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

Christoph Boehm Aaron Flaaen Nitya Pandalai Nayar

July 14, 2014

PRELIMINARY AND INCOMPLETE

Abstract

Using novel firm-level microdata and leveraging a natural experiment, this paper provides causal evidence for the role of trade and multinational firms in the cross-country transmission of shocks. Foreign multinational affiliates in the U.S. exhibit substantial intermediate input linkages with their source country. The scope for these linkages to generate cross-country spillovers in the domestic market depends on the elasticity of substitution with respect to other inputs. Using the 2011 Tōhoku earthquake as an exogenous shock, we estimate this elasticity for those firms most reliant on Japanese imported inputs: the U.S. affiliates of Japanese multinationals. These firms suffer large drops in U.S. output in the months following the shock, roughly one-for-one with the drop in imports. A structural model of input linkages, input inventories, and variable capacity utilization yields the firm-level estimate of the elasticity of substitution to be very low, nearly that implied by a Leontief production function.

JEL Codes: F10, F12, F23, F41, F44

Keywords: Multinational Firms, Production Elasticity of Substitution

*University of Michigan, Department of Economics, 238 Lorch Hall, Ann Arbor, MI 48109-1220. chrisboe@umich.edu, aflaaen@umich.edu and npnayar@umich.edu. We would like to thank Andrei Levchenko, Kyle Handley, Matthew Shapiro, and Linda Tesar for valuable comments, suggestions, and support. We also benefitted from conversations with Ruediger Bachmann, Javier Cravino, Chris House, Lutz Kilian, Maggie Levenstein, William Lincoln, Ryan Monarch, and Jeff Smith. This research has received research grants from the Michigan Institute for Teaching and Research in Economics (MITRE), for which we are very thankful. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure no confidential information is disclosed.

1 Introduction

One of the most important and encompassing topics in international economics in recent decades has been the spillover effects of trade and financial linkages across countries. The large expansions in trade and foreign direct investment (FDI) in the past twenty years have generated much discussion in whether they increase volatility (di Giovanni and Levchenko (2012)), increase comovement (Frankel and Rose (1998), Burstein, Kurz, and Tesar (2008)) or lead to less diversified production and specialization (Imbs (2004)). Identifying the micro-foundations underlying the role of these cross-country activities in the increased co-dependence of national economies is a complicated task, however. Advanced economies are highly connected, and most variables influenced by any candidate mechanism are often correlated with other developments in the source and host 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 has 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). There are two advances of this paper, however, that permit a rare quantitative evaluation of the nature and magnitude of these linkages. First, we utilize a novel dataset consisting of restricted Census Bureau microdata with new links to international ownership and affiliate information to narrow in on particular firms and their underlying behavior. In addition, a new methodology for isolating firm-level imports intended for further manufacture focuses attention on the channel of vertical linkages. Second, we utilize the Tōhoku earthquake and tsunami as a natural experiment of a large and exogenous shock disrupting the production linkages emanating from Japan in the months following March 2011.

Affiliates of Japanese multinationals are the firms most directly impacted by this shock in the U.S. economy. These affiliates display near universally large imports from their source country, a feature that is common to foreign multinational affiliates more generally (see Flaaen (2013b)). Substantial source-country imports do not by themselves point to these firms, or input-linkages more generally, as a key mechanism underlying the transmission of cross-country shocks. For example, changes to the characteristics of imported goods intended for final sale should have little effect on production or usage of factor inputs in the United States. Imported goods intended as intermediate inputs in domestic production could affect the production process directly, but only if these inputs are not easily substitutable with intermediates from alternative sources. In other words, the potential for the transmission of shocks via imported intermediate inputs relies crucially on 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 event. For the Japanese multinational affiliates on average, we find this elasticity to be very low, close to zero. Perhaps surprisingly, inventories do not appear to play a role in the firm responses to the shock as output shipments fall without a lag by a magnitude similar to the drop in imports. To explore the interaction of inventories, supply-chain disruptions, and other features such as capacity utilization, we develop a structural model which allows us to estimate this elasticity at the firm-level. Low inventories together with rigid production linkages imply the firm expects such disruptions to occur with such infrequency as to make the insurance mechanism of inventories not worth the cost.

The elasticity of substitution shows up in various forms in a wide span of models involving the exchange of goods across countries. And, as discussed by Backus, Kehoe, and Kydland (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 lack of available estimates for the elasticity of substitution, it has become a common practice to demonstrate the behavior of these models along a wide array of parameter values.

It is well known that a low value for this parameter (broadly stated to include 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) the relative degree of output vs consumption comovement are switched relative to the data. Due to the robust nature of these shortcomings,

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

Backus, Kehoe, and Kydland (1995) refer to them as the "price anomaly" and "quantity anomaly" respectively.

To understand the relationship between the elasticity and the quantity anomaly, it is helpful to recall that these models generate output comovement by inducing comovement in factor supplies, a source which 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, Kurz, and Tesar (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.[2] A recent paper by Johnson (2014) looks at a similar question, but applies greater input-output structure on the model. Vertical linkages 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.

It is also relatively straightforward to see how a lower elasticity increases volatility in the terms of trade. When 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): "greater complementarity is associated with a larger return to relative scarcity." However, it should be noted that Backus, Kehoe, and Kydland (1995) indicate that, in their model, a lower elasticity of substitution improves the variability of prices only to make the variability of quantities much smaller than is observed in the data. Perhaps defining the difference between imported intermediates used in domestic production and imported final goods would help to improve this tradeoff.

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)).[3] 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 4.

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

[3] Our results are based on a high-frequency shock, and therefore the presence of short run adjustment costs implies that they do not speak to the sort of labor volatility that is the subject of papers such as Kurz and Senses (2013).

This paper is also a contribution to the empirical evidence on the role of individual firms in aggregate fluctuations, emanating from the pioneering work of Gabaix (2011). Other related evidence comes from di Giovanni, Levchenko, and Méjean (2012), which uses 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 can arguably account for a quantitatively evident drop in U.S. industrial production following the Tōhoku earthquake (see Figure 3).

As is the case with most research based on an event-study, some care should be taken in extending the results to other situations. On the other hand, it should not be surprising that there are few substitutes for the specialized intrafirm trade undertaken between multinational affiliates. Other research appears to confirm the role of multinationals in business cycle comovement: Using the regional dispersion of multinationals within France, Kleinert, Martin, and Toubal (2012) find a positive correlation with regional GDP comovement with the origin of multinationals in that region. They also find a high share of import shipments for these multinational affiliates.

The rest of the paper proceeds as follows: Section 2 describes the empirical strategy and data sources used in this paper, and outlines evidence in favor of a low production elasticity of imported inputs for Japanese multinational affiliates in the U.S. We then incorporate the empirical features in a structural, single firm model of multinational production in Section 3. We estimate the parameters of this model and discuss their implications. Section 4 discusses a number of checks and robustness exercises. The final section offers concluding thoughts.

2 Empirical Strategy and Specification

This section outlines the empirical approach of the event-study framework surrounding the 2011 Tōhoku event as a way to estimate the production elasticity of imported inputs. We discuss the relevant details of this tragedy, document the aggregate effects, and then outline the empirical specification for the firm-level analysis we employ in subsequent sections.

2.1 Background

The Tōhoku earthquake and tsunami, which took place off the coast of Northeast Japan on March 11, 2011, offers a promising natural experiment to study the transmission of economic shocks. The Tōhoku event had a devastating impact on Japan, with estimates of almost twenty thousand dead or missing (see Schnell and Weinstein (2012)). The magnitude of the earthquake was recorded at 9.0 (Mw), which makes it the fourth largest earthquake event ever recorded in the modern era.[4] Most of the damage and casualties, however, 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 1 describes 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. 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 the US. It should be noted, however, that the largest Japanese ports (Yokohama, Tokyo, Kobe) which account for the considerable majority of Japanese 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 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 in magnitude to 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.

[4] 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)

Figure 3 demonstrates that there is indeed a drop in 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 seems clear.

Though tragic, the Tōhoku event offers a rare glimpse into the cross-country spillovers following an exogenous supply shock. The event satisfies much of what is required for it to be useful for purposes of identification. It was clearly exogenous with respect to any underlying features of the economy that could potentially influence the behavior of firms or aggregate variables. The shock was also specific to one country, which means any effects on other countries cannot be directly tied to the shock itself. The challenge associated with this using this event as an natural laboratory for cross-border spillovers comes from the remarkable rapidity with which Japan recovered from the disaster. Such a short time-span for study significantly limits the available datasets with information at the required frequency.

2.2 Firm-Level Analysis

Identification of the production elasticity of substitution will rely on the relative impacts on output and imported inputs following the shock. To be concrete, consider the production function:

$$x_t = \left[(1-\mu)^{\frac{1}{\psi}} \left[F_D(\cdot)\right]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} \left[M_t^J\right]^{\frac{\psi-1}{\psi}} \right]^{\frac{\psi}{\psi-1}}. \tag{1}$$

where output consists of combining a domestic bundle of inputs $F_D(\cdot)$ (an aggregate of domestic labor, capital, and materials), with a foreign imported input M^J. The parameter μ reflects the optimal input share of M^J in production.[5] The goal is to estimate the parameter ψ governing the degree of substitution between these inputs, using information on the output elasticity with respect to imported inputs ($\frac{dlnx}{dlnM}$) in the months following the shock.

Although it is fairly intuitive that there exists a mapping between the output elasticity and the elasticity of substitution after a shock to one of the inputs, we formalize this relationship with the following result:

Result 1. *Assuming i) constant relative input prices, ii) an aggregate import bundle M^J that*

[5] Specifying the prices of the inputs $F_D(\cdot)$ and M^J as P_F and P_M respectively, the firm's optimization problem yields the following: $\frac{(P_M)^{\psi} M^{J*}}{(P_M)^{\psi} M^{J*} + (P_F)^{\psi} F_D^*(\cdot)} \equiv \mu$.

reflects an optimal mix of subcomponents, and iii) M^J is (weakly) scarce, then the fraction $\frac{dlnx}{dlnM}$ is a monotone function of the parameter ψ.

Proof. See appendix A for details. □

Moreover, it is straightforward to show that $\lim_{\psi \to 0^+} \frac{dlnx}{dlnM^J} = 1$, which implies that a one-for-one drop in output following a drop in imported inputs implies a Leontief ($\psi = 0$) production function. We discuss the effects of relaxing the assumptions underlying this fact in later sections.

2.2.1 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 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.[6]

The core information used in this paper comes from 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 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 trade transactions in the LFTTD includes import clearance and export declaration 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, Jensen, and Schott (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 firm-level information on cross-country ownership and production linkages, and are a critical 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 Flaaen (2013b) and

[6] All of the analysis in this paper will be done at the firm-level.

appendix B.1. The second novel data resource used in this paper is a system to classify firm-level import transactions based on the expected use of the products. Although intermediate input trade represents as much as two-thirds of total trade (see Johnson and Noguera (2012)), one would expect trade intended for further production to have differing consequences for the transmission of shocks than that intended for final sale. In brief, 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. 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 Flaaen (2013b) or appendix B.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. Geologic 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 select the sample of firms plausibly affected by the shock. Please see Appendix B.3.2 for details. Figure 4 shows the USGS data along with the geocoded affiliate locations.

