Accounting for the Growth of MNC-Based Trade Using a Structural Model of U.S. MNCs

By SUSAN E. FEINBERG AND MICHAEL P. KEANE*

In recent decades, U.S. foreign trade grew much faster than GDP, but there is no consensus why. Notably lacking is an understanding of the role of multinational corporations (MNCs), which mediate over half of world trade. We use Bureau of Economic Analysis data on U.S. MNCs to study the rapid growth of MNC-based trade from 1983 to 1996. Using a model of U.S. MNCs and Canadian affiliates, we decompose this growth by source. Tariff reductions can largely explain increases in arms-length MNC-based trade. But intra-firm trade growth is attributed mostly to "technical change." We present additional evidence suggesting just-in-time production facilitated intra-firm trade. (*JEL* F13, F14, F23)

In recent decades, trade has grown much more rapidly than GDP. Indeed, between 1982 and 1994, the growth of U.S. foreign trade (exports plus imports) averaged more than 5 percent per year in real terms, while real GDP grew only 3 percent per year. Although the rapid growth of trade has been widely documented, there is no consensus on its source (see Raphael Bergoeing and Timothy J. Kehoe, 2001, and Kei-Mu Yi, 2003, for recent discussions of this issue).

In order to understand the sources of the growth in trade, it is important to understand the behavior of multinational corporations (MNCs). According to Alan M. Rugman (1988), over half of world trade involves MNCs. Such *MNC-based trade* includes two components: *arms-length trade*—shipments between a division of an MNC and unaffiliated buyers/suppliers in other countries—and *intra-firm trade*—internal shipments between divisions of an MNC. Between 1982 and 1994, the total trade of U.S.-based MNCs grew an average of 4.5 percent 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 percent to 45 percent of the total merchandise trade of U.S.-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 analyze the sources of the growth in trade activity by U.S. MNCs with affiliates in Canada, the largest trading partner of the United States. We estimate a structural

---

* Feinberg: Rutgers Business School, Rutgers University, Newark, NJ 07102 (e-mail: feinberg@rbsmail.rutgers.edu); Keane: Department of Economics, Yale University, 37 Hillhouse Avenue, New Haven, CT 06520 (e-mail: michael.keane@yale.edu). The statistical analysis of the confidential 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. The views expressed are those of the authors and do not necessarily reflect those of the Department of Commerce. Suggestions and assistance from William Zeile, Raymond Mataloni, Lorraine Eden, and Kevin Stiroh are gratefully acknowledged. We have benefited greatly from comments by participants at the NBER Universities Research Conference on Firm Level Responses to Trade Policy, especially Keith Head, Kala Krishna, and Bernie Yeung, and by participants at the NBER Conference on Firm Level Data, Trade and FDI (2003), and the European Research Workshop on International Trade. We have also received helpful comments from seminar participants at Yale University, Arizona State University, the University of Maryland, Princeton University, the U.S. International Trade Commission, and the University of Minnesota, especially Penny Goldberg, T. N. Srinivasan, Russ Hilberry, Jason Cummins, Richard Rogerson, Art Blakemore, Phil Haile, Sam Kortum, and Tom Holmes. All remaining errors are, of course, our own responsibility.

[1] For U.S. MNCs, these trade flows grew an average of 3.4 percent and 6.3 percent per year, respectively. We do not include in these statistics trade to and from U.S. affiliates of *foreign* MNCs. All our aggregate statistics on the growth of merchandise trade are derived from figures presented in William J. Zeile (1997).

