps of the algorithm here, using commonly seen relationships from a firm's profit maximization problem. For a detailed derivation of the equations that follow see *Imbs* (1999).

1. Construct a starting capital stock series using the perpetual inventory method from official investment series I_t and a quarterly depreciation rate of 0.025. For the initial value of the capital stock we chose $K_0 = \frac{I_1}{r+g_I}$, where g_I is the growth rate of investment

in Canada. We tested our results with other choices for the initial capital stock and found no substantive difference.

2. Construct an initial series for utilization u_t using the capital stock series K_t, output Y_t, and values for average depreciation $\bar{\delta}$ and the interest rate r from the equilibrium relationship $u_t = \left(\frac{Y_t/K_t}{Y/K}\right)^{\frac{\bar{\delta}}{\bar{\delta}+r}}$, where Y/K is the average period value.

3. Use the initial utilization series and the assumed relationship between depreciation and utilization $\delta_t = \bar{\delta} u_t^{1+r/\bar{\delta}}$ to construct a time-varying series for δ_t.

4. Together with the official series for investment and the time-varying δ_t, construct a new capital stock using the standard capital accumulation equation.

5. Using the new $\bar{\delta}$ and capital stock, return to step (1) and construct a new utilization series.

6. Iterate until the capital stock and $\bar{\delta}$ converge. Then construct the final implied u_t.

7. Construct a series for the household's labor effort e_t from $e_t = \left((1-\alpha)\frac{Y_t}{C_t}\right)^{f(w_t, N_t, Y_t)}$ using data on consumption C_t, wages w_t, and labor input N_t.[1]

8. Construct the utilization-adjusted TFP series from the production function $Y_t = X_t (u_t K_t)^{\alpha} (e_t N_t)^{(1-\alpha)}$.

The only additional data series required for this procedure are data on investment and wages. For consistency with the rest of our data, both series were taken from the *Ohanian and Raffo* (2012) dataset, which in turn uses data from the OECD Main Economic Indicators along with national databases.

[1] The derivation of this expression uses the household's optimization problem and can be found in *Imbs* (1999).

C.2 Chapter 3: Robustness Appendix

This section presents responses of core U.S. and Canadian variables to the three identified shocks, where the sentiment shock is identified using a measure of U.S. consumer confidence.

Figure C.1: The Impulse Responses to the US Surprise TFP Shock, Using US Consumer Confidence

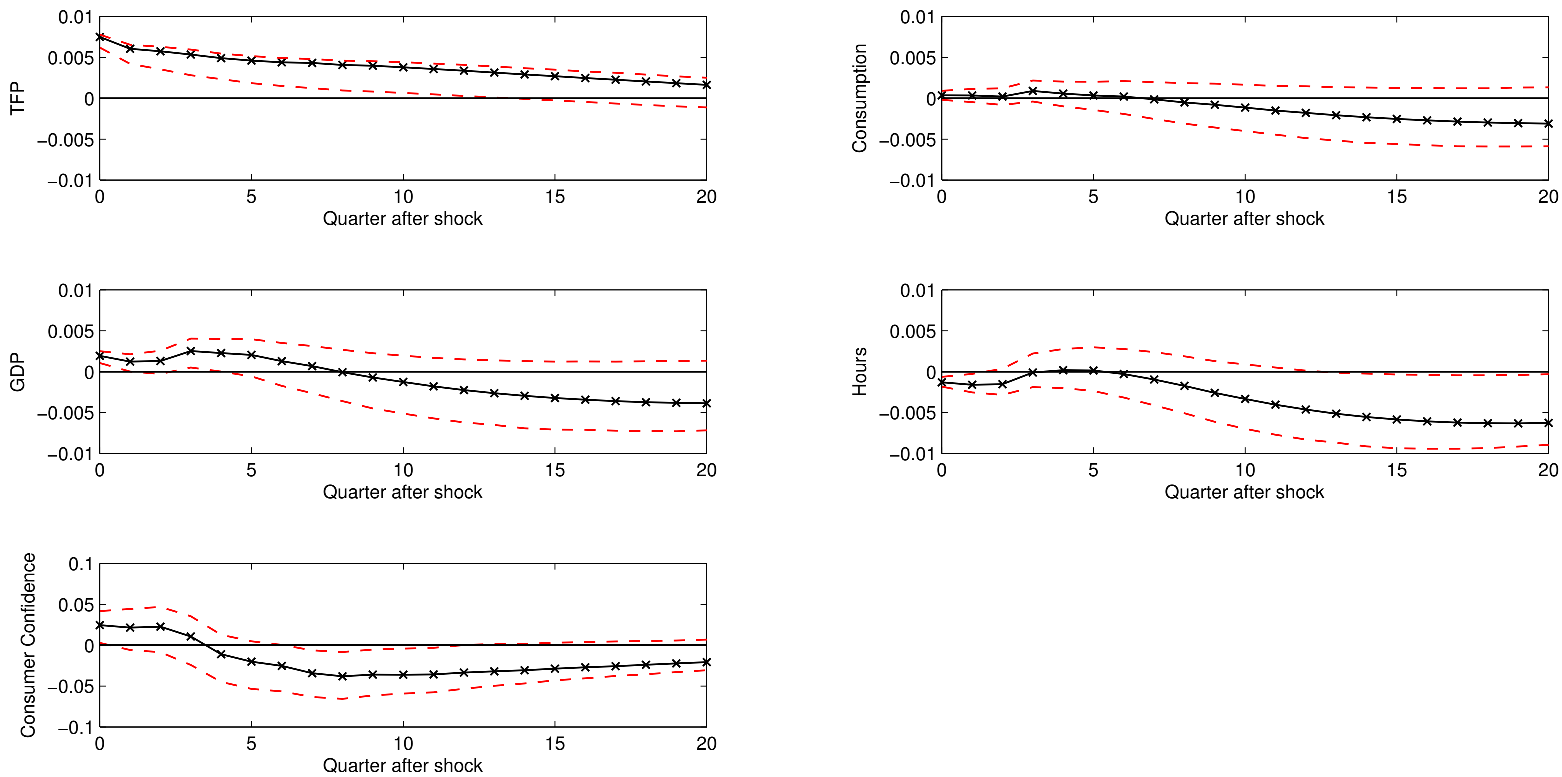

Notes: This figure plots the impulse responses of the core US variables to a surprise TFP innovation, identified in a VAR with the Michigan Consumer Confidence indicator ordered fifth. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure C.2: The Impulse Responses to the US News TFP Shock, Using US Consumer Confidence

