

Accounting for the Growth of MNC-based Trade Using a Structural Model of U.S. MNCs[†]

Susan E. Feinberg
Robert H. Smith School of Business
University of Maryland
3423 Van Munching Hall
College Park, MD 20742-1815
(301) 405-2251
feinberg@glue.umd.edu

Michael P. Keane
Department of Economics
Yale University
37 Hillhouse Avenue Room 29
New Haven, CT 06520-8264
(203) 432-3548
michael.keane@yale.edu

February 1, 2002
Revised March 29, 2002

[†] The statistical analysis of firm-level data on U.S. multinational corporations reported in this study was conducted at the International Investment Division, Bureau of Economic Analysis, U.S. Department of Commerce, under arrangements that maintained legal confidentiality requirements. Views expressed are those of the authors and do not necessarily reflect those of the Department of Commerce. Suggestions and assistance from William Zeile and Raymond Mataloni are gratefully acknowledged.

I. Introduction

It is well known that trade has been growing more rapidly than GDP in recent decades. Between 1982 and 1994, the growth of U.S. foreign trade (exports plus imports) averaged more than 5% per year in real terms, while real GDP grew 3% per year. According to Rugman (1988), over half of world trade involves large multinational corporations (MNCs). *MNC-based trade* includes two components: *Arms-length trade*, which means shipments between a division of an MNC and unaffiliated buyers/suppliers in other countries, and *intra-firm trade*, which means internal shipments between and MNC parent and its foreign affiliates. Between 1982 and 1984, the total trade of U.S. based MNCs grew an average of 4.5% per year. And the *intra-firm* component grew more rapidly than the *arms-length* component.¹ Indeed, between 1982 and 1994, intra-firm trade increased from 35.5% to 45% of the total trade of US-based MNCs. What explains the rapid growth in MNC-based trade in general, and the even faster growth of intra-firm trade between divisions of MNCs in particular?

In this paper, we estimate a structural model of the behavior of U.S. MNCs and their Canadian affiliates, and use the model to examine the sources of growth in both intra-firm and arms-length MNC based trade. We estimate the model using a confidential disaggregated panel data set from the Bureau of Economic Analysis on more than 500 U.S. MNCs and their Canadian affiliates from 1983-1996. We then use the model as a framework to decompose changes in intra-firm and arms-length trade flows into components due to changes in tariffs, changes in technology and changes in prices.²

There were significant bilateral tariff reductions between the U.S. and Canada during our sample period, arising both from the GATT/WTO and the Canada-U.S. Free Trade Agreement in 1989. And U.S. MNCs have a substantial presence in Canada. Thus, the U.S.-Canada context provides an excellent opportunity to examine the contribution of trade liberalization to the observed increases in MNC based trade.³

¹ For US MNCs, these trade flows grew an average of 3.4% and 6.3% per year, respectively. Note that we do not include in these statistics trade to and from U.S. affiliates of foreign MNCs.

² Our effort is somewhat analogous to the growth accounting literature (Solow, 1957), which is why we use the word “accounting” in the title.

³ As Treffer (2001) points out, most other trade liberalizations episodes that have been studied in the literature were accompanied by packages of economic reforms, so that the pure effect of tariffs is difficult to isolate. To a great extent this was not the case with the FTA, making it an “unusually clean trade policy exercise.”

As Yi (2001) points out, while most observers attribute the rapid growth of world trade since the early 1960s primarily to trade liberalization, standard trade models have difficulty explaining the data based on observed tariff reductions. The problem is that worldwide tariffs on manufactured goods have only declined by about 11 percentage points (on average) since the early 1960s.⁴ Yi argues that vertical specialization – the fragmentation of production processes into parts that are performed in several countries – can help resolve this quantitative puzzle. The idea is that a vertically specialized (or fragmented) production process may require goods in process to cross several international borders. Then, a 1 percentage point general tariff reduction can be magnified, leading to a greater than 1 percent fall in the cost of producing the good. Calibrating a two country (U.S. and Rest-of-world) Ricardian model with vertical specialization, Yi finds that it can explain about half of the growth of world trade since the 1960s based on trade liberalization, which is far more than is explained by standard models.

Inspection of the BEA data on U.S. MNCs and their Canadian affiliates suggests that fragmentation of production is pervasive. Table 1 reports descriptive statistics on 3385 firm year observations of U.S. parents and their Canadian affiliates from 1983-1996. The mean of parent's sales to the affiliate is more than one-third of mean affiliate total sales. And mean affiliate sales back to the parent are 39% of mean affiliate sales. Furthermore, as we'll see below, the extent of fragmentation (as indicated by intra-firm trade in intermediates) has been growing rapidly.

A striking aspect of the data (not shown in Table 1) is the number of cases in which intra-firm flows actually go in both directions - 69% of observations. In 39% of cases the parent exports to Canada and the affiliate also exports to the US. In only 11% of cases are the final sales of both the parent and the affiliate entirely in their own domestic market. Nearly one third of cases have intra-firm flows in both directions and arms-length flows in both directions! Clearly, goods in process often cross the border multiple times.

Unfortunately, existing theories of multinationals have difficulty explaining the multiplicity of MNC production arrangements we observe in the BEA data. Markusen and Maskus (1999a, b) provide an excellent discussion of current theories of the MNC. They divide the theories into those that generate vertical vs. horizontal MNCs. Classic examples of these two types of models are Helpman (1984) and Markusen (1984), respectively. Vertical MNCs

⁴ Elasticities of substitution between home and foreign goods need to be implausibly large for standard models to generate the observed growth in trade (by a factor of 3.4 as a share of world GDP) using an 11 point tariff reduction.

fragment the production process across countries to take advantage of factor price differentials, e.g., by locating unskilled labor intensive parts of the process in low wage countries. Horizontal MNCs basically replicate the entire production process in multiple countries, thus avoiding tariff and transport costs.

In both classes of model, a key reason for existence of the MNC is some firm level fixed cost that serves as a joint input (or “public good”) for MNC operations in all countries, thus creating multi-plant economies of scale. For example, this could be the fixed cost of inventing a production process, which can then be adopted in any number of plants. Or, it could be “headquarters services” as in Helpman (1984). Another example would be a fixed cost of advertising necessary to create brand equity, which is then internationally transferable.⁵

Theories of vertical MNCs have a hard time explaining extensive fragmentation of production across similar countries, like the U.S. and Canada, which presumably have similar factor prices.⁶ Thus, the high degree of fragmentation we observe in the BEA data is itself uncomfortable for these models. Of course, it is even more uncomfortable for horizontal models.

An even more fundamental problem is that the MNC forms we observe in the BEA data do not, for the most part, fall into the neat vertical/horizontal taxonomy that exists in the theories. Most striking is the fact that two-thirds of U.S. manufacturing MNCs have bilateral intra-firm

⁵ Theories of the MNC originated in the international business literature, with the so-called “Ownership-Location-Internalization (OLI) paradigm” – see Hymer (1976) and Dunning (1977). This approach informs the later work by Helpman, Krugman, Markusen, Venables, etc. A key idea in the OLI theory is that foreign operations are less cost efficient than domestic operations (e.g., due to communication and transport costs, language barriers, etc.), so a domestic firm will only be able to compete with foreign firms if it has some compensating advantage. Such an “ownership advantage” might be a high degree of brand equity that is recognized abroad, giving the owner of the brand name a degree of market power. This enables the domestic firm to compete in the foreign market.

Still, even with an ownership advantage, the firm could compete via exports. A “location advantage” gives the firm a reason to locate production abroad. This might be to jump the tariff wall and/or save on transport costs, to take advantage of lower factor prices, or because proximity to final customers is important (e.g., to provide service).

Yet, even this would not be sufficient for a firm to choose multinational operation. It could still contract out operations to foreign firms (e.g., licensing), thereby saving on the extra costs of foreign operation due to language barriers, dealing with local customs, unfamiliarity with local suppliers, etc. A key factor in MNC operations is the internalization advantage - see Rugman 1985. For example, if a firm’s ownership advantage comes from strong brand equity, it runs the risk that a licensee might shirk on quality and diminish the value of its brand name.

In the vertical and horizontal MNC models discussed by Markusen and Maskus, the firm level fixed cost (e.g., cost of inventing a better production process, or the advertising to create the brand name) creates the ownership advantage. The internalization advantage is left implicit, but it arises from the risk of a licensee stealing the process knowledge or diluting the brand name. In the vertical models the location advantage arises from factor price differentials, while in the horizontal models it arises from tariff jumping, the advantage of proximity to final customers, etc.

⁶ But this point is controversial. For instance, Davis and Weinstein (2001) argue that factor price differentials are indeed substantial even among wealthy OECD countries.

flows in intermediates. Furthermore, nearly one-third have bilateral intra-firm flows in intermediates and bilateral arms-length flows simultaneously. And in nearly all cases the parent and affiliate both have final sales to unaffiliated domestic buyers. In a pure vertical or horizontal model of the MNC we would not expect to see so many intra-firm and arms length trade flows going on at the same time.

Models of horizontal MNCs rule out substantial intra-firm flows of intermediates essentially by definition. The central idea of these models is that the common input that is located in the home country (e.g., “knowledge,” “brand equity”) can be transferred essentially costlessly across borders. In their simple form, horizontal models also do not generate arms length sales of final goods by the parent to the host country or by the affiliate back to the home country.⁷ But these models can be simply modified to generate such flows by assuming the final goods produced by the parent and affiliate are differentiated.⁸ Given that, about 12% of the firms in the BEA data appear to be “horizontal” MNCs.

The vertical MNC models in the literature, such as Helpman (1984, 1985) and Helpman and Krugman (1985) have two basic problems. First, they generate intra-firm trade in intermediates in (at most) one direction. Second, final goods are produced only by the parent or the affiliate (whichever receives the intermediates), but not both. For instance, in Helpman and Krugman, MNCs produce a differentiated product in three stages that require descending levels of capital intensity: headquarters services, production of intermediates and production of final goods. Suppose the factor prices are such that headquarters services and intermediate goods production are located in the home country and final assembly is done by the foreign affiliate. This leads to a one-way flow of intermediates from parent to affiliate, and sales of the final good by the affiliate in both the host and home countries.

⁷ This problem with horizontal MNC models has previously been noted in the aggregate industry/country level data. Rugman (1985) pointed out that in industry/country pairs where U.S. MNCs have substantial direct investment, the U.S. also tends to have substantial exports. And that, in these same industries, foreign firms tend to have both substantial direct investment and substantial exports to the U.S..

⁸ Recently, Baldwin and Ottaviano (1998) have extended the horizontal MNC model in a way that can generate bilateral arms length sales of final goods (and thereby explain the similarity between world investment and trade patterns noted earlier). Their model extends Brander’s (1981) model of intra-industry trade in similar goods to also generate intra-industry FDI. In their model, both the parent and the (horizontal) affiliate produce different varieties of a differentiated product. The key idea is that it is desirable to produce some varieties abroad because tariffs and transport costs partially shield home produced varieties from cannibalization by foreign produced varieties. Then, there will also be an incentive for both the parent and the affiliate to sell arms-length (i.e., export and import), as they are willing to accept lower price/cost margins on exports than on domestic sales. This still leaves the problem that the horizontal model does not generate intra-firm flows in intermediates however.

No firm in the BEA data literally follows this pattern. Almost all parents and affiliates are observed to sell final goods. But this problem is easy to fix. As Helpman (1985) notes, if one adds a distribution/marketing stage that must be tied to the point of sale, then both parents and affiliates will have final good sales. Simultaneous arms-length sales of final goods by both the parent and the affiliate seems more difficult to generate, but this could be achieved if the intermediate good can also be sold arms-length.

The more difficult problem for vertical models is to generate bilateral intra-firm flows of intermediates. Helpman and Krugman noted this problem. However, they also noted that in aggregate data, intermediates and final goods often share the same industry code, so their model would appear to generate two-way intra-industry trade. At the firm level, however, this does not resolve the problem. If we define vertical MNCs only as those that exhibit a one-way flow in intermediates, only 19% of the firms in the BEA data are vertical.⁹

An obvious way to generate bilateral intra-firm trade in intermediates in a vertical model is to fragment the production process into additional stages. Deardorff (1998) discusses this possibility.¹⁰ As long as there is not a strictly descending or ascending ordering of stages by factor intensity, this creates the possibility for bilateral flows of intermediates. This situation is illustrated in Figure 1.

The problem with such a multi-stage production process is that it would not be feasible to estimate with available (or prospectively available) data. For example, in the BEA we only know total employment for the parent and the affiliate. We would have no way to determine how much

⁹ The fact that existing theoretical models of MNCs are inconsistent with observed features of the individual MNCs we observe in the BEA data is not surprising, for at least three reasons. First, these data are confidential, and have only recently been made available to a handful of researchers. So it has only recently become possible to carefully document the production and sales patterns of individual MNCs. Second, primary the goal of existing theoretical work that models MNCs is not actually to predict the behavior of individual MNCs. Rather, the key goal of this work has been to predict patterns in aggregate intra-industry trade. For instance, Helpman and Krugman (1985) incorporate MNCs in their model specifically to help explain certain patterns in the aggregate composition of trade. Third, another goal of prior work on MNCs has been to mathematically formalize the conditions under which firms will choose to operate as MNCs as opposed to exporting or licensing. Models along these lines are, not surprisingly, rather stylized.

¹⁰ Deardorff also provides an excellent discussion of the causes and consequences of fragmentation in standard Ricardian and Heckscher-Ohlin trade models. He starts from the assumption that fragmentation is not itself a technical advance (i.e., geographic splitting up of parts of the production process does not make the process intrinsically more efficient). If fragmentation becomes technically feasible (say, due to improved communication or transportation), then it may be implemented to exploit factor price differentials. That is, if different parts of the production process involve different factor intensities, it may be optimal for a country to specialize in those fragments of the process where it has comparative advantage. But fragmentation of a production process will not occur in a world where factor price equalization has already been achieved.

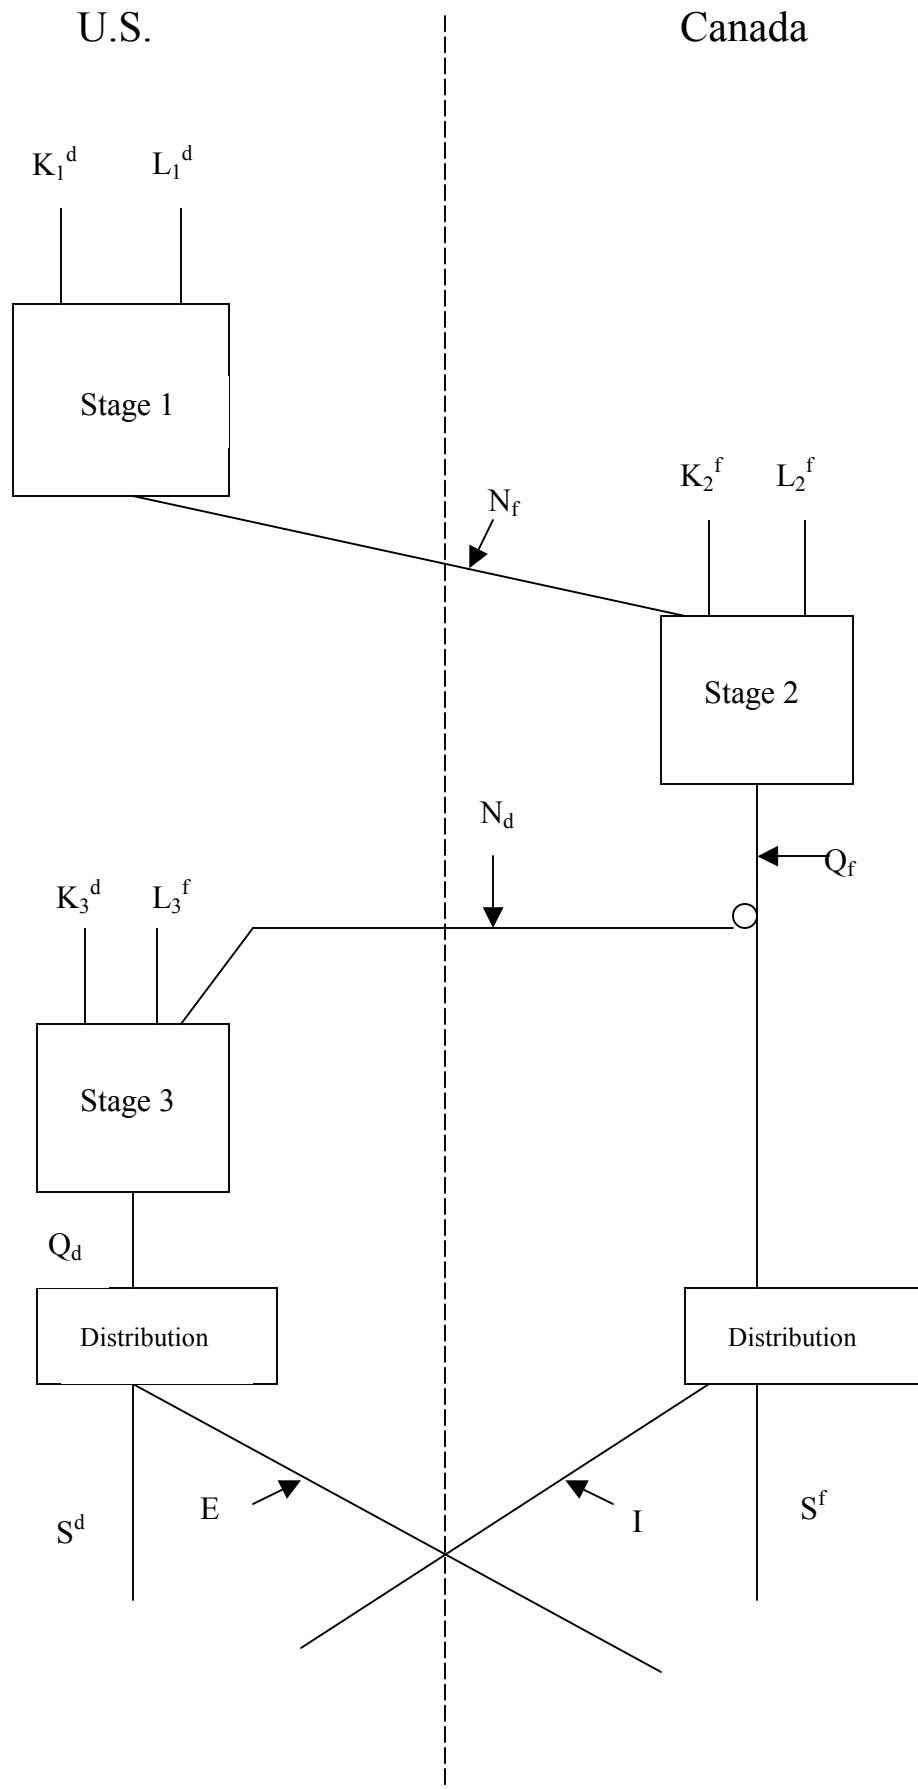


Figure 1: Multi-Stage Production Process

labor input the parent used in stage 1 vs. stage 3 of the production process in Figure 1. Furthermore, we only observe total values of the intra-firm flows, and not what stages of the production process these flows feed into. Thus, estimation of a multi-stage process appears to be hopeless.

