

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

Christoph Boehm

Aaron Flaaen

Nitya Pandalai-Nayar

February 27, 2015

PRELIMINARY

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 suffered large drops in U.S. output in the months following the shock, roughly one-for-one with the drop in imports and consistent with a Leontief relationship between imported and domestic inputs. Structural estimates of the production function for these firms yield disaggregated production elasticities that are similarly low. Our results suggest that global supply chains are sufficiently rigid to play an important role in the cross-country transmission of shocks.

JEL Codes: F10, F23, F44, E32

Keywords: Multinational Firms, International Business Cycles, 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 thank Ruediger Bachmann, Javier Cravino, Chris House, Robert Johnson, Lutz Kilian, Illenin Kondo, Chris Kurz, Margaret Levenstein, Logan Lewis, William Lincoln, Andrew McCallum, Ryan Monarch, Justin Pierce and Jeff Smith for helpful conversations. This research has received research grants from the Michigan Institute for Teaching and Research in Economics (MITRE), for which we are very thankful. Support for this research at the Michigan RDC from NSF (ITR-0427889) is also gratefully acknowledged. 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

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

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

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

respect to domestic factors of production.

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

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

The structural estimates are broadly in agreement with the results from our reduced form exercise. For Japanese multinationals, the elasticity of substitution across material inputs is 0.2 and the elasticity between material inputs and a capital/labor aggregate is 0.03. For non-Japanese firms using inputs from Japan, the elasticity of substitution across material inputs is somewhat higher at 0.6. While the high cost share and particularly low elasticity for Japanese affiliates explains their predominant contribution to the direct transmission of this shock to the U.S., the elasticity estimates for non-Japanese firms are still substantially lower than typical estimates used in the literature. We argue that the substantial share of intra-firm intermediate trade implies greater complementarities in aggregate trade than is currently recognized.

There are a number of important implications for such low values of the elasticity of substitution. This parameter appears in various forms in a wide span of models involving the exchange of goods across countries. As discussed by Backus, 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 uncertainty of available estimates for the elasticity of substitution, it is a common practice to evaluate the behavior of these models along a wide range of parameter values.

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

To understand the relationship between the elasticity and comovement, it is helpful to recall that these models generate output comovement by inducing synchronization in factor supplies, a mechanism that by itself generally fails to produce the degree of comovement seen in the data. Complementarities among inputs together with heterogeneous input shocks will generate direct comovement in production, augmenting the output synchronization based on factor movements. Burstein, 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.³ 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, page 621): “greater complementarity is associated with a larger return to relative scarcity.”

The estimates in this paper have implications for the role of trade in firm-level and aggregate volatility. Other research has argued that firms can diversify risk arising from country specific shocks by importing (Caselli et al. (2014)) or that firms with complex production processes of several inputs are less volatile as each input matters less for production (Koren and Tenreyro

¹For 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.

²Due to the robust nature of these shortcomings, Backus, Kehoe, and Kydland (1995) refer to them as the “price anomaly” and “quantity anomaly” respectively.

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

(2013)). On the other hand, there is a well-established fact that complementarities and multi-stage processing can lead to the amplification of shocks as in Jones (2011) and Kremer (1993). We discuss the potential for measured amplification in our context in Section 5.

This paper is also a contribution to the empirical evidence on the role of individual firms in aggregate fluctuations, emanating from the work of Gabaix (2011). Other related evidence comes from di Giovanni, Levchenko, and Méjean (2014), who use French micro-data to demonstrate that firm-level shocks contribute as much to aggregate volatility as sectoral and macroeconomic shocks combined. The so-called granularity of the economy is very much evident in our exercise; though the number of Japanese multinationals is small, they comprise a very large share of total imports from Japan, and are arguably responsible for a measurable drop in U.S. industrial production following the Tōhoku earthquake (see Figure 3).

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

As is the case with most research based on an event-study, some care should be taken in generalizing the results to other settings. Although we have already highlighted the aggregate implications of the effects we estimate, one might worry that the composition of Japanese trade or firms engaged in such trade is not representative of trade linkages more broadly. We believe the results we obtain are informative beyond the context of this particular episode for two reasons. First, the features of Japanese multinationals that are underlying the transmission of this shock are common to all foreign multinational affiliates in the United States.⁴ Second, estimates of a low elasticity of substitution should not be particularly surprising: they correspond to firms engaged in input trade that is largely intra-firm trade embodying firm-specific knowledge and capabilities, such that one would not expect substitution with other inputs or suppliers to be easy. This type

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

of trade accounts for a substantial share of total trade, particularly among advanced economies.

The next section describes the empirical strategy and data sources used in this paper and presents reduced form evidence in favor of a low production elasticity of imported inputs for Japanese multinational affiliates. In Section 4, we expand the scope of parameters we identify with a structural model of cross-country production linkages. We estimate the parameters of this model and discuss their implications. Section 5 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 using an event-study framework surrounding the 2011 Tōhoku event to estimate the production elasticity of imported inputs. We discuss the relevant details of this shock, document the aggregate effects, and then outline the empirical specification for the firm-level analysis.

2.1 Background

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

Figure 1 shows the considerable impact of the Tōhoku event on the Japanese economy. Japanese manufacturing production fell by roughly 15 percentage points between February and March 2011, and did not return to trend levels until July. Much of the decline in economic activity resulted

⁵Since 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).

from significant power outages that persisted for months following damage to several power plants – most notably the Fukushima nuclear reactor.⁶ Further, at least six Japanese ports (among them the Hachinohe, Sendai, Ishinomaki and Onahama) sustained significant damage and were out of operation for more than a month, delaying shipments to both foreign and domestic locations. It should be noted, however, that the largest Japanese ports (Yokohama, Tokyo, Kobe) which account for the considerable majority of Japanese 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 to that of Japanese manufacturing production: a 20 percentage point drop from March to April, with a full recovery by July 2011.

More striking is the response of U.S. industrial production in the months following the event. Figure 3 demonstrates that there is indeed a drop in U.S. manufacturing production in the months following the Japanese earthquake. Although the magnitudes are obviously much smaller — roughly a one percentage point drop in total manufacturing and almost two percentage points in durable goods — the existence of a measurable effect is clear.⁷

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

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

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

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

2.2 Data

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

An important component 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 the destination (or source) country, quantity and value shipped, the transport mode, and other details from point-of-trade administrative documents. Importantly for this study, the LFTTD includes import and export trade transactions at a *daily* frequency, which is easily aggregated to monthly-level trade flows. A number of important papers have utilized this resource, such as Bernard et al. (2007) and Bernard, 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 information, for the first time, to identify those U.S. affiliates part of a foreign parent company, as well as those U.S. firms with affiliate operations abroad. This information is 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 Appendix B.1.

The second novel data resource 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)), the LFTTD does not classify a trade transaction based on its intended use. To overcome this limitation, we use information on the products produced by U.S. establishments in a given industry to identify the set of products intended for final sale *for that industry*.⁹ The remaining products are presumably used by establishments in that industry

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

either as intermediate inputs or as capital investment. Details on this classification procedure are available in 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. Geological Survey (USGS). By geocoding the Japanese addresses of firms with U.S. operations, we construct an earthquake intensity measure for each Japanese affiliate location. We then apply such information to the U.S. operations as a way to further measure the sample of firms plausibly affected by the shock. Please see Appendix B.3.2 for details. Figure 4 shows the geographic distribution of one such USGS measure — the modified mercalli index (MMI) — along with the geocoded affiliate locations.

The ideal dataset to evaluate the transmission of the Tōhoku event on U.S. firms would consist of high frequency information on production, material inputs, and trade, separated out by geographic and ownership criteria. Unfortunately, Census data on production and material inputs at the firm-level is somewhat limited. The Annual Survey of Manufacturers (ASM) contains such information, but at an annual frequency and only for a subset of manufacturing firms. On the other hand, firm-level trade information is available at a nearly daily frequency, and covers the universe of firms engaged in exporting/importing. For the purposes of characterizing the shock to firm-level imports of intermediate goods, the LFTTD (and supplements identified above) is ideal. There remain significant gaps in information on a firm’s domestic input usage, a limitation we discuss in subsequent sections.

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

quality of this measure as a proxy for output in section 5.2.1.¹⁰

2.3 Basic Theory

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

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

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

The first order conditions imply that

$$\frac{F_D^*}{IM^*} = \frac{1 - \mu}{\mu} \left(\frac{p_M}{p_D} \right)^{\psi}, \quad (2)$$

where F_D^* and IM^* denote the optimal quantities of inputs. We would like to show the theoretical foundations underlying the intuitive result that a one-for-one drop in output with the fall in imported inputs implies an elasticity of zero. To do this, we make the following assumptions, all of which we will relax to some degree in the estimation framework in Section 4:

¹⁰Another consideration with the use of this proxy is whether it more accurately reflects production or sales, as the two are distinct in the presence of output inventories. In our case, this depends on whether the inventories are held in the U.S. or Canada/Mexico. Without further evidence, we interpret the proxy to be capturing some mix between production and sales. The structural estimation in section 4 will allow for such a mix.

1. Imported inputs shipments are disrupted, such that the firm receives a suboptimally low quantity of IM : $IM < IM^*$;
2. The firm is unable to adjust domestic inputs F_D^* or its price p_x after learning that it receives IM ;
3. The firm does not shut down.

Given these assumptions, the following result holds:

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

$$\frac{\partial \ln p_x x}{\partial \ln p_M IM} = \frac{1}{1 + \left(\frac{IM^*}{IM}\right)^{\frac{\psi-1}{\psi}} \left(\frac{1-\mu}{\mu}\right) \left(\frac{p_M}{p_D}\right)^{\psi-1}} \in (0, 1) \quad (3)$$

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

Proof. See Appendix A.1 for details. □

An immediate implication of this result is that the output elasticity is unity only when ψ approaches zero.¹¹ In this case $\left(\frac{IM^*}{IM}\right)^{\frac{\psi-1}{\psi}} \rightarrow 0$ (recall that $IM < IM^*$) and hence $\lim_{\psi \rightarrow 0} \frac{\partial \ln p_x x}{\partial \ln p_M IM} = 1$. Hence, observing a one-for-one drop in the value of output with the value of imported intermediates, we infer that ψ is close to zero. It is also straightforward to show that conditional on a value for $\psi \in (0, \infty)$, the output elasticity in (3) is increasing in the parameter μ . That is, conditional on a given drop in the imported input, a larger weight on this input leads to a larger percent response in output.