2.2.2 Empirical Specification

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 nearly a daily frequency, and covers the universe of firms engaged in exporting/importing. For the purposes of characterizing the shock to firm-level imports of intermediate goods, the LFTTD (and supplements identified above) is ideal. Of course, missing on the input side is then the firm's input consumption from domestic sources. We discuss this limitation in some detail in section 4.2.4.

High-frequency information on firms' U.S. production is also a considerable challenge. Once again, traditional sources of this information are limited to annual frequencies and for select samples of manufacturing firms. Thus, recognizing the key advantage of longitudinal frequency which the LFTTD affords, we use a proxy for a firm's U.S. production — namely the firm's exports of goods to North America (Canada and Mexico). The advantage of this approach is the ability to capture the flow of goods at a specific point in time. Moreover, there are few, if any, barriers to North American trade, and transport time is relatively short. Finally, as documented in Flaaen (2013b), exporting is a common feature of these firms, of which exports to North America is by far the largest component. The disadvantage of employing this approach is that it conditions on a positive trading relationship between Japanese affiliates in the U.S. and Canada/Mexico. The underlying assumption is that all firms export a fixed fraction of their U.S. output to neighboring countries. We will return to whether this is a useful proxy in section 4.2.1.[7]

We use this data to conduct an event study around the time of the Tōhoku event. While the primary analysis of the data will occur in the estimation of our structural model, we first construct some reduced form estimates of the effects on Japanese multinationals as a whole. This exercise will serve several purposes. First, it will provide a sense of the magnitude of this shock for these firms, in terms of both imported inputs and U.S. output. Second, it will give some intuition for how we achieve identification in the structural model. Finally, the estimates will serve as a baseline for measuring the production elasticity of substitution for imported inputs, and for discussing other relevant features that might affect this measurement.

Specifically, 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. Constructing the appropriate control group is not immediately clear, however, as the average Japanese affiliate is much different than the average firm in our sample.[8]

For simplicity and transparency, we use other multinational firms – both U.S. and non-Japanese foreign – as the natural control group. Table 5 reports the values for imported inputs and N.A. exports for these two groups in the months preceding the shock. The control group of firms should

[7]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.

[8]As shown in Flaaen (2013b), in terms of shipments the average foreign multinational firm is over 300 percent larger than a domestic manufacturing firm. Relative to an exporting (non-multinational) firm, the average foreign multinational is 200 percent larger.

also serve to soak up common seasonality patterns, and other demand-driven factors in the U.S. market.

To isolate the magnitude of the shock for a representative Japanese multinational, we look at the total imports of intermediate products for these firms. As is clear from Table 5, the share of imported inputs coming from Japan – 70% of the total for Japanese firms vs 2% for non-Japanese multinationals – is the critical feature driving the identification. Specifically, let M_{jt}^{INT} be intermediate imports of firm j in month t. Consider the following specification:

$$M_{jt}^{INT} = \alpha_j + \lambda_j t + \sum_{i=-4}^{9} \gamma_i E_i + \sum_{i=-4}^{9} \beta_i E_i D_{j,i} + u_{j,t} \tag{2}$$

where α_j are firm fixed-effects, λ_j removes a firm-specific linear trend (extrapolated post March 2011), 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.[9]

The β_i coefficients are of primary interest. The variables $D_{j,t}$ are dummy variables equal to one if the firm is owned by a Japanese parent company. Interacting these dummy variables with each month of the panel allows for a time-varying effect of Japanese ownership on a firm's overall basket of intermediate input imports, particularly during and after the Tōhoku event. The β_i 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 affiliates. To evaluate the differential impact on production for Japanese firms, we can simply replace the dependent variable in equation (2) with the firm's North American exports, denoted $X_{j,t}^{NA}$.

It is important to highlight that the specification in (2) 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. We are interested in calculating the mean effect of the Tōhoku event on some notion of a representative Japanese firm, rather than the effect on the mean Japanese firm. In a world of heterogeneous firms these effects are distinct. The former notion is captured with the levels specification that accounts for the absolute changes relative to a group-specific baseline. In section 3, we evaluate this framework with the results one would obtain when estimating the effect on a firm-by-firm basis.

[9] The baseline sample will consist of January 2009 to December 2011. We denote March 2011 as t=0.

Finally, there are several assumptions underlying this identification strategy. For the $\beta_i s$ to be valid estimates of the mean effect for Japanese affiliates, it must be the case that the control group is not itself "treated" by the shock. Specifically, for the non-Japanese multinationals comprising the control group, the share of Japanese inputs in the overall basket of imported inputs should be small enough to have no measurable effects on the whole. That is, there should be no direct effects on the control group of firms.[10] The low Japanese share of imported inputs for non-Japanese firms provides supporting evidence for this assumption.

2.2.3 Results: Total Manufacturing Sector

The top panel of Figure 5 plots the $\beta_i s$ from equation (2) for the months surrounding the Tōhoku event. As expected, there is a large drop in total intermediate input imports by Japanese firms in the months following the earthquake, relative to the control group. The drop in intermediate inputs bottoms out at 4 million USD in $t = 3$ (June 2011) and the 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 5, which looks for evidence for 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 5, the standard errors increase dramatically in the months following the shock, a feature we interpret to reflect heterogeneous 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 5 and divide by the average pre-shock level for these firms from Table 5. The results, plotted jointly in Figure 6, 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. In subsequent sections, we set up a structural model that will

[10]Additionally, this assumption also rules out any indirect effects of the Tōhoku event, such as general equilibrium effects on factor prices, on the control group of firms. Given the span of time we are considering here, we believe such effects to be negligible.

serve to use these empirical estimates to estimate features of the production elasticity of imported inputs for these firms.

The one-for-one drop in production following a drop in imported inputs shown in Figure 6 suggests a production function that is essentially Leontief in the imported input. On the other hand, one should be cautious in reaching this conclusion while abstracting away from many other practical features of firm behavior. In a later section, (see section 4) we both assess and discuss a number of factors that might influence our reduced-form results from section 2.2.3. Translating the reduced-form estimates into a notion of a production elasticity of substitution relies, naturally, on the magnitude of the production impact. For now, we simply point out that there are a number of features one could include to *dimish* the effect of such a shock on firm production; there are relatively few that would serve to artificially *magnify* the effect.

3 A Model of Cross Country Input Linkages

The results from the previous section lend strong support to a low elasticity of substitution for imported inputs in the production function of Japanese multinational affiliates. This section imposes more structure on this framework to further the analysis. Specifying a structural model has a number of advantages over the reduced-form analysis employed in this section. First, it allows us to directly account for factors that could plausibly influence our measurement, such as intermediate inventory holdings and variable utilization. Second, estimating the model can provide (arguably) a more precise quantitative range corresponding to a particular parameter. Third, the greater structure will allow us to expand the range of parameters for which we can extract information from the data. And finally, the structural approach allows for the distinction between a firm-level estimate of the production elasticity of substitution (a micro-elasticity) vs an aggregate, group-level estimate (macro-elasticity). As argued above, the appropriate aggregation of this parameter will depend on the modeling framework within which it is used.

In this section we build a structural model with the objective of estimating the elasticity of substitution between imported and domestically-sourced intermediate inputs. The model is intended to be simple and only incorporates the salient features of the scenario under study. We adopt a partial equilibrium specification of a firm with production facilities in the U.S. that sources differentiated intermediate inputs from Japan and other countries. Central to our analysis is the possibility of supply chain disruption. While a firm can choose the quantity of intermediate inputs under normal conditions, the deliveries become exogenous to the firm when its supply chain is disrupted. Recognizing this, the firm endogenously holds inventories of imported intermediate inputs. Time is discrete and the length of a period is one month.

3.1 The Model

3.1.1 Production function and demand:

We assume that the production function of the firm is a CES aggregate of the form

$$x_t = \phi \left[(1-\mu) \left[a_t (k_t)^{\alpha} (u_t)^{\beta} (l_t)^{1-\alpha-\beta} \right]^{\frac{\psi-1}{\psi}} + \mu \left[\left((1-\nu) (m_{-J,t})^{\frac{\omega-1}{\omega}} + \nu (m_{J,t})^{\frac{\omega-1}{\omega}} \right)^{\frac{\omega}{\omega-1}} \right]^{\frac{\psi-1}{\psi}} \right]^{\frac{\psi}{\psi-1}} . \quad (3)$$

Here $m_{J,t}$ are materials sourced from Japan, $m_{-J,t}$ are materials sourced from all countries other than Japan, a_t the firm's productivity, k_t capital, l_t labor, and u_t denotes a utilization variable. We think of u as capturing inputs that are easily adjustable such as electricity, gas, and water usage, changing the number of shifts of hourly workers, and so forth. The ϕ is a scaling parameter, μ the weight on the materials aggregate and ν the weight on specialized materials from Japan. The main objects of interest are ψ and ω . The ψ parameter represents the elasticity of substitution between materials and the capital-labor-utilization aggregate. The ω parameterizes the elasticity of substitution between specialized Japanese and non-Japanese intermediate inputs.

Let the price of the firm's good be $p_{x,t}$. Consistent with much of the literature on multinational production, we assume that the firm is monopolistically competitive and faces a downward sloping (CES) demand function

$$p_{x,t} = \left(\frac{Y_t}{x_t} \right)^{\frac{1}{\varepsilon}}, \quad (4)$$

where Y_t is a measure of aggregate demand and ε the associated elasticity.

3.1.2 Supply chain disruptions and inventories:

We model supply chain disruptions to the firm as a two-state Markov process. When the firm is in the normal state (n), the firm can choose how many intermediate inputs z_t to import from Japan.

Let s_t denote the stock of Japanese-sourced intermediate inputs at the beginning of the period. After importing z_t and using $m_{J,t}$ in production, the firm stores the remainder to have

$$s_{t+1} = (1-\delta)s_t + z_t - m_{J,t} \tag{5}$$

units available in the next period. The δ term denotes the depreciation rate of stored intermediate inputs.

If the firm enters the disrupted state, intermediate input shipments from Japan become exogenous to the firm and follow the AR(1) process

$$z_{t+1} = (1-\rho)\,\bar{z} + \rho z_t + \sigma_z \varepsilon_{z,t+1} \tag{6}$$

where $\varepsilon_{z,t+1}$ is a white noise error. We denote the probability of entering the disrupted state conditional on being in the normal state by q_n. The probability of staying in the disrupted state is then q_d.

In this model the only motive for inventory accumulation is to continue production in periods in which the supply chain is disrupted. An additional possibility is to allow the firm to hold inventories in its final good. In that case it can sustain a desired level of sales at times of supply chain disruption even while reducing production. Given our data the resulting behavior is indistinguishable. We focus on intermediate input inventories as they appear to be more directly relevant, both conceptually and in the data.

3.1.3 Simplifying assumptions and adjustability of capital and labor:

Since we model the behavior of firms over a time horizon of less than a year, we believe that a number of simplifying assumptions are appropriate. In particular, we assume that Y_t, characterizing overall demand, and a_t, the firm's productivity, are constant.[11] Furthermore, we assume a partial-equilibrium setting in which all factor prices are held constant. This corresponds to $\{R, w, p_u, p_{-J}, p_J\}$, denoting, respectively, the rental rate, the wage, the price of a unit of the utilization variable, the price of non-Japanese materials, and the price of Japanese materials. The

[11] This is consistent with those private-sector estimates of monthly U.S. GDP (i.e. constructed by Macroeconomic Advisers). Extracting a monthly "factor" series from a wide array of relevant data, as in Bernanke, Boivin, and Eliasz (2005), yields a similar conclusion.