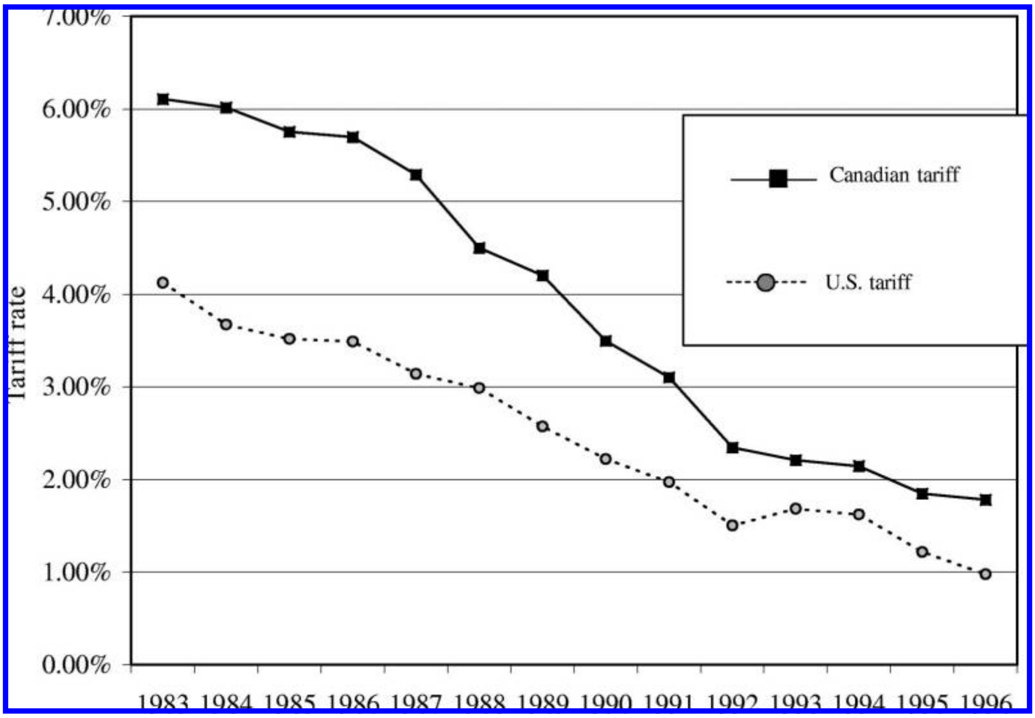

FIGURE 1. AVERAGE U.S. AND CANADIAN TARIFFS IN MANUFACTURING, 1983–1996

Notes: Average U.S. and Canadian tariffs are calculated for firms in the BEA sample in this study. Both tariffs are defined as duties paid in industry *j* divided by total sales in industry *j*. U.S. data were obtained from the U.S. Census Bureau and Canadian data from Statistics Canada.

model of the MNCs' production and trade decisions, using a confidential disaggregated panel dataset from the Bureau of Economic Analysis (BEA) on roughly 500 U.S. MNCs and their Canadian affiliates from 1983 to 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, changes in wages, and other factors (i.e., changes in prices, exchange rates, etc.).

There were significant bilateral tariff reductions between the United States and Canada during our analysis period, arising both from the General Agreement on Tariffs and Trade (GATT), the World Trade Organization (WTO), and the Canada-U.S. Free Trade Agreement (FTA) of 1989, as shown in Figure 1. This figure makes clear that tariff reductions were gradual, and not concentrated in the immediate aftermath of the FTA. As we will see, there was also substantial variation across industries in the extent of tariff reductions. 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.[2]

Using the same BEA firm-level data we examine in this study, Feinberg and Keane (2005) find that tariff reductions had no significant effect on MNC-based trade on the extensive margin. That is, tariff reductions did not cause MNCs to commence arms-length or intra-firm trade, or to establish new affiliates that engaged

---

[2] As Daniel Trefler (2004) points out, most trade liberalizations are accompanied by packages of economic reforms, making the pure effect of tariffs difficult to isolate. A notable example is the North American Free Trade Agreement (NAFTA), which included removal of restrictions on foreign direct investment (FDI) in Mexico. To a great extent, the U.S.-Canada FTA did *not* combine tariff reductions with other major policy changes, making it, according to Trefler, an "unusually clean trade policy exercise."

in such trade. Rather, the growth in MNC-based trade over the 1983–1996 period was almost entirely on the intensive margin, among the subset of firms already engaged in such activities.

Thus, in this paper we develop and estimate a structural model of the *marginal* production decisions of U.S. MNCs and their Canadian affiliates, conditioning on the MNCs' decisions to have an affiliate in Canada, and their decisions regarding whether to engage in intra-firm and arms-length trade. We estimate the first model of MNC behavior that incorporates all the production and trade decisions of MNCs (in a two-country setting): capital, labor and materials input in each country, output in each country, the volume of intra-firm flows, and the volume of arms-length imports and exports.[3]

Our main findings are as follows: first, we find that tariff reductions can explain a substantial fraction of the increase in arms-length MNC-based trade between the United States and Canada. Arms-length exports from the United States to Canada increased much more than arms-length imports, and the model attributes this to an increase in relative wages facing MNC affiliates in Canada as compared to the U.S. parents.