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

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

$$\frac{\partial \ln(p_x x)}{\partial \ln(p_M IM)} = \frac{\partial \ln(p_x x)}{\partial \ln(p_D F_D)} = \frac{\partial \ln(p_D F_D)}{\partial \ln(p_M IM)} = 1. \quad (4)$$

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

3 Reduced Form Evidence

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

3.1 Framework

Our analysis of the production function (1) above demonstrates that a natural measure to evaluate the potential conduits of the Tōhoku shock to the United States would be the degree of reliance on Japanese imported inputs. This is best expressed as the cost share of inputs from Japan, and can be constructed in a Census year by taking a firm’s Japanese imported inputs and dividing by all other inputs (which includes production worker wages and salaries, the cost of materials, and the cost of new machinery expenditures). Exposure to Japanese imported inputs is heavily concentrated among Japanese affiliates. In the year 2007, which is the closest available Census year, this cost share was nearly 22% on average for Japanese affiliates (see Table 1), compared to just 1% for other firms. For more detail on the heterogeneity across and within these firm groups, we construct a density estimate of such an exposure measure for the Japanese affiliates and non-Japanese multinationals. The results, shown in Figure 5, show little overlap between these distributions: there are few Japanese affiliates with low exposure to Japanese inputs, and few

non-Japanese firms have substantial exposure.¹²

We now estimate the relative impacts on imported inputs and output for the Japanese affiliates as a group. To do this, we implement a dynamic treatment effects specification in which a firm is defined as being treated if it is owned by a Japanese parent company.¹³ The effect on these firms can be inferred from the differential impact of the variable of interest relative to a control group, which soaks up common seasonal patterns and other demand-driven factors in the U.S. market. While there are a number of competing methodologies for this type of estimation, we use normalized propensity score re-weighting due to the relatively favorable finite-sample properties as discussed in Busso, DiNardo, and McCrary (2014), as well as for its transparent intuition. Consistent estimation of the average treatment effect on the treated requires the assumption of conditional independence: the treatment/control allocation is independent of potential outcomes conditional on a set of variables. As the average Japanese firm differs considerably from other firms in the data, we use other multinational firms – both US and non-Japanese foreign- as our baseline control group prior to reweighting. To compute the propensity scores for reweighting, we control for size and industry, which ensures the control group has a similar industrial composition and size distribution as our treated sample.¹⁴ Table 2 reports summary values for the sample, including statistics on the balancing procedure using the normalized propensity score.

The magnitude of the shock for a representative Japanese multinational is captured by the effect on total imported intermediate products at a monthly frequency.¹⁵ Including non-Japanese imported intermediates is important for applying the control group as a counterfactual, and the shares by source-country gives the necessary variation for identification: as shown in Table 2 the share of imported inputs from Japan is 70% of the total for Japanese firms and only 3.5% for non-Japanese multinationals. Let $V_{i,t}^M$ be the value of intermediate imports of firm i in month t , after removing a firm-specific linear trend through March 2011. We fit the following regression:

¹²The exposure measure used in Figure 5 is from 2010 and does not include the cost of domestic material usage.

¹³We could have also used a threshold of Japanese input usage for the classification of treatment status. Doing so yields estimates that are very similar, which is due to the patterns evident in Figure 5. We have also tried conditioning on our geographic information (i.e. the firm-level Japanese MMI index) in defining a Japanese firm as being treated. The results are largely unchanged from those we report here, and for the sake of clarity we report results pertaining to the full sample.

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

¹⁵We consider Japanese and non-Japanese intermediate imports separately in section 4.

$$V_{i,t}^M = \alpha_i + \sum_{p=-4}^9 \gamma_p E_p + \sum_{p=-4}^9 \beta_p E_p \text{JPN}_{i,p} + u_{i,t} \quad (5)$$

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

The β_p coefficients are of primary interest. The $\text{JPN}_{i,t}$ is an indicator variable equal to one if the firm is owned by a Japanese parent company. Interacting these indicator variables with each month of the panel allows for a time-varying effect of Japanese ownership on a firm's overall intermediate input imports, particularly during and after the Tōhoku event. The β_p coefficients will estimate the differential effect of the Tōhoku event on Japanese multinational affiliates in the U.S., compared to the control group of non-Japanese firms. To evaluate the differential impact on production for Japanese firms, we simply replace the dependent variable in equation (5) with the firm's North American exports, denoted $V_{i,t}^{NA}$.

It is important to highlight that the specification in (5) is in levels. There are several reasons for doing so, as opposed to using log differences or growth rates. First, allowing for the presence of zeros is important when the data are at a monthly frequency, particularly given the magnitude of the shock to imports for Japanese firms. The second reason is more conceptual. Because we are interested in calculating the average effect of these firms that represents (and can scale up to) the aggregate impact on the U.S. economy, it is appropriate to weight the firms based on their relative size. The levels specification does exactly this: the absolute deviations from trend will be greater for the bigger firms and hence will contribute disproportionately to the coefficient estimates.¹⁶ In section 4, we evaluate this framework with the results one would obtain when estimating the effect on a firm-by-firm basis.

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

¹⁶See Appendix C.1 for more discussion, as well as results obtained using other specifications. Importantly, in a reduced sample abstracting from zeros, a weighted regression using percentage changes directly yields estimates that are very close to those presented here.

shock is unlikely to have a measurable effect on imported inputs as a whole. We discuss strategic interaction in section 5.2.4.

3.2 Results: Total Manufacturing Sector

The top panel of Figure 6 plots the β_p coefficients from equation (5) for the months surrounding the Tōhoku event. Relative to the control group, there is a large drop in total intermediate input imports by Japanese firms in the months following the earthquake. The drop in intermediate inputs bottoms out at 4 million USD in $t = 3$ (June 2011) and the 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 6, which looks for evidence of the production/sales impact of this shock on Japanese firms via their North American exports. The differential time-path of N.A. exports also exhibits a substantial drop following the Tōhoku event, hitting a trough of 2 million USD below baseline in $t = 2$ (May 2011). The standard errors, which are clustered at the firm level, are themselves interesting. As made clear via the 95-percent confidence bands on the point estimates of Figure 6, the standard errors increase dramatically in the months following the shock, a feature we interpret to reflect heterogeneous incidence and timing of the shocks (as well as the recoveries) for the Japanese multinationals.

To gain a sense of the average percentage drops of these two data series for Japanese multinationals as a group, we take the two plots of the differential dollar amounts from Figure 6 and divide by the average pre-shock level for these firms (see Table 2). The results, plotted jointly in Figure 7, show the fraction below pre-shock trend levels for these firms, on average. There is a remarkable correlation between these two series – whereby there is essentially a one-for-one drop in output for a given drop in intermediate imports. Using the mapping from Result 1, these reduced form results suggest a production function that is essentially Leontief in the imported input.

One potential concern with the interpretation of these results is separating out the intermediate input channel with other channels, such as a direct “productivity shock” affecting the U.S. operations of Japanese affiliates. Separating an ownership channel from an imported input channel is difficult due to lack of substantial overlap we identified above: few Japanese firms have low input exposure and few non-Japanese firms have high input exposure. Nevertheless, to get a sense of the defining characteristics, we construct a simple probit model of a disruption in Japanese inputs and

U.S. output:

$$Pr(X_{ik}^D = 1) = \Phi [\beta_1 JPN_{ik} + \beta_2 Exposed_{ik} + \beta_3 MMI_{ik} + \beta_4 Port_{ik} + \gamma_k] \quad (6)$$

where the dependent variable (X_{ik}^D) is an indicator equal to one if the N.A. exports of firm i in industry k are on average 20% below trend during the five months following the Tōhoku event. The independent variables are also indicators: JPN_{ik} , for affiliates of Japanese multinationals; $Exposed_{ik}$, for firms with an exposure to Japanese inputs above 0.05 of total material; MMI_{ik} for firms with an elevated MMI value pertaining to their average Japanese manufacturing locations; and $Port_{ik}$ for firms that typically rely on imports via ports damaged by the tsunami.¹⁷ The γ_k term allows for industry-specific intercepts. To evaluate the determinants of an input disruption from Japan, we replace the dependent variable with J_{ik}^D , an indicator for a drop in Japanese imported inputs of 20% relative to trend.

Panel A of Table 3 evaluates firm characteristics predicting a drop in U.S. output (X_{ik}^D), as measured by our proxy. The columns (1)-(4) show the results from different specifications with various combinations of the covariates in equation (6). Both Japanese ownership and high exposure to Japanese inputs significantly increase the probability of an output disruption, as expected. In columns (3) and (4), we demonstrate that Japanese ownership is substantially more indicative of an output decline than high input exposure alone. In Panel B, we replace the dependent variable with the binary measure of a drop in Japanese intermediate inputs (J_i^D). The results from these regressions indicate, unsurprisingly, that high exposure to Japanese imports are highly predictive of a subsequent disruption following the Tōhoku event. Apart from their exposure to imports from Japan, the Japanese affiliates are no more likely to suffer a disruption to these imports (see column 8).¹⁸ While the results from Table 3 are somewhat inconclusive, they nevertheless point to unique features of the production function of Japanese affiliates that yields direct pass-through of Japanese shocks to the U.S. economy. Our estimation procedure that follows should help to clarify this point further.

¹⁷Specifically, the $MMI_{ik} = 1$ if the average Japanese manufacturing establishment corresponding to a U.S. firm is above the median (roughly an MMI of 5.2) of all firms with Japanese manufacturing locations. The affected ports are: Onahama, Hitachi, Kashima, Haramachi, Shiogama, Sendai, Shimizu, Ishinomaki, Hashinohe, Miya Ko, Kamaishi, Ofunato, and Kessennuma.

¹⁸The combined effect of the coefficients on Japan and $JPN*Exp$ is -0.16, and not significant.