R, w, p_u, and p_{-J} variables are determined in general equilibrium and are unlikely to be affected by an individual firm. We have discussed above the evidence in support of our assumption that p_J is constant.

Furthermore, we assume that disruptions are sufficiently short to make it unprofitable for the firm to adjust its capital stock or labor force. Instead, the labor force and the capital stock remain at their previous levels when the firm enters a disrupted period. This is once again consistent with the analysis of firm employment and payroll in Section 4.2.4.

3.1.4 Bellman equations:

We set up the firm's recursive problem using the notation of dynamic programming, such that we drop time subscripts and denote the next period's value of a variable a as a'. In the normal state, the value function is

$$V^n(s) = \max_{u,k.l,m_{-J},m_J \geq 0,z} \Big\{ p_x x - p_u u - Rk - wl - p_{-J}m_{-J} - p_J z + ... \\ \frac{1}{1+r}\left[(1-q_n)V^n(s') + q_n \mathbb{E}_n\left[V^d(s',z')\right]\right]\Big\} \quad (7)$$

where the maximization is subject to production function (3), demand function (4), and the transition equation of stored intermediate inputs from Japan (5). Notice that $\mathbb{E}_n[\cdot]$ denotes the expectations operator with respect to the *unconditional* distribution of z'.

While all inputs are subject to the firm's choice in normal periods, capital and labor are fixed once the firm enters the disrupted state. Additionally, intermediate input shipments from Japan are beyond the control of the firm and hence become a state variable. The value function in the disrupted state is then

$$V^d(s,z) = \max_{u \geq 0, m_{-J} \geq 0, m_J \geq 0} \Big\{ p_x x - p_u u - R\bar{k} - w\bar{l} - p_{-J}m_{-J} - \frac{c}{2}(m_{-J} - \bar{m}_{-J})^2 - p_J z + ... \\ \frac{1}{1+r}\left[(1-q_d)V^n(s') + q_d \mathbb{E}_d\left[V^d(s',z')\right]\right]\Big\} \quad (8)$$

and the maximization is subject to (3), (4), (5), and the AR(1) process for z, (6). To indicate that capital and labor are held fixed in the disrupted state, we denote them by $\bar{k}$ and $\bar{l}$. Finally, we add the term $\frac{c}{2}\left(m_{-J}-\bar{m}_{-J}\right)^2$ to penalize choices of m_{-J} that differ from its level in the normal state $\bar{m}_{-J}$. This term is intended to capture the fact that shipments of intermediate inputs are often fixed contractually several quarters in advance. The constant c parameterizes the severity of this penalty.

In contrast to the value function in the normal state (equation 7), expectations over next period's shipments z' are now formed based on the knowledge of this period's z. To make this distinction explicit, we denote this *conditional* expectations operator by $\mathbb{E}_d[\cdot]$. All technical details of the model solution are deferred to the appendix.

3.1.5 Model Behavior and Identification

Before we proceed with the estimation we briefly discuss and illustrate the model's behavior. This discussion also builds intuition for how the parameters (ψ, ω, q_n, c) that we estimate from the variation in the data are identified.

The right hand side of equation (7) represents a fairly standard problem of a neoclassical firm – with one exception. Because of the possibility of supply chain disruption it will often be optimal for the firm to accumulate intermediate input inventories. Consistent with intuition, it can be shown that under the assumption that $\mathbb{E}_d\left[V^d\left(s', z'\right)\right]$ is concave and differentiable in s', the choice of inventory holdings is larger, the smaller are p_J, δ, and r. That is, holding all else equal, lower costs of purchasing and storing the intermediate input lead to higher inventory accumulation. Under the same assumptions, it is true that a larger probability of supply chain disruption q_n, raises the desired level of inventory holdings in the normal state. We will use the fact that inventory holdings are informative about the probability of supply chain disruption (or the firm's belief thereof) to identify q_n in the structural estimation below.

Once the firm enters the disrupted state, it will optimally draw down its inventories. Since the duration of disruption is uncertain, the speed at which it uses up its Japanese-sourced materials inventory crucially depends on its expectations about how long the disruption will last and how many shipments z it will receive. The formation of these expectations depends on the parameters $\bar{z}$, ρ, σ_z, and q_d. If the firm expects a very transitory (low q_d) and mild (high $\bar{z}$) drop in shipments, it will use up its inventories rapidly. In contrast, if the firm expects a protracted period of disruption and low deliveries, then it will slowly draw down its inventories so as to avoid costly stockouts. Below we will estimate the parameters of the exogenous process (6) directly from the shipment data following the shock. Similarly, we will estimate q_d as the conditional probability of staying in the disrupted state. While the estimation of these parameters from the actual data somewhat limits the degree to which firms can be 'surprised' by the severity and duration of the disruption after the Tōhoku event, we view direct calibration of these parameters as too speculative. Notice also that we do not completely eliminate the possibility of surprise. An estimate of parameter q_n near zero would indicate that firms attributed a very low probability to a disruption of this magnitude and length.[12]

We next turn to an informal discussion of the identification of the two elasticities ψ and ω as well as the adjustment cost parameter c. Recall from Figure 6 that our proxy for production (firm exports to Canada and Mexico) and Japanese intermediate imports experienced simultaneous drops of similar magnitudes. This behavior rules out large values of the two elasticities. To see why, suppose ψ was large. Then the firm could increase utilization to prevent output from falling. Similarly, if ω was significantly greater than zero, then the firm could substitute towards intermediates from sources other than Japan. In that case the output proxy should fall by less than observed and non-Japanese intermediate input imports should not drop. Finally, note that in the absence of adjustment costs for non-Japanese-sourced intermediates ($c = 0$), small values of ω would make it optimal for the firm to adjust m_{-J} in lock-step with m_J. Positive values of c make instant adjustment suboptimal.

3.2 Estimation

In this section we proceed in three steps. First, we calibrate those parameters where either a conventional value in the literature exists, or where our data is unlikely to be informative. In the latter case we perform a variety of robustness checks. Second, we estimate the parameters ρ, $\bar{z}$, σ_z, and q_d directly from the data. Third, we estimate the remaining parameters structurally on the model described in the previous section.

[12] When we attempted to structurally estimate q_n and q_d jointly on simulated data, the estimates were very imprecise. It is possible that, given the data limitations we face, q_n and q_d are not separately identified. Intuitively, both parameters together determine the firm's inventory accumulation decision on which we have indirect information. However, q_n and q_d do not independently of one another affect any variables in our dataset in an obvious way.

Step 1: Calibration

First we set the three parameters which contain no embodied information such that $\phi = a = Y = 1$. Next, we set $\varepsilon = 8$, r such that the annual interest rate is 6 percent, and $\delta = 0.01$, i.e. one percent per month. The value for ε is subject to some debate, and so we test the robustness of our model to various values of this parameter.

Next we turn to the parameters governing the equilibrium cost shares of the various inputs used in production. For the ν parameter, we use the CM/LFTTD data on the cost share of Japanese imported inputs discussed in section 2. The μ parameter is then meant to capture the weight of the materials aggregate encompassing Japanese and non-Japanese (including domestic) materials inputs. The cost share of intermediate inputs as a whole is roughly 50 percent, according to recent estimates from Atalay (2014). Then, the remaining parameters governing the capital-labor-utilization aggregate are calibrated such that labor occupies roughly half of the cost share, with the remaining being split evenly between capital and utilization. If we assume a 10 percent profit share of revenue, then we can use the data on the cost shares for the costs of Japanese inputs, non Japanese inputs, labor, capital and utilization to back out values for the prices of each of these variables in the model.[13] Given this choice of units and the implied prices, we can then calculate the values for α, β, ν and μ. [14] Our calibrated parameters are summarized in Table 1.

Table 1: Calibration Details

Parameter	**Value**	**Source/Justification**
ε	8	various
r	0.06/12	6 percent p.a.
δ	0.01	1 percent per month
α	0.266	Cost of capital per Revenue = 0.12
β	0.222	Cost of utilization per Revenue = 0.1
ν	0.435	Cost of Japanese inputs per Revenue = 0.196
μ	0.514	Cost of non-Japanese inputs per Revenue = 0.254

[13] Note, here we are implicitly choosing the units of each type of intermediate, factor input and utilization such that their normal state levels are 1. The equations used here are $\frac{p_J m_J}{p_x x} = 0.196$, $\frac{p_{-J} m_{-J}}{p_x x} = 0.254$, $\frac{wl}{p_x x} = 0.23$, $\frac{Rk}{p_x x} = 0.12$ and $\frac{p_u u}{p_x x} = 0.10$

[14] The equations used here are $\alpha = \frac{R}{w+R+p_u}$, $\beta = \frac{p_u}{w+R+p_u}$, $\nu = \frac{p_J}{p_{-J}+p_J}$ and and $\mu = \frac{p_{-J}}{\left(1-\frac{1}{\varepsilon}\right)(1-\nu)}$

Step 2:

In the next stage of the calibration, we estimate the parameters $\bar{z}$, ρ, and σ_z in equation (6) directly on our data. Furthermore, we also estimate q_d as the conditional probability of staying in the disrupted state.

Step 3: Structural estimation

We are interested in estimating the remaining model parameters $\theta = (\psi, \omega, q_n, c)$. For a given value of θ, the model implies the policy functions $x_t = h(s_t, z_t, d_t; \theta)$ and $m_{-J,t} = g(s_t, z_t, d_t; \theta)$ which serve as the starting point for maximum likelihood estimation. Here, d_t is an indicator variable taking the value one if the firm is in the disrupted state. As indicated in previous sections, we use a proxy for x_t, the firm's exports to North America. The proximity of these neighboring countries combined with the fact that both are part of NAFTA suggests it is reasonable to assume that shipments occur at a regular basis over land without large accumulation of stocks before new shipments are made. Formally, we decompose total production into shipments to US consumers x_t^{US} and consumers in Canada and Mexico $\tilde{x}_t^{CM}$, that is

$$x_t = x_t^{US} + \tilde{x}_t^{CM}. \tag{9}$$

The key exogeneity assumption of the estimation procedure is the following relationship,

$$x_t^{US} = (1 - \kappa)\, x_t - \kappa \sigma_x \varepsilon_{x,t}. \tag{10}$$

In words, consumption of the firm's product is a constant fraction out of total production plus an error term $\varepsilon_{x,t}$ that is uncorrelated with the two state variables (s_t, z_t). Importantly, this assumption implies that the shipments z_t (and the level of inventories s_t) do not systematically change the composition of shipments between US and other North American consumers. Since we will only use data on firms that are in the disrupted state, a correlation between d_t (i.e. whether or not the firm is in the disrupted state), and the error $\varepsilon_{x,t}$ is permissible. The κ and σ_x terms are constants. Combining equations (9) and (10) yields

$$x_t^{CM} = x_t + \sigma_x \varepsilon_{x,t}.$$

That is, (scaled) exports to Canada equal production plus an error term.[15]

We assume that analogous argument holds for (scaled) non-US, non-Japanese imported materials which we denote by $m_{ROW,t}$. Hence, we can write

$$m_{ROW,t} = m_{-J,t} + \sigma_m \varepsilon_{m,t}.$$

For imported materials, the exogeneity assumption requires that the error term $\varepsilon_{m,t}$ does not depend on inventories of Japanese-sourced materials s_t or the current shipments z_t.