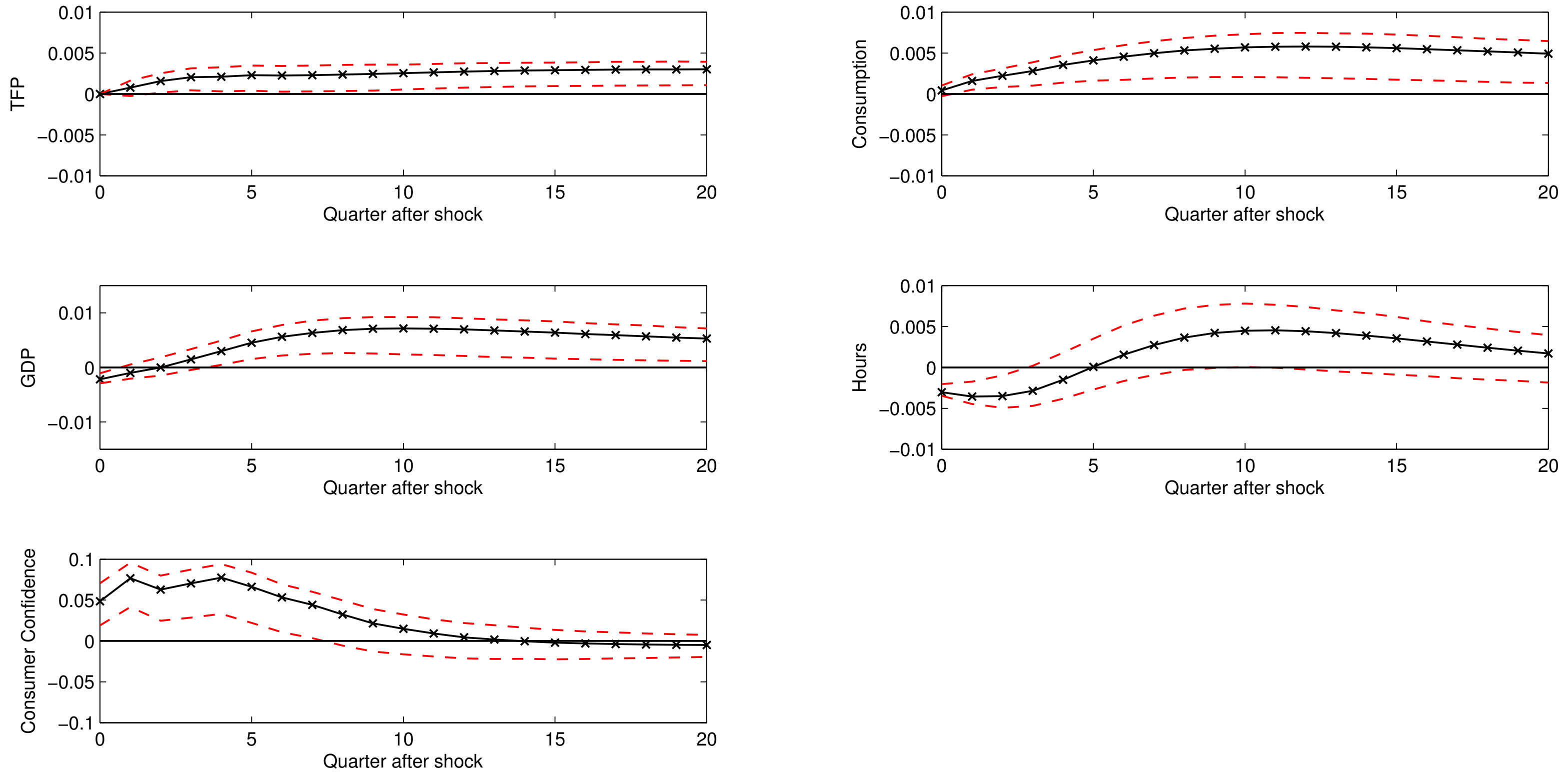

Notes: This figure plots the impulse responses of the core US variables to the news shock, identified in a VAR with the Michigan Consumer Confidence indicator ordered fifth. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure C.3: The Impulse Responses to the US Sentiment Shock, Using US Consumer Confidence