4 Structural Estimation of Cross Country Input Linkages

The relative movements of imported inputs and output of Japanese multinational firms point to little substitutability of Japanese intermediate inputs. In this section we expand our analysis by structurally estimating the production function of firms affected by the Tōhoku shock. We do this for three reasons. First, it is reassuring to find elasticities that are consistent with the heuristic evidence implied by our reduced-form results, when imposing a conventional production function framework. Second, by imposing additional structure, we are able to distinguish two elasticities: one between Japanese material inputs and other material inputs, and another between an aggregate bundle of material inputs and domestic capital and labor. Finally, by using an estimation procedure not relying on a control group we can obtain separate estimates for Japanese and non-Japanese firms. The results corroborate the claim that the supply chains of Japanese and non-Japanese exhibit different degrees of rigidity.

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

4.1 Framework

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

$$x_{i,t} = \phi_i \left[\mu_i^{\frac{1}{\zeta}} \left(K_{i,t}^\alpha L_{i,t}^{1-\alpha} \right)^{\frac{\zeta-1}{\zeta}} + (1 - \mu_i)^{\frac{1}{\zeta}} M_{i,t}^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}}, \quad (7)$$

where

$$M_{i,t} = \left(\nu_i^{\frac{1}{\omega}} \left(m_{i,t}^{-J} \right)^{\frac{\omega-1}{\omega}} + (1 - \nu_i)^{\frac{1}{\omega}} \left(m_{i,t}^J \right)^{\frac{\omega-1}{\omega}} \right)^{\frac{\omega}{\omega-1}}. \quad (8)$$

In this production function $x_{i,t}$, $K_{i,t}$, and $L_{i,t}$ denote the output, capital, and labor of firm i . The variable $M_{i,t}$ denotes an aggregate of intermediate inputs consisting of materials sourced from Japan ($m_{i,t}^J$) and materials sourced from all places other than Japan ($m_{i,t}^{-J}$), including domestic materials. We are interested in estimating ω and ζ , which parameterize the substitutability between Japanese

and non-Japanese materials and that between the capital-labor aggregate and the aggregate of intermediate inputs. The parameters μ_i and ν_i are firm-specific weights and ϕ_i parameterizes the firm's productivity, all of which we assume are constant over the short time horizon we consider. Further, we assume that the firm is monopolistically competitive and faces a CES demand function

$$p_{i,t}^x = \left(\frac{Y_{i,t}}{x_{i,t}} \right)^{\frac{1}{\varepsilon}}. \quad (9)$$

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

4.1.1 Pre-Tsunami period

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

$$p_{i,\tau-1}^x x_{i,\tau-1} - w_{\tau-1} L_{i,\tau-1} - R_{\tau-1} K_{i,\tau-1} - p_{i,\tau-1}^{-J} m_{i,\tau-1}^{-J} - p_{i,\tau-1}^J m_{i,\tau-1}^J$$

subject to (7), (8), and (9). The firm takes all factor prices as given. Material prices $p_{i,\tau-1}^J$ and $p_{i,\tau-1}^{-J}$ are firm-specific to indicate that different firms use different materials.

It is straightforward to show that this optimization problem implies

$$K_{i,\tau-1} = \frac{\alpha}{1-\alpha} \frac{w_{\tau-1} L_{i,\tau-1}}{R_{\tau-1}}, \quad (10)$$

$$\nu_i = \frac{\left(p_{i,\tau-1}^{-J} \right)^\omega m_{i,\tau-1}^{-J}}{\left(p_{i,\tau-1}^J \right)^\omega m_{i,\tau-1}^J + \left(p_{i,\tau-1}^{-J} \right)^\omega m_{i,\tau-1}^{-J}}, \quad (11)$$

$$\mu_i = \frac{\left(\left(\frac{R_{\tau-1}}{\alpha} \right)^\alpha \left(\frac{w_{\tau-1}}{1-\alpha} \right)^{1-\alpha} \right)^\zeta K_{i,\tau-1}^\alpha L_{i,\tau-1}^{1-\alpha}}{\left(P_{i,\tau-1}^M \right)^\zeta M_{i,\tau-1} + \left(\left(\frac{R_{\tau-1}}{\alpha} \right)^\alpha \left(\frac{w_{\tau-1}}{1-\alpha} \right)^{1-\alpha} \right)^\zeta K_{i,\tau-1}^\alpha L_{i,\tau-1}^{1-\alpha}}, \quad (12)$$

where

$$P_{i,\tau-1}^M = \left[\nu_i \left(p_{i,\tau-1}^{-J} \right)^{1-\omega} + (1-\nu_i) \left(p_{i,\tau-1}^J \right)^{1-\omega} \right]^{\frac{1}{1-\omega}}.$$

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

4.1.2 Post-Tsunami period

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

For these reasons we prefer an alternative approach, namely to estimate the production function without specifying the full optimization problem. We only assume that in period τ , firms operate the same technologies given by (7) and (8), and that no firm adjusts its capital stock such that $K_{i,\tau} = K_{i,\tau-1}$. Conditional on knowing the time-invariant features of the production function (ϕ_i, μ_i, ν_i), we next describe an estimation procedure that allows us to find the elasticity parameters most consistent with the observed input choices and output evident in the data.

4.2 Estimation

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

$$\kappa_i = \frac{V_{i,\tau-1}^{NA}}{p_{i,\tau-1}^x x_{i,\tau-1}}. \quad (13)$$

In words, κ_i is the fraction of firm i 's shipments exported to Canada and Mexico in the six months preceding the tsunami. We next make two assumptions that allow us to construct an estimation equation. First, we assume that a relationship analogous to (13) continues to hold in period τ , except for a log-additive error $u_{i,\tau}$. That is,

$$\ln V_{i,\tau}^{NA} = \ln \kappa_i p_{i,\tau}^x x_{i,\tau} + u_{i,\tau}. \quad (14)$$

The second assumption is that $\mathbb{E}[u_{i,\tau}|X_i] = 0$ where X_i is a vector of all right-hand-side variables. Setting the conditional mean of $u_{i,\tau}$ to zero is a standard exogeneity assumption requiring that, loosely speaking, the error is uncorrelated with all right-hand-side variables. It rules out, for example, that in response to a fall in Japanese intermediate imports firms export a fraction of their shipments to Canada and Mexico that systematically differs from κ_i . We provide evidence in section 5.2.1 that demonstrates that this is a reasonable assumption.

Using equation (7) we can rewrite (14) as

$$\ln(V_{i,\tau}^{NA}) = \ln(\kappa_i \phi_i) + \ln \left(p_{i,\tau}^x \left[\mu_i^{\frac{1}{\zeta}} (K_{i,\tau}^\alpha L_{i,\tau}^{1-\alpha})^{\frac{\zeta-1}{\zeta}} + (1 - \mu_i)^{\frac{1}{\zeta}} (M_{i,\tau})^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}} \right) + u_{i,\tau}. \quad (15)$$

Values for ν_i and μ_i are obtained from equations (11) and (12).¹⁹ Using (13), the intercept can be constructed from the previous period

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

Notice that κ_i and ϕ_i are not separately identified. Under standard assumptions, we can consistently estimate equation (15) using, e.g., nonlinear least squares. The only parameters to calibrate are the rental rate of capital R_τ and the capital share in the capital/labor aggregate α . We estimate the two elasticities, ζ and ω . Notice that ω appears in the intermediate aggregate $M_{i,\tau}$ as shown in equation (8). The estimates $(\hat{\zeta}, \hat{\omega})$ solve $\min_{\{\zeta, \omega\}} \sum_{i=1}^N (u_{i,\tau})^2$.

Why do we restrict the sample to the year surrounding the Tōhoku event? To understand this, recall that the principal difficulty of estimating production functions lies in unobserved inputs

¹⁹After constructing μ_i according to equation (12) we average by industry to reduce the level of noise.

and productivity. Since both are unobserved by the econometrician, they are absorbed into the error term. However, because they are known to the firm, other input choices depend on them. Hence, right-hand-side variables and the error term will generally be correlated, rendering estimates inconsistent.²⁰

By restricting the sample period to a single year, the assumption of constant firm productivity seems appropriate. If productivity is constant, it cannot be correlated with the error term, thereby ruling out one of the concerns.²¹ The fact that the Tōhoku event was an unexpected shock negates much of the concern about endogeneity arising from unobserved inputs. To see why, consider the case when the firm anticipates a supply chain disruption in a future period. Firm adjustment of unobserved inputs in expectation of this shock will impact input choices – leading to an endogeneity problem where inputs are correlated with the shock. Put simply, the unexpected nature of the Tōhoku event works towards equalizing the information sets between the econometrician and the firm because factor choices are not affected prior to the shock being realized.²²

Before turning to the data we briefly discuss the intuition of parameter identification. Unlike other approaches to estimating elasticities of substitution (e.g. Feenstra et al. (2014)), our method does not rely on the response of relative values to a change in relative prices.²³ In fact, in an econometric sense, our approach treats all inputs as *independent* variables.

A simple example illustrates how the parameters are identified. Consider the production function (7) and suppose that, for a particular firm, the initial period yields a value of $(1 - \mu) = 0.4$. The elasticity ζ determines how deviations from this measure of the optimal input mix between the intermediate aggregate $M_{i,\tau}$ and the capital labor aggregate translate into measured output. Thus, if we observe comparatively fewer intermediates $M_{i,\tau}$, reflecting a different mix of inputs than that given by 0.4, we obtain an elasticity estimate for ζ that best matches the response in output. Because the estimates for μ , ν , and $\kappa_i \phi_i$ are themselves functions of the elasticities, this procedure must iterate across the parameter space to find the estimate most consistent with the data. Similar reasoning applies for the identification of the ω elasticity based on relative movements in Japanese

²⁰This problem is discussed in greater detail in, for example, Akerberg, Caves, and Frazer (2006).

²¹Of course, the size and exogeneity of the shock also helps with this concern: any idiosyncratic productivity movements during this time are surely subsumed by the earthquake/tsunami.

²²An unobserved input that could remain operative in our case is that of factor utilization. Since the scope for substantial adjustment along this dimension seems quite limited, we remain confident that our estimates would be robust to the inclusion of this missing ingredient.