Under the assumption that $\varepsilon_{x,t}$ and $\varepsilon_{m,t}$ are jointly normal with zero mean, it follows that

$$\begin{pmatrix} x_t^{CM} \\ m_{ROW,t} \end{pmatrix} | \left(s_t, z_t, d_t = 1\right) = N\left(\begin{pmatrix} h\left(s_t, z_t, 1; \theta\right) \\ g\left(s_t, z_t, 1; \theta\right) \end{pmatrix}, \begin{pmatrix} \sigma_x^2 & \rho_{xm} \\ \rho_{xm} & \sigma_m^2 \end{pmatrix}\right),$$

where ρ_{xm} is the correlation between the two error terms.

Finally, let $L\left(\theta, \sigma_x^2, \sigma_m^2, \rho_{xm}\right)$ denote the likelihood function. Then the maximum likelihood estimator solves

$$\left(\hat{\theta}, \hat{\sigma}_x^2, \hat{\sigma}_m^2, \hat{\rho}_{xm}\right) = \arg \max_{\theta, \sigma_x^2, \sigma_m^2, \rho_{xm}} \ln L\left(\theta, \sigma_x^2, \sigma_m^2, \rho_{xm}\right).$$

When we implement the estimation, we additionally impose that the two shocks $\varepsilon_{x,t}$ and $\varepsilon_{m,t}$ are uncorrelated so that $\rho_{xm} = 0$.

To construct the sample of firms on which we estimate this model, we first take the baseline sample of firms used in section 2.2.3. Next, in order to be able to make valid inferences based on movements in the monthly values of trade, we condition on consistent Japanese imports and N.A. exports, such that we drop any firm that has monthly zeros of either of these transactions in the 12 months preceding the shock. We then calculate the year-on-year growth rates of each series using the measure advanced by Davis, Haltiwanger, and Schuh (1996).[16] For example, the year-on-year growth in imported Japanese intermediates is given by:

$$g_{M^J,i,t} = \frac{(M_{i,t}^J - M_{i,t-12}^J)}{0.5(M_{i,t}^J + M_{i,t-12}^J)} \tag{11}$$

[15] Note that we continue to assume that prices of exported goods are constant.

[16] See Davis, Haltiwanger, and Schuh (1996) or Tornqvist, Vartia, and Vartia (1985) for a discussion of the properties of this growth rate.

Finally we normalize the growth rates by their pre-shock average to get a stationary series at the onset of the shock $\tilde{g}_{M^J,i,t} = \frac{g_{M^J,i,t}}{\sum_{j=t-7}^{j=t-1} g_{M^J,i,j}}$.[17] We define a firm to be disrupted if $\tilde{g}_{M^J,i,t} < 0.7$ in the months following the Tōhoku event, and will remain disrupted until $\tilde{g}_{M^J,i,t} > 0.85$ or 6 months have elapsed.[18] We keep only those firms that have experienced a disruption, and then keep the number of disrupted periods per firm, and the three observables $\left(x_{i,t}^{CM}, m_{ROW,i,t}, z_{i,t}\right)$ where i indexes the firm.

Specific estimation results and robustness checks are withheld until we obtain approval for disclosure from the U.S. Census Bureau. The following section will provide a broad qualitative summary of the results.

3.3 Summary of Results

Discussion of the results: *Coming soon...*

The combination of a low production elasticity of substitution together with very low observed materials inventories can be reconciled with a very low expected probability of supply chain disruption on the part of the firm. Thus, the insurance motive for inventory holdings does not generally cover the holding cost of such inventories.

The estimates from this section correspond to the individual firm, and therefore can be thought of as a micro-level estimates. On the other hand, the results from section 2.2.3 correspond to the group of multinationals as a whole, and therefore can be interpreted as the "macro" elacticity estimate. It is often the case that elasticities calculated at different levels of aggregation can be quite different, even when using the same data. A recent paper 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.[19] We do not find such an aggregation bias in this context, even though with a non-linear model, there is no ex-ante reason for the "micro" and "macro" estimates to be similar.

Because these elasticity estimates play an important role when used to calibrate theoretical

[17] We trim the sample at the 90th and 10th percentiles of this pre-shock average growth rate to remove firms with extreme average growth rates in these input shipments.

[18] When calculating the disruption indicator, we use a 3-month moving average of $\tilde{g}_{M^J,i,t}$ to remove the effects of temporary improvements to import shipments.

[19] Using U.S. data, they find an elasticity of between 2.5 and 3 when imposing homogeneity across sectors. The weighted aggregate elasticity based on sectoral estimates, however, is between 6 and 7.

models, it is important to apply the estimates from the appropriate level of aggregation. The "macro" elasticity we calculate is the most appropriate model analogue for an aggregate DSGE model populated by a representative multinational firm, where the focus is on understanding aggregate features such as business cycle dynamics. The "micro" elasticity, on the other hand, is more relevant for models in which the focus is on the behavior of the firm, such as estimating firm level labor adjustment costs.

4 Discussion

4.1 Implications

The structural estimates of the model are broadly in agreement with what could be observed casually from section 2.2.3: imported inputs are strong complements with other inputs in the production function. On reflection, the rigidity of the production function for multinational firms is not particularly surprising, due in large part to the substantial presence of intra-firm trade in what is presumably highly-specialized inputs. Nonetheless, this result has a number of important implications for how we think about business cycle co-movement, firm volatility, and multinational firms more broadly.

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, Flaaen, and Nayar (2014)), we evaluate the quantitative importance of such complementarities of source-country inputs by mulitnational affiliates. When separately accounting for intermediate input trade by multinationals with traditional trade in final goods, the model separates the production elasticity of imported inputs with the traditional "Armington" elasticity used to bundle together international goods for consumption. The complementarities in import linkages by multinationals increases value-added comovement in the model by 17 percentage points relative to a benchmark without such firms.

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 Tenreyro (2013)). Kurz and Senses (2013) establish that firms with substantial import and export behavior have lower employment volatility than domestic firms in the medium to long term, which they attribute partly to the diversification of risk.[20] 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. That is, a firm with a complex production structure spanning many inputs can have increased volatility if each input is key to the production process. The increase in supply chain risk is a considerable counterweight to diversification, particularly when the production elasticity is low.

In the model, we consider the aggregate bundle of imported intermediates, abstracting away from product-level detail. In reality, the firms in our dataset often import many distinct intermediate inputs. The structure of a CES production function implies that if each of these inputs was non-substitutable with another (a 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.[21] We evaluate this possibility, and the influence it may have on the estimation of the other elasticity parameters, in section 4.2.3 below.

Finally, the rigid production networks of multinational firms influences our understanding of why firms segment production across country borders. In a related paper, Flaaen (2013b) shows that while the motive for multinational production appears to be to serve the domestic market (consistent with the horizontal framework of FDI), substantial import linkages also exist (consistent with a vertical framework). The result could be called "horizontal FDI with production sharing."[22] The evidence for strong complemenarities 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 that is not easily transferable when the firm moves production abroad.[23] Understanding the dynamics behind

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

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

[22] Ramondo, Rappoport, and Ruhl (2007) is another recent example arguing for a more nuanced interpretation of multinational production. The model of knowledge sharing by multinational firms in Keller and Yeaple (2013), is also consistent with this modified framework of multinational activity.

[23] Alternatively, domestic content requirements may provide incentives to produce specified inputs in one location over another.

these sourcing decisions is an area in need of further work.

4.2 Robustness and Extensions

4.2.1 Mis-measurement of Firm Production

The most natural concern would be with respect to using 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 easily obtainable production, sales, and inventory data 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.[24] The baseline specification is the same as in equation (2), where the dependent variable is 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.[25]

Figure 7 shows the results. Relative to their U.S. counterparts, Japanese firms operating in the United States experienced a large drop in production following the Tōhoku event. The fall in production bottomed out in May of 2011 — two months after the event — at just over -6400 units a month.[26]The point estimates return to a level near zero in September of 2011, implying that the shock affected production for nearly 6 months.[27]

We interpret these results to be largely supportive of the results obtained using the exports-based proxy for production. In fact, the percentage drops in the two series are remarkably similar: a trough of 52% at $t = 2$ in the automotive data vs 56% 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

[24] Appendix C.3 details further features of this data and explains how the sample was constructed.

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

[26] The average monthly plant-level production at these firms during December 2010 through February 2011 was about 12,200 units a month. This implies a 52% drop relative to before the event.

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

U.S. production behavior.[28]

4.2.2 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 any measurable impact on the production process.[29] 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 intermediate input inventories in the production impacts of Figures 5 (Panel B) or 7. 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 through 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.[30] Panel A of Table 5 calculates the production-weighted averages over a select set of firm groups.[31] 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 this data corresponds to 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 5 also

[28] 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 4.2.1 should also serve to allay such concerns.

[29] 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.

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

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

reports the corresponding estimates for output inventories.[32]

4.2.3 Multi-Products and Coordination Failures

As is clear in figure 4, the Tōhoku event had a heterogeneous impact across Japan. As a result, one might wonder whether there is considerable dispersion in the impact on the *products* imported by a particular U.S. firm. 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 product mix at all times.

To better understand the implications of these assumptions on the inference with regard to the elasticity calculation, we modify the production function from equation 1 to incorporate multiple imported products. We now have:

$$x_t = \left[(1-\mu)^{\frac{1}{\psi}} \left[F_D(\cdot)\right]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} \left[M_t^J\right]^{\frac{\psi-1}{\psi}} \right]^{\frac{\psi}{\psi-1}} \tag{12}$$

where the term M_t^J is now composed of J products:

$$M_t^J = \left(\sum_{j=1}^{J} \nu_j^{\frac{1}{\zeta}} (m_{j,t}^J)^{\frac{\zeta-1}{\zeta}} \right)^{\frac{\zeta}{\zeta-1}} \tag{13}$$

As implied implicitly before, M_t^J refelcts the optimal mix of products, as identified by the weights v_j. 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 $\widetilde{M}_t^J = \sum_{j=1}^{J}(m_{j,t}^J) \geq M_t^J$. And the lower the elasticity of substitution among products, the more severe the disconnect between the measured imports and the "effective" imports — that which is actually useful in downstream production.

A simple example will help to clarify. Suppose there are two products m_1 and m_2 whose optimal weights are $\nu_1 = 0.8$ and $\nu_2 = 0.2$. These two inputs are strong complements, with an elasticity of

[32] 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.

$\zeta = 0.2$. Imports from product m_2 fall from a steady state value of 20 to 10 (a 50 percent drop) following the shock, whereas the product m_1 falls from 80 to 70 (a 12.5 percent drop). The total measured imports falls from 100 to 80, a 20 percent drop. However, the mix is now suboptimal relative to the steady state (a 0.875 /0.125 mix). Thus, the effective imports according to equation 12 is now roughly 68.4, which is a 31.5 percent drop relative to before the shock.[33] This information is summarized in the table below:

Table 1. A Simple Example With 2 Inputs

	Input 1 (m_1)	Input 2 (m_2)	Measured Imports ($\widetilde{M}^J$)	Effective Imports[1] (M^J)
Before	80	20	100	100
After	70	10	80	68.4
Percent Drop	0.125	0.5	0.2	0.315

[1] Using equation 13 with two products where $\nu_1 = 0.8$, $\nu_2 = 0.2$, and $\zeta = 0.2$.