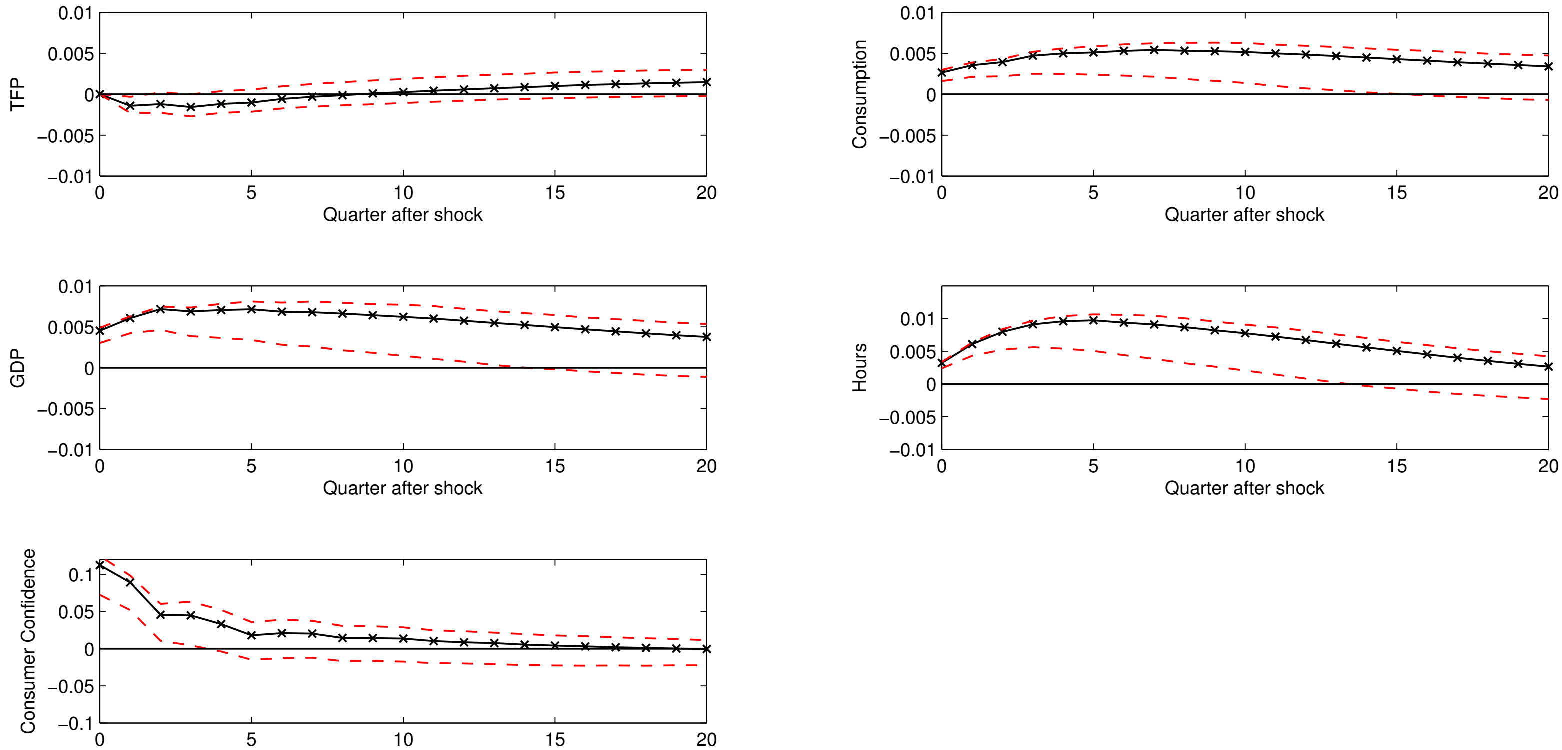

Notes: This figure plots the impulse responses of the core US variables to the sentiment shock, identified in a VAR with the Michigan Consumer Confidence indicator ordered fifth. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure C.4: The Impulse Responses of Canadian TFP to the Three US Shocks, Using US Consumer Confidence

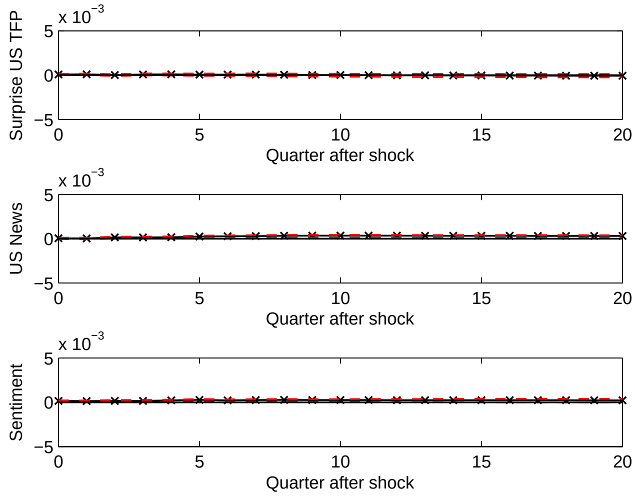

Notes: This figure plots the impulse responses of Canadian utilization-adjusted TFP to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure C.5: The Impulse Responses of Canadian GDP to the Three US Shocks, Using US Consumer Confidence

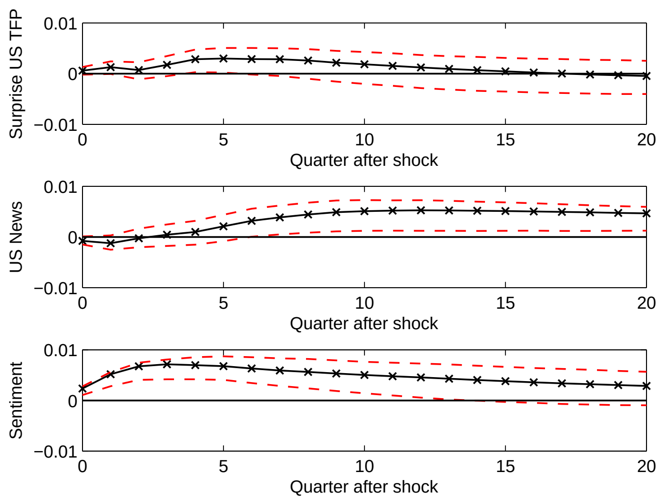

Notes: This figure plots the impulse responses of Canadian GDP to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure C.6: The Impulse Responses of Canadian Exports to the US to the Three US Shocks, Using US Consumer Confidence