²³Given that we observe little systematic variation in prices (see section 5.2.4), we believe that our approach is more appropriate in this setting.

materials, non-Japanese materials, and output. The estimates we obtain are the best fit across the firms in each sample.

4.3 Connecting Model and Data

Estimation of the model requires data on employment, Japanese and non-Japanese material inputs, as well as on exports to North America and output prices for periods $\tau - 1$ and τ . Since data on firm-specific capital stocks are hard to obtain and likely noisy, we use equation (10) to construct it from firm payroll and a semi-annual rental rate of 7 percent for period $\tau - 1$.²⁴ Recall that the capital stock is not adjusted over this time horizon so that $K_{i,\tau} = K_{i,\tau-1}$. The parameter α is calibrated to $1/3$.²⁵

Quarterly employment information comes from the Business Register, which we adjust to reflect the average value over the 6 month periods we study, as they do not align with the quarters defined within a calendar year.²⁶ As discussed in earlier sections, the LFTTD contains firm-level data of Japanese imports and North American exports. For non-Japanese material inputs, we would ideally combine the non-Japanese imported materials with information on domestic material usage for these firms. As information on domestic material inputs is not available in Census data at this frequency, we utilize information on the *total* material expenditures from the Census of Manufacturers (CM) to construct a firm-level scaling factor to gross up non-Japanese intermediate imports. Put differently, we impute non-Japanese material inputs from non-Japanese input imports. For each firm, we construct the scaling factor as

$$\frac{P_i^M M_i - p_i^J m_i^J}{p_i^{-J} m_i^{-J}} \quad (16)$$

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

Regarding information on prices, the LFTTD records the value and quantity of each trade

²⁴This comes from assuming a real interest rate of 4 percent, combined with an annual depreciation rate of 10 percent, and then adjusting for a semi-annual frequency. The estimates are insensitive to alternative values of the rental rate.

²⁵In principle it is possible to construct a firm-specific value for α , using value-added information available in a census year. We are currently exploring the feasibility of this option.

²⁶Specifically: $L_{\tau-1} = \frac{1}{6}Emp_{2010Q3} + \frac{1}{2}Emp_{2010Q4} + \frac{1}{3}Emp_{2011Q1}$ and $L_{\tau} = \frac{1}{2}Emp_{2011Q2} + \frac{1}{2}Emp_{2011Q3}$.

transaction (at the HS10 level), and thus it is possible to construct the associated price, or “unit-value” of each shipment directly.²⁷ Aggregating up these shipments into a firm-month observation is complicated, of course, by the differing quantity units. Lacking any better alternative, we simply average the transaction prices using the dollar value of each transaction as weights.

Finally, we restrict the sample of firms to those that have regular imports from Japan and non-Japan over the periods we study, as well as regular North American exports.²⁸ While this substantially limits the number of firms in each sample, the shares of trade represented by these firms in each category remains very high (see Table 4).

We obtain standard errors using bootstrap methods, which also allow us to account for the uncertainty implied by the imputation of non-Japanese material inputs. We draw randomly with replacement from our set of firms to construct 5000 different bootstrap samples. For each of these samples, the non-Japanese materials share is imputed as described above before the estimation proceeds.

4.4 Summary of Results

The results of the estimation are shown in Table 4. The elasticity between material inputs for Japanese affiliates is 0.2, while the elasticity between the aggregate material input and capital/labor is 0.03. Together, these estimates are indeed consistent with the reduced-form evidence for the (ψ) elasticity from section 3.2. The relative magnitudes are also intuitive: while Japanese imported inputs are strong complements with other material inputs — consistent with the high share of intra-firm transactions comprising this trade — there is even less scope for substitution between material inputs and domestic capital/labor.

The estimation procedure also allows us to estimate these elasticities for a sample of non-Japanese multinationals. While the estimates for the ζ elasticity are indeed very close, the elasticity estimate corresponding to material inputs for non-Japanese multinationals is considerably higher, at 0.6. The lower share of intra-firm imports from Japan for this group is consistent with the argument that this type of trade is the key source of non-substitutability in the short-run. We should note, however, that the estimates for both parameters are significantly lower than what is

²⁷Those transactions with missing or imputed quantity information are dropped. Future efforts will evaluate whether it is possible to recover the quantity values from prior transaction details.

²⁸Specifically, we drop any firm that has more than 3 months of zeros for any of these values, over the period $\tau - 1$ or the period τ .

commonly assumed (typically unity or higher) in the literature.

Other details from this estimation make intuitive sense. The weight on Japanese materials in the aggregate materials bundle (ν) is 0.17 for Japanese multinationals, and 0.04 for non-Japanese firms. The weight on domestic capital/labor (μ) is considerably lower for Japanese firms — 0.22 compared to 0.51 for non-Japanese firms — which could point to either i) a more assembly-oriented production or ii) a different industry mix.

Although the number of firms included in this estimation is small (105 Japanese affiliates and 304 non-Japanese multinationals), they account for a large share of economic activity in the United States. Looking at their combined share of total trade, these firms account for over 80% of Japanese intermediate imports, 68% of non-Japanese intermediate imports, and well over 50% of North American exports. Such high concentration of trade among relatively few firms is consistent with other studies using this data (see Bernard et al. (2007)).

5 Discussion

The structural estimates of the model are broadly in agreement with the evidence in section 3.2: imported inputs are strong complements with other inputs in the production function. The rigidity of the production function for multinational firms is likely due to the substantial presence of intra-firm trade in what is presumably highly specialized inputs.²⁹ This result has a number of important implications for how we think about business cycle co-movement, firm volatility, and multinational firms more broadly.

5.1 Implications

The rigid production networks of foreign-owned multinationals will have direct consequences on the destination (host) economy. Previous literature has hypothesized that input linkages could generate business-cycle comovement, but supportive empirical evidence has been difficult to find. This paper can be seen as a first step in establishing empirical evidence for a causal relationship between trade, multinational firms, and business cycle comovement. In a companion paper (Boehm, Flaaen, and Pandalai Nayar (2014)), we evaluate the quantitative importance of such complementarities

²⁹The vertical integration of production across countries, within the firm, has shown to be a key driver of the decline in joint ventures (see Desai, Foley, and Hines (2004)).

of imported inputs by multinational affiliates. When separately accounting for intermediate input trade by multinationals and traditional trade in final goods, the model distinguishes between the production elasticity of imported inputs and the traditional “Armington” elasticity used to bundle together international goods for consumption. The complementarities in import linkages by multinationals increases value-added comovement in the model by 11 percentage points relative to a benchmark without such firms.

This model share similarities with several other existing models, particularly Burstein, Kurz, and Tesar (2008). A key advantage of Boehm, Flaaen, and Pandalai Nayar (2014), however, is a tight link to Census data for matching other features of multinationals and trade. Johnson (2014) also looks at the role of vertical linkages on comovement, but applies greater input-output structure on the model. Such features will generate increases in value-added comovement in his model, the magnitude of which becomes significant only when the elasticity of substitution among inputs is sufficiently low. Other work also identifies multinationals as a key source of the transmission of shocks: Cravino and Levchenko (2014) demonstrates that foreign multinational affiliates can account for about 10 percent of aggregate productivity shocks.³⁰

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

Another branch of literature on the diversification of risk has studied whether firms using complex production structures with several intermediates could be less volatile (Koren and Tenreyro (2013)). Kurz and Senses (2013) establish that firms with substantial imports and exports have lower employment volatility than domestic firms in the medium to long term, which they attribute

³⁰Of course, shocks can be passed through to affiliates through other means as well. See Peek and Rosengren (1997) and Peek and Rosengren (2000) for the case of U.S. affiliates from Japan.

partly to the diversification of risk.³¹ The key result in this paper points to a possibly overlooked fact: the extent of the benefits from diversification depends heavily on the substitutability of inputs. Conditional on a given number of inputs used in production, a firm will likely experience greater volatility if each input is key to the production process and inputs are subject to heterogeneous shocks.³² Conceptually, an increase in the use of imported inputs should not be viewed necessarily as diversification. A fragmentation of production can lead to an increased supply chain risk that is an important counterweight to whatever efficiencies such complex input sourcing might afford, particularly when the production elasticities are low.

The rigid production networks of multinational firms also influences our understanding of why firms segment production across country borders. In a related paper, Flaaen (2013b) shows that despite the presence of substantial and complex import linkages with the source country (consistent with a vertical framework of FDI), the motive for multinational production appears to be to serve the domestic market (consistent with the horizontal framework of FDI). The result could be called “horizontal FDI with production sharing.”³³ The evidence for strong complementarities in this production sharing, however, presents a puzzle. Why does the firm replicate only select portions of the supply chain, considering the penalties for disruptions and mismatched inputs are so great? It is perhaps the case that the segments of the production chain that remain in the source country have a location-specific component that is not easily transferable when the firm moves production abroad.³⁴ Understanding the dynamics behind these sourcing decisions is an area in need of further research.³⁵

Finally, the estimates from section 4 correspond to an average across the firms in each group (Japanese multinationals and non-Japanese firms), the results from section 3.2 correspond to the aggregated group of Japanese multinationals. Because the level of aggregation differs, it is important for the researcher to understand the appropriate model analogue when using these estimates

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

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

³³Ramondo, Rappoport, and Ruhl (2014) is another recent example arguing for a more nuanced interpretation of multinational production.

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

³⁵For a recent example of how such investment and sourcing decisions can alter a country’s comparative advantage over time, see Alvarez (2014).

for the calibration of theoretical models.³⁶ For example, a researcher may decide that a group-level elasticity is the most appropriate for an aggregate DSGE model populated by a representative multinational firm, where the focus is on understanding aggregate features such as business cycle dynamics.