Suboptimal product mix will imply that measured imports ($\widetilde{M}_t^J$) are greater than the effective imports (M_t^J). As a result the measured output response to the import shock will be larger than otherwise, resulting in a downwarad bias in the elasticity estimate. Such a bias is decreasing in the elasticity parameter itself, as complementarity itself is the driving force between differences in $\widetilde{M}_t^J$ and M_t^J. In addition, the bias is increasing in the distance from the optimal product mix.

Does this exert a quantitatively large effect on our point estimates? First, recall that the deviation from the optimal input mix must extend beyond Japanese production and into the measured imports into the United States. Given the emphasis on low inventories and lean production processes in downstream operations, it is likely that across-product adjustment may take place before sending the inputs abroad. To analyze this empirically, 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. Then, using the LFTTD data we can construct for each firm j, the share of total imports from Japan for a given product code p. We define this share to be $s_{p,t}^j$. Then we can define the distance from optimum DO_t^j as:

[33] Notice that before the shock $\widetilde{M}_t^J = \sum_{j=1}^{J}(m_{j,t}^J) = M_t^J$ since the bundle is optimal.

$$DO_t^j = \sum_{p=1}^{P} \left(|s_{p,t}^j - s_{p,s}^j| \right)$$

where P is the total number of products. If we normalize $DO_s^j = 100$ — which could be defined as the April-May period of 2010 — then we can evaluate how this measure moves at a monthly (or bi-monthly) frequency, with particular interest in the time following the shock. One can calculate this at various levels of product aggregation (i.e. HS4, HS6, HS8, HS10) as well.

4.2.4 Other Considerations

Strategic Behavior: Another possibility that could bias the interpretation of the results from Figure 6 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 $\beta_i s$ from the equation with $X_{j,t}^{NA}$ as the dependent variable.[34] To evaluate this 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 below:

Prices: Traditionally, estimating the elasticity of substitution is accomplished via price and quantity data for products over extended periods of time. There are several reasons why we do not utilize prices and quantities separately in this paper. First, and perhaps most obviously, the duration of this shock is short enough such that 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, Tables 2 and 5 make clear 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

[34] Specifically, the $\gamma_i s$ would be higher than would be expected without the shock, and hence the $\beta_i s$ artificially low.

remained relatively stable during this period. The results in Appendix Table C.3 confirm that there are few significant price movements on import or export transactions for either Japanese or non-Japanese multinationals surrounding the Tōhoku event.[35]

Domestic Inputs and Utilization: So far, we have been restricted to information on imported inputs and firm output – we have ignored firm usage of locally sourced inputs. Whether orders of other inputs increased or decreased would be valuable not only in terms of the elasticity of substitution, but also for the effects of spillovers on other sectors of the economy. Unfortunately, data on domestic input usage at the frequency necessary for evaluating this shock is relatively scarce. The only information at our disposal is firm employment and payroll, at a quarterly frequency, which is taken from the Census Bureau's Business Register (BR).[36] We consider the evidence in Appendix C.1 and find no significant effects on either employment or payroll for Japanese firms in the quarter(s) following the shock (see Table C.3).

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.

4.3 External Validity

Finally, we discuss issues surrounding 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.

When viewed in light of the substantial fraction of intrafirm imports comprising multinational affiliate trade, the low elasticity of substitution should not appear as a surprise. One would not expect close substitutes for the sort of specialized firm-specific products that likely represent this trade. Moreover, such low estimates 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.[37]

[35] Further details on the construction of the data underlying the analysis of unit values is available in Appendix C.2.

[36] 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).

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

The pattern of strong intermediate input linkages with the source country is not restricted to Japanese affiliates only. As shown in Flaaen (2013b), 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.

Any elasticity estimate is somewhat tied to the time-period to which it corresponds. Work by Ruhl (2008) 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, a similar argument could be made that the firm would be more willing to substitute between inputs if the shock was extremely persistent. For a shock believed to last only a few months, the firm might be less willing to restructure existing supplier relationships, and the observed production response would indicate a lower elasticity of substitution than otherwise. Part of our underlying hypothesis, however, is that the inputs sourced by multinational firms in their source country are of particular importance to their production, whether it is because they embody unique firm-specific knowledge, or because they are specialized inputs that are tied to some degree to the location of the source country. If this is true, then we may not expect estimates of these elasticities to be much higher even for a more persistent shock, as the firm cannot easily change the underlying cause of the non-substitutability.

5 Conclusions

Using a novel firm-level dataset to analyze the firm behavior at the time of a large exogenous shock, this paper reveals the mechanisms underlying cross-country spillovers. We find strong rigidities in the international production networks of the U.S. affiliates of Japanese firms, such that U.S. output declines dramatically following the Tōhoku earthquake which caused an equally large decline in imported input shipments. 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. A structural model of imported inputs, variable capacity utilization and material inventories confirms this finding. The model demonstrates that a low elasticity production function together with few material inventories can be rationalized by a very infrequent percieved probability of supply chain disruption by the firm.

This elasticity plays 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, we document the behavior of a model with such complementarities in imported intermediates, and discuss how these two elasticity parameters interact in the model. Calibrating this model to the share of multinational affiliate trade in intermediates yields an increase in value-added comovement by 17 p.p., and reduces the gap between consumption and value-added comovement by an even larger margin.

Such rigid production networks will 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.

References

Atalay, Enghin. 2014. "How Important Are Sectoral Shocks?" Working Paper.

Backus, David K., Patrick J. Kehoe, and Finn E. Kydland. 1994. "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?" *American Economic Review* 84 (1):84–103.

———. 1995. "International Business Cycles: Theory and Evidence." In *Frontiers of Business Cycle Research*, edited by Thomas F. Cooley. Princeton: Princeton University Press.

Bernanke, Ben S., Jean Boivin, and Piotr Eliasz. 2005. "Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach." *Quarterly Journal of Economics* 120 (1):387–422.

Bernard, Andrew, J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott. 2007. "Firms in International Trade." *Journal of Economic Perspectives* 21 (3):105–130.

Bernard, Andrew, J. Bradford Jensen, and Peter K. Schott. 2006. "Transfer Pricing by U.S.-Based Multinational Firms." Working paper, Yale School of Management.

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

Boehm, Christoph, Aaron Flaaen, and Nitya Pandalai Nayar. 2014. "The Role of Multinational Firms in Business Cycle Comovement." Working paper, University of Michigan.

Burstein, Ariel, Christopher Kurz, and Linda Tesar. 2008. "International Trade, Production Sharing and the Transmission of Business Cycles." *Journal of Monetary Economics* 44 (4):775–795.

Caselli, Francesco, Míklos Koren, Milan Lisicky, and Silvana Tenreyro. 2014. "Diversification through Trade." Working paper.

Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction.* Cambridge, Massachusetts: MIT Press.

di Giovanni, Julian and Andrei Levchenko. 2012. "Country Size, International Trade, and Aggregate Fluctuations in Granular Economies." *Journal of Political Economy* 120 (6):1083–1132.

di Giovanni, Julian, Andrei Levchenko, and Isabelle Méjean. 2012. "Firms, Destinations, and Aggregate Fluctuations." Working Paper 9168, CEPR Discussion Paper.

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

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

Frankel, Jeffrey A. and Andrew K. Rose. 1998. "The Endogeneity of the Optimum Currency Area Criteria." *Economic Journal* 108 (449):1009–1025.

Gabaix, Xavier. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica* 79:733–772.

Goldberg, Penelope, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova. 2010. "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India." *Quarterly Journal of Economics* 125 (4):1727–1767.

Halpern, László, Miklós Koren, and Adam Szeidl. 2011. "Imported Inputs and Productivity." Working Paper.

Heathcote, Jonathan and Fabrizio Perri. 2002. "Financial Autarky and International Business Cycles." *Journal of Monetary Economics* 49:601–627.

Hirschmann, A.O. 1958. "Interdependence and Industrialization." In *The Strategy of Economic Development*. New Haven: Yale University Press.

Imbs, Jean. 2004. "Trade, Finance, Specialization, and Synchronization." *Review of Economics and Statistics* 86 (3):723–734.

Imbs, Jean and Isabelle Méjean. 2011. "Elasticity Optimism." Working Paper 7177, CEPR Discussion Paper.

Jarmin, Ron and Javier Miranda. 2002. "The Longitudinal Business Database." Mimeo, available at http://www.vrdc.cornell.edu/info7470/2007/Readings/jarmin-miranda-2002.pdf.

Johnson, Robert C. 2014. "Trade in Intermediate Inputs and Business Cycle Comovement." *American Economic Journal: Macroeconomics* forthcoming.

Johnson, Robert C. and Guillermo Noguera. 2012. "Accounting for Intermediates: Production Sharing and Trade in Value Added." *Journal of International Economics* 86 (2):224–236.

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

Keller, Wolfgang and Stephen Yeaple. 2013. "The Gravity of Knowledge." *American Economic Review* 103 (4):1414–1444.

Kleinert, Jörn, Julien Martin, and Farid Toubal. 2012. "The Few Leading the Many: Foreign Affiliates and Business Cycle Comovement." Working Paper DP 9129, Center for Economic Performance.

Koren, Miklós and Silvana Tenreyro. 2013. "Technological Diversification." *American Economic Review* 103 (1):378–414.

Kremer, Michael. 1993. "The O-Ring Theory of Economic Development." *Quarterly Journal of Economics* 108 (3):551–575.

Kurz, Christopher Johann and Mine Zeynep Senses. 2013. "Importing, Exporting and Firm-Level Employment Volatility." Working Paper 2013-44, Board of Governors of the Federal Reserve System.

Leontief, Wassily. 1936. "Quantitative Input and Output Relations in the Economic Systems of the United States." *The Review of Economics and Statistics* 18 (3):105–125.

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

Ramey, Valerie A. 1989. "Inventories as Factors of Production and Economic Fluctuations." *American Economic Review* 79 (3):338–354.

Ramondo, Natalia, Veronica Rappoport, and Kim J. Ruhl. 2007. "Horizontal vs Vertical FDI: Revisiting Evidence from U.S. Multinationals." Working paper, Columbia Business School.

Ruhl, Kim J. 2008. "The International Elasticity Puzzle." Working paper, NYU Stern School of Business.

Schnell, Molly and David E. Weinstein. 2012. "Evaluating the Economic Response to Japan's Earthquake." Working Paper 12-P-003, RIETI.

Stern, Robert, Jonathan Francis, and Bruce Schumacher. 1976. "Price Elasticities in International Trade." In *The Strategy of Economic Development*. London: Macmillan.

Tornqvist, Leo, Pentti Vartia, and Yrjo Vartia. 1985. "How Should Relative Change Be Measured?" *American Statistician* 39 (1):43–46.

Wald, Davis J, Bruce C. Worden, Vincent Quitoriano, and Kris L. Pankow. 2006. "ShakeMap Manual: Technical Manual, User's Guide and Software Guide." Tech. rep., U.S. Geological Survey. URL `http://pubs.usgs.gov/tm/2005/12A01/`.

Walker, E. 1997. "The Census Bureaus Business Register: Basic Features and Quality Issues." Paper presented at the Joint Statistical Meetings, Anaheim, CA.