Instead we assume what might be called “holistic” or “circular” (as opposed to vertical) specialization. That is, we assume the final good produced by the parent may be needed as an intermediate input into the affiliate’s production process, and that simultaneously, the final good produced by the affiliate may be needed as an intermediate input into the parent’s production process.

An example of this type of circular structure is provided by the petroleum industry. Finished products from a petroleum refinery include fuel oil, lubricants and Naphtha. In turn, these products are used in oil drilling. Consider Mobil’s off shore drilling operation in the Hibernia field off Newfoundland. Since Mobil had little refining capacity in Canada, the crude was primarily shipped to Mobile refineries in Texas and Louisiana. Lubricants and fuel oil were then shipped back to Hibernia as inputs to run the rig. Fuel oil requirements are substantial in off shore oil drilling, since all power must be generated at the rig. Similarly, Venezuelan crude oil from the Oronoco flow is extremely heavy. Naphtha produced at Mobil refineries in Louisiana is shipped to Venezuela and pumped into the crude oil flow to facilitate extraction.¹¹

While we do not wish to claim that such circular linkages are a major source of observed intra-firm flows - we have no way to quantify their importance given available data. But this assumption on technology appears to be the simplest, most direct and most empirically tractable way to generate bilateral intra-firm flows in intermediates, which we observe for 69% of cases in the BEA data.¹²

¹¹ I thank Bob Thomson (Ph.D. Ch.E. Notre Dame ’88) for these examples. Another example of this type of circular structure is in the design of computer hardware and software. The architectural components of a computer are designed using simulation software that is run on a computer (see, e.g., the JSim software for Course 6.004 at MIT). More generally, it is rather common in the electrical machinery industry that various types of electronic equipment are needed to build the components that make up that equipment.

¹² Suppose that the true data generating process is a multi-stage fragmented production process, with multiple border crossings by intermediates at different stages. Or that the parent and affiliate are multi-product firms, and the intermediates they ship intra-firm and the final goods they sell to consumers (or other firms) are completely different products. We cannot identify these situations from the data, so in a sense our production process is a reduced form representation of these underlying production structures. This may cause the model to break down for certain policy experiments where the underlying structure inside the black box of the aggregated production function matters. But that is generally true when one aggregates.

Another key problem we confront is that, while the BEA data on MNCs is in some ways very rich, many basic quantities that one would like to use to estimate a structural model of MNC behavior are not measured. For instance, the BEA data does not allow one to separate quantities of production from prices. Similarly, while one can observe values of intra-firm flows, imports and exports – the prices and quantities for these cannot be observed separately. Data on the capital stock, and even the capital share, appears to be quite unreliable as well. Given the lack of information on quantities of inputs and outputs, it would be very difficult to estimate plant level returns to scale, the role of intangible latent inputs, or, for that matter, any very general production technology.

In this paper, we take a very different approach from the prior theoretical and empirical literature on MNCs. First, we specify a production function in which all productive inputs are (at least in principle) observable. Second, we specify a simple, constant returns to scale, Cobb-Douglas production technology. Given these assumptions, we show that all the parameters of a structural model of MNC behavior (i.e., parameters of the production technology) can be identified without separate data on quantities and prices of output, intra-firm flows, imports and exports, or data on capital, provided we make the following additional strong assumptions:

- 1) There are no adjustment costs for capital, and in each period (year) the supply of capital is allocated across firms in such a way that the rate of profit (i.e., the ratio of profit to value of capital) is equalized in expected value.
- 2) The demand functions that confront firms are characterized by constant price elasticity of demand.

Given these assumptions, we estimate the first model of MNC behavior that is rich enough to account for most of the major decisions of MNCs (in a two country setting): capital and labor input in each country, output in each country, the volume of intra-firm flows, and the volume of arms-length imports and exports (i.e., direct sales from the U.S. parent to the foreign market and direct sales from the foreign affiliate to the U.S. market). To our knowledge, the only prior structural modeling of MNCs (Cummins, 1995, Ihrig 2000) is limited to modeling investment or repatriation decisions, and does not consider the production and trade aspects of MNC behavior.¹³

¹³ On the other hand, Cummins' model of investment decisions is much richer – including capital adjustment costs that are allowed to be interrelated across the parent and affiliate. Note that while the focus of Ihrig's work is on

Our main findings are as follows: First we find there is a clear bifurcation of MNCs into two types: those that have intra-firm flows and those that do not. The increase in the volume of intra-firm flows over the sample period does not occur because more MNCs organized production in such a way that they had intra-firm flows. Rather, it occurred for those firms that already had positive flows to begin with. In other words, the increase in intra-firm trade has occurred almost entirely on the intensive rather than the extensive margin.

Second, we find that tariff reductions explain a substantial fraction of the increase in intra-firm flows, as well as in arms-length trade, but they cannot explain nearly all of the increase. Technical change, in the form of changes in the share parameters of the Cobb-Douglas production function, accounts for a large portion of the increases in trade.

According to our estimates, those MNCs whose initial technological state was such that they were organized to have intra-firm flows, also experienced technical change such that it was optimal for them to substantially increase intra-firm flows. In contrast, for those MNCs who do not have intra-firm flows, our model is able to explain their behavior well using essentially fixed share parameters. Our estimates imply that all MNCs experienced substantial general technical progress (about 4.5 percent per year).

Third, steady state simulation of our model implies that the size of MNCs and the volume of intra-firm flows is being restrained by adjustment costs. Our model implies that, holding tariffs, technology and prices fixed, intra-firm flows will increase on the order of 30 percent when steady state is reached.

Finally, an interesting policy experiment is to compare the baseline levels of intra-firm and arms-length trade flows and domestic and foreign labor in the last year of our sample (1995), given actual tariffs in that year, to the levels of labor and flows MNCs would have chosen had tariffs remained at their first year (1984) level. Similarly, we compare the baseline levels of tariffs and trade flows with levels that would have obtained if tariffs were eliminated completely. Using steady state simulations, we find that under the first scenario (in which tariffs are not reduced), MNCs would have hired 8.5% less foreign labor and 1.21% less domestic labor and sent 19% fewer arms-length exports from the US to Canada and 30.7% fewer arms-length exports from Canada to the U.S. Similarly, MNCs would have chosen to ship significantly

decisions about repatriation of profits, this is one major decision we abstract from (i.e., we assume all profits are repatriated each year).

smaller volumes of goods intra-firm under this scenario. If tariffs were eliminated completely, MNCs would have chosen to ship 2.1% more goods intra-firm from the Canadian affiliate to the U.S. parent, and 6.6% more intra-firm shipments from the U.S. parent to the Canadian affiliate.

Our work also makes four methodological/econometric contributions. Dynamics in our model arise from labor force adjustment costs. Thus, solution of our model using a full solution method would require us to specify how firms form expectations of future labor force size (which in turn depends on their forecasts of future demand, technology, prices, tariffs, exchange rates, etc.). To avoid this, we adopt the common approach of working with the Euler conditions of the firm's optimization problem. However, rather than estimating these using GMM, we instead use simulated maximum likelihood. This approach requires us to specify a distribution for forecast errors, which we assume are normally distributed (subject to a Box-Cox transformation). Such an approach has typically been avoided because there is no obvious basis for assuming a distribution on forecast errors. A similar approach was recently successfully adopted by Krussell, Ohanian, Rios-Riull and Violante (2000). We go beyond their econometric contribution in a number of ways.

Our first innovation is to show how to test the distributional assumption by simulating the posterior distributions of forecast errors conditional on our estimated model and the data. We find that these distributions appear strikingly close to normality. Statistical tests do not reject normality even at very low levels of significance.

The second innovation is the introduction of a new probability simulator. Estimation of our model requires a new recursive probability simulator that is the discrete/continuous analogue to the GHK method for simulation of discrete choice probabilities.

Third, our model generates a likelihood in which the Jacobian of the transformation from the stochastic terms to the data is intractably complex (i.e., we cannot even write the likelihood). Nevertheless, we show how to construct a numerical approximation to the Jacobian, which also requires simulation.

Fourth, we develop yet another new recursive simulator for purely continuous distributions that is required to generate the posterior distribution of the stochastic terms in our model (i.e., shocks to technology and demand) conditional on the data.

Because we do not use a full solution method (i.e., we estimate the model using first order conditions) we cannot simulate data from the model. But, since we can simulate the

posterior distribution of the stochastic terms in the model, we can evaluate how well the model fits by evaluating how well the shape of these posterior distributions match up with the shapes we assume. In this sense the model appears to fit extremely well, in that in only a few cases are our distributional assumptions rejected. Also, although we cannot simulate data from our dynamic model, we can simulate the steady state behavior predicted by the model (i.e., simply by setting the adjustment cost parameters to zero we obtain a static model that we can also solve for optimal firm behavior).

The remainder of this paper is organized as follows. We review some prior empirical work on MNCs and on the U.S.-Canada FTA in section 2. The model and estimation techniques are discussed in section 3, and the construction of our dataset is described in section 4. We describe the results of our structural estimation in section 5 and the policy experiments from steady-state simulations of our model in section 6. Section 7 concludes.

II. Review of Empirical Work on MNCs and the U.S. Canada FTA

The body of empirical work on MNCs and on the U.S. Canada FTA is rather extensive, so this review is quite selective, and limited to some recent studies of direct relevance to the present investigation.

Much recent empirical work on MNCs has focused on the issue of whether the vertical or horizontal model of MNCs is supported by the data. Authors who have studied data on foreign direct investment and international trade flows at the industry/country level generally conclude that horizontal MNCs are the dominant form. Markusen and Maskus (1999a, b) argue that: “the overwhelming proportion of world direct investment is from high-income developed countries to other similar high-income developed countries. This suggests that horizontal investment is much more important ... than vertical investment ... motivated by factor endowment differences,” the point being that endowment differences are presumably small among the high-income countries. However, Davis and Weinstein (2001) argue that endowment and factor price differences are substantial even within wealthy OECD countries, and state that “classical comparative advantage is likely to be quite important in the North.” Nevertheless, Brainard (1993) finds little relation between factor endowments and patterns of FDI.

Brainard (1997) examines sales of U.S. affiliates abroad, and reaches conclusions similar to Markusen and Maskus. Using the BEA industry level data for 1989, she notes that 64% of the

affiliates' sales are in the host country (on average), while only 13% are back to the U.S. She states that: "affiliate production destined for export ... back to the home country ... is the activity ... associated with vertical integration ... while local sales in the foreign market are more likely to be associated with horizontal expansion."¹⁴

On the other hand, Hanson, Mataloni and Slaughter (2001) challenge this view, arguing that a closer look at the BEA data on U.S. affiliates abroad provides strong evidence that vertical fragmentation is important and growing. For instance, in manufacturing, affiliate value added as a fraction of total affiliate sales fell from 37% in 1982 to 29% in 1994, suggesting that affiliates are becoming less like stand alone (horizontal) operations. Similarly, affiliate imports of intermediate goods as a share of total sales rose from 9.8% in 1982 to 12.2% in 1994. Interestingly, the figures for Canada are much higher (21.6% and 33.5%, respectively). And manufacturing affiliate exports as a share of total sales rose from 33.9% in 1992 to 44.4% in 1998. Most obviously, using the firm level data, they observe that more than a third of manufacturing affiliates are not in the same industry as the parent.

Our work follows the line of Hanson, Mataloni and Slaughter (2001) in that we provide an even more detailed analysis of the BEA firm level data. But we differ in that our analysis is focused exclusively on Canadian affiliates.

There have been a number of prior studies of the impact of the U.S.-Canada FTA on Canadian industry. Gaston and Trefler (1997) and Trefler (2001) use industry level data to identify the effects of tariff reductions on Canadian manufacturing industries. Their identification strategy exploits the fact that tariff decreases differed substantially across industries. Trefler (2001) concludes that the FTA reduced employment of production workers, increased earnings of production workers, reduced output, and raised labor productivity. He concludes that the FTA did not affect either the number or scale of Canadian manufacturing plants. Head and Ries (1997) show that theoretical predictions regarding effects of tariffs on the number and scale of manufacturing plants are quite ambiguous. Using a similar basic identification strategy to that of Trefler (i.e., comparison of changes across industries with different degrees of tariff reduction) they estimate that the FTA had little effect on the scale of Canadian manufacturing plants.

¹⁴ In the Canadian case, however, she reports that 28% of affiliate sales were back to the U.S., so the case is not so clear.

The role of tariffs in determining the scale of Canadian manufacturing plants has been a key issue for many years. Many authors have argued that, as a result of tariff protection, Canadian manufacturing plants operated at an inefficiently small scale – see Eastman and Stykolt (1967), Caves (1984, 1986), Baldwin and Gorecki (1986), Caves (1990). This led to concern that the FTA would lead to “hollowing out” of Canadian industry, because U.S. MNCs could most efficiently serve both markets from large U.S. plants. In our own previous work using the BEA firm level data – Feinberg, Keane and Bognanno (1998), Feinberg and Keane (2001) we found no evidence of such effects. In fact, we found evidence that assets, employment and value added of Canadian affiliates of U.S. MNCs actually increased following tariff reductions. We also found evidence that intra-firm flows, particularly flows from the affiliate to the parent, were increased following the tariff reductions. This suggests that the fragmentation of production was increased.

At first our earlier results may appear to contradict those of Gaston and Trefler, since they find (small) negative employment effects while we find positive effects. But it is important to bear in mind that their results are for all of Canadian manufacturing while our results are only for affiliates of U.S. MNCs. Since tariffs are a tax on internal flows within MNCs, it would not be surprising if tariff reductions benefited MNCs relative to national firms.

III. The Model

We do not model the MNCs decision to place an affiliate in Canada, which we take as exogenous. We shall simply assume there exists some firm level fixed cost that generates multi-plant economies of scale, and that this fixed cost does not affect the MNCs marginal production decisions. For instance we could assume that firms have market power by virtue of a fixed investment they have made in establishing brand names or inventing a proprietary production process, and that licensing out of foreign operations would potentially dissipate brand equity or risk revelation of proprietary knowledge.¹⁵

To motivate bilateral arms-length trade in finished goods, and domestic sales by both the parent and the affiliate, we assume that the parent and the affiliate produce two different goods.

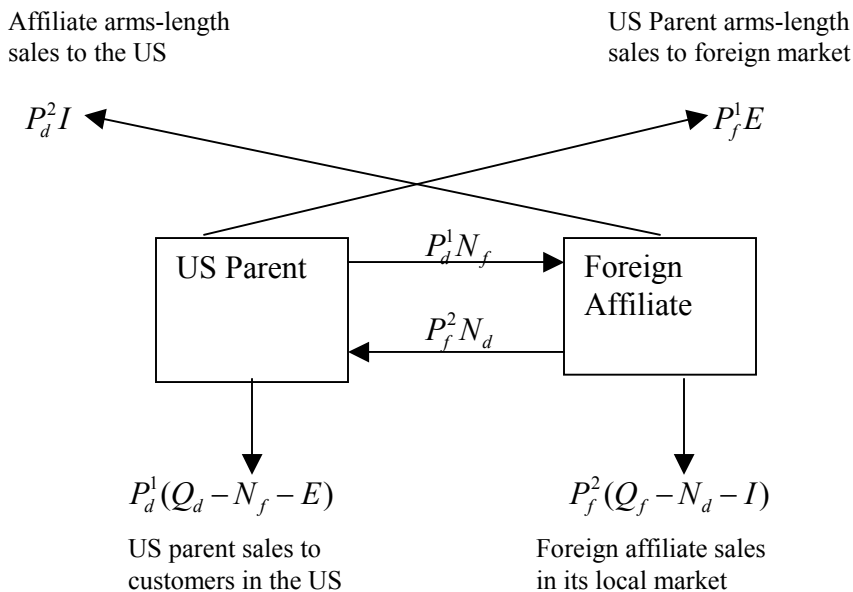
¹⁵ This assumption is consistent with Markusen’s (1995) discussion of characteristics of industries with significant multinational production. Specifically, multinationals tend to be important in industries with high R&D/sales, technically complex products, high levels of product differentiation and advertising, and a large share of skilled workers in their workforces.

Each is a variety of a non-rival differentiated product, so that both the parent and affiliate have market power. To motivate bilateral intra-firm trade in intermediates between the MNC parent and its foreign subsidiary, we assume that the parent may use final goods from the affiliate's production process as an input into its own production process, and vice-versa.

We will first specify the most general case, in which a single MNC exhibits all 4 of the potential trade flows. Instances where an MNC has only a subset of these 4 flows are special cases. The MNC will decide on which of the 4 flows to utilize. We estimate reduced form decision rules for the utilization of each flow, and estimate these jointly with the structural model of the production process.

III.1. Mathematical Structure of the Model

Figure 2: MNC trade flows in the model



Let I (imports) denote the quantity of goods sold arms-length by the Canadian affiliate to consumers in the U.S., and let E (exports) denote arms-length exports from the U.S. parent to consumers in Canada. Next, let Q_d and Q_f denote output of the parent and affiliate, respectively. Of affiliate output, let N_d denote the part transferred for use as intermediates by the parent. Similarly, let N_f denote intermediates transferred from the parent to the affiliate. Then, $(Q_d - N_f - E)$ is the quantity of its output the parent sells in the U.S., and $(Q_f - N_d - I)$ is the quantity of its output the affiliate sells in Canada. Finally, the P denote prices, with the superscript $j=1,2$ denoting the

good (i.e., that produced by the parent or the affiliate) and the subscript $c=d,f$ denoting the point of sale.

Note that the price on N_d , the intermediate good shipped from the affiliate to the parent, is set equal to the price on the same good if sold to final customers in Canada (and vice versa for N_f). Thus, we ignore issues of transfer pricing. We hope that the corporate tax structures of the U.S. and Canada are sufficiently similar that this assumption will not do too much violence to the data.¹⁶

We do not observe prices and quantities separately in the data. Hence, we will work with six MNC firm-level trade and domestic sales flows, which are: $P_d^2 I, P_f^1 E, P_d^1 N_f, P_f^2 N_d,$

$P_d^1(Q_d - N_f - E)$ and $P_f^2(Q_f - N_d - I)$. These are shown in Figure 2.

We assume that the MNC's domestic and foreign production functions are Cobb-Douglas, given by:

$$(1) \quad Q_d = A_d K_d^{\alpha_{Kd}} L_d^{\alpha_{Ld}} N_d^{\alpha_{Nd}} M_d^{\alpha_{Md}}$$

$$(2) \quad Q_f = A_f K_f^{\alpha_{Kf}} L_f^{\alpha_{Lf}} N_f^{\alpha_{Nf}} M_f^{\alpha_{Mf}}$$

Note that there are four factor inputs: capital (K), labor (L), intermediate goods (N) and materials (M), where the intermediate goods are the quantities of product shipped intra-firm to the parent from the foreign affiliate (in 2) and vice versa (in 1).