5.2 Robustness and Extensions

5.2.1 Mis-measurement of Firm Production

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

To evaluate this concern, we narrow our study to the automotive sector, which has data on production, sales, and inventory at a monthly frequency. Using the Ward’s electronic databank, which reproduces the published series in the annual Automotive Yearbook, we obtain plant-level information on production, and model-line information on inventory, sales, and incentives.³⁷ The baseline specification is the same as in equation (5), where the dependent variable is now Q_{jit} : production of plant j of firm i in month t . The Japanese multinational firms are, in this case, those automakers with plants located within North America but whose parent company is headquartered in Japan.³⁸

Figure 8 shows the results, where we once again divide by pre-shock levels to gain a sense of the percentage effects of these changes. Relative to their U.S. counterparts, Japanese automakers in the United States experienced large drops in production following the Tōhoku event. Production bottomed out in May of 2011 — two months after the event — at almost 60 percent below trend.³⁹ The point estimates return to a level near zero in September of 2011, implying that the shock affected production for nearly 6 months.⁴⁰ We interpret these results to be largely supportive of

³⁶A 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.

³⁷Appendix C.5 details further features of this data and explains how the sample was constructed.

³⁸These firms are Honda, Mitsubishi, Nissan, Toyota, and Subaru.

³⁹The average monthly plant-level production at these firms during December 2010 through February 2011 was about 12,200 units a month. The magnitude of the drop in May was -7200 units.

⁴⁰We describe additional results on the behavior of inventories, sales, incentives, and production in Japan in an appendix.

the results obtained using the exports-based proxy for production. The percentage drops in the two series are remarkably similar: a trough of 59% at $t = 2$ in the automotive data vs 53% at $t = 2$ using the proxy. We conclude that, at least for this exercise, the proxy appears to be providing valuable information on a firm’s U.S. production behavior.⁴¹

5.2.2 Intermediate Input Inventories

Inventories are another obvious feature that should influence the relationship between input shipments, production, and the elasticity of substitution. In particular, inventories of intermediate inputs allow the firm to absorb unforeseen shocks to input deliveries without an impact on the production process.⁴² 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 Figure 6 (Panel B) or Figure 8. The effect on production appears to be almost immediate, indicating that the stock of inventories of imported intermediates is low (less than one month’s supply) for these firms.

We obtain a rough sense of the degree of inventory holdings from the Census of Manufacturers micro-data. Combining information on the beginning period stock of materials inventories with the annual usage of materials, we calculate the average monthly supply of inventories for each firm.⁴³ Panel A of Table 1 calculates the production-weighted averages over a select set of firm groups.⁴⁴ 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

⁴¹In 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 5.2.1 should also serve to allay such concerns.

⁴²The 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.

⁴³Unfortunately, the CM data does not report imported materials inventory separately.

⁴⁴These numbers are broadly similar, though somewhat lower than other estimates in the literature. See Ramey (1989) for one example.

oft-cited “lean” production processes made famous by Japanese firms in previous decades. Though these data are for the year 2007, there is little reason to believe these relative magnitudes have changed substantially over a period of a few years. For completeness, Panel A of Table 1 also reports the corresponding estimates for output inventories.⁴⁵

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

5.2.3 Multi-Products and Sub-Optimal Mix

In the frameworks used in sections 3.1 and 4, we consider the aggregate bundles of imported intermediates, abstracting away from product-level detail. In reality, the firms in our dataset often import many distinct intermediate inputs from Japan. The structure of a CES production function implies that if each of these within-country inputs was non-substitutable with one another (a further, nested Leontief structure), the production impact of a disruption in the supply of just one input could be amplified relative to the value of that input.⁴⁶ We evaluate this possibility below.

This is particularly true given the heterogeneous impact of the Tōhoku event across Japan (see Figure 4). This could translate into considerable dispersion in the impact on the *products* imported by a particular U.S. firm or Japanese affiliate. With product-level shocks, considering the effect on the aggregate import bundle amounts to assuming either 1) perfect substitutability among products, or 2) that the firm maintains an optimal within-country product mix at all times.

To be concrete, it may be more accurate to view the M_t in equation (1) as a further C.E.S.

⁴⁵At 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.

⁴⁶This point has been made in somewhat differing contexts, by Kremer (1993) and Jones (2011).

function of multiple products. Thus, we can define the proper measurement of this variable as

$$V_{i,t}^M = P_{i,t}^M \left(\sum_{n=1}^N \eta_n^{\frac{1}{\chi}} (m_{n,i,t}^J)^{\frac{\chi-1}{\chi}} \right)^{\frac{\chi}{\chi-1}}, \quad (17)$$

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

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

A suboptimal product mix indicates that measured imports ($\widehat{V}_{i,t}^M$) are greater than the effective imports ($V_{i,t}^M$). As a result the measured output response to the import shock will be larger than otherwise, resulting in a downward “bias” in the elasticity estimates from section 3.1 and 4.⁴⁷ Such an effect is decreasing in the product-level elasticity parameter χ , as complementarity itself is the driving force between differences in $\widehat{V}_{i,t}^M$ and $V_{i,t}^M$. In addition, the effect is also increasing in the degree of deviation from the optimal product mix.

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

⁴⁷Because the source of this downward pressure on the estimate for ψ (or ω) is itself a very low product-level elasticity, it is unclear whether this should be considered a bias in the traditional sense.

$$DO_{i,t} = \frac{1}{N^i} \sum_{n=1}^{N^i} (|s_{n,i,t} - s_{n,i,s^*}|) \quad (18)$$

where N^i is the total number of products imported by firm i . We define the period s^* to be the months of April-June of 2010, and then evaluate DO_i at a monthly frequency, with particular interest in the months following the Tōhoku event. While there may be natural movements in the bundle of products imported from Japan, evidence for substantial coordination failure in product composition or heterogeneity in product-level shocks would come from any abnormal jumps in this index in the months of the disruption. One can calculate this at various levels of product aggregation (i.e. HS4, HS6, HS8, HS10), though we report results using the HS6 level.⁴⁸

The results of this exercise are shown in Figure 9. We plot the average DO_i across Japanese firms for each month (the figure shows a 3-month moving average) during the period 2009-2011. Mechanically, this measure should be relatively close to zero in the months consisting of the benchmark (April-June 2010). While there is a secular rise in this measure on either side of this benchmark period, there do not appear to be any large jumps in the months directly following the Tōhoku event. More interesting, perhaps, are the considerably larger values for this measure during early 2009, which might reflect the effects of the trade collapse associated with the Great Recession. We interpret Figure 9 as evidence that the potential for suboptimal mix across products from Japan does not pose a serious problem to our measurement in previous sections.

5.2.4 Other Considerations

Strategic Behavior: Another possibility that could affect the interpretation of the results from Figure 7 might be strategic behavior, particularly on the part of the competitors of Japanese firms in the United States. These firms could raise production or prices following the negative supply shock affecting their competitors, which would serve to bias downward the β_p coefficients from the equation with $X_{i,t}^{NA}$ as the dependent variable.⁴⁹ 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

⁴⁸The level of aggregation we use attempts to balance concerns along two dimensions. With less product aggregation (i.e. HS10 level), one might be concerned with the inherent lumpiness of product-level firm imports. More product aggregation, on the other hand, could mask important product differences within a particular product grouping.

⁴⁹Specifically, in equation (5) the γ_p 's would be higher than would be expected without the shock, and hence the β_p 's artificially low.

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

Prices: Traditionally, estimating the elasticity of substitution is accomplished via price and quantity data for products over extended periods of time. For the short horizon we consider in this paper, there are several reasons why prices may not have the scope to adjust. Many supplier relationships negotiate prices for longer periods of time than one or two months. Second, and perhaps more importantly, Table 2 demonstrates that the large majority of imported intermediate inputs are intra-firm. The observed prices of these transactions are transfer-prices (within firm) and not likely to change reflecting any short-term disturbance. However, because the LFTTD contains both quantity and price information, we can confirm whether or not prices remained relatively stable during this period. The results in Appendix Table A5 confirm that there are few significant price movements on import or export transactions for either Japanese or non-Japanese multinationals surrounding the Tōhoku event.⁵⁰

Domestic Inputs: It is also possible to evaluate the response of domestic inputs directly, using the limited information we have on quarterly firm-level employment and payroll information, taken from the Census Bureau’s Business Register (BR).⁵¹ We consider the evidence in Appendix C.3 and find no significant effects on either employment or payroll for Japanese firms in the quarter(s) following the shock (see Table A4). Of course, there are a number of reasons — principally labor adjustment costs — why one would expect little, if any, impacts on employment following this short-lived shock. Press releases dispatched by the Japanese automakers during this time indicated that no layoffs would occur. Rather, the firms indicated that they would use the production stoppages for employee skill and safety training.

⁵⁰Further details on the construction of the data underlying the analysis of unit values is available in Appendix C.4.

⁵¹The 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).

5.3 External Validity

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

When viewed in light of the substantial fraction of intra-firm imports comprising multinational affiliate trade, the low elasticity of substitution should not come as a surprise. One would not expect close substitutes for the sort of specialized products reflecting firm-specific knowledge that likely comprises this trade. Moreover, such a low estimate for an elasticity of this nature is not without precedent. Using different methodologies, recent work by Atalay (2014) highlights strong complementarities between intermediate inputs, using industry-level data for the United States.⁵²

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. The cost share of all imported inputs is actually quite close: 35 percent for Japanese affiliates vs 32 percent for all foreign affiliates.

A related concern is whether the estimates for Japanese affiliates are driven solely by the automotive sector. The ideal check would be to run industry-by-industry subgroup estimates for the elasticities. Unfortunately, the small number of firms applicable for this analysis, combined with disclosure requirements associated with the Census Bureau data usage, prevents this degree of detail. Using the published data from the B.E.A., the automotive sector is a large but not overwhelming percentage of total Japanese manufacturing affiliates in the U.S. The entire motor vehicle sector as a whole comprises significantly less than half of value-added (roughly 40 percent) for the Japanese manufacturing affiliates.

Any elasticity estimate is tied to the time-horizon to which it corresponds. Ruhl (2004) emphasizes the difference between elasticities implied by responses to temporary vs permanent shocks.

⁵²The point estimate for the elasticity of substitution among intermediate inputs from Atalay (2014) is 0.03.