Table 2: Composition of Japanese Imports by Japanese Multinationals

	Japanese Multinationals
Share Intermediate	
	53.8
Share Related-Party	
in intermediate imports	85.7
in final imports	97.1

Source: CMF, LFTTD, DCA, and UBP as explained in the text. The data are for year 2007. This table reports the composition of Japanese imports from Japanese multinationals.

Table 3: 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

Table 4: Baseline Sample: Summary Statistics

	Japanese Firms	Other Multinationals
N.A. Exports	3,815,744	4,019,298
share intra-firm	73.0	52.2
Intermediate Input Imports	8,792,903	12,272,996
share from Japan	70.0	2.3
share intra-firm	86.0	28.8

Figure 1: Index of Japanese Industrial Production: Manufacturing Jan.2010 - Jan-2012

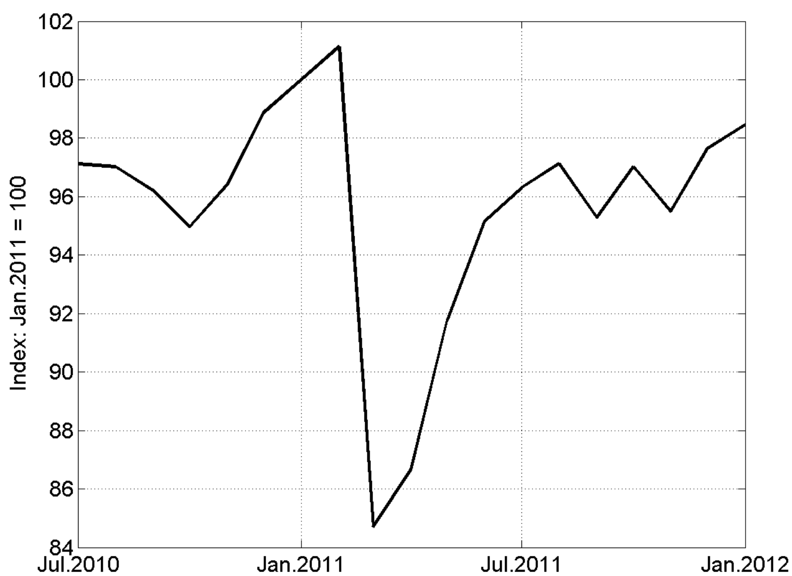

Source: Japanese Ministry of Economy, Trade, and Industry (METI). Series is Seasonally Adjusted.

Figure 2: U.S. Imports from Japan and Rest of World, Jan.2010 - Jan-2012

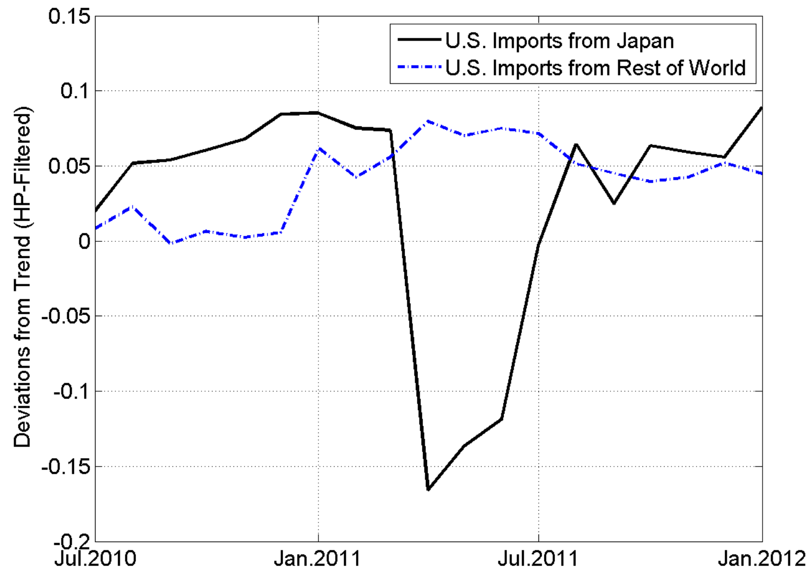

Source: U.S. Census Bureau, based on Published Totals. The series are logged, HP-Filtered, after seasonally adjusting.

Figure 3: U.S. Industrial Production: Manufacturing and Durable Goods

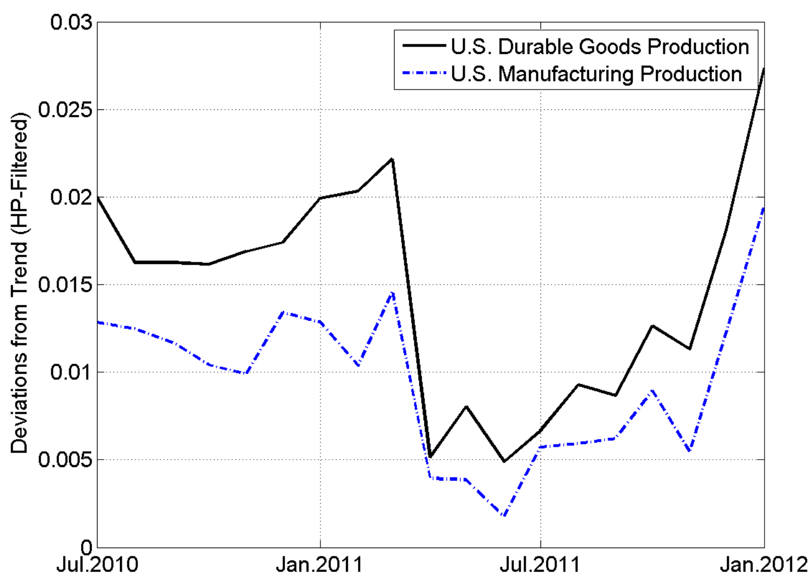

Source: Federal Reserve Board. Series is Seasonally Adjusted.

Figure 4: Geographic Distribution of Earthquake Intensity and Affiliate Locations

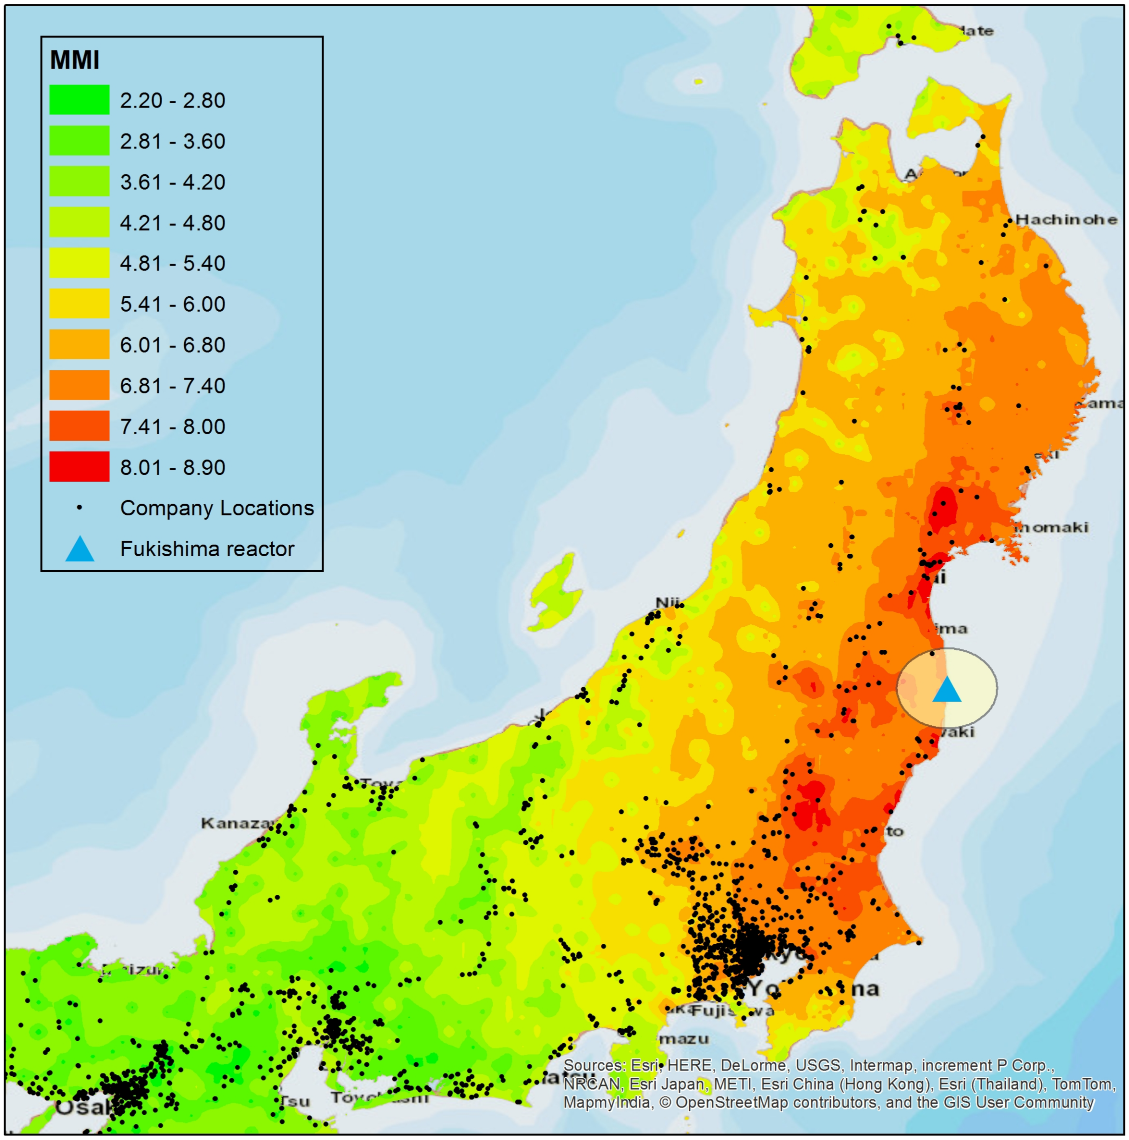

Source: USGS and DCA/Uniworld Directories
This figure plots the geographic distribution of the Tōhoku earthquake, based on recorded measurements taken directly after the event. The "Modified Mercalli Intensity" (MMI) scale is constructed based on a relation of survey response and measured peak acceleration 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.3.2.