If $N_d = \phi_{Nd} Q_f$ and $N_f = \phi_{Nf} Q_d$ then:

$$Q_d = A_d K_d^{\alpha_{Kd}} L_d^{\alpha_{Ld}} M_d^{\alpha_{Md}} \left(\phi_{Nd} A_f K_f^{\alpha_{Kf}} L_f^{\alpha_{Lf}} M_f^{\alpha_{Mf}} N_f^{\alpha_{Nf}} \right)^{\alpha_{Nd}}.$$

Substituting for N_f and solving for Q_d we can express (1) and (2) as:

$$(3) \quad Q_d = \left\{ \left[A_d A_f^{\alpha_{Nd}} \phi_{Nd}^{\alpha_{Nd}} \phi_{Nf}^{\alpha_{Nd} \alpha_{Nf}} \right] \left(K_d^{\alpha_{Kd}} L_d^{\alpha_{Ld}} M_d^{\alpha_{Md}} \right) \left(K_f^{\alpha_{Kf}} L_f^{\alpha_{Lf}} M_f^{\alpha_{Mf}} \right)^{\alpha_{Nd}} \right\}^{\frac{1}{1 - \alpha_{Nf} \alpha_{Nd}}}$$

Thus, the ‘‘aggregate’’ technology for the MNC as a whole exhibits constant returns to scale (RTS) provided that ϕ_{Nd} and ϕ_{Nf} are fixed as capital, labor and material inputs vary. However, the aggregate technology may exhibit increasing or decreasing RTS if ϕ_{Nd} and/or ϕ_{Nf} vary along

¹⁶ Since we do not observe prices and quantities separately, it does not seem possible to model transfer pricing.

with capital, labor and material when input prices, output prices and/or tariffs vary. We would expect tariff reductions to both lead to increases in the size of MNCs (increases in K, L, M inputs) and lead to increases in the ϕ . Thus, these may make it appear MNCs have increasing RTS if an aggregate production function rather than separate parent and affiliate production functions are estimated.

For the domestic- and foreign-produced goods, the MNC faces the following constant price elasticity demand functions:

$$\begin{aligned} P_d^1 &= P_{\circ d}^1 S_d^{-g_1} & \therefore \ln S_d &= \frac{1}{g_1} \ln P_{\circ d}^1 - \frac{1}{g_1} \ln P_d^1 \\ P_f^2 &= P_{\circ f}^2 S_f^{-g_2} & \therefore \ln S_f &= \frac{1}{g_2} \ln P_{\circ f}^2 - \frac{1}{g_2} \ln P_f^2 \end{aligned}$$

where S_d denotes the quantity of the U.S. produced good sold in the U.S., S_d denotes the quantity of affiliate sales in Canada, and g_1 and g_2 are the inverses of the price elasticities of demand for the domestic and foreign produced good, respectively.

The domestic and foreign produced good can also be exported and imported.

$P_f^1 = P_{\circ f}^1 E^{-g_1}$ is the price of the domestic good (good #1) in Canada. We assume that the price elasticity of demand for the good $\left(-\frac{1}{g_1}\right)$ is the same in Canada, but that the demand function intercept differs $\left(P_{\circ f}^1\right)$. Similarly, the price of the foreign good (good #2) in the US is represented by $P_d^2 = P_{\circ d}^2 I^{-g_2}$.

We assume the MNC faces labor force adjustment costs. It is often assumed such costs are quadratic, e.g.: $AC_{dt} = \delta_d [L_{dt} - L_{d,t-1}]^2$, where $\delta_d > 0$. However, we found that a generalization of this function led to a substantial improvement in fit and could accommodate many reasonable adjustment cost processes¹⁷:

$$AC_{dt} = \delta_d \left((L_{dt} - L_{d,t-1})^2 \right)^\mu / L_{d,t-1}^\Delta \quad \text{where } \delta_d > 0, \mu > 0, \Delta \geq 0.$$

¹⁷ For example, setting $\mu = 1$ and $\Delta = 0$ produces $\delta_d [L_{dt} - L_{d,t-1}]^2$. Similarly, $\mu = \frac{1}{2}$ and $\Delta = 1$ gives $\delta_d |(L_{dt} - L_{d,t-1}) / L_{d,t-1}|$.

Having specified the domestic and foreign production functions, we can write the MNC's period specific profits as:

$$\begin{aligned}\Pi &= P_d^1(Q_d - N_f - E) - P_d^1 N_f(T_f + C_f) + P_f^1 E(I - T_f - C_f) \\ &+ P_f^2(Q_d - N_d - I) - P_f^2 N_d(T_d + C_d) + P_d^2 I(I - T_d - C_d) \\ &- w_d L_d - w_f L_f - \phi_d M_d - \phi_f M_f - \gamma_d K_d - \gamma_f K_f \\ &- AC_d(L_d, L_d^{(-1)}) - AC_f(L_f, L_f^{(-1)})\end{aligned}$$

Here, T_f and C_f are the foreign tariff and transportation costs the MNC faces when shipping products from the U.S. to Canada (and similarly for T_d and C_d).

The MNC maximizes $E \sum_{\tau=1}^{\infty} \beta^\tau \Pi_{t+\tau}$ by choice of eight variables $\{L_{dt}, M_{dt}, K_{dt}, N_{dt}, L_{ft}, M_{ft}, K_{ft}, N_{ft}\}$.

The first order conditions for parent factor inputs and parent's exports to Canada are:

$$L_d : \alpha^{Ld} \left(\frac{P_d^1 Q_d}{L_d} \right) - g_1 \alpha^{Ld} \left(\frac{P_d^1 Q_d}{L_d} \right) \left(\frac{P_d^1(Q_d - N_f - E) - P_d^1 N_f(T_f + C_f)}{P_d^1(Q_d - N_f - E)} \right) - w_d - \left[\frac{\partial AC_d}{\partial L_d} - \beta E \frac{\partial AC_d^{(+1)}}{\partial L_d} \right] = 0$$

$$K_d : \alpha^{Kd} \left(\frac{P_d^1 Q_d}{K_d} \right) - g_1 \alpha^{Kd} \left(\frac{P_d^1 Q_d}{K_d} \right) \left(\frac{P_d^1(Q_d - N_f - E) - P_d^1 N_f(T_f + C_f)}{P_d^1(Q_d - N_f - E)} \right) - \gamma_d = 0$$

$$M_d : \alpha^{Md} \left(\frac{P_d^1 Q_d}{M_d} \right) - g_1 \alpha^{Md} \left(\frac{P_d^1 Q_d}{M_d} \right) \left(\frac{P_d^1(Q_d - N_f - E) - P_d^1 N_f(T_f + C_f)}{P_d^1(Q_d - N_f - E)} \right) - \phi_d = 0$$

$$\begin{aligned}N_d : \alpha^{Nd} \left(\frac{P_d^1 Q_d}{N_d} \right) - g_1 \alpha^{Nd} \left(\frac{P_d^1 Q_d}{N_d} \right) \left(\frac{P_d^1(Q_d - N_f - E) - P_d^1 N_f(T_f + C_f)}{P_d^1(Q_d - N_f - E)} \right) \\ + g_2 P_f^2 \left(\frac{P_f^2(Q_f - N_d - I) - P_f^2 N_d(T_d + C_d)}{P_f^2(Q_f - N_d - I)} \right) - (1 + T_d + C_d) P_f^2 = 0\end{aligned}$$

$$E : P_d^1 + g_1 P_d^1 \left(\frac{P_d^1(Q_d - N_f - E) - P_d^1 N_f(T_f + C_f)}{P_d^1(Q_d - N_f - E)} \right) + (1 - g_1) P_f^1 (1 - T_f - C_f) = 0$$

The first order conditions for L_f , K_f , M_f , N_f and I are similar. By defining:

$$A = \left(\frac{P_d^1(Q_d - N_f - E) - P_d^1 N_f(T_f + C_f)}{P_d^1(Q_d - N_f - E)} \right) = \left(\frac{P_d^1 S_d - P_d^1 N_f(T_f + C_f)}{P_d^1 S_d} \right)$$

$$B = \left(\frac{P_f^2(Q_f - N_d - I) - P_f^2 N_d(T_d + C_d)}{P_f^2(Q_f - N_d - I)} \right) = \left(\frac{P_f^2 S_f - P_f^2 N_d(T_d + C_d)}{P_f^2 S_f} \right)$$

$$FD = \partial \mu ((L_{dt} - L_{d,t-1})^2)^{\mu-1} (L_{dt} - L_{d,t-1}) / L_{d,t-1}^{\Delta} \\ - \beta \mu ((L_{dt+1} - L_{dt})^2)^{\mu-1} (L_{dt+1} - L_{dt}) / L_{dt}^{\Delta} - \beta \Delta ((L_{dt+1} - L_{dt})^2)^{\mu} (L_{dt+1} - L_{dt}) / L_{dt}^{\Delta+1}$$

we can write the first order conditions more compactly as:

$$L_d : \alpha^{Ld} (1 - g_1 A) \left(\frac{P_d^1 Q_d}{L_d} \right) - w_d - E(FD) = 0$$

$$K_d : \alpha^{Kd} (1 - g_1 A) \left(\frac{P_d^1 Q_d}{K_d} \right) - \gamma_d = 0$$

$$M_d : \alpha^{Md} (1 - g_1 A) \left(\frac{P_d^1 Q_d}{M_d} \right) - \phi_d = 0$$

$$N_d : \alpha^{Nd} (1 - g_1 A) \left(\frac{P_d^1 Q_d}{N_d} \right) + g_2 P_f^2 B - (1 + T_d + C_d) P_f^2 = 0$$

$$E : (1 - g_1) P_f^1 (1 - T_f - C_f) - (1 - g_1 A) P_d^1 = 0$$

The first order condition for N_d equates the marginal revenue product from increasing the input of N_d in domestic production, to the effective cost of importing N_d , which is

$(1 - g_2 B) P_f^2 + (T_d + C_d) P_f^2$. The first term is lost revenue from failing to sell N_d in Canada, the second term is the tariff and transport cost. Similarly, the first order condition for E equates the marginal revenue from exports of the domestically produced good and the marginal revenue of selling it domestically.

Since prices and quantities are not separately observed, we cannot take these FOCs directly to the data. We now describe how the technology parameters can be identified using only data on the nominal values of the six trade and sales flows. If we multiply each first order condition by the associated control variable, we obtain:

$$L_d : \alpha^{Ld} (1 - g_1 A) (P_d^1 Q_d) + -w_d L_d - \delta_d E(FD) L_d = 0$$

$$K_d : \alpha^{Kd} (1 - g_1 A) (P_d^1 Q_d) - \gamma_d K_d = 0$$

$$M_d : \alpha^{Md} (1 - g_1 A) (P_d^1 Q_d) - \phi M_d = 0$$

$$N_d : \alpha^{Nd} (1 - g_1 A) (P_d^1 Q_d) + g_2 (P_f^2 N_d) B - (1 + T_d + C_d) (P_f^2 N_d) = 0$$

$$E : (1 - g_1) (P_f^1 E) (1 - T_f - C_f) - (1 - g_1 A) (P_d^1 E) = 0$$

$$\begin{aligned}
&= (1 - g_1)(P_f^1 E)(1 - T_f - C_f) - (1 - g_1 A) \left(\frac{P_d^1}{P_f^1} \right) (P_f^1 E) = 0 \\
&= (1 - g_1)(P_f^1 E)(1 - T_f - C_f) - (1 - g_1 A) \left(\frac{P_{od}^1}{P_{of}^1} \right) \left(\frac{P_d^1 S_d}{P_f^1 E} \right)^{-g_1} \cdot (P_f^1 E) = 0
\end{aligned}$$

Note that in the first order condition for E, the endogenous quantity ($P_d^1 E$) is not observable.¹⁸

Thus, in the second and third lines of the equation for E, we have exploited the fact that:

$$(P_d^1 E) = \left(\frac{P_d^1}{P_f^1} \right) (P_f^1 E) = \left(\frac{P_{od}^1}{P_{of}^1} \right) \left(\frac{P_d^1 S_d}{P_f^1 E} \right)^{-g_1} \cdot (P_f^1 E)$$

to express the first order condition for E in terms of observable quantities and the ratio of

$$\text{demand function parameters: } \left(\frac{P_{od}^1}{P_{of}^1} \right).$$

Similarly, in the FOCs for the factor inputs, the quantity ($P_d^1 Q_d$) is also not observed. We

can rewrite this quantity as $P_d^1 Q_d = P_d^1 S_d + P_d^1 N_f + \left(\frac{P_d^1}{P_f^1} \right) (P_f^1 E)$ but, again, $P_d^1 E$ is not

observed. We therefore repeat the same type of substitution to obtain:

$$L_d : \alpha^{Ld} (1 - g_1 A) \left[P_d^1 S_d + P_d^1 N_f + \left(\frac{P_{od}^1}{P_{of}^1} \right) \left(\frac{P_d^1 S_d}{P_f^1 E} \right)^{-g_1} (P_f^1 E) \right] - w_d L_d - \delta_d E (FD) L_d = 0$$

and similarly for K_d , M_d and N_d .

We now have a total of eight first order conditions (L_d , K_d , M_d , N_d , L_f , K_f , M_f , E , I) expressed in terms of observable quantities, unknown parameters, and an unmeasured expectation term. We could in principle use nonlinear GMM to estimate our model using these eight first order conditions and a sufficiently large set of instruments to identify the complete set of model parameters. But first we need a stochastic specification for the model. To illustrate, focusing on L_d we could write:

$$\begin{aligned}
\alpha_{it}^{Ld} &= \exp \left\{ \underline{\alpha}^{Ld} + \varepsilon_{(it)}^{Ld} \right\} \\
g_{lit} &= \exp \left\{ \underline{g}_l + g^l_{(it)} \right\}
\end{aligned}$$

¹⁸ Note: $P_d^1 E$ is the physical quantity of exports times their domestic (not foreign) price – an object we cannot construct since we do not observe prices and quantities separately

$$\left(\frac{P_{od}^l}{P_{of}^l} \right)_{it} = \exp \{ \underline{PR}_l + pr_{l(it)} \}$$

Note that only the ratio of the demand parameters P_{od}^l and P_{of}^l is identified from these FOCs, an issue we return to later. Finally, invoking a rational expectations assumption:

$$E_t(FD_{it})L_{dit} = FD_{it}L_{dit} - \eta_{it}^d$$

we obtain the following equation for L_d , incorporating forecast errors, η_{it}^d :

$$L_d : e^{\underline{\alpha}^{Ld}} (e^{\underline{\varepsilon}_{it}^{Ld}}) (1 - e^{-\underline{g}_1} e^{\underline{g}_1} A) \left[P_d^l S_d + P_d^l N_f + e^{\underline{PR}_1} e^{pr_{l(it)}} \left(\frac{P_d^l S_d}{P_f^l E} \right)^{-e^{-\underline{g}_1} e^{\underline{g}_1}} \cdot (P_f^l E) \right] - w_d L_d - \delta_d FD_{it} L_{dit} \eta_{it}^d = 0$$

and similarly for the other first order conditions. We would then need a set of instruments that enables us to identify the model parameters (of which, e.g., $\underline{\alpha}^{Ld}$, \underline{g}_1 , \underline{PR}_1 , and δ_d appear in the L_d equation).

There appear to be a number of serious problems with a GMM approach in this context. First, the stochastic terms enter the first order condition in such a highly nonlinear way that it is difficult to see any way to linearize the equation (the same problem was confronted by Krusell et al. (2001)). Second, the choice of instruments is not at all obvious. Usual candidates like input and output prices might well be correlated with other firm specific technology parameters if technology changes over time in response to price changes. Third, nonlinear GMM is known to often suffer from severe numerical instability, even in much simpler models than this.

More fundamentally, even if we could implement a GMM approach successfully, it would not be adequate the present context. The usual argument for GMM over ML is that one avoids making distributional assumptions – in this case on the technology, market power and demand shift parameters that are assumed heterogeneous across firms, and on the forecast errors – and thereby obtains more robust estimates of model parameters. But we need to assume (and then estimate) distributions of the heterogeneous parameters in order to be able to simulate the response of the whole population of firms to changes in tariffs or other features of the environment.

On the other hand, a complete stochastic specification of the model would also require us to specify how firms form expectations of future labor inputs, which would require us to specify how they forecast future demand and technology shocks, tariffs, exchange rates, etc. This would go well beyond the scope of the present investigation.

Our approach is a compromise in which we make parametric assumptions on the distributions of the demand and taste shocks, as in a full ML approach. But rather than specifying stochastic processes for the dynamics of all the forcing variables (e.g., tariffs, exchange rates, etc.) we simply substitute the realization of the $t+1$ labor demand terms for their expectation, as in a typical GMM. However, rather than implementing GMM, we make a parametric assumption on the forecast errors, as in Krusell et al (2001), and estimate the model by (simulated) maximum likelihood.

In terms of what we can do with the model once it is estimated, our approach also represents an intermediate case or compromise between the full solution ML approach and the GMM approach. Since we do estimate the complete distribution of technology and demand parameters for the MNCs, we can do *steady state* simulations of the responses of the whole population of firms to changes in the tariffs and other features of the environment. But since we do not model the evolution over time of all the forcing processes, we cannot simulate transition paths to a new steady state.

III.2 (Inverse) Solution of the model

In order to understand the construction of the likelihood function, it is important to first understand the mapping from the observed data for a firm, to the set of firm specific parameters that rationalize the data.

To start, we note that we do not feel there are reliable measures of the payments to capital $\gamma_d K_d$ and $\gamma_f K_f$ for parents and affiliates in the BEA data. To deal with this problem, we assume that the expected rate of profit is equalized across firms (a condition that should hold in equilibrium if there are no capital adjustment costs – which we assume – and ignoring risk adjustments). We can then back out the payment to capital and profit from data on total revenues and payments to the other factors as follows: First define:

$$\text{Domestic revenue: } RD = P_d^l S_d + P_d^l N_f + (1 - T_c - C_f)(P_f^l E)$$

Domestic costs: $CD = w_d L_d + \delta_d K_d + \phi_d M_d + P_f^2 N_d (1 - T_d + C_d) + AC_d$

Domestic costs excluding capital: $CDI \equiv CD - \delta_d K_d$

Now, let R_K denote the fraction operating profit that is pure profit, leaving $(1 - R_K)$ as the fraction that is the payment to capital. This gives $\Pi_d = R_K \cdot [RD - CDI]$ and thus:

$$\delta_d K_d = (1 - R_K) \cdot [RD - CDI]$$

Thus, the rate of profit for domestic operations is $R = \frac{\Pi_d}{\delta_d K_d} = \frac{R_K}{(1 - R_K)}$. We treat R as a common parameter across firms and countries that we estimate (we also assume it is equal for the parent and the affiliate). That is, for the affiliate we have the analogous equation:

$$\delta_f K_f = (1 - R_K) \cdot [RF - CF1]$$

Now, given the common parameters, $R_K, \beta, \delta_d, \delta_f, \mu, \Delta$ and a draw for the forecast errors, η^d, η^f , we can calculate the firm specific technology and demand parameters, and then construct their joint density. We detail the steps involved in calculating the firm-specific parameters from the eight first order conditions in Appendix 1. Here we give a brief summary:

The key insight used in solving the model up is that by observing the fraction of sales of the domestically produced good in the US versus exports of the domestically produced good to Canada, we can infer a *ratio* of domestic to foreign demand parameters. This means that without separate data on prices and quantities (other than wages and labor input, which we need because we estimate the labor force adjustment cost function)¹⁹ we are able to identify the market power parameters (g_1, g_2), the technology parameters $\alpha_d^L, \alpha_d^M, \alpha_d^K, \alpha_d^N, \dots$ and the ratios of the demand shift parameters for good 1, $PR_1 = \left(\frac{P_{od}^I}{P_{of}^I} \right)$, and good 2, $PR_2 = \left(\frac{P_{of}^2}{P_{od}^2} \right)$. To identify the *levels* of the demand shift parameters, we need to bring in extra information that enables us to separate price from quantity. The extra information we need are price indices for capital and materials inputs.