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

6 Conclusions

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

These elasticities play a critical role in the way international trade impacts both source and destination economies. Such complementarities between domestic and foreign goods have been shown to improve the ability of leading theoretical models to fit key moments of the data. We emphasize here the distinction between substitutability between domestic and foreign final goods (a “consumption” elasticity of substitution, or the so-called Armington elasticity) and substitutability between domestic and foreign intermediate goods (a “production” elasticity of substitution). In a companion paper (Boehm, Flaaen, and Pandalai Nayar (2014)), we document the behavior of a model with such complementarities in imported intermediates, and discuss how these elasticity

parameters interact. Calibrating this model to the share of multinational affiliate trade in intermediates yields an increase in value-added comovement of 11 p.p., and reduces the gap between consumption and value-added comovement by a considerable margin.

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

References

- Akerberg, Daniel A., Kevin Caves, and Garth Frazer. 2006. “Structural Identification of Production Functions.” Working paper, University of Michigan.
- Alviarez, Vanessa. 2014. “Multinational Production and Comparative Advantage.” Working paper, University of Michigan.
- Atalay, Englin. 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.
- 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. “Complementarities in Multinational Production and Business Cycle Dynamics.” 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.
- Busso, Matias, John DiNardo, and Justin McCrary. 2014. “New Evidence on the Finite Sample Properties of Propensity Score Matching and Reweighting Estimators.” *Review of Economics and Statistics* forthcoming.
- Carvalho, Vasco M., Makoto Nirei, and Yukiko U. Sato. 2014. “Supply Chain Disruptions: Evidence from the Great East Japan Earthquake.” Working paper.
- Caselli, Francesco, Miklos Koren, Milan Lisicky, and Silvana Tenreyro. 2014. “Diversification through Trade.” Working paper.
- Cravino, Javier and Andrei Levchenko. 2014. “Multinational Firms and International Business Cycle Transmission.” Working paper, University of Michigan.
- Desai, Mihir A., C. Fritz Foley, and James R. Hines. 2004. “The Costs of Shared Ownership: Evidence from International Joint Ventures.” *Journal of Financial Economics* 73 (2).

- 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. 2014. "Firms, Destinations, and Aggregate Fluctuations." *Econometrica* 82 (4):1303–1340.
- Feenstra, Robert C., Philip A. Luck, Maurice Obstfeld, and Katheryn N. Russ. 2014. "In Search of the Armington Elasticity." Working 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* 6 (4):39–83.
- 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.
- 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.
- Peek, Joe and Eric S. Rosengren. 1997. "The International Transmission of Financial Shocks: The Case of Japan." *American Economic Review* 87 (4):495–505.
- . 2000. "Collateral Damage: Effects of the Japanese Bank Crisis on Real Activity in the United States." *American Economic Review* 90 (1):30–45.
- 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.
- Pravin, Krishna and Andrei Levchenko. 2014. "Comparative Advantage, Complexity, and Volatility." *Journal of Economic Behavior and Organization* 94:314–329.
- 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. 2014. "Horizontal vs Vertical FDI: Revisiting Evidence from U.S. Multinationals." Working paper, Columbia Business School.
- Ruhl, Kim J. 2004. "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.
- 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 Bureau Business Register: Basic Features and Quality Issues." Paper presented at the Joint Statistical Meetings, Anaheim, CA.
- Wasi, Nada and Aaron Flaaen. 2014. "Record Linkage Using Stata: Pre-Processing, Linking and Reviewing Utilities." Working paper, University of Michigan.

Table 1: Summary Statistics: Imported Inputs and Inventories by Firm Type

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

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

Table 2: Baseline Sample: Summary Statistics

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

Source: LFTTD, DCA, and UBP as explained in the text.

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

Table 3: Predicting Japanese Import and U.S. Output Disruption by Firm Characteristics

Panel A: Disruption to U.S. Output (proxy)				Panel B: Disruption to Japanese Imports				
$X_i^D = 1$				$J_i^D = 1$				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Japan	0.443*** (0.0921)		0.352*** (0.117)	0.347** (0.152)	0.707*** (0.0917)		0.310*** (0.115)	0.686*** (0.150)
Exposed		0.351*** (0.0886)	0.145 (0.112)	0.140 (0.149)		0.814*** (0.0880)	0.636*** (0.110)	0.991*** (0.144)
JPN*Exp				-0.00771 (0.228)				-0.848*** (0.222)
MMI	-0.176*** (0.0676)	-0.121* (0.0646)	-0.178*** (0.0676)	-0.178*** (0.0683)	0.346*** (0.0691)	0.389*** (0.0667)	0.341*** (0.0694)	0.306*** (0.0704)
Ports	-0.174 (0.224)	-0.144 (0.225)	-0.197 (0.226)		0.248 (0.211)	0.217 (0.212)	0.168 (0.213)	0.174 (0.213)
Constant	-0.674 (0.681)	-0.674 (0.681)	-0.674 (0.681)	-0.674 (0.681)	-4.672 (85.78)	-4.672 (85.78)	-4.672 (85.78)	-4.668 (85.00)
Industry Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2451	2451	2451	2451	2451	2451	2451	2451

*** p < 0.01, ** p < 0.05, * p < 0.1

Source: LFTTD, DCA, UBP, and USGS as explained in the text. This table reports the results of a probit model prediction of JPN import and N.A. exports (output) disruption based on firm characteristics. See section 3.1 for a definition of the variables.


Table 4: Firm-Level Estimation: Summary Statistics and Results

Panel A: Calibration		
Parameter	Value	Source/Justification
R_t	0.07	
α	1/3	
Panel B: Estimation Results		
	Japanese Multinationals	Non-Japanese Multinationals
ω	0.201 [0.02 0.43]	0.624 [0.16 0.69]
ζ	0.032 [0.030 0.673]	0.038 [0.035 0.508]
Sample Details		
Weight on K/L Aggregate ($\bar{\mu}$)	0.223	0.514
Weight on JPN Materials ($1 - \bar{\nu}$)	0.173	0.044
Number of Firms	105	304
Share of Total Trade		
JPN int imports	0.60	0.23
Non-JPN int imports	0.02	0.66
N.A. exports	0.08	0.47

Source: CM, LFTTD, DCA, and UBP as explained in the text.


This table reports the results from the firm-level estimation detailed in section 4. Panel A outlines the parameters that are calibrated prior to estimation. The top two rows of Panel B reports the point estimates of the elasticities, and the corresponding 95 percent confidence intervals using a bootstrapping procedure. (See Appendix C.2 for more details on the measurement of dispersion for these estimates.) Rows 3 and 4 report other estimates related to the calculated production functions. The final rows of Panel B describe features of the two estimation samples.

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




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

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




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


This figure displays density estimates of the (log) exposure measure to Japanese imported inputs, separately for Japanese affiliates and non-Japanese multinational firms. The estimates correspond to year 2010. The exposure measure is defined as the ratio of Japanese imported inputs to total imported inputs plus U.S. salaries and wages. It is equivalent to the cost share measure shown in table 1 with the omission of domestic material usage and investment expenditures. Density estimates at either tail are suppressed for purposes of confidentiality.

Figure 6: Dynamic Treatment Effects: Japanese Firms

A. Relative Intermediate Input Imports of Japanese Firms




B. Relative North American Exports of Japanese Firms



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

These figures report the intermediate imports and North American exports of the U.S. affiliates of Japanese firms relative to a control group of other multinational firms. The values are coefficient estimates taken from an interaction of a Japanese-firm dummy with a monthly dummy – additional baseline monthly dummies remove seasonal effects. See equation 5 in the text. Standard errors are clustered at the firm level.


Figure 7: Relative Imported Inputs and Output (Proxy) of Japanese Firms: Fraction of Pre-Shock Level



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

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


Figure 8: Assessing the Output Proxy Using Monthly Automotive Production



Source: Ward's Automotive Database

This figure reports the production levels of Japanese auto plants relative to a control group of non-Japanese auto plants. The values are percent changes from a pre-shock level, defined as the average of the months December 2010, January 2011, and February 2011. See equation A13 in the text. For purposes of comparison, we also include the equivalent measure corresponding to total manufacturing of Japanese affiliates using the output proxy from Census data (from Figure 7). The Japanese automakers are Honda, Mazda, Mitsubishi, Nissan, Toyota, and Subaru. For the sake of clarity, we suppress the standard errors for the automotive series, though there are 4 months with below zero production based on a 95 percent confidence interval. See section C.5 for more details.

Figure 9: Japanese Products: Average Distance from Benchmark Cost Shares: JPN Multinationals



Source: LFTTD-DCA-UBP as explain in the text

Underlying this figure is the calculation of the average total (absolute) deviations from a benchmark measure of a firm's cost shares across input products from Japan. See equation 18 in the text. The figure reports the mean across the Japanese multinationals used in the section 4.

A Basic Theory Appendix

A.1 Proof of Result 1

Suppose that the firm solves

$$\max p_x x - p_D F_D - p_M IM$$

subject to

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

and

$$p_x = \left(\frac{Y}{x} \right)^{\frac{1}{\varepsilon}}$$

The first order conditions are

$$\begin{aligned} \left(1 - \frac{1}{\varepsilon} \right) (Y)^{\frac{1}{\varepsilon}} (x)^{\frac{1}{\psi} - \frac{1}{\varepsilon}} (1 - \mu)^{\frac{1}{\psi}} [F_D]^{-\frac{1}{\psi}} &= p_D \\ \left(1 - \frac{1}{\varepsilon} \right) (Y)^{\frac{1}{\varepsilon}} (x)^{\frac{1}{\psi} - \frac{1}{\varepsilon}} \mu^{\frac{1}{\psi}} [IM]^{-\frac{1}{\psi}} &= p_M \end{aligned}$$

Dividing one by the other gives

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

The same equation can be obtained under perfect competition.