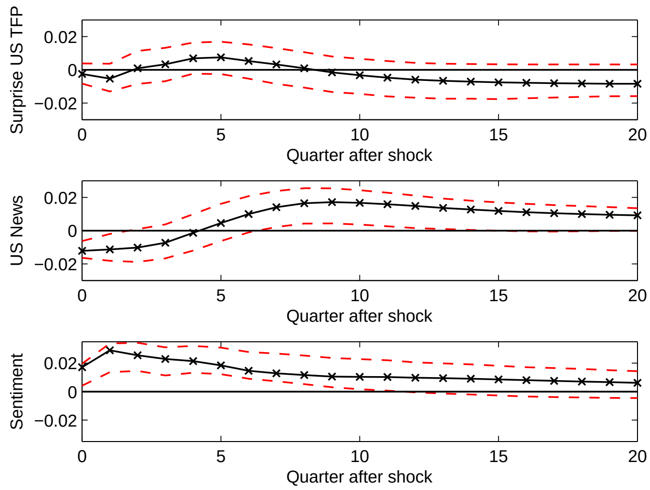

Notes: This figure plots the impulse responses of Canadian exports to the US to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure C.7: The Impulse Responses of Canadian Imports from the US to the Three US Shocks, Using US Consumer Confidence

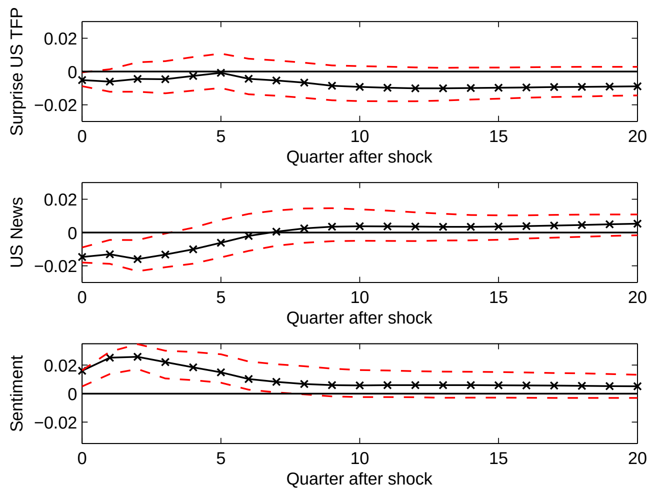

Notes: This figure plots the impulse responses of US exports to Canada to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure C.8: The Impulse Responses of Canadian Hours to the Three US Shocks, Using US Consumer Confidence

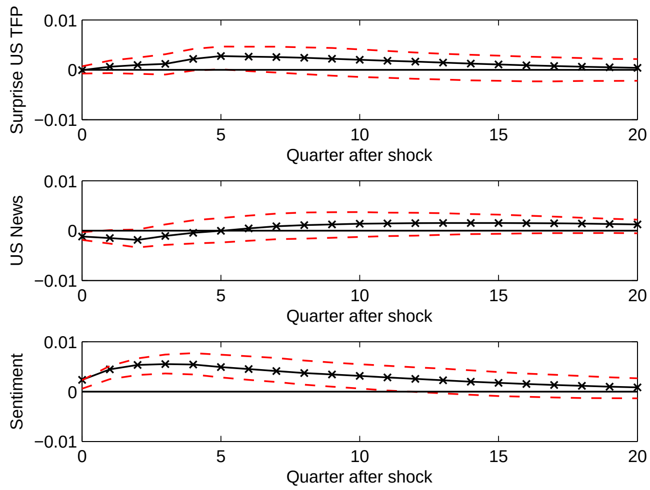

Notes: This figure plots the impulse responses of Canadian hours per worker to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure C.9: The Impulse Responses of Canadian Utilization to the Three US Shocks, Using US Consumer Confidence

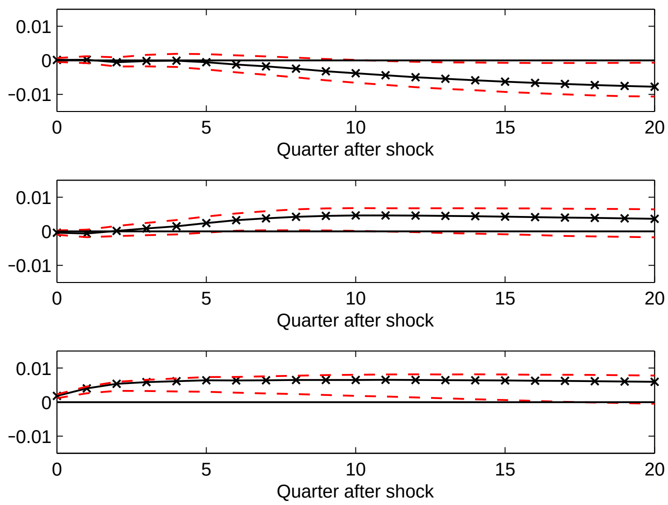

Notes: This figure plots the impulse responses of Canadian capital utilization to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure C.10: The Impulse Responses of Canadian Consumption to the Three US Shocks, Using US Consumer Confidence