Figure 5: Dynamic Treatment Effects: Japanese Firms

A. Relative Intermediate Input Imports of Japanese Firms

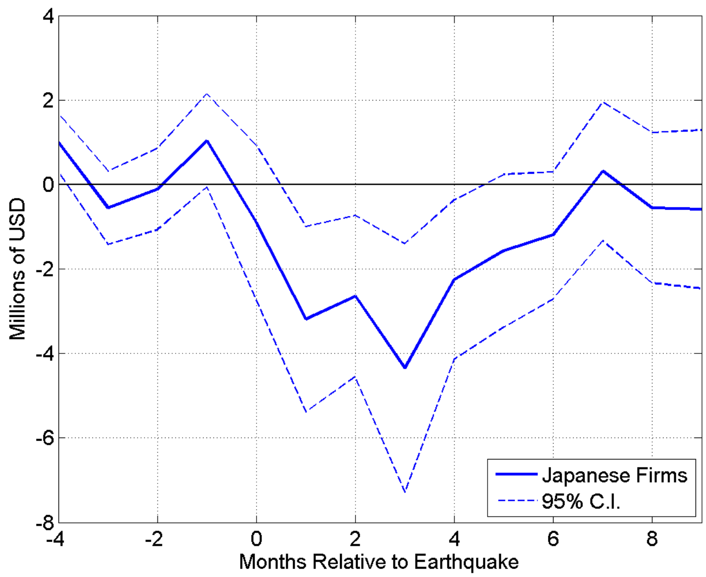

B. Relative North American Exports of Japanese Firms

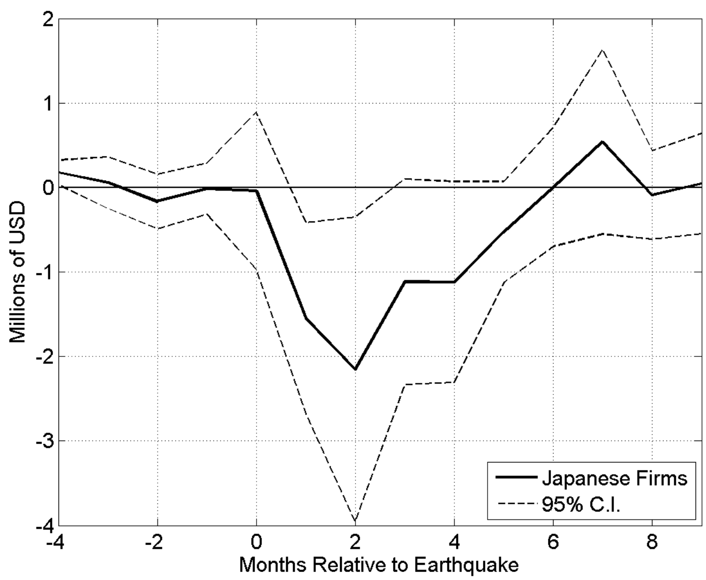

Source: LFTTD-DCA as explained in text.
These figures report the intermediate imports and North American exports of the U.S. affiliates of Japanese firms relative to a control group of other multinational firms. The values are coefficient estimates taken from an interaction of a Japanese-firm dummy with a monthly dummy – additional baseline monthly dummies remove seasonal effects. See equation 2 in the text. Standard errors are clustered at the firm level.

Figure 6: Relative Imported Inputs and Output (Proxy) of Japanese Firms: Fraction of Pre-Shock Level

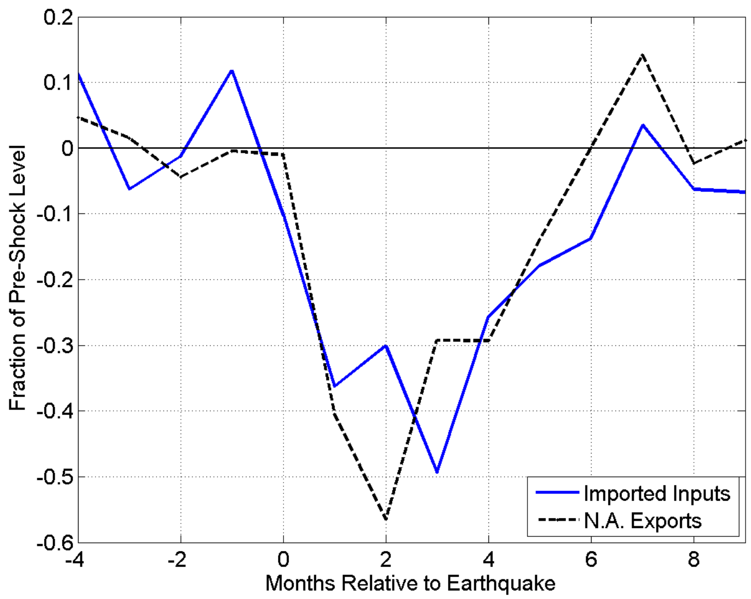

Source: LFTTD-DCA as explained in text.
This figure reports the intermediate imports and output proxy (North American exports) of the U.S. affiliates of Japanese firms relative to a control group of other multinational firms. The values are percent changes from the pre-shock level of each series, defined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

Figure 7: North America: Relative Production of Japanese Automotive Plants

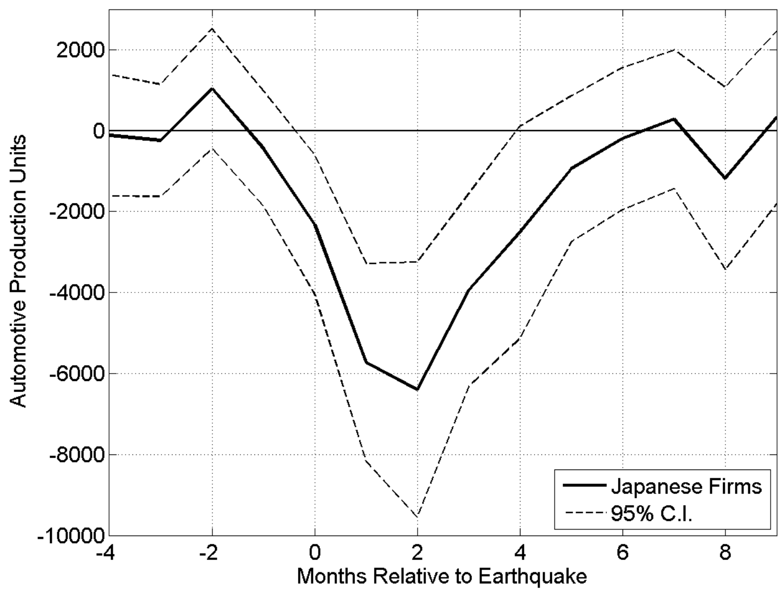

Source: Ward's Automotive Database
This figure reports the production levels of Japanese auto plants relative to a control group of non-Japanese auto plants. The values are coefficient estimates taken from an interaction of a Japanese-firm dummy with a monthly dummy – additional baseline monthly dummies remove any seasonal effects. See equation A8 in the text. Standard errors are clustered at the plant level. The Japanese automakers are Honda, Mazda, Mitsubishi, Nissan, Toyota, and Subaru. See section C.3 for more details.

A Mapping the Output Elasticity into the Elasticity of Substitution

Proof of Result 1
The first assumption allows us to equate the movements in dollar values that we observe in the data with movements in quantities. The second assumption allows us to separate out the ψ parameter from a product-level elasticity, the failure of which we evaluate in section 4.2.3. The final assumption implies that the firm receives less M^J than it would optimally require. We will observe $\frac{dlnx}{dlnM^J}$ in the data:

$$\frac{dlnx}{dlnM^J} = \frac{dlnx}{dM^J}\frac{dM^J}{dlnM^J} \tag{A1}$$

It is easy to show that $\frac{dM^J}{dlnM^J} = M^J$, and so we are left with $\frac{dlnx}{dM^J}$

$$\ln x = \frac{\psi}{\psi - 1} \ln\left(\left[(1-\mu)^{\frac{1}{\psi}} \left[F_D(\cdot)\right]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} \left[M^J\right]^{\frac{\psi-1}{\psi}}\right]\right) \tag{A2}$$

$$\frac{dlnx}{dM^J} = \frac{\mu (M^J)^{\frac{\psi-1}{\psi}-1}}{(1-\mu)^{\frac{1}{\psi}} \left[F_D(\cdot)\right]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} \left[M^J\right]^{\frac{\psi-1}{\psi}}} \tag{A3}$$

Combining we have the following expression:

$$\frac{dlnx}{dlnM^J} = \frac{\mu^{\frac{1}{\psi}} (M^J)^{\frac{\psi-1}{\psi}}}{(1-\mu)^{\frac{1}{\psi}} \left[F_D(\cdot)\right]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} \left[M_t^J\right]^{\frac{\psi-1}{\psi}}} \tag{A4}$$

To continue, we establish the sign of the derivative of equation (A4) with respect to ψ, and then check the bounds of the parameter space $(0, \infty)$ separately.

Output Elasticity for $\psi \in (0,\infty)$
We evaluate the following expression:

$$\frac{\partial \frac{dlnx}{dlnM^J}}{\partial \psi}$$

For ease of exposition let $\gamma \equiv \frac{M^J}{F_D(\cdot)}$. Then:

$$
\begin{aligned}
\frac{\partial \frac{dlnx}{dlnM^J}}{\partial \psi} &= \frac{\mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} \frac{1}{\psi^2} (\ln\gamma - \ln\mu) \left((1-\mu)^{\frac{1}{\psi}} + \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}}\right) - \left(-(1-\mu)^{\frac{1}{\psi}} \ln(1-\mu) + \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} (\ln\gamma - \ln\mu)\right) \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} \frac{1}{\psi^2}}{\left((1-\mu)^{\frac{1}{\psi}} + \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}}\right)^2} \\
&= \frac{\mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} \left[(\ln\gamma - \ln\mu) \left((1-\mu)^{\frac{1}{\psi}} + \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}}\right) - \left(-(1-\mu)^{\frac{1}{\psi}} \ln(1-\mu) + \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} (\ln\gamma - \ln\mu)\right)\right]}{\psi^2 \left((1-\mu)^{\frac{1}{\psi}} + \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}}\right)^2} \\
&= \frac{\mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} \left[(1-\mu)^{\frac{1}{\psi}} \ln\gamma - (1-\mu)^{\frac{1}{\psi}} \ln\mu + \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} \ln\gamma - \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} \ln\mu + (1-\mu)^{\frac{1}{\psi}} \ln(1-\mu) - \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} (\ln\gamma - ln\mu)\right]}{\psi^2 \left((1-\mu)^{\frac{1}{\psi}} + \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}}\right)^2} \\
&= \frac{\mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} \left[(1-\mu)^{\frac{1}{\psi}} ln\gamma - (1-\mu)^{\frac{1}{\psi}} \ln\mu + (1-\mu)^{\frac{1}{\psi}} \ln(1-\mu)\right]}{\psi^2 \left((1-\mu)^{\frac{1}{\psi}} + \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}}\right)^2} \\
&= \frac{\mu^{\frac{1}{\psi}} (1-\mu)^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}} \left[\ln\gamma - \ln\mu + ln(1-\mu)\right]}{\psi^2 \left((1-\mu)^{\frac{1}{\psi}} + \mu^{\frac{1}{\psi}} \gamma^{\frac{\psi-1}{\psi}}\right)^2}
\end{aligned}
$$

Since $\mu^{\frac{1}{\psi}} > 0$ and $(1-\mu)^{\frac{1}{\psi}} > 0 \ \forall \psi$, the sign of this expression is determined by the term in brackets within the numerator:

$$
\left(\ln\left(\frac{M^J}{F_D(\cdot)}\right) - \ln\mu + \ln(1-\mu)\right)
$$

According to assumption iii), we are in the case where $\frac{M^J}{F_D(\cdot)} < \frac{\mu}{1-\mu}$. For ease of exposition, we then re-write $\frac{M^J}{F_D(\cdot)} = x\frac{\mu}{1-\mu}$ where $x < 1$. Then:

$$
\left(\ln\left(\frac{M^J}{F_D(\cdot)}\right) - \ln\mu + \ln(1-\mu)\right) = \ln(x) < 0 \quad \text{since} \quad x < 1
$$