Now take the production function and multiply it by p_x

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

Taking logs gives

$$\ln(p_x x) = \frac{\psi}{\psi-1} \ln \left(p_x \left[(1 - \mu)^{\frac{1}{\psi}} [F_D]^{\frac{\psi-1}{\psi}} + (p_M)^{-\frac{\psi-1}{\psi}} \mu^{\frac{1}{\psi}} [p_M IM]^{\frac{\psi-1}{\psi}} \right] \right) \quad (\text{A1})$$

$$= \frac{\psi}{\psi-1} \ln \left(p_x \left[(1 - \mu)^{\frac{1}{\psi}} \exp \left(\frac{\psi-1}{\psi} \ln [F_D] \right) + (p_M)^{-\frac{\psi-1}{\psi}} \mu^{\frac{1}{\psi}} \exp \left(\frac{\psi-1}{\psi} \ln [p_M IM] \right) \right] \right) \quad (\text{A2})$$

Before differentiating, recall the assumption that the firm takes prices p_M as given and that it cannot change p_x after learning about the shock. Then

$$\frac{\partial \ln p_x x}{\partial \ln p_M M} = \frac{\psi}{\psi - 1} \frac{p_x (P_M)^{-\frac{\psi-1}{\psi}} \mu^{\frac{1}{\psi}} \exp\left(\frac{\psi-1}{\psi} \ln [p_M IM]\right) \frac{\psi-1}{\psi}}{p_x \left[(1 - \mu)^{\frac{1}{\psi}} \exp\left(\frac{\psi-1}{\psi} \ln [F_D]\right) + (p_M)^{-\frac{\psi-1}{\psi}} \mu^{\frac{1}{\psi}} \exp\left(\frac{\psi-1}{\psi} \ln [p_M IM]\right) \right]} \quad (\text{A3})$$

$$= \frac{1}{1 + \left(\frac{1-\mu}{\mu}\right)^{\frac{1}{\psi}} \left[\frac{F_D}{IM}\right]^{\frac{\psi-1}{\psi}}} \quad (\text{A4})$$

We evaluate this elasticity at

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

so that

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

A.2 On Flexibility in Domestic Inputs

Under the assumption of perfect competition, the first order conditions are:

$$\begin{aligned} x(1 - \mu) &= (p_D)^\psi F_D \\ x\mu &= (p_M)^\psi IM \end{aligned}$$

If the firm takes prices p_x , p_M , and p_D as given, the following elasticities are immediate:

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

The above equations demonstrate that a constant returns to scale production function combined with these assumptions on market structure imply that the output elasticity will equal one for all values of the elasticity of substitution. For this reason, we require some assumptions limiting the flexibility of domestic inputs following the import disruption.

Below we show an alternative way of understanding the interaction of competitive factor markets, changes in domestic inputs, and the mapping of the output elasticity into parameter values for the elasticity of substitution. Consider the total derivative of $\ln(x)$:

$$d \ln x = \frac{\partial \ln x}{\partial IM} d \ln IM + \frac{\partial \ln x}{\partial F} d \ln F \quad (\text{A5})$$

$$d \ln x = \frac{\mu^{\frac{1}{\psi}} (IM)^{\frac{\psi-1}{\psi}} d \ln IM}{(1-\mu)^{\frac{1}{\psi}} [F_D]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} [IM]^{\frac{\psi-1}{\psi}}} + \frac{(1-\mu)^{\frac{1}{\psi}} (F_D)^{\frac{\psi-1}{\psi}} d \ln F_D}{(1-\mu)^{\frac{1}{\psi}} [F_D]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} [IM]^{\frac{\psi-1}{\psi}}} \quad (\text{A6})$$

Dividing by $d \ln IM$ yields:

$$\frac{d \ln x}{d \ln IM} = \frac{\mu^{\frac{1}{\psi}} (IM)^{\frac{\psi-1}{\psi}}}{(1-\mu)^{\frac{1}{\psi}} [F_D]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} [IM]^{\frac{\psi-1}{\psi}}} + \frac{(1-\mu)^{\frac{1}{\psi}} (F_D)^{\frac{\psi-1}{\psi}}}{(1-\mu)^{\frac{1}{\psi}} [F_D]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} [IM]^{\frac{\psi-1}{\psi}}} \frac{d \ln F_D}{d \ln IM}$$

Now, as before, combining the first order conditions from the profit maximization problem, we have:

$$\frac{F_D(\cdot)}{IM} = \frac{1-\mu}{\mu} \left(\frac{p_D}{p_M} \right)^{-\psi} \quad (\text{A7})$$

Log-differentiating this expression:

$$\begin{aligned} d \ln \left(\frac{F_D}{IM} \right) &= -\psi d \ln \left(\frac{p_D}{p_M} \right) \\ d \ln F_D - d \ln IM &= -\psi d \ln \left(\frac{p_D}{p_M} \right) \\ \frac{d \ln F_D}{d \ln IM} &= 1 - \psi \frac{d \ln \left(\frac{p_D}{p_M} \right)}{d \ln IM} \end{aligned} \quad (\text{A8})$$

Finally, we have:

$$\frac{d \ln x}{d \ln IM} = \frac{\mu^{\frac{1}{\psi}} (IM)^{\frac{\psi-1}{\psi}}}{(1-\mu)^{\frac{1}{\psi}} [F_D]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} [IM]^{\frac{\psi-1}{\psi}}} + \frac{(1-\mu)^{\frac{1}{\psi}} (F_D)^{\frac{\psi-1}{\psi}}}{(1-\mu)^{\frac{1}{\psi}} [F_D]^{\frac{\psi-1}{\psi}} + \mu^{\frac{1}{\psi}} [IM]^{\frac{\psi-1}{\psi}}} \left[1 - \psi \frac{d \ln \left(\frac{p_D}{p_M} \right)}{d \ln IM} \right] \quad (\text{A9})$$

Thus, if there is no change in the relative input price following the disruption in IM of the firm: $\frac{d \ln \left(\frac{p_M}{p_D} \right)}{d \ln IM} = 0$, then the output elasticity will be equal to one regardless of the value of ψ . On the other hand, any assumptions that yield a non-zero change in the relative input prices will then yield the result that $\frac{d \ln x}{d \ln IM} = 1$ provided $\psi \rightarrow 0$.

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 directories to the Census data infrastructure.

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 a set of record linking utilities described in Wasi and Flaaen (2014). This program uses a bigram string comparator algorithm on multiple variables with

differing user-specified weights.⁵³ 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.

B.1.3 The Unit of Matching

The primary unit of observation in the DCA, UBP, and BR datasets is the business establishment. Hence, the primary unit of matching is the establishment, and not the firm. However, there are a number of important challenges with an establishment-to-establishment link. First, the DCA (UBP) and BR may occasionally have differing definitions of the establishment. One dataset may separate out several operating groups within the same firm address (i.e. JP Morgan – Derivatives, and JP Morgan - Emerging Markets), while another may group these activities together by their common address. Second, the name associated with a particular establishment can at times reflect the subsidiary name, location, or activity (i.e. Alabama plant, processing division, etc), and at times reflect the parent company name. Recognizing these challenges, the primary goal of the matching will be to assign each DCA (UBP) establishment to the most appropriate business location of the parent firm identified in the BR. As such, the primary matching variables will be the establishment name, along with geographic indicators of street, city, zip code, and state.

B.1.4 The Matching Process: An Overview

The danger associated with probabilistic name and address procedures is the potential for false-positive matches. Thus, there is an inherent tension for the researcher between a broad search criteria that seeks to maximize the number of true matches and a narrow and exacting criteria that eliminates false-positive matches. The matching approach used here is conservative in the sense that the methodology will favor criteria that limit the potential for false positives at the potential expense of slightly higher match rates. As such, the procedure generally requires a match score exceeding 95 percent, except in those cases where ancillary evidence provides increased confidence in the match.⁵⁴

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

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

⁵³The term bigram refers to two consecutive characters within a string (the word *bigram* contains 5 possible bigrams: “bi”, “ig”, “gr”, “ra”, and “am”). The program is a modified version of Blasnik (2010), and assigns a score for each variable between the two datasets based on the percentage of matching bigrams. See Flaaen (2013a) or Wasi and Flaaen (2014) for more information.

⁵⁴The 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 the parent firm is headquartered outside of the United States, we designate this establishment as part of a Foreign multinational firm. We also retain the nationality of parent firm.⁵⁵

There can be a number of issues when translating the DCA-based indicators through the DCA-BR bridge for use within the Census Bureau data architecture. First, there may be disagreements between the DCA and Census on what constitutes a firm, such that an establishment matches may report differing multinational indicators for the same Census-identified firm. Second, such an issue might also arise due to joint-ventures. Finally, incorrect matches may also affect the degree to which establishment matches agree when aggregated to a firm definition. To address these issues, we apply the following rules when using the DCA-based multinational indicators and aggregating to the (Census-based) firm level. There are three potential cases:⁵⁶

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.

⁵⁵The multinational status of firms from the UBP directories are more straightforward.

⁵⁶Some of these cases also apply to the UBP-BR bridge.

4. If one particular type of link yields an employment share above 0.5 and total firm employment is below 10,000, then apply that link to all establishments within the firm.
5. All other cases require manual review.

Potential 3: A Census-identified firm in which one establishment is matched to a non-multinational firm, and another establishment is matched to a foreign-country parent firm (or U.S. multinational firm).

Apply same steps as in Potential 2.

B.2 Classifying Firm-Level Trade

The firm-level data on imports available in the LFTTD does not contain information on the intended use of the goods.⁵⁷ 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.⁵⁸ 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

⁵⁷This is one advantage of the survey data on multinational firms available from the Bureau of Economic Analysis. There are, however, a number of critical disadvantages of this data source, as outlined in Flaaen (2013b).

⁵⁸To be more precise, this set will include a combination of intermediate and capital goods.

are typically products that are self-identified by the firm but do not match any of the pre-specified products identified for that industry by Census.) Rather than ignoring the value of shipments corresponding to these codes, we attempt to match at a more aggregated level. Specifically, we iteratively try to find a product code match at the 6, 5, and 4 digit product code level, and use the existing set of 7-digit matches as weights to allocate the product value among the 7-digit product codes encompassing the more aggregated level.

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

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

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

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

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

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

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

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

There are several advantages to using the CM-Products file rather than using an input-output table.⁵⁹ First, within a given CM year, the classification can be done at the firm or establishment level rather than aggregating to a particular industry. This reflects the fact that the same imported product may be used as an input by one firm and sold to consumers as

⁵⁹ 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 concurring to the trade data is considerably more difficult.

a final product by another. Second, the CM-Products file is one of the principal data inputs into making the input-output tables, and thus represents more finely detailed information. Related to this point, the input-output tables are produced with a significant delay – the most recent available for the U.S. is for year 2002. Third, the input-output tables for the U.S. are based on BEA industry classifications, which imply an additional concordance (see below) to map into the NAICS-based industries present in the Census data.