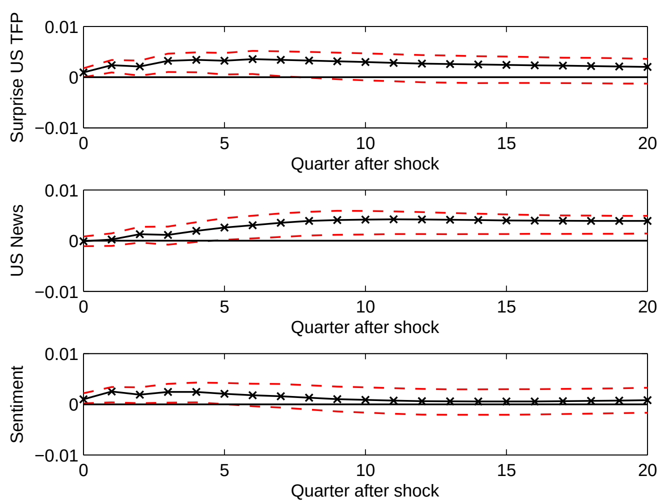

Notes: This figure plots the impulse responses of Canadian consumption to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

BIBLIOGRAPHY

BIBLIOGRAPHY

Acemoglu, D., D. Autor, D. Dorn, G. H. Hanson, and B. Price (2014), Import Competition and the Great U.S. Employment Sag of the 2000s, *NBER Working Papers 20395*, National Bureau of Economic Research, Inc.

Ackerberg, D. A., K. Caves, and G. Frazer (2006), Structural identification of production functions, *Working paper*, University of Michigan.

Alfaro, L., and A. Charlton (2009), Intra-industry foreign direct investment, *American Economic Review*, *99*(5), 2096–2119.

Alviarez, V. (2014), Multinational production and comparative advantage, *Working paper*, University of Michigan.

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

Angeletos, G.-M., and J. La'O (2013), Sentiments, *Econometrica*, *81*(2), 739–779.

Angeletos, G.-M., F. Collard, and H. Dellas (2014), Quantifying confidence, nBER Working Paper 20807.

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

Antràs, P., and D. Chor (2013), Organizing the Global Value Chain, *Econometrica*, *81*(6), 2127–2204.

Antràs, P., and E. Helpman (2004), Global Sourcing, *Journal of Political Economy*, *112*(3), 552–580.

Antràs, P., T. C. Fort, and F. Tintelnot (2014), The Margins of Global Sourcing: Theory and Evidence from U.S. Firms, *Working Paper 154856*, Harvard University OpenScholar.

Atalay, E. (2014), How important are sectoral shocks?, working Paper.

Autor, D., D. Dorn, G. Hanson, and J. Song (2014), Trade adjustment: Worker-level evidence, *American Economic Review*, *129*(4), 1799–1860.

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

Backus, D. K., P. J. Kehoe, and F. E. Kydland (1992), International real business cycles, *Journal of Political Economy*, *100*(4), 745–75.

Backus, D. K., P. J. Kehoe, and F. E. Kydland (1994), Dynamics of the trade balance and the terms of trade: The j-curve?, *American Economic Review*, *84*(1), 84–103.

Backus, D. K., P. J. Kehoe, and F. E. Kydland (1995), International business cycles: Theory and evidence, in *Frontiers of Business Cycle Research*, edited by T. F. Cooley, Princeton University Press, Princeton.

Barba-Navaretti, G., D. Castellani, and A.-C. Disdier (2010), How does investing in cheap labour countries affect performance at home? firm-level evidence from france and italy, *Oxford Economic Papers*, *62*, 234–260.

Barsky, R. B., and E. R. Sims (2011), News shocks and business cycles, *Journal of Monetary Economics*, *58*(3), 273–289.

Barsky, R. B., and E. R. Sims (2012), Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence, *American Economic Review*, *102*(4), 1343–77.

Barsky, R. B., S. Basu, and K. Lee (2014), Whither news shocks?, *NBER Macroeconomics Annual*, *29*(1), 225–264.

Basu, S., J. G. Fernald, and M. S. Kimball (2006), Are technology improvements contractionary?, *American Economic Review*, *96*(5), 1418–1448.

Beaudry, P., and F. Portier (2006), Stock prices, news, and economic fluctuations, *American Economic Review*, *96*(4), 1293–1307.

Beaudry, P., and F. Portier (2014), News-driven business cycles: Insights and challenges, *Journal of Economic Literature*, *52*(4), 993–1074.

Beaudry, P., M. Dupaigne, and F. Portier (2011), Modeling news-driven international business cycles, *Review of Economic Dynamics*, *14*(1), 72–91.

Benhabib, J., X. Liu, and P. Wang (2015a), Sentiments, financial markets, and macroeconomic fluctuations, forthcoming, *Journal of Financial Economics*.

Benhabib, J., P. Wang, and Y. Wen (2015b), Sentiments and aggregate demand fluctuations, *Econometrica*, *83*, 549–585.

Bernard, A., and J. B. Jensen (2007), Firm structure, multinationals, and manufacturing plant deaths, *Review of Economics and Statistics*, *89*(2), 193–204.

Bernard, A., J. B. Jensen, and P. K. Schott (2006), Transfer pricing by u.s.-based multinational firms, *Working paper*, Yale School of Management.

Bernard, A., J. B. Jensen, S. J. Redding, and P. K. Schott (2007), Firms in international trade, *Journal of Economic Perspectives*, *21*(3), 105–130.

Bernard, A., J. B. Jensen, S. J. Redding, and P. K. Schott (2010), Intra-firm trade and product contractibility, *Working Paper 15881*, NBER.

Blanchard, O. J., and D. Quah (1989), The Dynamic Effects of Aggregate Demand and Supply Disturbances, *American Economic Review*, *79*(4), 655–73.

Blanchard, O. J., J.-P. L'Huillier, and G. Lorenzoni (2013), News, noise, and fluctuations: An empirical exploration, *American Economic Review*, *103*(7), 3045–70.

Blasnik, M. (2010), Reclink: A stata module to probabilistically match records, *Working paper*.

Blaum, J., C. LeLarge, and M. Peters (2015), The Gains from Input Trade in Firm-Based Models of Importing, *NBER Working Papers 21504*, National Bureau of Economic Research, Inc.