¹⁹ In an earlier working paper, we showed that in the absence of separate information on wages and the labor input we could still identify the technology, price elasticity of demand, and *ratio* of demand shift parameters in a *static* model. So in the dynamic model, the separate data on labor input and wages enables us to identify the adjustment cost function.

Appendix 2 augments Appendix 1 by giving an intuitive explanation of how the more subtle model parameters (A , R_k and the levels of the demand shift parameters P_0) are identified.

III.3 Stochastic Specification

The production function parameters must be positive and sum to one (given CRTS). To impose that they are positive and sum to one, we use a logistic transformation; e.g.:

$$\frac{\alpha^{Ld}}{1 - \alpha^{Ld} - \alpha^{Md} - \alpha^{Kd}} \equiv \alpha_R^{Ld}$$

(where α_R^{Nd} is normalized to equal 1) and then specify that α_R^{Ld} is stochastic and positive. A natural specification would be log normality:

$$\ln \alpha_R^{Ld} = x\beta^{Ld} + \varepsilon^{Ld} \quad \varepsilon^{Ld} \sim N(0, \sigma_{Ld}^2)$$

Then $\alpha_R^{Ld} > 0$ for all ε^{Ld} . Now given a vector α_R^{Ld} , α_R^{Md} , α_R^{Kd} we can solve for α_R^{Ld} , α_R^{Md} and α_R^{Kd} and obtain:

$$\begin{aligned} \alpha^{Ld} &= \frac{\alpha_R^{Ld}}{1 + \alpha_R^{Ld} + \alpha_R^{Md} + \alpha_R^{Kd}} \\ &= \frac{\exp\{x\beta^{Ld} + \varepsilon^{Ld}\}}{1 + \exp\{x\beta^{Ld} + \varepsilon^{Ld}\} + \exp\{x\beta^{Md} + \varepsilon^{Md}\} + \exp\{x\beta^{Kd} + \varepsilon^{Kd}\}} \end{aligned}$$

Note that:

$$\alpha^{Nd} = \frac{1}{1 + \exp\{\} + \exp\{\} + \exp\{\}}.$$

So α^{Nd} is the “base alternative.”

This specification insures that given any values for the $x\beta$ and any draw for the ε , the technology parameters are guaranteed to be positive and sum to 1. Originally we estimated the model assuming log normality for the ε , and found that this was severely rejected for some stochastic terms. So instead we turned to a Box-Cox transformation. We describe the specification of the Box-Cox parameters in Appendix 3.

Since we have a panel, it is important to allow for serial correlation of the errors within each firm over time. We specify a permanent-transitory (random effects) structure, e.g.:

$$\varepsilon^{Ld}(it) = \mu^{Ld}(i) + V^{Ld}(it)$$

and similarly for the other parameters.

$$\mu^{Ld}(i) \sim N(0, \sigma_{\mu^{Ld}}^2) \quad \text{and} \quad v^{Ld}(it) \sim N(0, \sigma_{v^{Ld}}^2).$$

$V(\varepsilon_{it})$ has a basic block-diagonal structure, and covariances across periods are given by

$Cov(\varepsilon_{it} \cdot \varepsilon_{i,t-j})$. Defining $\varepsilon_i = (\varepsilon_{i1}, \dots, \varepsilon_{iT_i})$ we have:

$$V(\varepsilon_i) = \begin{pmatrix} V_1 & & \\ C_{12} & & \\ C_{1T_i} & V_{T_i} & \end{pmatrix}$$

So far, we have only considered the most general case where a firm has all 4 potential trade flows. If a firm has $N_{dt}=0$ ($N_{ft}=0$) at time t , then there is no value for $\varepsilon^{Kd}(it)$ (or $\varepsilon^{Kf}(it)$).

Similarly, if $E_t=0$, there is no value for $P_{of}^1(it)$, and if $I_t=0$ there is no value for $P_{od}^2(it)$. In such cases, $\Sigma^{\alpha d}$, $\Sigma^{\alpha f}$, and Σ^p are collapsed in the obvious way (by removing the relevant rows and columns) and $V(\varepsilon_i)$ is also collapsed. Also, some of our firms are not observed for consecutive years. In the case of a missing year, the relevant rows and columns of $V(\varepsilon_i)$ are removed.

Of course, the likelihood is the joint density of the data, not of the stochastic terms. If y_i denotes the vector of data elements for firm i , then:

$$f(y_i) = \left| \frac{\partial \varepsilon_i}{\partial y_i} \right| f(\varepsilon_i)$$

where $\frac{\partial \varepsilon_i}{\partial y_i}$ is the Jacobian of the transformation from the data to the stochastic terms. In our

case, the 12 data items observed for the firm (or as few as 8 if $N_d = N_f = E = I = 0$) at time t are:

$$y_{it} = \{P_d^1 S_d, P_f^2 S_f, L_d, L_f, w_d L_d, w_f L_f, P_f^2 N_d, P_d^1 N_f, P_f^1 E, P_d^2 I, \phi_d M_d, \phi_f M_f\}$$

Observe that the Jacobian is not block diagonal within periods, because ε_{it} is affected by $L_{d,t-1}$, $L_{f,t-1}$, and $L_{f,t+1}$. It is not possible (as far as we can tell) to write out an analytic expression for the Jacobian, because the mapping from the data elements to the stochastic terms is so highly nonlinear. Furthermore, the mapping depends on the values of the forecast errors (which we are conditioning on here, but which will have to be integrated out by simulation). Therefore, we construct the Jacobian numerically. This is done by bumping one element of the data vector y , recalculating all the elements of ε_i , and then forming numerical derivatives of the elements of ε_i with respect to the elements of y . Then we use these to fill in the column of the Jacobian that corresponds to that element of y :

$$J = \begin{pmatrix} \frac{\partial \varepsilon_{i11}}{\partial y_{i11}} & \frac{\partial \varepsilon_{i11}}{\partial y_{i12}} & \dots & \frac{\partial \varepsilon_{i11}}{\partial y_{i1T_1,12}} \\ \frac{\partial \varepsilon_{i12}}{\partial y_{i11}} & & & \cdot \\ \cdot & & & \cdot \\ \cdot & & & \cdot \\ \cdot & & & \cdot \\ \cdot & & & \cdot \\ \frac{\partial \varepsilon_{iT_1,12}}{\partial y_{i11}} & \dots & \dots & \frac{\partial \varepsilon_{iT_1,12}}{\partial y_{i1T_1,12}} \end{pmatrix}$$

For example, we could bump $y_{i11} = P_d^1 S_{d(i1)}$. and then recalculate the ε_i elements. In this case only elements from period 1 will be affected. We would then construct numerical derivatives and use these to fill in the first column of the Jacobian. Finally, we can now consider forming a likelihood expression by integrating over the normal distribution of the forecast errors:

$$\begin{aligned} \mathcal{L}(\theta) &= \prod_{i=1}^N \int_{\eta_i} f(y_i | \eta_i, \theta) f(\eta_i) d\eta_i \\ &= \prod_{i=1}^N \int_{\eta_i} \left| \frac{\partial \varepsilon_i(\eta_i)}{\partial y_i} \right| f(\varepsilon_i(\eta_i | \theta)) f(\eta_i) d\eta_i \end{aligned}$$

Here, $\eta_i \equiv (\eta_{i1}^d, \eta_{i1}^f, \dots, \eta_{iT_i}^d, \eta_{iT_i}^f)$, $f(\eta_i)$ denotes the joint density of the forecast errors for firm i , and the notation $\varepsilon_i(\eta_i)$ shows that the ε_i vector we construct for a firm is a function of the η_i draw. We assume that the forecast errors are independent over time, as implied by rational expectations, and that within a period:

$$\begin{pmatrix} \eta^d \\ \eta^f \end{pmatrix} \sim N(0, \Sigma_\eta)$$

Let $\Sigma_\eta = CC'$ where C is the lower triangular Cholesky decomposition $\begin{pmatrix} C_{11} & 0 \\ C_{12} & C_{22} \end{pmatrix}$.

Naively, we could simulate the likelihood function by taking iid draws from the distribution of η_i . An important complication arises here, however. Firm behavior cannot be rationalized given any arbitrary choice for the forecast errors. Conditional on (η^d, η^f) , we must have $g_1 > 0$, $g_2 > 0$, all technology parameters positive, and all demand shift parameters positive. From Appendix 1 we have:

$$g_1 \left[AYE - \frac{AB(P_d^1 N_f)(P_f^2 N_d)}{BVI} \right] = RD - CD - \delta_d E(FD)L_d + \frac{(RF - CF - \delta_f E(FF)L_f)}{BVI} (P_f^2 N_d)B$$

In our data, the term $AYE - AB(P_d^1 N_f)(P_f^2 N_d) / BVI$ is always positive.²⁰ So the right-hand side the above equation must be positive in order for g_1 to be positive. But if $E(FD) = FD - \eta^d$ or $E(FF) = FF - \eta^f$ is too large then the right-hand side can be driven negative. Thus, firm behavior implies bounds on the forecast errors. We describe the derivation of the bounds on (η^d, η^f) in Appendix 4.

For any draw (η_d, η_f) that does not satisfy the bounds, no values of the firm specific parameters can rationalize firm behavior. But we can still write the likelihood as:

²⁰ A and B are usually close to 1. Therefore, AYE is close to sales of the domestically produced good, and BVI is close to sales of the foreign produced good. Thus, the second term includes $P_f^2 N_d / BVI$, a number < 1 , times $P_d^1 N_f$, which is just one component of sales of the domestically produced good. So the entire term is approximately sales of good 1 minus a fraction thereof.

$$\mathcal{L}(\theta) = \prod_{i=1}^N \int_{\eta_i} |J(\eta_i)| f(\varepsilon_i(\eta_i)|\theta) f(\eta_i) d\eta_i$$

provided we define $f(\varepsilon_i(\eta_i)|\theta) = 0$ if $\eta_i \notin B_i$, and as equal to the joint normal density of the $\varepsilon_i(\eta_i)$ vector otherwise. We can rewrite the likelihood as:

$$\begin{aligned} \mathcal{L}(\theta) &= \prod_{i=1}^N \int_{\eta_{i1}} \dots \int_{\eta_{iT_i}} |J(\eta_i)| f(\varepsilon_i(\eta_i)|\theta) f(\eta_{iT_i}) \dots f(\eta_{i1}) d\eta_{iT_i} \dots d\eta_{i1} \\ &= \prod_{i=1}^N \int_{\eta_{i1} \in B_{i1}} \dots \int_{\eta_{iT_i} \in B_{iT_i}} |J(\eta_i)| f(\varepsilon_i(\eta_i)|\theta) f(\eta_{iT_i}) \dots f(\eta_{i1}) d\eta_{iT_i} \dots d\eta_{i1} \\ &= \prod_{i=1}^N \int_{\eta_{i1} \in \beta_{i1}} \dots \int_{\eta_{iT_i} \in \beta_{iT_i}} |J(\eta_i)| f(\varepsilon_i(\eta_i)|\theta) I[\eta_{iT_i} \in B_{iT_i}] \dots I[\eta_{i1} \in B_{i1}] \\ &\quad f(\eta_{iT_i}) \dots f(\eta_{i1}) d\eta_{iT_i} \dots d\eta_{i1} \\ &= \prod_{i=1}^N \int_{\eta_{i1}} \dots \int_{\eta_{iT_i}} |J(\eta_i)| f(\varepsilon_i(\eta_i)|\theta) I[\eta_{iT_i} \in B_{iT_i}] \dots I[\eta_{i1} \in B_{i1}] \\ &\quad f(\eta_{iT_i}) \dots f(\eta_{i1}) d\eta_{iT_i} \dots d\eta_{i1} \end{aligned}$$

Clearly, then, the likelihood contribution for firm i can be simulated using a frequency simulator as follows:

$$(F) \quad \hat{\mathcal{L}}_i(\theta) = M^{-1} \sum_{m=1}^M |J(\eta_i^m)| f(\varepsilon_i(\eta_i^m)|\theta) I[\eta_{i1}^m \in B_{i1}] \dots I[\eta_{iT_i}^m \in B_{iT_i}]$$

There are two fundamental problems with this approach, however:

- 1) Observe that for some draws m , the likelihood contribution is zero. For instance, if $P(\eta_{it} \in B_{it}) = .95 \forall t$, this means that only 54% of draws would belong to B_i , on average. This implies a serious numerical inefficiency, since a high percentage of draws are of no use in evaluating the $|J(\eta_i)| f(\varepsilon_i(\eta_i)|\theta)$ term, and are only used to evaluate the event probabilities $P(\eta_i \in B_i)$, which could be much more accurately evaluated using other means (see below).

2) The simulator (F) is not a smooth function of the model parameters θ . The value of the simulator takes discrete jumps at values of θ such that one of the draws η_i^m is exactly on the boundary of B_i . This means that gradient-based search algorithms cannot be used to maximize the likelihood function, and derivatives are not available to calculate standard errors of parameter estimates.

To deal with these problems, we implement a more efficient smooth simulator of the likelihood using a (new) importance sampling algorithm that is a discrete/continuous data analogue of the GHK algorithm for simulating event probabilities in discrete choice models. As in GHK, the idea is to draw the η_i from the “wrong” density, chosen so that all η_i will be consistent with firm behavior. The likelihood is then simulated using a weighted average over these draws where the weights are ratios of the likelihood of a draw under the correct density $f(\eta_i)$ vs. the likelihood of a draw under the incorrect density $f(\eta_i | \eta_i \in B_i)$. Like GHK, the algorithm is recursive. We describe the implementation of the algorithm for period t in Appendix 5. Appendix 6 describes process by which we simulate the posterior distribution of the model parameters in order to test the distributional assumptions of our model.

Implementing the algorithm, we obtain the following smooth unbiased simulator for the likelihood contribution of firm i :

$$(G) \quad \hat{L}_i(\theta) = M^{-1} \sum_{m=1}^M \left| J(\eta_i^m) \right| f(\varepsilon_i(\eta_i^m) | \theta) P(\eta_{d1} \in B_1) P(\eta_{f1} \in B_1 | \eta_{d1}^m) \cdot \dots \cdot P(\eta_{dT_i} \in B_{T_i}) P(\eta_{fT_i} \in B_{T_i} | \eta_{dT_i}^m)$$

Since many economic models have a structure where certain stochastic terms must be in particular ranges in order for a continuous outcome to be observed, the basic approach used here can be useful in many contexts. Note that this approach can be extended trivially to accommodate serial correlation of the η_t . The bounds for η_2 would be a function of η_1^m , and so on. We rule out such correlation here because of the forecast error interpretation of η .

To complete the specification of the structural model, we assume the following structure for the parameters δ_d , δ_f , σ_d and σ_f :

$$\delta_{dt} = \exp\{\delta_d(1) + \delta_d(2) \ln w_{dt} + \delta_d(3) \cdot t + \delta_d(4) \cdot I[N_{dt} > 0]\}$$

$$\begin{aligned}\delta_{ft} &= \exp\{\delta_f(1) + \delta_f(2)\ln w_{ft} + \delta_f(3)\cdot t + \delta_f(4)\cdot I[N_{ft} > 0]\} \\ \sigma_{dt} &= \exp\{\tau_{od} + \tau_1 L_{dt} + \tau_2 w_{dt}\} \\ \sigma_{ft} &= \exp\{\tau_{of} + \tau_1 L_{ft} + \tau_2 w_{ft}\}\end{aligned}$$

Recall that forecast errors were specified as:

$$\begin{aligned}E(FD_{it})L_{dit} &= FD_{it}L_{dit} + \eta_{it}^d \\ &= FD_{it}L_{dit} + \sigma_{dt}\eta_{dt}^*\end{aligned}$$

where η_{dt}^* is standard normal. Clearly then, we expect that σ_{dt} will be an increasing function of L_{dit} . Also, we allow the δ to be functions of the wage rate, on the premise that cost of adjusting labor force size by a particular amount may depend on the wages of the workers involved (e.g. search costs for high wage / highly skilled labor are higher).

As with the other structural parameters, we allow for the possibility that adjustment costs vary over time through the time trends in the δ equations, and we allow for the possibility that the adjustment costs differ between firms that do and do not have intra-firm flows.

Finally, we turn to the problem that not all firms have all four trade flows (N^d , N^f , E , I) in each period. It is beyond the scope of this paper to structurally model the MNC's decisions with regard to how to organize its operations to have intra-firm flows and to import or export. Our focus is on structurally modeling the MNC's production decisions *conditional* on its choice of whether to have intra-firm flows, exports and imports. However, at the same time, we are not comfortable in treating the discrete decisions $I[N^d > 0]$, $I[N^f > 0]$, $I[E > 0]$, and $I[I > 0]$ as exogenous with respect to the firms' decisions regarding continuous quantities of inputs, outputs and flows.