Second, our model implies that tariffs can explain only a small portion of the increase in intra-firm MNC-based trade between the United States and Canada. Our estimates imply that "technical change" (i.e., shifts in intermediate input share parameters), accounts for most of the increase.[4] To explore the source of the technical change, we regress the intermediate input share parameters on a wide range of firm/industry characteristics. We find that growth in intra-firm trade is highly correlated with improvements in the inventory-to-sales ratio. We take this as evidence that advances in logistics management, such as the just-in-time production system, have reduced the inventory carrying cost associated with any given level of intra-firm trade in intermediates.

The remainder of the paper is organized as follows. We describe some key features of the BEA data, and their implications for modeling MNC behavior, in Section I. The model and estimation techniques are discussed in Sections II and III. The construction of our dataset is described in Section IV. We describe our estimation results in Section V. In Section VI we use simulations of our model to quantify the impact of tariffs and other factors on MNC-based trade. In Section VII we study the sources of technical change in our model. Section VIII concludes.

I. A First Look at the BEA Firm-Level Data in Light of Theories of the MNC

Before we settle on a model of MNC behavior to estimate, we examine the BEA data at a descriptive level, to see which, if any, of the extant theoretical models of MNCs might be able to fit their observed behavior. One contribution of our work is simply to use the confidential firm-level BEA data to provide a description of the production and trade arrangements of U.S. MNCs and their foreign affiliates. Few such detailed descriptions of the data at the firm level have previously been available, presumably because of the restrictions on access to the data.[5]

James R. Markusen and Keith E. Maskus (1999, 2002) provide an excellent discussion of the existing theories of the MNC. As they note, these theories can be divided into those that generate vertical versus horizontal MNCs. Classic examples of these two types of models are Elhanan Helpman (1984) and Markusen (1984),

---

[3] Prior structural modeling of MNCs by Jason G. Cummins (1998), Jane Ihrig (2000), and Sanghamitra Das et al. (2001) has examined investment, repatriation, or export decisions in isolation, and has not considered the production and trade aspects of MNC behavior. Cummins (1998) models marginal investment decisions of U.S. parents with affiliates in Canada. He allows for capital adjustment costs that are interrelated across the parent and affiliate. Ihrig (2000) models MNC decisions about repatriation of profits, which we abstract from (i.e., we assume all profits are repatriated each year). Nadia Soboleva (2000) models MNC decisions to locate foreign production facilities across a portfolio of potential host countries, and also models production in each location. But she does not examine intra-firm trade or bilateral arms-length trade.

[4] In a sense, our effort is analogous to the growth accounting literature (Robert M. Solow, 1957). We obtain "technological change" as a residual not accounted for by other factors, and we do not structurally model its ultimate source.

[5] Gordon H. Hanson et al. (2002) also provide a useful description of the BEA data, but they mostly focus on data at the industry/country level, while we focus on manufacturing production arrangements at the firm level.

respectively. Vertical MNCs 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.

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. Models of horizontal MNCs rule out substantial intra-firm flows of intermediates essentially by definition.[6] But inspection of the BEA firm-level data on U.S. MNC parents and their Canadian affiliates suggests that vertical specialization—the fragmentation of production processes into parts that are performed in different countries—is indeed pervasive. We find that the value of intermediates shipped from U.S. parents to their Canadian affiliates is, on average, more than one-third of affiliate total sales. And the value of intermediates shipped by affiliates to parents is, on average, 39 percent of affiliate total sales. Only about 12 percent of the firms in the BEA data are pure "horizontal" MNCs, in the sense of having negligible intra-firm flows of intermediates.

But standard vertical MNC models do not describe the data either. The classic models of Helpman (1984, 1985) and Helpman and Paul R. Krugman (1985) do generate intra-firm trade in intermediates, but it goes in only *one* direction. For instance, in Helpman and Krugman, MNCs produce a differentiated product in three stages that requires descending levels of capital intensity: headquarters services, production of intermediates, and production of final goods. Suppose 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.[7] But a striking aspect of the BEA data is the number of cases in which intra-firm flows actually go in *both* directions—69 percent of observations.[8] In fact, if we define vertical MNCs only as those that exhibit a one-way flow in intermediates, only 19 percent of the firms in the BEA data are purely vertical. To our knowledge, prior empirical work has not documented the great extent of intra-firm flows in intermediates, and the fact that fewer than a third of U.S. MNCs fall into the neat vertical-versus-horizontal dichotomy.[9]

Furthermore, this prevalence of bilateral intra-firm flows of intermediates is not a special feature of U.S.-Canada trade. While U.S. MNC affiliates in Canada tend to be more vertically specialized than affiliates in other countries, we find broadly similar patterns worldwide. For example, among all foreign manufacturing affiliates, nearly 34 percent have some two-way trade with their U.S. parents.[10] Nor is the prev-

---

[6] Obviously, simple horizontal models also fail to generate arms-length sales of final goods (either by the parent to the host country or by the affiliate back to the home country), since FDI is used to avoid such trade in these models. But James A. Brander (1981) showed how horizontal models can be modified to generate bilateral arms-length trade in similar goods by assuming the final goods produced by the parent and affiliate are differentiated. Recently, Richard E. Baldwin and Gianmarco I. P. Ottaviano (2001) have extended Brander's model to also generate intra-industry FDI. We note that Alan M. Rugman (1985) has quantified the importance of bilateral intra-industry trade and FDI at the aggregate industry/country level.