Boehm, C., A. Flaaen, and N. Pandalai-Nayar (2014a), Input linkages and the transmission of shocks: Firm level evidence from the 2011 tōhoku earthquake, *Working paper*, University of Michigan.

Boehm, C., A. Flaaen, and N. Pandalai-Nayar (2014b), Complementarities in multinational production and business cycle dynamics, *Working paper*, University of Michigan.

Braconier, H., and K. Ekholm (2000), Swedish multinationals and competition from high and low-wage locations, *Review of International Economics*, *8*(3), 448–461.

Burstein, A., C. Kurz, and L. Tesar (2008), International trade, production sharing and the transmission of business cycles, *Journal of Monetary Economics*, *44*(4), 775–795.

Busso, M., J. DiNardo, and J. McCrary (2014), New evidence on the finite sample properties of propensity score matching and reweighting estimators, *Review of Economics and Statistics*, *forthcoming*.

Caliendo, L., and F. Parro (2015), Estimates of the Trade and Welfare Effects of NAFTA, *Tech. Rep. 1*.

Cameron, C., J. Gelbach, and D. Miller (2011), Robust inference with multiway clustering, *Journal of Business and Economic Statistics*, *29*(2), 238–249.

Canova, F. (2005), The transmission of us shocks to latin america, *Journal of Applied Econometrics*, *20*(2), 229–251.

Canova, F., and G. de Nicoló (2003), On the sources of business cycles in the g-7, *Journal of International Economics*, *59*(1), 77–100.

Carvalho, V. M., M. Nirei, and Y. U. Sato (2014), Supply chain disruptions: Evidence from the great east japan earthquake, *Working paper*.

Caselli, F., M. Koren, M. Lisicky, and S. Tenreyro (2014), Diversification through trade, *Working paper*.

Christiano, L., M. Eichenbaum, and C. Evans (1999), Monetary policy shocks: What have we learned and to what end?, in *Handbook of Macroeconomics*, vol. 1, edited by J. Taylor and M. Woodford, chap. 2, pp. 65–148, Elsevier.

Corsetti, G., L. Dedola, and S. Leduc (2014), The international dimension of productivity and demand shocks in the us economy, *Journal of the European Economic Association*, *12*(1), 153–176.

Costinot, A., J. Vogel, and S. Wang (2013), An Elementary Theory of Global Supply Chains, *Review of Economic Studies*, *80*(1), 109–144.

Cravino, J., and A. Levchenko (2014), Multinational firms and international business cycle transmission, *Working paper*, University of Michigan.

Cushman, D. O., and T. Zha (1997), Identifying monetary policy in a small open economy under flexible exchange rates, *Journal of Monetary Economics*, *39*(3), 433–448.

Davis, S. J., and J. Haltiwanger (1992), Gross job creation, gross job destruction and employment reallocation, *Quarterly Journal of Economics*, *107*(3), 819–863.

Davis, S. J., and J. Haltiwanger (2001), Sectoral job creation and destruction responses to oil price changes, *Journal of Monetary Economics*, *48*(3), 465–512.

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

Davis, S. J., J. Haltiwanger, K. Handley, R. Jarmin, J. Lerner, and J. Miranda (2014), Private Equity, Jobs, and Productivity, *American Economic Review*, *104*(12), 3956–90.

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

Debaere, P., H. Lee, and J. Lee (2010), It matters where you go: Outward foreign direct investment and multinational employment growth at home, *Journal of Development Economics*, *91*(2), 301–309.

Decker, R., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2014), The Secular Decline in Business Dynamism in the U.S., *Tech. rep.*

Desai, M. A., C. F. Foley, and J. R. Hines (2004), The costs of shared ownership: Evidence from international joint ventures, *Journal of Financial Economics*, *73*(2).

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

di Giovanni, J., and A. Levchenko (2012), Country size, international trade, and aggregate fluctuations in granular economies, *Journal of Political Economy*, *120*(6), 1083–1132.

di Giovanni, J., A. Levchenko, and I. Méjean (2014), Firms, destinations, and aggregate fluctuations, *Econometrica*, *82*(4), 1303–1340.

Eaton, J., and S. Kortum (2002), Technology, Geography, and Trade, *Econometrica*, *70*(5), 1741–1779.

Ebenstein, A., A. Harrison, M. McMillan, and S. 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.

Ebenstein, A., A. Harrison, and M. McMillan (2015), Why are American Workers getting Poorer? China, Trade and Offshoring, *NBER Working Papers 21027*, National Bureau of Economic Research, Inc.

Eickmeier, S., and T. Ng (2015), How do credit supply shocks propagate internationally? a gvar approach, *European Economic Review*, *74*, 128–145.

Engel, C., and J. Wang (2011), International trade in durable goods: Understanding volatility, cyclicality, and elasticities, *Journal of International Economics*, *83*(1), 37–52.

Faust, J., and E. M. Leeper (1997), When do long-run identifying restrictions give reliable results?, *Journal of Business and Economic Statistics*, *15*(3), 345–53.

Feenstra, R., and G. Hanson (1996), Foreign investment, outsourcing and relative wages, in *The Political Economy of Trade Policy: Papers in Honor of Jagdish Bhagwati*, edited by G. G. Robert Feenstra and D. Irwin, MIT Press.

Feenstra, R. C., P. A. Luck, M. Obstfeld, and K. N. Russ (2014), In search of the armington elasticity, *Working paper.*

Fernald, J. (2014), A quarterly, utilization-adjusted series on total factor productivity, federal Reserve Bank of San Francisco Working Paper 2012-19.

Fernández-Villaverde, J., J. F. Rubio-Ramírez, T. J. Sargent, and M. W. Watson (2007), ABCs (and Ds) of Understanding VARs, *American Economic Review*, *97*(3), 1021–1026.