Hence, we specify reduced form approximations to the firm's decision rule regarding $N^d > 0$, $N^f > 0$, $E > 0$ and $I > 0$ and estimate these jointly with the structural model. We specify the approximate decision rule as a heterogeneous multinomial logit with alternative value functions (latent indices) given by:

$$\begin{aligned}V_{it}^{ND} &= \psi_{10} + \psi_{11}\alpha_{it}^{Kd} + \psi_{12}\alpha_{it}^{Kf} + \psi_{13}[\tau_{dit} + c_{dit}] + \psi_{14}g_{lit} + \psi_{15}w_{dit} + \psi_{16}w_{fit} + \psi_{17}t + \psi_{18}I[N_{it-1}^d > 0] \\ &\quad + \psi_{19}I[N_{it-1}^d > 0] + \psi_{1,10}[\tau_{di} + c_{di}] + \mu_i^{ND} + v_{it}^{ND} \\ V_{it}^{NF} &= \psi_{20} + \psi_{21}\alpha_{it}^{Kf} + \psi_{22}\alpha_{it}^{Kd} + \psi_{23}[\tau_{fit} + c_{fit}] + \psi_{24}g_{2it} + \psi_{25}w_{fit} + \psi_{26}w_{dit} + \psi_{27}t + \psi_{18}I[N_{it-1}^f > 0] \\ &\quad + \psi_{29}I[N_{it-1}^f > 0] + \psi_{2,10}[\tau_{fi} + c_{fi}] + \mu_i^{NF} + v_{it}^{NF}\end{aligned}$$

$$V_{it}^E = \psi_{30} + \psi_{31} \alpha_{it}^{Kd} + \psi_{32} [\tau_{fit} + c_{fit}] + \psi_{33} g_{lit} + \psi_{34} w_{dit} + \psi_{35} w_{fit} + \psi_{36} t + \psi_{37} I [E_{i,t-1} > 0] \\ + \psi_{38} I [E_{i1} > 0] + \psi_{39} [\tau_{f1l} + c_{f1l}] + \mu_i^E + v_{it}^E$$

$$V_{it}^I = \psi_{40} + \psi_{41} \alpha_{it}^{Kf} + \psi_{42} [\tau_{dit} + c_{dit}] + \psi_{43} g_{2it} + \psi_{44} w_{dit} + \psi_{45} w_{fit} + \psi_{46} t + \psi_{47} I [I_{i,t-1} > 0] \\ + \psi_{48} I [I_{i1} > 0] + \psi_{49} [\tau_{di1} + c_{di1}] + \mu_i^I + v_{it}^I$$

The name “heterogeneous” logit refers to the heterogeneity terms μ_i , which we assume have a joint normal distribution:

$$\begin{pmatrix} \mu_i^{ND} \\ \mu_i^{NF} \\ \mu_i^E \\ \mu_i^I \end{pmatrix} \sim N(0, \Sigma_\mu^V) \quad \begin{pmatrix} v_{it}^{ND} \\ v_{it}^{NF} \\ v_{it}^E \\ v_{it}^I \end{pmatrix} \sim \text{iid extreme value .}$$

To accommodate state dependence, we include the lagged choice in each value function. This introduces an initial conditions problem (Heckman (1991)) which we accommodate using a recent suggestion by Wooldridge (2000). The idea is to specify the distribution of the random effects conditional on the initial condition. For example, letting $\mu_i^{ND,*}$ denote the firm effect in the V_{it}^{ND} equation, we assume:

$$\mu_i^{ND,*} \sim N(Z_{i1}^{ND}, \Gamma_{ND}, \sigma_{ND}^2)$$

Here Z_{i1}^{ND} contains covariates that characterize the initial condition – including both $I [N_{i1}^d > 0]$ and other covariates dated at $t=1$. In our case, we included both $I [N_{i1}^d > 0]$ and the $t=1$ tariff and transport cost variable in Z_{i1}^{ND} . The reasoning is as follows: Obviously, the distribution of the firm effect should depend on $I [N_{i1}^d > 0]$, since the latter is a function of the firm effect. Also, it is often argued, based on political economy considerations, that tariffs placed on a particular industry may be affected by the importance of trade flows to that industry (e.g., if $\mu_i^{ND,*}$ is large, then that firm may lobby for low tariffs, leading to a negative covariance between $\mu_i^{ND,*}$ and τ_{di1}).

Writing:

$$Z_{it}^{ND} \Gamma_{ND} = \psi_{1,9} I[N_{it}^d > 0] + \psi_{1,10} [\tau_{di} + c_{di}]$$

we obtain the V_{it}^{ND} equation. (Similarly for the other equations). We expect to find $\psi_{1,9} > 0$ and $\psi_{1,10} < 0$.

The intuition behind this approach is as follows: If there is no true state dependence, then $I[N_{it-1}^d > 0]$ may still appear significant just because $I[N_{it}^d > 0]$ is correlated with the random effect, $\mu_i^{ND,*}$, and because the lagged choice proxies for $I[N_{it}^d > 0]$. But if the lagged choice is significant even after controlling for $I[N_{it}^d > 0]$, then there is clear evidence of dynamics.

Note that our choice model is fully simultaneous with the continuous structural model. For instance, we let V_{it}^{ND} depend on α_{it}^{Kd} and α_{it}^{Kf} , the capital intensities of both the parent and the affiliate. We also let V_{it}^{ND} depend on g_{lit} , the market power parameter of the parent. (Similar parameters appear in the other choice equations).

Because these unobserved structural parameters appear in the choice model, it and the continuous component of our model must be estimated jointly. Thus, e.g., if $\psi_{11} > 0$, our model would implement a selection correction for the case that more capital intensive parents are more likely to have $ND > 0$.

IV. Data

IV.1 Construction of the Panel

The data set used in this paper is from the Benchmark and Annual Surveys of U.S. Direct Investment Abroad administered by the Bureau of Economic Analysis. These surveys contain the most comprehensive data available on the activities of U.S.-based MNC and their foreign affiliates. For this study, we use the BEA data disaggregated at the individual foreign affiliate level for each MNC from 1983-1996. The BEA collects data separately from parents and affiliates, which can be merged using an identifier for the parent.

Several alterations were made to the original BEA sample to construct the panel used in this research. First, in the BEA data, firms are classified into approximately 120 2-3 digit ISI industry codes, of which 50 are manufacturing industries. Because many non-manufacturing industries include non-tradeables (such as retailing, real estate and hotels), we use only data on affiliates in manufacturing industries in our panel. This initial screen reduced the total number of affiliate-year observations from 24,313 to 12,241.

Since we aim to capture firm-level production and trade decisions in response to changes in industry tariff and transportation costs, we are indifferent to the organizational form in which MNCs choose to structure these activities. For example, two U.S.MNC parents in the semiconductor industry might choose to structure their Canadian operations in one large affiliate (which might consolidate several plants), or several different affiliates, in which each affiliate is effectively a single plant. For our purposes, what matters is how much semiconductor production the U.S. parent locates in the U.S. (for both the U.S. and Canadian markets) and in Canada (for both the U.S. and Canadian markets), and how much of the goods produced in both markets for export are shipped intra-firm versus arms-length. It is irrelevant whether the U.S. parent chooses to consolidate its same-industry affiliates or report them as separate entities. We therefore merge together affiliates in the same industry with the same parent. Before we get to this step, however, we make several important alterations to the BEA sample, keeping in mind that we required a minimum of three consecutive firm-year observations for each MNC.²¹

- 1) We examine the full set of observations (up to 14 years) for all affiliates that “change” industries over the time period in which they are observed to determine whether the change is due to mis-reported data, or to an actual change in the activity of the affiliate.²² This step was necessary for two reasons. First, to assure smoothness in merging same-industry affiliates together, we had to avoid merging (or not merging) particular affiliates arbitrarily because of mis-reported industry data. Otherwise, an affiliate could double in size from one

²¹ Note that a “firm-year” observation is a parent-affiliate pair observed in a given year.

²² The determination of “mis-reported data” versus actual change was gauged by examining whether, for example, the industry change occurred in a single year and was bracketed by consistent observations, or whether the industry changed in a given year and did not change back. If the latter were the case, we also examined key variables such as employment and sales to ascertain whether a change in the affiliate’s activity occurred. Most “real” industry changes were due to affiliates changing from manufacturing to sales (or vice versa), or changing from being diversified to undiversified (or vice versa). The full set of (14 potential) observations had to be examined to ascertain whether there was a “real” change in the affiliate’s operations. Although only 691 observations were flagged (the flag was created when the change occurred, so if the industry changed and did not change back, the affiliate had only had one flagged observation), a total of 2086 affiliate-year observations represented affiliates with at least one flagged observation.

year to the next, then shrink back again due to noise in the data. Second, we combine parent and affiliate data with tariff and transportation cost data measured in the same 3-digit industries used by the BEA (see below), and we wanted to avoid noisy changes in the exogenous data resulting from mis-reported industry changes. We also examined parents that “changed” industries, but did so primarily for the second reason, since we do not merge together parent data.

- 2) Second, we ideally wanted to restrict our sample to only single-industry parents and affiliates in order to match the firm data with the relevant exogenous data. Since affiliates are almost completely undiversified²³ this screen did not pose a problem with the affiliate data. Again, however, because we combine same-industry affiliates, we could not simply remove observations from the sample in which affiliates report sales in more than one industry in any particular year. We first needed to check the full set of observations (up to 14 years) for all affiliates that reported sales in more than one industry to determine whether the affiliate(s) were diversified over the entire reporting window. This was done in concert with step #1 above. If affiliates were not consistently diversified, they would not be consistently merged with other same-industry affiliates, creating a problem similar to the industry change issue discussed above. In general, the decision rule was to drop affiliates that sold less than 80% of sales in one industry over the entire time period in which they were observed. This step and step #1 above resulted in dropping 1677 affiliate-year observations.²⁴

U.S. parents tend to be considerably more diversified than affiliates. Since parents are not merged together, this was not as severe a problem as it would have been for affiliates. However, the classification of large diversified parents into a single “industry” for the purposes of matching BEA data with exogenous tariff and transportation cost data was not always practicable. We address this issue later in this section.

- 3) Finally, because the BEA conducts two different surveys (the Benchmark and Annual Surveys of U.S. Direct Investment Abroad) with different reporting requirements in terms of affiliate size, reported data are not available for all the affiliates throughout the 14-year period (see Feinberg and Keane, 2001). In particular, the Benchmark Surveys, conducted in 1977, 1982, 1989 and 1994 include the whole population of MNCs and their foreign affiliates, and smaller affiliates are required to report. But in the Annual Surveys, many of the small affiliates that reported data in a Benchmark survey are exempt from filing. In cases in which affiliates report data in a Benchmark Survey but are exempt from the Annual Surveys, the BEA carries them forward by estimating data. As a result of this sampling

²³ “Diversification” refers to the proportion of total (parent or affiliate) sales in only one industry. Affiliates can report sales in up to five industries, and parents can report sales in up to eight industries. In the sample of 12,241 manufacturing affiliate-year observations, an average of 91% affiliate sales is in one industry. The median affiliate sells all its output in one industry. By contrast, only 65% of parent sales in this sub-sample are in one industry.

²⁴ Out of the 12,241 manufacturing affiliate-year observations, 2474 affiliate-year observations reported less than 80% of sales in one industry. “Diversified” affiliate observations were kept if the observations were bracketed by non-diversified observations; or if the reported industries were the same at the 2-digit level; or if the affiliates reported sales in both the manufacturing and sales categories for the same products. In the final sample, 94% of affiliate sales are in one industry. The final sample contains 1100 affiliate-year observations in which less than 100% of affiliate sales are in one industry, but the median affiliate in this sub-sample sells 84% of total sales in one industry.

procedure, most of the observations for smaller affiliates contained estimated data in non-Benchmark years. Ideally, we would remove all estimated data, but similar to steps #1 and #2 above, care had to be taken to avoid removing observations for affiliates that would be merged with other affiliates. For example, an affiliate might have ten reported observations and four estimated observations, but the estimated observations are scattered across different years. If we simply removed the estimated observations and merged the affiliate together with another same-industry affiliate (with the same parent), the consolidated affiliate would appear to change size frequently due to pure noise in the data. So, the decision rule we used in this step was to remove ALL observations for affiliates that did not have at least three consecutive reported observations. We kept estimated observations if they were bracketed by reported observations. If an affiliate had estimated observations at the beginning and end of the 14-year window and reported data in the middle, we eliminated the estimated data after checking to be sure the affiliate was not going to be merged with another same-industry affiliate. Altogether, we removed 4247 affiliate-year observations that contained estimated data.

In steps 1-3 above, we removed a total of 5858 affiliate-year observations from the original sample, leaving 6358 observations.²⁵ Data were removed several more times to arrive at the final sample. First, affiliates in the same industry with the same parent were combined, reducing the sample to 5583 affiliate-year observations. Second, data were removed because of missing affiliate or parent data for sales, employment or employee compensation. Third, non-consecutive observations were removed. These steps left a total of 5175 firm-year observations (on 551 parents and 716 affiliates) in our sample.²⁶ This constituted our initial data set.

We made one final change to our working data set. Specifically, we decided the final panel should contain only one parent observation per year. This necessitated removing from the sample affiliates with the same parent, since only 3877 out of 5175 affiliate-year observations contained *unique* parent-affiliate pairs. (i.e., many parents in the sample had multiple affiliates so parent data was repeated as many as five times in the same year). In this step, we decided to keep in the data set each parent's largest affiliate, since larger affiliates are more likely to have reported observations, and hence, more complete coverage over the 14-year observation window.

²⁵ Although we discuss the three primary issues that arose in creating our panel (changes in affiliate/parent industry, diversification and estimated data) separately, the steps taken to remove observations were made concurrently. In other words, for each of the 12,241 affiliate-year observations, we systematically examined the following information: (1) whether the observation had a flag for diversification or industry change; (2) whether the observation contained estimated or reported data; (3) whether the observation was flagged as being a "merge-candidate" in the sense that there were affiliate observations in the same industry in the same year with the same parent. If none of the above obtained, we did not look at any additional data for a particular affiliate. If we found any of the above flags, we looked at the affiliates' reported industries, sales, employment, employee compensation, imports and exports to ascertain whether there was a change in the affiliates' activities and/or where the flagged observations occurred in the 14-year window. If relevant, we also looked at data on the other affiliate(s) with which the focal affiliate would be merged. In that sense, we give a somewhat stylized account of the creation of our panel by breaking the process into three steps. Indeed, at least two of the above factors necessarily came into consideration before we dropped an observation from the sample. Similarly, in some cases, affiliates that were removed had more than one flag (e.g., 68 deleted affiliate observations had both estimated data and were diversified).

²⁶ Note that our initial "working" data set contained duplicate parent-year observations which we remove in a later step discussed below.

After this step, our final data set contained 3385 unique affiliate-year observations on 543 unique parents and affiliates.

IV.2 Production function variables

We calculated the following production function variables for the parents and affiliates in our sample: labor and labor cost, capital, materials and intermediate goods. Note that labor and labor costs were the only variables for which we observed separate data on prices and quantities. As discussed in section III.2, we construct capital inside the estimation, since we did not feel there were reliable measures of payments to capital for parents and affiliates in the BEA data. In addition to the production function input variables, we constructed two trade flows for parents and affiliates: local sales and arms-length sales to the foreign market. Since different data were available for parents and affiliates, we describe the construction of their production function and trade flows separately.

For affiliates, the BEA collects more comprehensive data on costs and sales. We therefore had to calculate only affiliate raw materials inputs – which were not reported in the BEA data. We calculated affiliate raw materials inputs by subtracting employee compensation, imports from U.S. parents, and current depreciation from cost of goods sold.²⁷

U.S. parents do not report cost of goods sold (CGS), so we needed to obtain this item from Compustat to calculate raw materials inputs for parents. Since cusip numbers are not regularly reported by US parents in the BEA data, we generally used the name of the parent firm to match up BEA and Compustat data. To merge the datasets, a Compustat file containing firm names and ticker symbols was created for the entire population of Compustat firms. This file was brought into the BEA, and a file was created which contained only the ticker symbols for firms that could be matched with BEA data. Using this file, Compustat data was collected for the full 14-year time window and re-merged with BEA data. We were able to find some Compustat data for 384 out of 551 U.S. parents (3494 out of 4433 total parent-year observations)²⁸. However, upon examination, we found the BEA and Compustat data were not always comparable on items reported in both datasets, such as employment. Further examination of the two datasets revealed that although reported levels of the same data items were not always comparable, the ratios of employment-to-sales and employment-to-assets in both datasets had nearly identical distributions. Figure 2 shows the kernel density estimates for the ratio (US Parent) employment/assets in both the BEA and Compustat data.


Using this property of both datasets, we could construct parent cost of goods sold and materials costs using the formula: [Materials cost = (CGS + SG&A) – wages – imports from Canadian affiliates]. To construct (CGS + SG&A) in the BEA data, we calculated the ratio [(CGS + SG&A)/Sales] in Compustat and took the average by parent (BEA) primary industry for the subset of BEA parents that had Compustat data. We then multiplied BEA parent sales by the relevant industry average ratios to construct the sales cost measure. We could also use the ratios to construct parent sales cost for firms with no Compustat data, since there was no reason to

²⁷ Note that current depreciation and depletion charges are included in the BEA's cost of goods sold item. Note also that the BEA cost of goods sold item for affiliates also includes selling, general and administrative expenses (SG&A).

²⁸ The number of parents and parent-year observations reflect the fact that we were using our “working” data set of 5175 affiliate-year observations when we collected the Compustat data.

assume the two subsets of firms (with and without Compustat data) should differ systematically in their production technologies.

Figure 3: Kernel Density Estimates: BEA and Compustat Employment/Assets



The final production function input variable constructed for parents was arms-length sales to Canada. (This item was reported for affiliates, but not parents). To construct this variable, we assumed that a given parent's ratio of arms-length/intra-firm exports was the same for Canada as for other countries.²⁹ We therefore calculated US parent arms-length sales to Canada by multiplying US parent (intra-firm) exports to Canadian affiliates (in the affiliate data) by the (parent-level) ratio of arms-length/intra-firm exports.³⁰ This calculation required refinement in several specific cases. First, in 508 cases, parents reported zero arms-length sales, so no calculation was necessary. Second, in 637 cases, affiliates reported zero purchases from U.S. parents, so the formula could not be used. (Note that 154 of these 637 observations also had zero arms-length sales from parents. Similarly, 71 of the 637 observations were small parents for which all arms-length sales were assigned to Canada. Thus, arms-length trade had to be re-calculated for 402 parent-year observations). For these 402 cases, US-parent arms-length sales to Canada were calculated by multiplying the Canadian share of total (all countries) arms-length exports from all U.S. parents (obtained from the Benchmark Survey published data) by each individual parent's total arms-length sales.³¹

Trade flows, wages and sales were deflated using the U.S. GDP deflator. U.S. and Canadian raw materials prices were not observed, changes in the U.S. and Canadian raw materials price index were used to capture price changes. The U.S. price index is the BLS PPI for intermediate supplies, materials and components for manufacturing, and the Canadian price index, obtained from Statistics Canada, is the PPI for manufacturing intermediates. The cost of capital was assumed to be the same in the US and Canada, and similar to the procedure for

²⁹ Note that after removing transportation equipment from total US-parent intra-firm trade to Canada, the ratio of total intra-firm/arms-length trade between the US and Canada was approximately 84%, as compared to 73% for the US with all other countries in 1989.

³⁰ Note that in the parent data, we observe total parent arms-length and intra-firm exports, but we do not observe the destination. The destination of trade is reported in the Benchmark, but not Annual Surveys.

³¹ The Canadian share of total arms-length trade was 7.7% in 1982 and 11% in 1989. A growth rate was calculated using these data points and applied to each firm's first observation. Observations after the first year were assumed to grow at the same growth rate of each parent's total arms-length exports.

inferring price changes for materials, we used changes in the capital investment price index to reflect changes in cost of capital. Specifically, we used the price index for gross private domestic (non-residential) investment in producer durables obtained from the BLS.

U.S. and Canadian tariffs were measured as annual ratios of the value of duties paid in the United States (Canada) on imports of Canadian (U.S.) goods in industry j at time t divided by the total value of imports to the U.S. (Canada) from Canadian (U.S.) importers in industry j at time t (see Feinberg and Keane, 2001 for a detailed discussion of these measures). A similar measure of transportation costs was constructed by dividing the industry-level cost of insurance and freight by the total value of imports in each industry j at time t ^{32,33}. Although the tariff and transportation cost measures used here are considerably more aggregated than the level at which tariffs are actually imposed, they are more disaggregated than measures often used in empirical work (see Grubert and Mutti, 1991). Indeed, US and Canadian tariffs drop significantly over the 14-year period (1983-1996), with Canadian tariffs falling from an average of nearly 6% to 1.75%, and U.S. tariffs falling from 4% to less than 1%. There is also considerable cross-sectional variation in both sets of tariffs, with U.S. tariffs highest in tobacco (average 13%) and lowest in motor vehicles and pulp and paper (average less than 0.2%). Tobacco and apparel are the highest Canadian tariff industries (both averaging over 17%), and agricultural chemicals, autos and farm machinery are among the lowest Canadian tariff industries (all averaging approximately 1%).