Thus, the output elasticity is decreasing in ψ provided that $\frac{M^J}{F_D(\cdot)} < \frac{\mu}{1-\mu}$. The magnitude of this correspondence depends on ψ itself as well as the deviation from the optimal input mix. For completeness, we now show the behavior of $\frac{dlnx}{dlnM^J}$ as $\psi \to 0$ and as $\psi \to \infty$

Output Elasticity as $\psi \to 0$

Re-organizing and taking the limits of equation (A4) as $\psi \to 0$:

$$
\lim_{\psi \to 0^+} \frac{dlnx}{dlnM^J} \quad = \quad \frac{\lim_{\psi\to 0^+} \mu^{\frac{1}{\psi}} \left(\frac{M^J}{F_D(\cdot)}\right)^{\frac{\psi-1}{\psi}}}{\lim_{\psi\to 0^+} (1-\mu)^{\frac{1}{\psi}} + \lim_{\psi\to 0^+} \mu^{\frac{1}{\psi}} \left[\frac{M^J}{F_D(\cdot)}\right]^{\frac{\psi-1}{\psi}}}
$$

Once again we rely on the assumption that $\frac{M^J}{F_D(\cdot)} < \frac{\mu}{1-\mu}$. Re-writing $\frac{M^J}{F_D(\cdot)} = x\frac{\mu}{1-\mu}$ where $x < 1$ we have:

$$\lim_{\psi \to 0^+} \frac{dlnx}{dlnM^J} = \frac{\lim_{\psi \to 0^+} \mu x^{\frac{\psi-1}{\psi}}}{\lim_{\psi \to 0^+}(1-\mu) + \lim_{\psi \to 0^+} \mu x^{\frac{\psi-1}{\psi}}} = \frac{\infty}{(1-\mu)+\infty} = 1$$

Output Elasticity as $\psi \to \infty$
Intuitively we know that this approximates the relationship of perfect substitutes, and the two inputs can be used interchangeably.

$$\lim_{\psi \to \infty} \frac{dlnx}{dlnM^J} = \frac{\lim_{\psi \to \infty} \mu^{\frac{1}{\psi}} \left(\frac{M^J}{F_D(\cdot)}\right)^{\frac{\psi-1}{\psi}}}{\lim_{\psi \to \infty}(1-\mu)^{\frac{1}{\psi}} + \lim_{\psi \to \infty} \mu^{\frac{1}{\psi}} \left[\frac{M^J}{F_D(\cdot)}\right]^{\frac{\psi-1}{\psi}}}$$

Because $\mu < 1$, then $\lim_{\psi \to \infty} \mu^{\frac{1}{\psi}} = 1$ and $\lim_{\psi \to \infty}(1-\mu)^{\frac{1}{\psi}} = 1$. And because $\lim_{\psi \to \infty} \frac{\psi-1}{\psi} = 1$ we know $\lim_{\psi \to \infty} \left(\frac{M^J}{F_D(\cdot)}\right)^{\frac{\psi-1}{\psi}} = \frac{M^J}{F_D(\cdot)}$. So, putting this together we have:

$$\lim_{\psi \to \infty} \frac{dlnx}{dlnM^J} = \frac{\frac{M^J}{F_D(\cdot)}}{1 + \frac{M^J}{F_D(\cdot)}} = \frac{M^J}{F_D(\cdot) + M^J}$$

This analysis makes sense. If there is a suboptimal mix of the two inputs under a (nearly) Leontief relationship (as $\psi \to 0$), the non-scarce input has essentially no impact on output. On the other hand, movements in the scarce input will transfer one for one to output. In addition, one can see that the effect is increasing in the distance from μ, the optimal allocation between $F_D(\cdot)$ and M^J.

The case when $\psi \to \infty$ is similarly intuitive: Beginning from a steady state where there exists an optimal allocation between the two inputs, then the effect on output will simply be μ, the share of the imported input. Here the effect will depend on the relative proportions between the two inputs, whereas on the other side of the parameter space of ψ, the effect depended simply on the relative scarcity.

B Data Appendix

B.1 Matching Corporate Directories to the Business Register

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

B.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 information.

B.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 the program *reclink* created by Michael Blasnik. This program uses a bigram string comparator algorithm on multiple variables with differing user-specified weights.[38] 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 bigram character matches. Further information on this program is available in Flaaen (2013a) or Blasnik (2010).

B.1.3 The Unit of Matching

The primary unit of observation in both the DCA and BR datasets is the business establishment. Hence, the primary unit of matching for this bridge will be the establishment, and not the firm. However, there are a number of important challenges with an establishment-to-establishment link. First, the DCA 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 bridge will be to assign each DCA 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.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 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.[39]

This is an iterative process, in which a series of matching procedures are applied with decreasingly restrictive sets of matching requirements. In other words, the initial matching attempt uses the most stringent standards possible, after which the non-matching records proceed to a further matching iteration, often with less stringent standards. In each iteration, the matching records are assigned a flag that indicates the standard associated with the match.

See Table A1 for a summary of the establishment-level match rate statistics by year and type of firm.

[38] 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 assigns a score for each variable between the two datasets based on the percentage of matching bigrams.

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

B.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 hte 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.

There can be a number of issues when translating these 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 DCA-based multinational indicators and aggregating to the (Census-based) firm level. There are three potential cases:

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. 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 Classifying Firm-Level Trade

The firm-level data on imports available in the LFTTD does not, unfortunately, contain any information on the intended use of the goods.[40] 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.[41] 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.

B.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.

[40] 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 **?**.

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

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.3.2 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.[42] First, within a given CM year, the classification can be done at the firm or establishment level rather than aggregating to a particular industry. (Hence, we could build S_{pi} rather than S_{pj}.) 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

[42] 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.

using the industry-level product classifications calculated above.

B.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. [43] While 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:

$$\left.\begin{aligned} M_{ij}^{\text{int}} &= \sum_{p \notin P_j} M_{ipj} \\ M_{ij}^{\text{fin}} &= \sum_{p \in P_j} M_{ipj} \end{aligned}\right\} \qquad \text{where} \qquad P_j = \{p \mid S_{pj} \geq W\}\,. \tag{A5}$$

The bottom section of Table B.3.2 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 1 (about 3 percentage points), the number is essentially unchanged when raising it further to 0.2.

B.3 Sample Selection

B.3.1 Constructing the Baseline Dataset

This section will discuss the steps taken to construct the sample used in section 2.2.2.

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

[43] 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 the that firm.

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 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. Finally, the firms' monthly trade can 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 that is 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.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.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.1.

Table A1: Match Statistics: 2007-2011

	# of 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	Coming Soon...		
2011	Coming Soon...		
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	Coming Soon...		
2011	Coming Soon...		
Foreign Multinationals			
2007	10,331	7,555	0.73
2008	9,351	6,880	0.74
2009	11,142	8,193	0.74
2010	Coming Soon...		
2011	Coming Soon...		

Table A2: 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

C Appendix: Other Results

C.1 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.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 2.2.2, we run the following regression:

$$\Delta \text{emp}_{j,t} = \sum_{i=-3}^{3} \gamma_i E_i + \sum_{i=-3}^{3} \beta_i E_i D_{j,i} + u_{j,t} \tag{A6}$$

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 A6 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 C.3.

C.2 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.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 ($\mathrm{m}_{p,j,t} = log(\mathrm{M}_{p,j,t})$:

$$\mathrm{m}_{p,j,t} = \alpha_{pj} + \sum_{i=-19}^{9} \gamma_i E_i + \sum_{i=-19}^{9} \beta_i E_i D_{j,i} + u_{j,t} \tag{A7}$$

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 C.3.

C.3 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 2.2.2:

$$Q_{jt} = \alpha_j + \lambda_j t + \sum_{i=-4}^{9} \gamma_i E_i + \sum_{i=-4}^{9} \beta_i E_i D_{j,i} + u_{j,t} \tag{A8}$$

where here the variable $Q_{j,t}$ is auto production by plant j in month t. 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 2.2.3. 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 C.3.

Table A3: 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 A6 in the text. The dependent variable is the four-quarter log difference of employment (payroll).

Table A4: 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 A7 in the text.

Figure A1: Automotive Production, Inventory, Sales by Firm Type, Distributed Lag Model

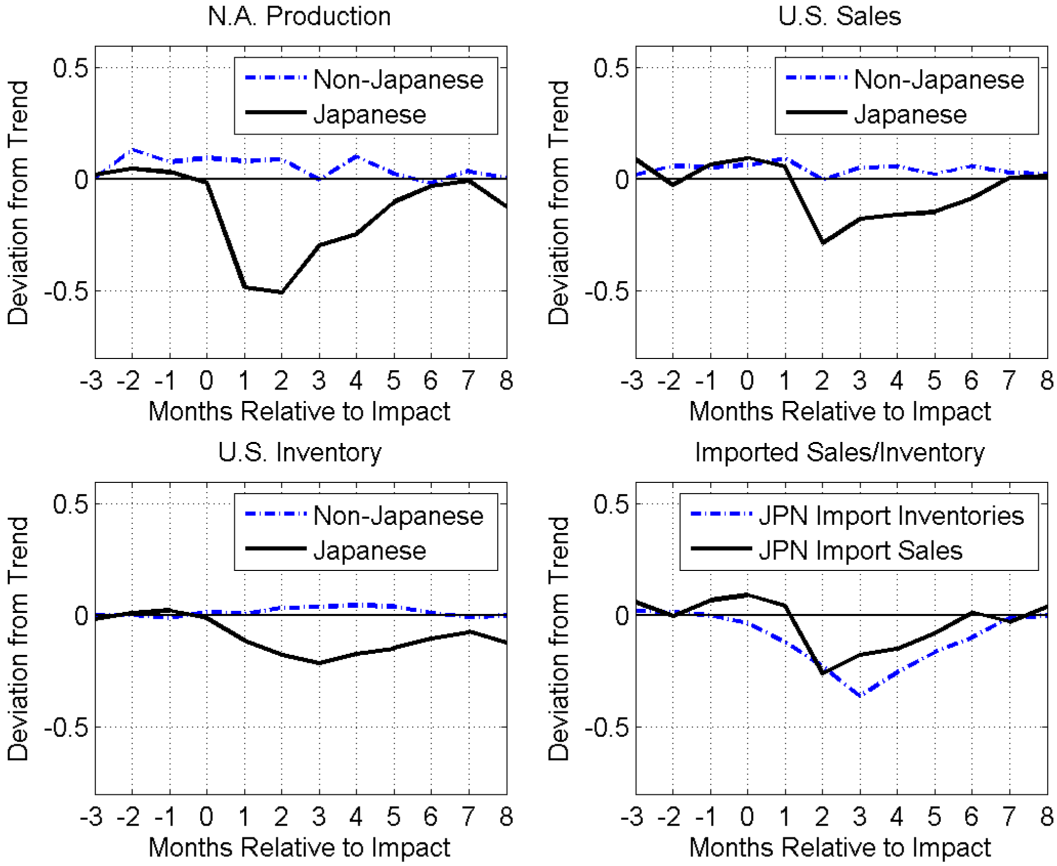

Source: Ward's Automotive Database
This figure reports North American production, and U.S. sales and inventory data according to firm type: Japanese and non-Japanese firms. The values are coefficient estimates taken from a distributed lag model, exploiting time-series variation only. The underlying series 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 A5: 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 .

D Model Appendix

Coming soon ...