Fisher, J. D. M. (2006), The Dynamic Effects of Neutral and Investment-Specific Technology Shocks, *Journal of Political Economy*, *114*(3), 413–451.

Flaaen, A. (2013a), Multinational firms in context, *Working paper*, University of Michigan.

Flaaen, A. (2013b), Constructing a bridge between the lexisnexis directory of corporate affiliations and the u.s. business register, *Working paper*, University of Michigan.

Forni, M., and L. Gambetti (2014), Sufficient information in structural VARs, *Journal of Monetary Economics*, *66*, 124–136.

Forni, M., L. Gambetti, and L. Sala (2014), No news in business cycles, *The Economic Journal*, *124*(581), 1168–1191.

Fort, T., and S. Klimek (2015), What's in a name? the effect of industry classification changes on u.s. employment composition, *Working paper*, U.S. Census Bureau.

Francis, N., M. T. Owyang, J. E. Roush, and R. DiCecio (2014), A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks, *The Review of Economics and Statistics*, *96*(3), 638–647.

Frankel, J. A., and A. K. Rose (1998), The endogeneity of the optimum currency area criteria, *Economic Journal*, *108*(449), 1009–1025.

Gabaix, X. (2011), The granular origins of aggregate fluctuations, *Econometrica*, *79*, 733–772.

Galí, J. (1999), Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations?, *American Economic Review*, *89*(1), 249–271.

Galí, J., and P. Rabanal (2005), Technology shocks and aggregate fluctuations: How well does the rbs model fit postwar u.s. data?, *NBER Macroeconomics Annual*, *19*, 225–288.

Goldberg, P., A. K. Khandelwal, N. Pavcnik, and P. Topalova (2010), Imported intermediate inputs and domestic product growth: Evidence from india, *Quarterly Journal of Economics*, *125*(4), 1727–1767.

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

Halpern, L., M. Koren, and A. Szeidl (2011), Imported inputs and productivity, working Paper.

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

Hamilton, J. D. (1994), *Time Series Analysis*, Princeton University Press, Princeton, NJ.

Hanson, G. H., R. J. Mataloni, and M. J. Slaughter (2005), Vertical Production Networks in Multinational Firms, *The Review of Economics and Statistics*, *87*(4), 664–678.

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

Heathcote, J., and F. Perri (2002), Financial autarky and international business cycles, *Journal of Monetary Economics*, *49*, 601–627.

Helbling, T., R. Huidrom, M. A. Kose, and C. Otrok (2011), Do credit shocks matter? a global perspective, *European Economic Review*, *55*(3), 340–353.

Helpman, E., M. Melitz, and S. Yeaple (2004), Exports vs fdi with heterogeneous firms, *American Economic Review*, *94*, 300–316.

Hijzen, A., S. Jean, and T. Mayer (2011), The effects at home of initiating production abroad: Evidence from matched french firms, *Review of World Economics*, *147*(3), 457–483.

Hirschmann, A. (1958), Interdependence and industrialization, in *The Strategy of Economic Development*, Yale University Press, New Haven.

Hummels, D., R. J. rgensen, J. Munch, and C. Xiang (2014), The wage effects of offshoring: Evidence from danish matched worker-firm data, *American Economic Review*, *104*(6), 1597–1629.

Huo, Z., and N. Takayama (2015), Higher order beliefs, confidence, and business cycles, mimeo, University of Minnesota.

Ilzetzki, E., and K. Jin (2013), The puzzling change in the international transmission of u.s. macroeconomic policy shocks, mimeo, LSE.

Imbs, J. (2004), Trade, finance, specialization, and synchronization, *Review of Economics and Statistics*, *86*(3), 723–734.

Imbs, J., and I. Méjean (2011), Elasticity optimism, *Working Paper 7177*, CEPR Discussion Paper.

Imbs, J. M. (1999), Technology, growth and the business cycle, *Journal of Monetary Economics*, *44*(1), 65 – 80.

Jarmin, R., and J. Miranda (2002), The longitudinal business database, mimeo, available at http://www.vrdc.cornell.edu/info7470/2007/Readings/jarmin-miranda-2002.pdf.

Johnson, R. C. (2014), Trade in intermediate inputs and business cycle comovement, *American Economic Journal: Macroeconomics*, *6*(4), 39–83.

Johnson, R. C., and G. Noguera (2012), Accounting for intermediates: Production sharing and trade in value added, *Journal of International Economics*, *86*(2), 224–236.

Jones, C. I. (2011), Intermediate goods and weak links in the theory of economic development, *American Economic Journal: Macroeconomics*, *3*(2), 1–28.

Justiniano, A., and B. Preston (2010), Can structural small open-economy models account for the influence of foreign disturbances?, *Journal of International Economics*, *81*(1), 61–74.

Keller, W., and S. Yeaple (2013), The gravity of knowledge, *American Economic Review*, *103*(4), 1414–1444.

Kilian, L. (1998), Small-Sample Confidence Intervals For Impulse Response Functions, *The Review of Economics and Statistics*, *80*(2), 218–230.

Kim, S. (2001), International transmission of u.s. monetary policy shocks: Evidence from var's, *Journal of Monetary Economics*, *48*(2), 339–372.

King, R. G., and S. T. Rebelo (1999), Resuscitating real business cycles, in *Handbook of Macroeconomics*, vol. 1, edited by J. B. Taylor and M. Woodford, chap. 14, pp. 927–1007, Elsevier.

Kollmann, R. (2013), Global banks, financial shocks, and international business cycles: Evidence from an estimated model, *Journal of Money, Credit and Banking*, *45*(s2), 159–195.

Konings, J., and A. P. Murphy (2006), Do multinational enterprises relocate employment to low-wage regions, *Weltwirtschaftliches Archiv (Review of World Economics)*, *142*(2), 267–286.