Since affiliates in the sample are predominantly single-industry, U.S. tariff and transportation cost data were converted from 3-digit SIC codes to the Canadian affiliates' 3-digit ISI codes used by the BEA. For diversified U.S. parents, we constructed sales-weighted average Canadian tariff and transportation cost measures across the (up to) eight industries in which U.S. parents report sales. In cases where parents reported sales in non-manufacturing industries, no tariffs were assigned.

Finally, in initial estimations, we found systematic differences in technological change between MNCs with and without intra-firm trade. We therefore created two different time trends – for firms with and without intra-firm trade – to use in subsequent estimations.

V. Results

Table 1 gives descriptive statistics for the firms in our sample. In our structural model, we needed one-period lagged and lead labor force data to construct current and expected future labor adjustment costs. Thus, although our panel extends from 1983-1996, we only use data from 1983 and 1996 only to model labor adjustment costs. Table 1 reports summary statistics on the “big” data set (which includes data from 1983 and 1996 – if the firms appear in the sample in 1984 and 1995). Columns 4 and 5 in Table 1 give the number of zero observations for the four trade flows in the big and “small” data set. (The small set includes only observations from 1984-1995.) Note that most affiliates have zero arms-length sales to the US. Indeed, out of 3385 firm-year observations in the “big” data set, 2061 have zero affiliate arms-length sales to the US. By

³² U.S. tariff and transportation cost data were obtained from the United States Census Bureau, and Canadian tariff data were obtained from Statistics Canada. Canadian tariffs were reported in three-digit SIC codes and were converted into US SIC codes, then BEA ISI codes.

³³ Because similar transportation cost data was not available from Statistics Canada and no systematic transportation cost differences were assumed to exist, we also used the Census data to construct transportation cost measures from Canada to the US (calculated at the industry of the affiliates).

contrast, 953 affiliate-year observations have no intra-firm sales to parents. Note that Canadian affiliates' sales to U.S. parents are larger, on average, than U.S. parents' sales to Canadian affiliates. The bulk of the very large flows from affiliates to US parents are accounted for by the transportation sector.

Our estimates of the full model, obtained by estimating the structural model jointly with the heterogeneous logit models for the trade flows, are shown in Table 2. The logit estimates are shown in Table 3.

The first portion of the table shows the parameter estimates of the MNC's technology. Note that the domestic and foreign labor and materials cost parameter estimates contain several elements. The first term is the intercept. The second term is a time trend for the subset of firms with positive intra-firm flows. For domestic labor and materials, the relevant flow is intra-firm imports from the foreign affiliate. Note that when this flow is positive, U.S. parents have four elements in their production technology (labor, capital, raw materials, and imports of intermediates from the foreign affiliate). When the flow is zero, the parent's technology has only three elements, and similarly for the affiliate technology. The third term is a time trend for the subset of firms with no intra-firm imports of intermediates. The fourth term is an intercept shift parameter for the subset of firms with positive intra-firm flows. Finally, the fifth term is the Box-Cox parameter.

In initial estimations, we found that the restriction of common time trends on the technology parameter for the entire set of firms – those with positive intra-firm trade and those with no trade – was clearly rejected. By incorporating heterogeneous time trends, we substantially improved the fit of the model. A fascinating aspect of the results is that the time trends are small and insignificant for firms without intra-firm flows. Thus, the behavior of these firms is well described by a fixed CRTS Cobb-Douglas technology.

In contrast, in the subset of MNCs with positive intra-firm trade, the time trends for the technology parameters are all highly significant. The direction of the trends points to reduced share of labor, materials and capital for both parents and affiliates. Thus, the estimates imply that, conditional on MNCs having positive intra-firm flows, technical change was driving up the share of intra-firm flows. (We return to this critical point below when we discuss the estimates of the decision rules for positive flows.)

For all the technology parameters, the Box-Cox parameters are (slightly) less than zero, implying that a transformation function (slightly) more strongly concave than the log is necessary to bring them into line with normality.

The next panel of Table 2 (bottom of first page) contains estimates of the two scale parameters, sc_d and sc_f which account for the fact that the scale of the parameters and stochastic terms is not comparable across the three-outcome ($N_d = 0, N_f = 0$) and four-outcome ($N_d > 0, N_f > 0$) specifications. As expected both parameters are less than one, since the scale should be reduced in the three-outcome case. In the same box as the scale parameters are the other common parameters R_k , the rate of profit; A , the rate of disembodied technological change (4.5% across all firms); and μ and β -- parameters in the MNCs' generalized adjustment cost function. The estimates of μ and β imply that adjustment costs are not well described by the common linear-quadratic specification.

The next panels of Table 2 (top of second page) contain estimates of additional labor force adjustment cost parameters. The estimates for $\delta_d W$ and $\delta_f W$ are both close to one, implying that the cost of a given change in the labor force increases one-for-one with the wage rate. It is reasonable that search costs are higher (as are severance costs) for high-wage / highly

skilled workers. The intercept shift for $N_f > 0$ is significant and positive for foreign affiliate labor force adjustment, suggesting that labor force adjustment costs are greater for affiliates that ship intermediates to parents. (A dummy for $N_d > 0$ was not significant in this equation, and similarly, $N_f > 0$ was not significant in the parent adjustment cost equation). Adjustment costs also increased over time, as is evidenced by the significant positive time trends in the domestic and foreign labor force adjustment cost functions.

The next panel contains estimates for the parameters that determine the variance of labor adjustment forecast errors (η). Not surprisingly, the forecast error variances are an increasing function of labor force size. Interestingly, the domestic forecast error variance is higher, and domestic and foreign forecast errors are positively correlated.

Below the labor force adjustment parameters are the parameter estimates for domestic and foreign market power (g_1 and g_2). The market power parameters are estimated with a common time trend and an intercept shift for firms with positive trade flows. Note that the intercept shift is insignificant in estimates of both the domestic and foreign market power parameters. Also, domestic market power does not vary significantly over time, but foreign market power does. The Box-Cox parameter for g_1 implies a transform close to the log, but the transform for g_2 is closer to the square root.

Finally, the last four panels of table 2 contain estimates of the four demand function parameters (i.e. domestic and foreign demand for the domestic and foreign produced goods). All four demand functions exhibit negative time trends.

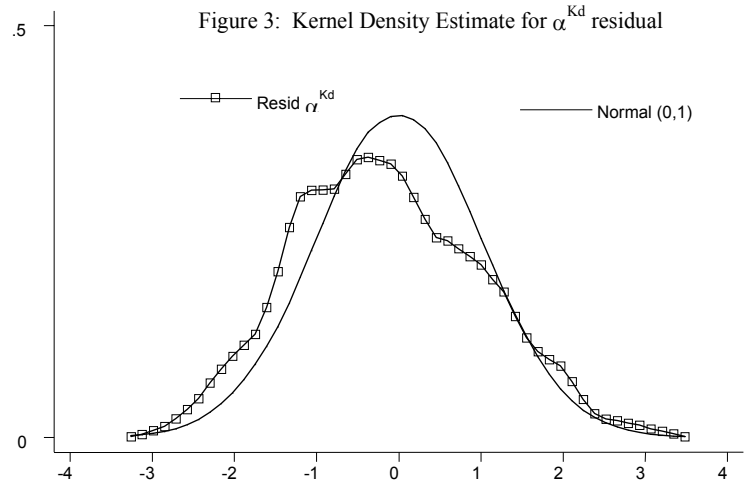
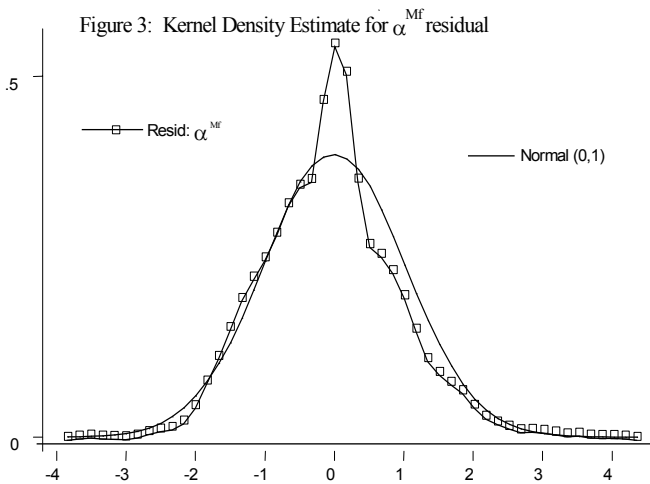
Each parameter is also estimated with a random effect, μ_i , and a transitory error, ε_{it} . Estimates of the covariance matrices are given in Appendix 7. Also in Appendix 7 are correlation matrices for the composite error: $\mu_i + \varepsilon_{it}$, and the composite error at time t and $(t+1)$. Although all the correlations are highly significant, the smallest correlations are between the residuals for demand parameters P_{od}^1 and P_{od}^2 , and between P_{of}^1 and P_{of}^2 . By contrast, the correlation between the residuals for domestic demand for the domestically produced good (P_{od}^1) and foreign demand for the domestically produced good, (P_{of}^1) is 0.97! Similar results obtain for the correlations between errors at time t and $(t+1)$.

Our model appears to fit the data extremely well, in the sense that our assumptions about the distributions of the stochastic terms appear (for the most part) to be very strongly supported. As discussed in section 1.3, a critical modeling assumption we make is that the errors are distributed normally. Tables 6, 6A and 7 in Appendix 7 show that this assumption is substantially supported. Indeed, we reject normality at the 5% level for only two residuals—domestic capital and foreign materials input, and at the 1% level, we reject normality only for foreign materials input.³⁴ Since the test is based upon the deciles of the normal distribution, we show the deciles of each residual in Appendix 7 Table 7, and compare these to the deciles of the standard normal distribution function. Starred entries in Table 7 denote deciles of the residuals that depart significantly from normality. Only the residuals for domestic capital and foreign materials input show any significant differences from the normal distribution function. We reject normality for four of the deciles of the α^{Kd} residual and two of the deciles of the α^{Mf} residual.

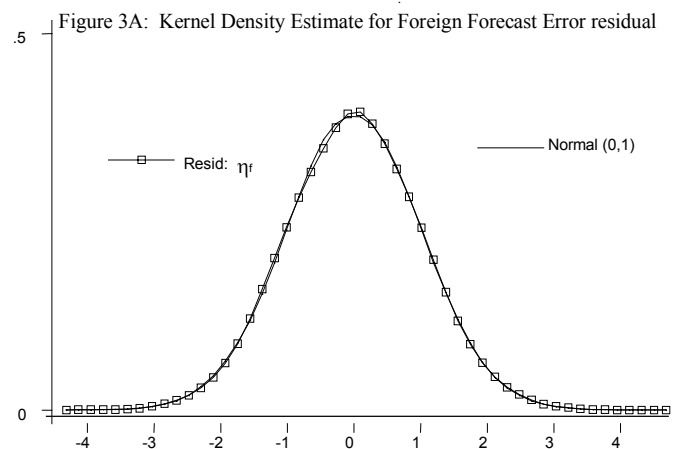
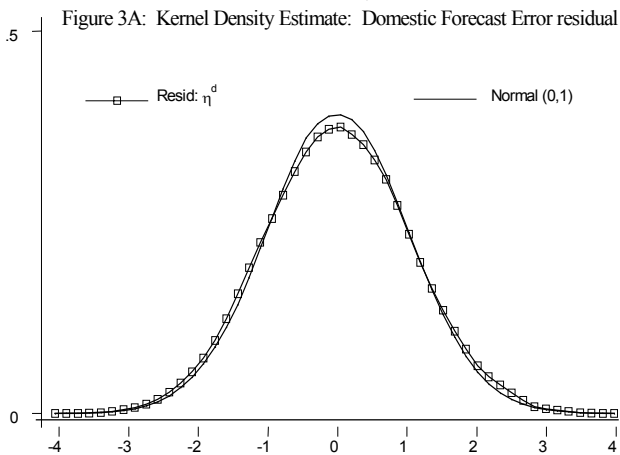
To further examine the fit of our model, we simulated values (and associated importance sampling weights) for each residual using the importance sampling algorithm discussed in section 1.3. We then plotted kernel density estimates for the residuals and compared these to the

³⁴ The test statistic is $\chi^2(9)$, based upon the deciles of the normal distribution. $\chi^2_{.05}(9)=16.92$, and $\chi^2_{.01}(9)=21.67$.

(assumed) normal density function. Figure 3 shows the kernel density estimates for the residual of parameter α^{Mf} , for which normality was rejected at the 1% level, and the residual of parameter α^{Kd} , for which normality was rejected at the 5% level.



Recall that a key assumption of our model is that forecast errors are normal. The statistics in Appendix 7 indicate strong support for this assumption. In fact, the simulated distributions of the η_d and η_f are essentially indistinguishable from normality. This is particularly important given that researchers have typically wished to avoid normality assumptions for forecast errors. Figure 3A shows the kernel density estimates for the residuals of the domestic and foreign forecast errors, η^d and η^f .



Next we turn to our estimates of the decision rules for whether the MNC will have positive values of each of the four trade flows. Table 3 shows the heterogeneous multinomial logit results for the four trade flows: N_d – Canadian affiliate intra-firm sales to U.S. parents; N_f – US parent intra-firm sales to Canadian affiliates; E – US parent arms-length sales to Canada; and I – Canadian affiliate arms-length sales to the US. This logit model aims to endogenize the MNC's decision rule with respect to whether or not to export intra-firm or at arms-length to and

from the US and the foreign affiliate / country. The state variables in the specifications for the two intra-firm flows are: (US or Canadian) tariffs and transportation costs in the parent or affiliate's industry j ³⁵ at time t , the domestic and foreign capital share parameters, domestic (foreign) market power, domestic and foreign wages at time t , a time trend, the lagged choice with regard to the particular flow, and an indicator for the choice at $t=1$ and tariff level at $t=1$.

Two remarkable features of these models are evident. First, (contemporaneous) tariffs and transportation costs are insignificant to the MNC's decision to have intra-firm flows—in either direction. Second, the initial tariff level (time $t=1$) is also irrelevant to the current choice. So despite the fact that US tariffs declined from an average of 4% to less than 1%, and Canadian tariffs fell from nearly 6% to 1.75%, neither the initial high tariffs, nor the current reduced tariff levels were important factors in the MNC's decision to ship goods intra-firm. Not only are the initial tariffs uncorrelated with the MNC's current choice, they are uncorrelated with the random effects (μ_i), further indicating that the decision to have trade flows is unrelated to the tariff level.

The most important predictors of the decision to ship goods intra-firm from the foreign affiliate to the US parent (N_d) are the lagged choice, the choice at $t=1$, and the affiliate's capital share (this parameter being negatively related to the decision to export intermediates from the affiliate to the parent!).

Similar results obtain for the U.S. parent's decision to ship goods intra-firm to the foreign affiliate. Canadian tariffs (contemporaneous and initial) are insignificant, as are transportation costs. Lagged and initial export choices are significantly correlated with the U.S. parent's decision to export goods intra-firm to the affiliate at time t . Parent and affiliate capital share are both significantly associated with the intra-firm export choice, with affiliate capital share again negative. Note that although capital shares are related to intra-firm trade decisions for both parents and affiliates, the firm's domestic and foreign wages are not. However, in both specifications, the coefficient on the foreign wage is negative and the coefficient on the domestic wage is positive. The affiliate's market power parameter is positively associated with the parent's decision to ship goods intra-firm to the affiliate. It seems intuitively reasonable that U.S. parents, wanting to assure global quality control with regard to particular brands or products, might opt to produce some important components centrally and ship them to their foreign subsidiaries for further processing. This finding is similar to Helpman's (1985) suggestion that intra-firm shipments of intermediates from the parent to the subsidiary are likely when the intermediate products require "general know-how" that is typically more abundant in the "entrepreneurial unit."

Essentially, our model says that changes in tariffs and other exogenous factors do not alter whether an MNC has intra-firm flows. There are MNCs that have intra-firm flows, and MNCs that do not, and firms don't change type over time in response to tariffs. The increase in intra-firm trade among the universe of MNCs is accounted for entirely by increased intra-firm trade by MNCs that had positive intra-firm trade to begin with.

As we saw above in Table 2, our estimates of time trends in the technology parameters imply that technical change was one factor inducing the MNCs that started the sample with positive intra-firm flows to increase those flows. Of course, another factor was reduced tariffs. In the next section, we estimate the fraction of the increase in intra-firm trade that is accounted for by technology change vs. tariff reductions.

³⁵ Recall that for diversified U.S. parents, we used sales-weighted average measures of tariffs and transportation costs for the industries in which the parents report sales.

Next we turn to the decision rules for whether to engage in arms length trade. We included the same set of predictor variables in the models for arms-length trade, with the exception of the other-unit capital share. In other words, we include parent capital share in the model for U.S. parent arms-length exports to Canada, but not affiliate capital share (and similarly for affiliate arms-length exports to the U.S.). These results are similar to the results for intra-firm trade in that the lagged and initial export choices are important predictors of current export decisions. Again, similar to the intra-firm trade results, current and initial tariffs are not associated with parent or affiliate export choices. However, several interesting differences are evident in the arms-length trade models.

First, neither the parent nor the affiliate capital share is significantly associated with the parent's (or affiliate's) decision to export goods arms-length to Canada (the U.S.). As our structural model showed, MNCs that trade intra-firm configure their technology differently from firms with no intra-firm trade. This result appears in the significant association between parent and affiliate capital share and the choice to trade intra-firm. It seems that arms-length trade does not require the same sort of technological adaptation as intra-firm trade. This makes considerable sense, as the intermediate goods shipped intra-firm are inputs into the recipient's production technology, while goods shipped arms-length go to entirely different end users.

Second, parent wages are positive and significantly associated with the parent's decision to export goods arms-length to Canada. Evidently, high-wage U.S. MNC parents are more likely to export goods arms-length to Canada. Finally, transport costs are negatively correlated with the decision to export arms-length to Canada, although this relationship is significant at only the 10-15% level.

In the model for Canadian affiliate arms-length exports to the U.S., we find, interestingly, that affiliate market power is significant and negatively associated with the decision to ship goods arms-length to the U.S. While this may seem counterintuitive and contrasts with the positive coefficient on affiliate market power in the intra-firm trade model, we discuss in an earlier paper differences in the composition of the intra-firm versus arms-length trade flows (Feinberg and Keane, 2001). In particular, there is a greater concentration of manufacturing goods such as farm and construction machinery and refined chemicals being traded intra-firm as contrasted with a higher concentration of resource-based goods (such as primary metal and pulp and timber) being sold arms-length. It seems likely that the manufactured goods shipped intra-firm are more easily differentiated (e.g., confer more market power) than the commodity products shipped at arms-length.