[7] Note that, in this example, all final goods are produced by the affiliate, which it sells in both the host and home countries. In the BEA data on U.S. MNCs and Canadian affiliates, it is always the case that both the parent and affiliate have final sales. As Helpman (1985) notes, vertical models can generate final sales by both parent and affiliate if one adds a distribution/marketing stage that must be tied to the point of sale. Alternatively, one could simply assume that intermediates can be sold as final goods to third parties (which is what we assume in our model).

[8] It is also common for arms-length flows to go in both directions—39 percent of cases in the U.S.-Canada case. Nearly one-third of the MNCs we examine have bilateral intra-firm flows and bilateral arms-length flows simultaneously.

[9] Helpman and Krugman noted that failure to generate bilateral intra-firm flows of intermediates was a problem for their model. But 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.

[10] Much recent empirical work on MNCs has focused on whether the vertical or horizontal model is supported by the data. Authors who have studied FDI at the *industry/country* level generally conclude that horizontal MNCs are the dominant form, because most FDI occurs among high-income developed countries, and this is hard to rationalize as vertical specialization based on endowment differences (see Markusen and Maskus, 1999, 2002; S. Lael Brainard, 1993, 1997). But Donald R. Davis and David E. Weinstein (2001) argue that endowment and factor price differences are substantial even within wealthy OECD countries. On the other

alence of bilateral intra-firm trade in the U.S.-Canada context limited to a small number of industries. While intra-firm flows are large in the automotive industry, there is also significant bilateral intra-firm trade in many other industries, such as computers, aerospace, chemicals, and pulp and paper goods.

An obvious way to generate bilateral intra-firm trade in intermediates in a vertical model is to fragment the production process into additional stages. Avinash K. Dixit and Gene M. Grossman (1982) and Alan V. Deardorff (2001) discuss this possibility. Absent a strictly descending or ascending ordering of stages by factor intensity, this creates the possibility for bilateral flows of intermediates. Unfortunately, the BEA Annual Survey data do not contain information on stage of processing of intermediates (nor, for that matter, do they break down other inputs like labor by stage of processing). Thus, estimation of an explicit multistage production process appears impossible. As a consequence, there is no "off-the-shelf" model of MNC behavior that we can estimate structurally that would fit the BEA data at the firm level.

Thus, a key challenge we face is how to specify a model that generates bilateral intra-firm trade, yet that is estimable given available data. We chose to do this in the most direct possible way. We simply assume a production process in which the final good produced by the parent may be required as an intermediate input in the affiliate's production process, and, simultaneously, the final good produced by the affiliate may be required as an intermediate input in the parent's production process.[11] We describe our model in detail in the next section.

II. Our Model of MNC Behavior

A. *Overview*

This subsection provides a nonmathematical overview of our model. At the outset, we stress that we model only the *marginal* production and trade decisions of the MNC, conditioning on its decision to place an affiliate in Canada, and its decisions regarding whether to import, export, and have intra-firm flows of intermediates. We do not view this as a serious limitation, because results in Feinberg and Keane (2005) indicate that essentially all the changes in MNC-based trade over the 1983–1996 period occurred on the intensive, and not the extensive, margin. Loosely speaking, the MNCs engaged in exporting, importing, and trading intra-firm at the start of our analysis period were the same ones engaged in these activities at the end. Hence, a model of MNC marginal decisions is what is most important for understanding changes in trade flows.

An MNC may potentially utilize four trade flows: arms-length imports and exports, and intra-firm trade in intermediates from parent to affiliate, or vice versa, giving $2^4 = 16$ possible MNC configurations.[12] While we condition on the MNC configuration, we do not assume it is exogenous. That could bias estimates of our structural model if MNC configuration is influenced by firm-specific unobservables. Hence, we estimate our structural model jointly with the reduced-form decision rules for utilization of each flow, presented in Feinberg and Keane (2005).