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

B.2.2 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.⁶⁰ 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:

$$\left. \begin{aligned} M_{ij}^{\text{int}} &= \sum_{p \notin P_j} M_{ipj} \\ M_{ij}^{\text{fin}} &= \sum_{p \in P_j} M_{ipj} \end{aligned} \right\} \quad \text{where} \quad P_j = \{p \mid S_{pj} \geq W\}. \quad (\text{A10})$$

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

⁶⁰It 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 would not serve to bias the firms’ classification of imported goods, as the upstream products are not actually “final” goods for that firm.

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 3.1.

Beginning with the raw files of the LFTTD export/import data, we drop any transactions with missing firm identifiers, and those pertaining to trade with U.S. territories. Next, we merge the LFTTD files with the HS-NAICS6 product concordance from Pierce and Schott (2012); if there is no corresponding NAICS6 code for a particular HS code, then we set NAICS6 equal to XXXXXX. We then aggregate up to the level of Firm-Country-Month-NAICS6, and then create extracts according to three sets of destinations/sources: Japan, Non-Japan, and North America (Canada and Mexico). Then, assigning each firm to an LBD-based industry (see below), we run the NAICS-based trade codes through the intermediate/final goods filter discussed in Appendix B.2. The firms' monthly trade can then be split into intermediate and final goods components. We repeat this step for years 2009, 2010, and 2011.

Using the Longitudinal Business Database, we drop inactive, ghost/deleted establishments, and establishments that are not in-scope for the Economic Census. To create the sample of manufacturing firms in the U.S., we first create a firm industry code defined as the industry encompassing the largest share of firm employment. We then drop non-manufacturing firms. Next, we merge the LBD for each year with the DCA-Bridge (see section B.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: DCA Match Statistics: 2007-2011

	# of DCA Establishments	Matched to B.R.	Percent Matched
Total			
2007	112,346	81,656	0.73
2008	111,935	81,535	0.73
2009	111,953	81,112	0.72
2010	111,998	79,661	0.71
2011	113,334	79,516	0.70
U.S. Multinationals			
2007	22,500	16,396	0.73
2008	23,090	16,910	0.73
2009	22,076	16,085	0.73
2010	21,667	15,785	0.73
2011	21,721	15,557	0.72
Foreign Multinationals			
2007	10,331	7,555	0.73
2008	9,351	6,880	0.74
2009	11,142	8,193	0.74
2010	11,308	8,181	0.72
2011	11,619	8,357	0.72

Table A2: Uniworld Match Statistics: 2006-2011

	# of Uniworld Establishments	Matched to B.R.	Percent Matched
Foreign Multinationals			
2006	3,495	2,590	0.74
2008	3,683	2,818	0.76
2011	6,188	4,017	0.65
U.S. Multinationals ¹			
2007	4,043	3,236	0.80
2009	4,293	3,422	0.80

¹U.S. multinationals include only the establishment identified as the U.S. headquarters.

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

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

C Appendix: Other Results

C.1 Alternate Specifications for Treatment Effects Regressions

Our results from section 3.2 are based on a sample including all Japanese multinationals in manufacturing, and therefore uses a levels specification to allow for zeros in the firm-month observations. Because larger firms exhibit greater absolute deviations from trend, this roughly amounts to weighting firms based on size, such that the results correspond to a representative firm based on the aggregate effect of the group.

To see this, and to explore how the levels specification influences our interpretation, we repeat the analysis on a subset of the firms for which we can view the percentage changes directly. Specifically, we drop any firms with zeros in any month for intermediate imports or N.A. exports during the sample, and then take logs and HP-filter each series to obtain percentage deviations from trend for each firm.⁶¹ The results of this exercise are shown in Panel A of Figure A1. We suppress standard errors for the sake of clarity; the drops are significant at the 95% level for between 2-4 months following the shock. If we rerun these regressions while also weighting according to the pre-shock size of firms, we obtain a picture that looks much closer to Figure 7, see Panel B of Figure A1.

These results indicate that the larger firms appear to be affected the most from this shock. This could be partly a result of our proxy being less effective for smaller firms that may not engage in consistent exports to North America.

C.2 Bootstrapping Standard Errors

We use bootstrapping methods to compute measures of the dispersion of our point estimates. Using random sampling with replacement within each group of firms, we create 5000 new artificial samples and re-run the estimation procedure. The standard deviation of the point estimates across these bootstrap samples is shown in Table 4. To gain a more complete picture of the dispersion, we create density estimates for each sample of firms across the parameter space for the elasticities. These densities are shown in Figure A2.

C.3 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.

⁶¹We re-weight the control group as described in section 3.1.

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 3.1, we run the following regression:

$$\Delta \text{emp}_{j,t} = \sum_{i=-3}^3 \gamma_i E_i + \sum_{i=-3}^3 \beta_i E_i D_{j,i} + u_{j,t} \quad (\text{A11})$$

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 A11 using payroll $\text{pay}_{j,t}$ 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 A4.

C.4 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 ($m_{p,j,t} = \log(M_{p,j,t})$):

$$m_{p,j,t} = \alpha_{pj} + \sum_{i=-19}^9 \gamma_i E_i + \sum_{i=-19}^9 \beta_i E_i D_{j,i} + u_{j,t} \quad (\text{A12})$$

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 vari-

ables $D_{j,t}$ 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 A5.

C.5 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 3.1:

$$Q_{i,t} = \alpha_0 + \alpha_i + \sum_{p=-14}^9 \gamma_p E_p + \sum_{p=-14}^9 \beta_p E_p \text{JPN}_{i,p} + u_{i,t} \quad (\text{A13})$$

where here the variable $Q_{i,t}$ is auto production by plant i in month t , after removing a plant-specific trend through March 2011. Because these plants can be tracked with some confidence back in time, it is reasonable here to remove seasonality directly, rather than assume a shared seasonal component between the treated and control groups as in section 3.2. We use the X12-ARIMA model, provided by the National Bank of Belgium, and apply it to each series before correcting for trend. The results for the Japanese plants are mostly similar, as shown in table A6.

Figure A1: Relative Inputs and Output (Proxy) of Japanese Firms (Reduced Sample) Logged, HP-Filtered




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


These figures report the relative percentage deviations from trend of Japanese affiliates relative to a control group of other multinational firms. The values are coefficient estimates taken from an interaction of a Japanese-firm dummy with a monthly dummy – additional baseline monthly dummies remove seasonal effects. These results reflect a reduced sample with no firm-month zeros in imported inputs or N.A. exports. The data is logged, and HP-filtered using a monthly smoothing parameter.

Figure A2: Density Estimates of Elasticities Across Bootstrap Samples

A. Materials Elasticity (ω)



B. Materials-Capital/Labor Elasticity (ζ)



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

Table A4: Dynamic Treatment Effects: Quarterly Employment/Payroll Surrounding Tōhoku Event

Independent Variables	Log 4-Quarter Difference	
	Employment (1)	Payroll (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 A11 in the text. The dependent variable is the four-quarter log difference of employment (payroll).

Table A5: Dynamic Treatment Effects: Unit Values of Trade Surrounding Tōhoku Event


Independent Variables	Log Unit-Value of:			
	JPN Imports: Related Party (1)	JPN Imports: Non-Related Party (2)	N.A. Exports Related Party (3)	N.A. Exports Non-Related Party (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 A12 in the text.

Figure A3: 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 A6: 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 (649.9)	17.78 (608.8)	JPN x Nov 2010 (t=-4)	-195.8 (841.9)	-341.7 (799.2)
Dec 2010 (t=-3)	-1,973*** (467.5)	310.3 (497.5)	JPN x Dec 2010 (t=-3)	-385.0 (736.5)	-408.3 (706.4)
Jan 2011 (t=-2)	-611.5 (637.3)	1,083* (618.7)	JPN x Jan 2011 (t=-2)	781.0 (792.1)	-1,092 (804.6)
Feb 2011 (t=-1)	694.9* (401.9)	756.3* (394.7)	JPN x Feb 2011 (t=-1)	-1,142 (696.2)	-1,210* (666.8)
Mar 2011 (t=0)	4,356*** (524.9)	1,483*** (389.1)	JPN x Mar 2011 (t=0)	-3,515*** (812.0)	-2,592*** (842.7)
Apr 2011 (t=1)	-216.2 (707.7)	305.5 (620.4)	JPN x Apr 2011 (t=1)	-6,239*** (1,303)	-6,099*** (1,282)
May 2011 (t=2)	1,584*** (525.4)	799.1 (511.3)	JPN x May 2011 (t=2)	-7,244*** (1,651)	-6,625*** (1,740)
Jun 2011 (t=3)	1,366** (623.6)	-499.3 (594.9)	JPN x Jun 2011 (t=3)	-4,564*** (1,248)	-3,423** (1,320)
Jul 2011 (t=4)	-4,512*** (878.4)	123.3 (606.2)	JPN x Jul 2011 (t=4)	-2,143 (1,430)	-3,723*** (1,045)
Aug 2011 (t=5)	685.6 (744.0)	-1,323** (648.1)	JPN x Aug 2011 (t=5)	-1,275 (970.8)	-1,108 (1,012)
Sep 2011 (t=6)	-836.5 (663.7)	-1,895*** (641.5)	JPN x Sep 2011 (t=6)	-359.4 (930.7)	40.37 (959.8)
Oct 2011 (t=7)	-338.0 (662.3)	-1,434** (632.4)	JPN x Oct 2011 (t=7)	93.27 (885.6)	-265.4 (785.8)
Nov 2011 (t=8)	-1,393** (582.8)	-1,443** (601.2)	JPN x Nov 2011 (t=8)	-1,318 (1,159)	-2,059* (1,183)
Dec 2011 (t=9)	-4,511*** (774.4)	-1,619** (655.5)	JPN x Dec 2011 (t=9)	759.1 (1,105)	24.95 (803.9)
Constant				-1,535*** (89.30)	-1,683*** (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 .