Finally, although significant at only the 15-20% level, the initial U.S. tariff is negatively associated with the affiliate's decision to ship goods arms-length to the U.S. The contemporaneous tariff is not significant. We explore the relationship between tariffs and trade in further detail in the policy simulations we present next.

VI. Simulations

In this section we use simulations of our model to examine effects of changes in tariffs on MNC behavior – including the levels of intra-firm and arms-length trade and foreign and domestic labor input. Specifically, using a steady-state simulation (discussed in section 1 and Appendix 6), we compare the baseline levels of intra-firm and arms-length trade flows and domestic and foreign labor in the last year of our sample (1995), given actual tariffs in that year,

to the levels that would have obtained under two different counterfactual scenarios: (1) no tariff reductions (i.e., keeping tariffs at their 1984 levels) and (2) complete tariff elimination.

Recall that since we have not implemented a full-solution algorithm (i.e., we have not specified how expected future labor force size is determined), we cannot simulate the short-run outcomes generated by our dynamic model. However, we can use a steady-state model (which assumes no labor force adjustment costs) to obtain draws from the distribution of firm-specific parameters in order to simulate the long-run response of the population of firms to changes in the policy environment.

Table 4: Simulated Responses to Tariff Changes

	US Parent Sales in US	CA Affil. Sales in Canada	US Parent Sales to CA Affil.	US Parent Arms-Length Sales to CA	CA. Affil Sales to US Parent	CA. Affil Arms-Length Sales to US	US Labor	CA Affil Labor
Base (tariffs = 1995): Steady State	3546430	210123.8	236867.2	37326.7	319125	34426.2	23810	1549
Tariffs = 1984 (no liberalization)	3522975.6	198446.6	227528.8	30238.4	312081.5	23871.4	23523	1417
% difference from Base	-0.7%	-5.6%	-3.9%	-19.0%	-2.2%	-30.7%	-1.2%	-8.5%
Tariffs = 0 (elimination)	3576732.3	227546.6	252606	47938.7	325698.2	41608.4	24115	1645
% difference from Base	0.9%	8.3%	6.6%	28.4%	2.1%	20.9%	1.3%	6.2%
% difference from Tariffs=1984	1.5%	14.7%	11.0%	58.5%	4.4%	74.3%	2.5%	16.1%

We report the results of our steady state simulations in Table 4. The first row shows the levels of domestic sales, intra-firm and arms-length trade flows and domestic and foreign labor that MNCs are predicted to choose (in the long-run steady state) under the actual policy regime that prevailed in 1995, and given the technology and demand parameters that existed in 1995. The second row shows the levels that would have obtained if US and Canadian tariffs were not reduced from their 1984 levels. Not surprisingly, U.S. parent domestic sales and labor change least under the high-tariff regime, with the former dropping by 0.7% and the latter dropping by 1.2%. By contrast, Canadian affiliate domestic sales and labor fall by 5.6% and 8.5%, respectively, under the high-tariff regime.

Also interesting to note are the differences in sensitivity of the intra-firm and arms-length trade flows to changes in tariff levels. For instance, US parent intra-firm sales to Canadian affiliates would have been 3.9% lower if tariffs were kept at their 1984 levels, while U.S. arms-length sales to Canada would have decreased 19%! Similarly, Canadian affiliate intra-firm flows to the US would have decreased only 2.2%, while affiliate arms-length exports to the US would have dropped 30.7% if tariffs had not been cut.

This corresponds with two previous results discussed earlier. First, recall that MNCs with intra-firm trade underwent significant technological change during the period, as reflected by the highly significant time trends for their technology parameters. The technology parameters of MNCs with no intra-firm trade remained close to fixed over this period. So MNCs that trade intra-firm are technologically different, and the decision to trade intra-firm appears to be more a reflection of firm-specific technological considerations than tariff rates. This is also borne out in the logit results. Recall that domestic and foreign capital shares are significant in the intra-firm trade models, but not the arms-length trade models. Similarly, transportation costs and initial U.S. tariffs are marginally significant in the arms-length trade models, but no “policy” variables are significant in the intra-firm trade models.

Figure 4 illustrates the levels of intra-firm and arms-length trade under the actual tariff levels and the no-tariff-reduction scenario. In both charts, the dashed lines represent the base and the solid lines represent the high-tariff (i.e., 1984 level) case. Note the considerably larger increases for US arms-length sales to Canada than for US parent intra-firm trade to Canadian affiliates.

Figure 4: Intra-firm and Arm-length trade from U.S. Parents

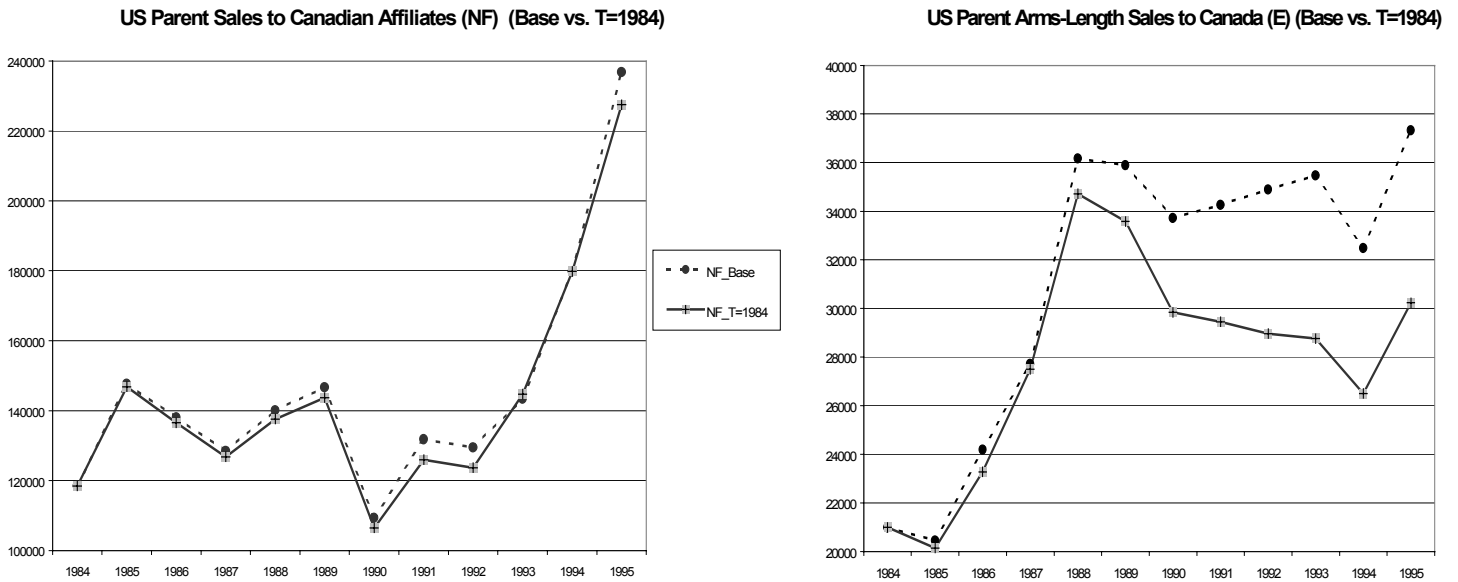


Figure 5: Intra-firm and Arms-Length Sales from Canadian Affiliates

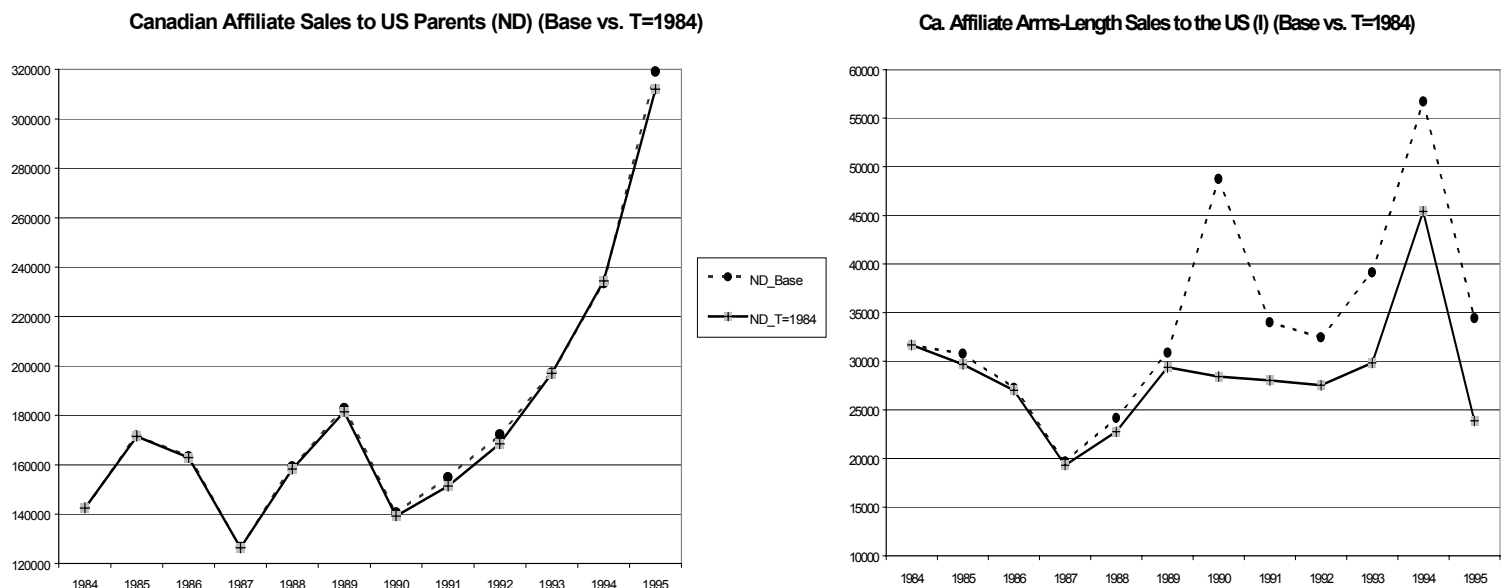


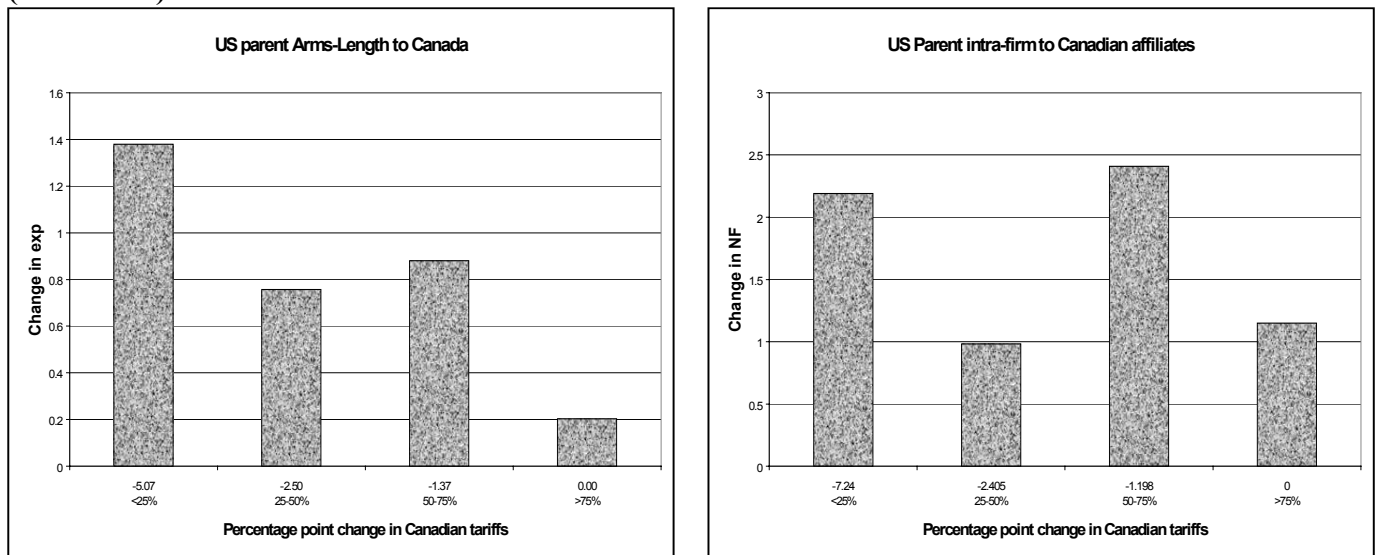
Figure 5 shows how intra-firm and arms-length flows from Canadian affiliates change when we move from the actual to the high-tariff regime. Similar to the U.S. parent flows, affiliate arms-length trade is considerably more sensitive to tariff changes than intra-firm trade to U.S. parents. Note also the significant upward trend in both intra-firm flows, as compared to the more variable nature of exports. Recall that in the previous section, we found that the time trends for the technology parameters of firms with positive trade flows indicated reduced shares of labor, materials and capital for both parents and affiliates. For those firms with intra-firm flows, the share of these flows is increasing as a component of their domestic and foreign production.

Also apparent in all the graphs are bumps in 1989-90 and 1994. These periods correspond with the BEA's Benchmark Survey, so the composition of the panel changes somewhat as a result of the sampling procedure.

The greater sensitivity of arms-length flows to tariff changes, versus intra-firm flows is illustrated in Figure 6 for U.S. Parent sales to Canada (Canadian affiliates) and Figure 7 for Canadian Affiliate sales to the US (US parents). The charts in Figures 6 and 7 show the average change in intra-firm and arms-length trade flows between 1984 and 1989, grouped by changes in US or Canadian tariffs. Changes in the flows are calculated from the posterior estimates of the trade flows predicted for each firm, given the actual tariff regime.

The horizontal axis in Figure 6 gives the quantiles for the tariff changes for firms observed in both years. The vertical axis gives the average percent change in the trade flow. Obviously, we expect larger changes in flows in the larger tariff-change quantiles (left side of the graphs). This is clearly borne out for US parent arms-length sales to Canada (except for the last tariff quantile). We see very large changes in US parent arms-length sales to Canada with larger Canadian tariff reductions, and we see steadily smaller increases in this flow when Canadian tariffs decline less.

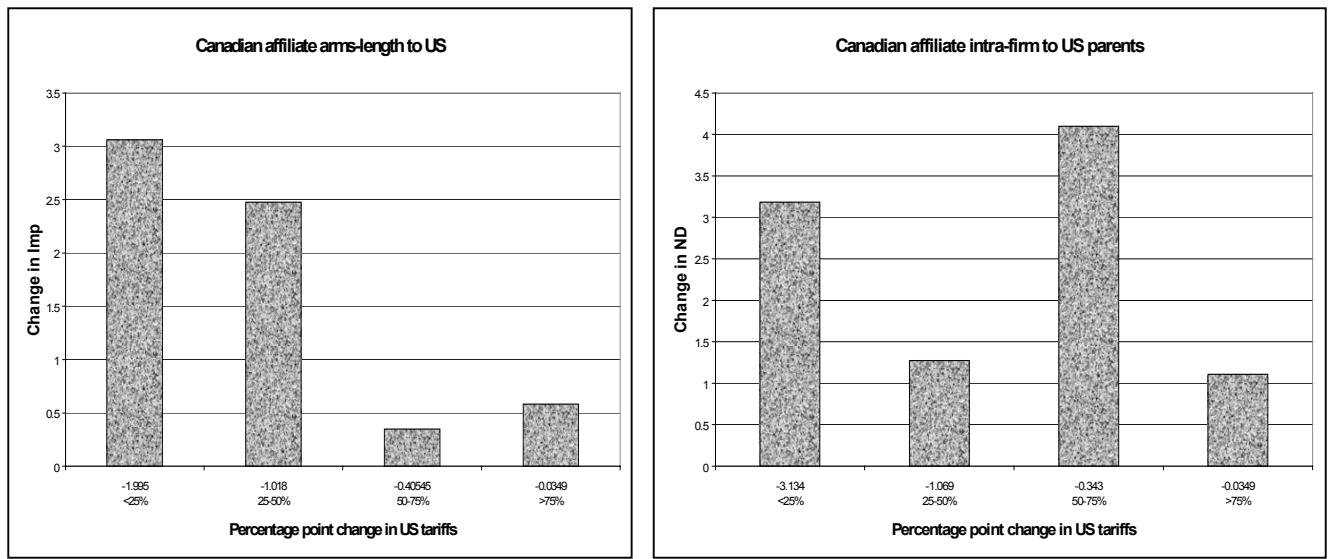
Figure 6: Median Changes in US parent flows to Canada by changes in Canadian Tariffs (1989-1995)



*Notes: All tariff changes are calculated conditional on each flow being positive and greater than 5 in both periods. Tariff changes on the horizontal axis are percentage point changes. Flow changes (vertical axis) are on the vertical axis.

By contrast, we see no evident pattern of changes in US parent intra-firm sales to Canada. This flow does not appear to be affected systematically by changes in Canadian tariffs. Rather, technological change, as evident in the large and significant time trends in the production technology of MNCs that trade intra-firm, appears to be the factor responsible for changes in this flow.

Figure 7: Median Changes Canadian Affiliate flows to the US by changes in US Tariffs (1989-1995)



*Notes: All tariff changes are calculated conditional on each flow being positive and greater than 5 in both periods. Tariff changes on the horizontal axis are percentage point changes. Flow changes (vertical axis) are on the vertical axis.

Figure 7 shows a very similar pattern for the two trade flows into the US. Increases in affiliate arms-length sales to the US (left chart) decline considerably with smaller US tariff reductions. However, the growth in Canadian affiliate intra-firm sales to US parents is nearly the same across all tariff quartiles.

Returning to the simulations, the last three rows of Table 4 give the levels of flows and labor that would have obtained if tariffs were reduced completely. Below the level is the percent difference between the tariff elimination scenario and the base level and the percent difference between the tariff elimination scenario and the no liberalization scenario.

For US parent sales in the US and US parent labor, the difference between complete tariff elimination and actual tariff levels (the base) are nearly the same as the differences between no liberalization and the base. US parent sales in the US would increase by 0.9% and labor would increase by 1.3% if tariffs were eliminated entirely.

The interesting differences that do emerge here may stem from the fact that Canadian tariffs remain higher than US tariffs in most industries. For example, if tariffs were eliminated completely, Canadian affiliates would ship only 2.1% more intermediates to their parents. Similarly, under the no-liberalization scenario, Canadian affiliate shipments to US parents would have only declined by 2.2%. So the change is the same under the two cases.

However, under the tariff elimination case, US parents would ship 6.6% more products to their Canadian affiliates, contrasted with 3.9% less shipments under the no liberalization scenario. This asymmetry could imply that tariffs continue to be a factor in some industries.

Similar results obtain for the arms-length flows. These are again more sensitive to tariff changes than intra-firm flows as depicted in Figures 4-7. We also find a large relative increase in US flows into Canada (and vice versa for Canadian flows into the US) under complete tariff elimination. For example, with no liberalization, US parents would have shipped 19% less goods at arms-length goods Canada, whereas under complete liberalization, they would ship 28.4% more goods into Canada. Canadian tariffs, although clearly lower in 1995 than in 1984, still impede flows into Canada.