Koren, M., and S. Tenreyro (2013), Technological diversification, *American Economic Review*, *103*(1), 378–414.

Kosaka, M. S. (2013), News-driven international business cycles, *The BE Journal of Macroeconomics*, *13*(1), 751–793.

Kovak, B., L. Oldenski, and N. Sly (2015), The labor market effects of offshoring by u.s. multinational firms: Evidence from changes in global tax policies, *Working paper*.

Kremer, M. (1993), The o-ring theory of economic development, *Quarterly Journal of Economics*, *108*(3), 551–575.

Kurmann, A., and E. Mertens (2014), Stock prices, news, and economic fluctuations: Comment, *American Economic Review*, *104*(4), 1439–1445.

Kurz, C. J., and M. Z. Senses (2013), Importing, exporting and firm-level employment volatility, *Working Paper 2013-44*, Board of Governors of the Federal Reserve System.

Leontief, W. (1936), Quantitative input and output relations in the economic systems of the united states, *The Review of Economics and Statistics*, *18*(3), 105–125.

Lorenzoni, G. (2009), A theory of demand shocks, *American Economic Review*, *99*(5), 2050–84.

Lütkepohl, H. (2005), *New Introduction to Multiple Time Series Analysis*, Springer-Verlag, Berlin Heidelberg.

Maćkowiak, B. (2007), External shocks, u.s. monetary policy and macroeconomic fluctuations in emerging markets, *Journal of Monetary Economics*, *54*(8), 2512–2520.

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

Milani, F. (2014), Sentiment and the u.s. business cycle, mimeo, UC Irvine.

Miyamoto, W., and T. L. Nguyen (2014), Understanding the cross country effects of u.s. technology shocks, mimeo, Bank of Canada and Santa Clara University.

Monarch, R., J. Park, and J. Sivadasan (2013), Gains from offshoring? evidence from us microdata, *US Census Bureau Center for Economic Studies Paper, No. CES-WP-13-20*.

Muendler, M.-A., and S. O. Becker (2010), Margins of Multinational Labor Substitution, *American Economic Review*, *100*(5), 1999–2030.

Nam, D., and J. Wang (2015), The effects of surprise and anticipated technology changes on international relative prices and trade, *Journal of International Economics*, *97*(1), 162–177.

Ohanian, L. E., and A. Raffo (2012), Aggregate hours worked in oecd countries: New measurement and implications for business cycles, *Journal of Monetary Economics*, *59*(1), 40–56.

Peek, J., and E. S. Rosengren (1997), The international transmission of financial shocks: The case of japan, *American Economic Review*, *87*(4), 495–505.

Peek, J., and E. S. Rosengren (2000), Collateral damage: Effects of the japanese bank crisis on real activity in the united states, *American Economic Review*, *90*(1), 30–45.

Pierce, J., and P. 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, J., and P. Schott (2013), The Surprisingly Swift Decline of U.S. Manufacturing Employment, *Working Papers 13-59*, Center for Economic Studies, U.S. Census Bureau.

Pravin, K., and A. Levchenko (2014), Comparative advantage, complexity, and volatility, *Journal of Economic Behavior and Organization*, *94*, 314–329.

Ramey, V. A. (1989), Inventories as factors of production and economic fluctuations, *American Economic Review*, *79*(3), 338–354.

Ramondo, N., V. Rappoport, and K. J. Ruhl (2014), Horizontal vs vertical fdi: Revisiting evidence from u.s. multinationals, *Working paper*, Columbia Business School.

Ruhl, K. J. (2004), The international elasticity puzzle, *Working paper*, NYU Stern School of Business.

Schmitt-Grohé, S. (1998), The international transmission of economic fluctuations: Effects of u.s. business cycles on the canadian economy, *Journal of International Economics*, *44*(2), 257–287.

Schnell, M., and D. E. Weinstein (2012), Evaluating the economic response to japan's earthquake, *Working Paper 12-P-003*, RIETI.

Simpson, H. (2012), Investment abroad and labour adjustment at home: Evidence from uk multinational firms, *Canadian Journal of Economics*, *45*(2), 698–731.

Sims, E. R. (2012), News, non-invertibility, and structural vars, *Advances in Econometrics*, *28*, 81–136.

Slaughter, M. (2000), Production transfer within multinational enterprises and american wages, *Journal of International Economics*, *50*(2), 449–472.

Stern, R., J. Francis, and B. Schumacher (1976), Price elasticities in international trade, in *The Strategy of Economic Development*, Macmillan, London.

Stockman, A. C., and L. L. Tesar (1995), Tastes and technology in a two-country model of the business cycle: Explaining international comovements, *American Economic Review*, *85*(1), 168–85.

Uhlig, H. (2003), What drives gnp?, mimeo, Euro Area Business Cycle Network.

Uhlig, H. (2004), Do technology shocks lead to a fall in total hours worked?, *Journal of the European Economic Association*, *2*(2-3), 361–371.

Wald, D. J., B. C. Worden, V. Quitoriano, and K. L. Pankow (2006), Shakemap manual: Technical manual, user's guide and software guide, *Tech. rep.*, U.S. Geological Survey.

Walker, E. (1997), The census bureau business register: Basic features and quality issues, paper presented at the Joint Statistical Meetings, Anaheim, CA.

Wasi, N., and A. Flaaen (2014), Record linkage using stata: Pre-processing, linking and reviewing utilities, *Working paper*, University of Michigan.

Wen, Y. (2007), By force of demand: Explaining international comovements, *Journal of Economic Dynamics and Control*, *31*(1), 1–23.

Whelan, K. (2000), A guide to the use of chain aggregated nipa data, mimeo, Federal Reserve Board of Governors.