Under no liberalization, Canadian affiliates would have shipped 30.7% less goods to the US than under the actual tariff regime. By contrast, if tariffs were eliminated completely, affiliates would have increased arms-length sales to the US by 20.9%. Since US tariffs are lower, the complete elimination scenario gives a smaller “bump-up” in production to Canadian affiliates.

VII. Conclusions

In this paper, we estimate a structural model of MNC production and trade to gain insight into factors causing the significant observed increase in MNC-based trade in general, and intra-firm trade in particular. We find that in an environment characterized as a “unique window” into an “unusually clean trade policy exercise” (Trefler, 2001), tariffs had surprisingly little impact on intra-firm trade. By contrast, MNCs’ arms-length sales to unaffiliated buyers and suppliers were considerably more sensitive to tariff reductions in both the US and Canada.

We find remarkable persistence in the MNCs that had initially organized such that they produced and sold goods intra-firm and the MNCs that had no intra-firm sales. Nearly all the observed increase in intra-firm trade came from firms already organized to sell intra-firm. Technological change, rather than trade liberalization, seemed to be the driving force in the increases in intra-firm trade between US MNCs and their Canadian affiliates.

Organizing to trade intra-firm involves complex international decisions about the organization of production, and the decision to either stop trading intra-firm (when tariffs are increased) or significantly increase intra-firm trade (when tariffs are eliminated entirely) is considerably less sensitive to changes in tariffs than the decision to export goods to unaffiliated buyers. Effectively, changes in intra-firm trade decisions are analogous to changes in the location of production.

An executive at a Canadian affiliate of a US auto paint manufacturer, interviewed as part of this study, summed up the factors involved in this decision³⁶.

When duty rates existed, or tariffs got in the way, when you looked at all the considerations, you really didn’t save any money [by moving production]. Unless there was a shift in technology which made sense in addition to something like the tariff or the duty. You need more than one component to make it a worthwhile move. And, for example, there’s a geographic factor. We service a number of plants in Mexico... and some of them have gotten into waterbase [coatings]. The problem is that the waterbase has certain heat sensitivities and for a place that gets as hot as Mexico... when you figure out the cost to chill all this product and get it there, a lot of the economies disappear (1996).

³⁶ More detail on the interview data is available from the authors.

In estimating our structural model, we also make several significant econometric contributions. First, we introduce a new probability simulator that is the discrete / continuous analogue to the GHK method for simulating discrete choice probabilities. Since many economic models share similar structural characteristics to the one estimated here, the smooth simulator used here can be useful in many contexts. Second, we show how to test our distributional assumptions by simulating the posterior distributions of forecast errors conditional on our estimated model and data. We find the distribution of these errors to be indistinguishable from normality. Finally, we develop another recursive simulator to generate the posterior distributions of the stochastic terms in our model (conditional on the data). This allows us to examine the fit of our model by comparing the shape of these posterior distributions with the shapes we assume a priori. The recursive simulator also allows us to simulate the complete distribution of technology and demand parameters for MNCs so we can do steady-state simulations of their responses to changes in tariffs.

References

- Baldwin, J.R., and P. Gorecki. 1986. *The role of scale in Canada-U.S. productivity differences*. Toronto: University of Toronto Press.
- Baldwin, R. E. and G.I.P. Ottaviano, Gianmarco I. P. 1998. Multiproduct Multinationals and Reciprocal FDI Dumping. *NBER Working Paper #6483*.
- Brainard, S.L. 1993. An Empirical Assessment of the Factor Proportions Explanation of Multinational Sales. *NBER Working Paper # 4583*.
- Brainard, S.L. 1997. An Empirical Assessment of the Proximity-Concentration Trade-off Between Multinational Sales and Trade. *American Economic Review*, 87: 520-544.
- Brander, J.A. 1981. Intra-Industry Trade in Identical Commodities. *Journal of International Economics*, 11(1): 1-14.
- Caves, R. E. 1982. *Multinational enterprise and economic analysis*. New York : Cambridge University Press.
- Caves, R.E. 1990. *Adjustment to International Competition*. Ottawa: Canadian Government Publishing Centre.
- Cummins, J. G. and R. G. Hubbard. 1995. The Tax Sensitivity of Foreign Direct Investment: Evidence from Firm-Level Panel Data, in Feldstein, M., Hines, J.R., and R. G. Hubbard (eds.) *The Effects of Taxation on Multinational Corporations*. Chicago, IL: University of Chicago Press.
- Davis, D. R. and D. E. Weinstein. 2001. An Account of Global Factor Trade. *American Economic Review*, 91(5): 1423-53.
- Deardorff, A.V. 1998. Fragmentation in Simple Trade Models. University of Michigan Department of Economics *working paper 98-11*.
- Dunning, J. H. 1975. *Economic analysis and the multinational enterprise*. New York, Praeger.
- Eastman, H.C., and S. Stykolt. 1967. *The Tariff and Competition in Canada*. Toronto: Macmillan.
- Feinberg, S. E., Keane, M. P. and M. F. Bognanno. 1998. Trade Liberalization and 'Delocalization': New Evidence from Firm-Level Panel Data. *Canadian Journal of Economics* 31(4): 749-77 .
- Feinberg, S. E. and M.P. Keane. 2001. US-Canada Trade Liberalization and MNC Production Location. *The Review of Economics and Statistics*, 83(1): 118-132.

- Gaston, N., and D. Trefler. 1997. The labour market consequences of the Canada-US Free Trade Agreement. *Canadian Journal of Economics*, 30: 18-41.
- Grubert, H. and J. Mutti. 1991. Taxes, Tariffs and Transfer Pricing in Multinational Corporate Decision Making. *The Review of Economics and Statistics*, 73(2): 285-293.
- Hanson, G. H., Mataloni, R. J. Jr. and M. J. Slaughter. 2001. Expansion Strategies of US Multinational Firms. *NBER Working Paper #8433*.
- Head, K. and J. Ries. 1997. Market-Access Effects of Trade Liberalization: Evidence from the Canada-U.S. Free Trade Agreement. In Feenstra, R. C., ed., *The effects of U.S. trade protection and promotion policies*. Chicago: University of Chicago Press.
- Heckman, J. J. 1991. Identifying the Hand of the Past: Distinguishing State Dependence from Heterogeneity. *American Economic Review*, 81: 2, p. 75.
- Helpman, E. 1984. A Simple Theory of International Trade with Multinational Corporations. *Journal of Political Economy*, 92(3): 451-71.
- Helpman, E. 1985. Multinational Corporations and Trade Structure. *The Review of Economic Studies*, 52: 442-458.
- Helpman, E. and P.R. Krugman. 1985. *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*. Cambridge, MA: MIT Press.
- Hymer, S. 1976. *The international operations of national firms : a study of direct foreign investment*. Cambridge, Mass. : MIT Press.
- Ihrig, J. 2000. Multinationals' response to repatriation restrictions. *Journal of Economic Dynamics & Control*, (24): 9, p. 1345.
- Krusell, P., Ohanian, L. E., Ríos-Rull, J.V., and G. L. Violante, 2000. Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica*, 68:5, p.1029.
- Markusen, J. R. 1984. Multinationals, Multi-plant Economies, and the Gains from Trade. *Journal of International Economics*, 16(3-4): 205-26.
- Markusen, J. R. 1995. The Boundaries of Multinational Enterprises and the Theory of International Trade. *Journal of Economic Perspectives*, 9(2): 169-89.
- Markusen, J.R. 2001. A Unified Approach to Intra-Industry Trade and Direct Foreign Investment. *NBER Working Paper #8335*.
- Markusen, J. R. and K. E. Maskus. 1999. Discriminating Among Alternative Theories of the Multinational Enterprise. *NBER Working paper #7164*.

- Markusen, J. R. and K. E. Maskus. 1999. Multinational Firms: Reconciling Theory and Evidence. *NBER Working paper #7163*.
- Markusen, J. R. and A. J. Venables. 2000. The Theory of Endowment, Intra-industry and Multi-national Trade. *Journal of International Economics*, 52(2): p. 209-234.
- Rugman, A. M. 1985. *The Determinants of Intra-industry Direct Foreign Investment*. New York: St. Martin's Press.
- Solow, R. M. 1957. Technical Change and the Aggregate Production Function. *The Review of Economics and Statistics*, 39(3): 312-320.
- Trefler, D. 2001. The Long and Short of the Canada-U.S. Free Trade Agreement. *NBER Working Paper #8293*.
- Yi, K. M. 2001. Can Vertical Specialization Explain the Growth of World Trade?. Forthcoming, *Journal of Political Economy*.

Table 1: Descriptive Statistics

	Mean	St.Dev	Obs=0	Obs=0 (small set)
$P_f^2(Q_f - N_d - I)$ Ca. Affil. Sales in Canada	178349	676671		
$P_f^2 N_d$ Ca. Affil. Sales to US Parent	114981	875130	953	579
$P_d^2 I$ Ca. Affil. Arms-Length Sales to US	13819	49822	2061	1357
$P_d^1 N_f$ US Parent Sales To Ca. Affil.	101671	778194	637	411
$P_f^1 E$ US Parent Arms-Length Sales to Canada	32800	156119	508	325
$P_d^1(Q_d - N_f - E)$ US Parent Sales in US	3392343	9175289		
Ca. Affil. Employee Compensation	42753	142745		
Ca. Affil. Employment	1110	3315		
Ca. Affil. Materials Cost	98625	416527		
US Parent Employee Compensation	777997	2200047		
US Parent Employment	19111	43198		
US Parent Materials Cost	1399165	4035402		

Notes: There are 3385 firm-year observations in the “big set” (which includes 1983 and 1996) and 2335 observations in the “small set” (1984-1995). “Obs=0” refers to the number of firm-year observations for which the specified trade flow is 0. All variables except US and Canadian materials cost are in 1992 US PPI dollars. US Materials cost and Canadian materials cost are in 1992 dollars deflated with US and Canadian deflators for manufacturing intermediates (described in section 2). Variables in dollars are in 000’s.

Table 2: Technology, Adjustment Cost, Market Power and Demand Parameter Estimates

Parameter	Estimate	Std. Error
α^{Ld}_o	-0.2288	(0.0925) ^b
$\alpha^{Ld}_T[Nd>0]$	-0.1135	(0.0089) ^a
$\alpha^{Ld}_T[Nd=0]$	-0.0149	(0.0199)
α^{Ld}_{shift}	5.4032	(0.1096) ^a
Box-Cox (α^{Ld}_o)	-0.0182	(0.0035) ^a
α^{Md}_o	0.2311	(0.0834) ^a
$\alpha^{Md}_T[Nd>0]$	-0.0922	(0.0083) ^a
$\alpha^{Md}_T[Nd=0]$	-0.0044	(0.0168)
α^{Md}_{shift}	4.8860	(0.1142) ^a
Box-Cox (α^{Md}_o)	-0.0398	(0.0038) ^a
α^{Kd}_o	4.5094	(0.1125) ^a
$\alpha^{Kd}_T[Nd>0]$	-0.0764	(0.0058) ^a
Box-Cox (α^{Kd}_o)	-0.0635	(0.0045) ^a
α^{Lf}_o	-0.2373	(0.1172) ^b
$\alpha^{Lf}_T[Nf>0]$	-0.0448	(0.0058) ^a
$\alpha^{Lf}_T[Nf=0]$	0.0167	(0.0134)
α^{Lf}_{shift}	1.5053	(0.0964) ^a
Box-Cox (α^{Lf}_o)	-0.0771	(0.0091) ^a
α^{Mf}_o	0.4295	(0.1920) ^b
$\alpha^{Mf}_T[Nf>0]$	-0.0301	(0.0103) ^a
$\alpha^{Mf}_T[Nf=0]$	0.0259	(0.0186)
α^{Mf}_{shift}	1.3688	(0.1694) ^a
Box-Cox (α^{Mf}_o)	-0.0277	(0.0052) ^a
α^{Kf}_o	1.2639	(0.2013) ^a
$\alpha^{Kf}_T[Nf>0]$	-0.0808	(0.0084) ^a
Box-Cox (α^{Kf}_o)	-0.0260	(0.0062) ^a
SCALE1 ^{‡(sd)}	0.4065	(0.0117) ^a
SCALE2 ^{‡(sf)}	0.6334	(0.0141) ^a
R _K	0.1652	(0.0033) ^a
A	0.0453	(0.0080) ^a
μ	0.7161	(0.0015) ^a
β	0.2216	(0.0047) ^a

Notes: a = significant at 1% level; b=significant at 5% level; c=significant at 10% level. "Shift" refers to intercept shift for N (d or f) >0, T refers to the time trend and Scale 1 and Scale 2 adjust for the differences in the scale of technology parameters and stochastic terms in the 3 vs. 4 outcome case (scale 1: 3 outcome Nd=0, scale 2: 3 outcome Nf=0)

Table 2: Continued

Parameter	Estimate	Std. Error
δ_{d0}	-5.1078	(0.0995) ^a
δ_{dW}	1.0194	(0.0267) ^a
δ_{dT}	0.0240	(0.0038) ^a
δ_{dshift}	-0.0151	(0.0287)
δ_{f0}	-5.4817	(0.0506) ^a
δ_{fW}	0.9667	(0.0073) ^a
δ_{fT}	0.0884	(0.0049) ^a
δ_{fshift}	0.1563	(0.0500) ^a
η_{od}	0.3186	(0.1409) ^b
η_{of}	-0.2072	(0.1247) ^c
η_L	1.1368	(0.0143) ^a
CORR^{##}	0.2974	(0.0808) ^a
g_{10}	-1.3610	(0.0960) ^a
g_{1T}	0.0001	(0.0005)
g_{1shift}	0.0055	(0.0061)
Box-Cox (g_{10})	0.6084	(0.0674) ^a
g_{20}	-1.0878	(0.0435) ^a
g_{2T}	-0.0021	(0.0004) ^a
g_{2shift}	-0.0002	(0.0037)
Box-Cox (g_{20})	0.8424	(0.0433) ^a
$P_{od}^1 o$	2.6898	(0.0908) ^a
$P_{od}^1 T$	-0.0526	(0.0085) ^a
$P_{od}^1 shift$	0.0647	(0.0346) ^c
Box-Cox ($P_{od}^1 o$)	-0.0049	(0.0157)
$P_{od}^2 o$	2.5330	(0.1875) ^a
$P_{od}^2 T$	-0.0740	(0.0123) ^a
$P_{od}^2 shift$	0.3912	(0.1018) ^a
Box-Cox ($P_{od}^2 o$)	0.1013	(0.0235) ^a
$P_{of}^1 o$	2.4520	(0.0816) ^a
$P_{of}^1 T$	-0.0539	(0.0083) ^a
$P_{of}^1 shift$	0.0583	(0.0269) ^b
Box-Cox ($P_{of}^1 o$)	-0.0135	(0.0136)
$P_{of}^2 o$	2.8451	(0.2123) ^a
$P_{of}^2 T$	-0.0888	(0.0128) ^a
$P_{of}^2 shift$	0.4714	(0.1129) ^a
Box-Cox ($P_{of}^2 o$)	0.1568	(0.0195) ^a

Notes: a = significant at 1% level; b=significant at 5% level; c=significant at 10% level. "Shift" refers to intercept shift for N (d or f) >0, T refers to the time trend.

^{##} parameter for unexpected change in labor

Table 3: Heterogeneous Logit Estimates for Trade Flows

Nd: Foreign Affiliate Intra-firm sales to U.S. parent

Parameter	Estimate	Std. Error
Trans.Cost _t	-1.8682	(1.3094)
α^{Kd}	2.6540	(3.5285)
α^{Kf}	-3.9909	(1.3219) ^a
U.S. Tariff _t	-3.5521	(9.4507)
g ₁	14.8839	(24.2639)
W _d	0.0025	(0.0155)
W _f	-0.0040	(0.0175)
Time trend	-0.0640	(0.0404)
I[N ^d _{it-1} >0]	2.7082	(0.2613) ^a
I[N ^d _{it} >0]	2.0432	(0.3885) ^a
U.S. Tariff _{t=1}	0.6535	(8.1493)

Nf: U.S. parent Intra-firm sales to Foreign Affiliate

Parameter	Estimate	Std. Error
Trans.Cost _t	-1.2835	(1.6649)
α^{Kf}	-17.6723	(3.9726) ^a
α^{Kd}	5.3974	(2.9113) ^c
Canadian Tariff _t	-3.8424	(9.6310)
g ₂	62.1002	(19.1856) ^a
W _d	0.0057	(0.0227)
W _f	-0.0276	(0.0264)
Time trend	-0.0923	(0.0732)
I[N ^f _{it-1} >0]	2.6841	(0.3699) ^a
I[N ^f _{it} >0]	3.1723	(0.6660) ^a
Canadian Tariff _{t=1}	-0.3242	(7.7293)

E: U.S. Parent Arms-Length Sales to Canada

Parameter	Estimate	Std. Error
Trans.Cost _t	-2.3931	(1.5575)
α^{Kd}	4.0199	(5.2116)
Canadian Tariff _t	6.7880	(9.7591)
g ₁	-33.2785	(25.1416)
W _d	0.0763	(0.0269) ^a
W _f	-0.0295	(0.0323)
Time trend	0.0022	(0.0633)
I[E _{it-1} >0]	2.9516	(0.4054) ^a
I[E _{it} >0]	3.3328	(0.7916) ^a
Canadian Tariff _{t=1}	-7.6606	(7.8508)

a=Significant at 1%
 b=Significant at 5%
 c=Significant at 10%

Table 3: Heterogeneous Logit Estimates (cont')

I: Canadian Affiliate Arms-Length Sales to the U.S.

Parameter	Estimate	Std. Error
Trans.Cost _t	-0.3746	(0.7916)
α^{Kf}	2.6306	(1.9747)
U.S. Tariff _t	-0.0552	(9.9939)
g_2	-30.3235	(10.8977) ^a
W_d	-0.0165	(0.0162)
W_f	-0.0018	(0.0156)
Time trend	-0.0430	(0.0364)
$I[I_{it-1} > 0]$	2.5033	(0.2408) ^a
$I[I_{it} > 0]$	2.0843	(0.3913) ^a
U.S. Tariff _{t=1}	-12.0584	(8.9857)

Covariance Matrix: Logit Random Effects

Parameter	Estimate	Std. Error
σ_{uND}	1.8843	(0.7022) ^a
$\sigma_{uNF}\sigma_{uND}$	1.9975	(0.5412) ^a
σ_{uNF}	2.9511	(1.2200) ^b
$\sigma_{uE}\sigma_{uND}$	1.0310	(0.4852) ^b
$\sigma_{uE}\sigma_{uNF}$	1.0927	(0.5503) ^b
σ_{uE}	3.1797	(1.3698) ^b
$\sigma_{uI}\sigma_{uND}$	1.4941	(0.4337) ^a
$\sigma_{uI}\sigma_{uNF}$	1.0196	(0.4629) ^b
$\sigma_{uI}\sigma_{uE}$	0.9690	(0.5185) ^c
σ_{uE}	2.2521	(0.7104) ^a

a=Significant at 1%

b=Significant at 5%

c=Significant at 10%