A key data problem that influences the setup of our model is that the BEA data do not allow us to separate quantities of production from prices. This problem is certainly not specific to our application. In fact, it is common to most production function estimation. It has been

---

hand, Hanson et al. (2002) argue that the BEA data on U.S. affiliates abroad provide strong evidence that vertical fragmentation is important and growing.

[11] As an example of this type of process, consider the petroleum industry. Finished products from a petroleum refinery include fuel oil, lubricants, and Naphtha. In turn, these products are also used as intermediates in oil drilling. Consider Mobil's (later Exxon/Mobil's) offshore drilling operation in the Hibernia field off Newfoundland. Since Mobil had little refining capacity in Canada, the crude was shipped primarily to Mobil 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 offshore oil drilling, since all power must be generated at the rig.

[12] These decisions should be based on expected present values of profits under each possible MNC configuration. Thus, estimation of a complete structural model would require solving a discrete/continuous stochastic dynamic programming problem with 16 discrete alternatives each period. It is well beyond available computational capacity to implement such a model. In fact, the current state of the art in this literature is dynamic modeling of the export decision *alone,* as in Das et al. (2001) or Andrew Bernard and J. Bradford Jensen (2004).

typical in the literature to simply use industry-level price indices to deflate nominal sales revenue data in order to construct real output. But Zvi Griliches and Jacques Mairesse (1995) and Tor Jakob Klette and Griliches (1996) have pointed out that this procedure is valid only in perfectly competitive industries, so that price is exogenous to the firms. This condition is obviously violated for MNCs, since they have market power. This problem has received a great deal of attention recently in the industrial organization (IO) literature (see, e.g., Hajime Katayama et al., 2003, and James A. Levinsohn and Marc J. Melitz, 2002).

The only general solution to the problem of endogenous output prices is to estimate the production function jointly with an assumed demand system. But in our case the problem is further exacerbated by the fact that, while the BEA data report nominal values of intra-firm flows, imports, and exports, the prices and quantities for these flows cannot be observed separately. Nor can we separate price and quantity for capital and materials inputs, or for intermediate inputs shipped intra-firm. Furthermore, the price of such intermediate inputs is endogenous, since it depends on the MNC's other input and trade decisions.

The only general solution to the problems created by the inability to observe prices and quantities of outputs or intermediate inputs separately is to assume: (a) constant returns-to-scale (CRTS) Cobb-Douglas production functions for both the parent and affiliate; and (b) that both parent and affiliate face isoelastic demand in the market for final goods. These two assumptions enable us to identify the price elasticities of demand faced by parents and affiliates using only information on revenues and costs (i.e., by exploiting Lerner-type conditions). Then, given the elasticities of demand, we can pin down the Cobb-Douglas share parameters using only information on factor shares of revenues (appropriately modified to account for market power).

The solutions proposed by Katayama et al. (2003) and Levinsohn and Melitz (2002) allow estimation of more general production functions, but they assume that real input quantities are observed. In our case, generalizations of Cobb-Douglas seem infeasible, since we lack the input price variation needed to identify substitution elasticities. In Section VII, however, we show that the unit elastic assumption does not drive our main results.

The CRTS assumption may seem objectionable, since theories of the MNC stress multi-plant economies of scale as a key reason for multinational organization. These economies of scale typically arise from "intangible" firm-level fixed inputs, which serve as joint inputs for MNC operations in all countries. When we take MNC structure as given and estimate a model of the MNC's marginal decisions, we implicitly assume there exists a firm-level fixed cost of obtaining the intangible input,[13] and that this fixed cost does not affect marginal production decisions. We stress, however, that increasing returns to scale for the MNC as a whole (due to intangible fixed inputs) is perfectly compatible with CRTS in the *tangible* inputs of both the parent and affiliate.

Given this background we now lay out the main assumptions of our model as follows:

1. The parent and affiliate each produce a different good.
2. The good produced by the affiliate may serve a dual purpose: it can be sold as a final good to third parties (in Canada or the United States), or it can be used as an intermediate input by the parent. We make a symmetric assumption for the good produced by the parent.
3. The parent and affiliate each produce a variety of a differentiated product, and have market power in final goods markets. The two final goods they produce are unrelated (i.e., they have zero cross-price elasticity of demand).
4. The parent and affiliate both produce output using a CRTS Cobb-Douglas production function that takes labor, capital, and materials as inputs. In addition, intermediates produced by the affiliate may be a required input in the parent's production process, and intermediates produced by the parent may be a required input in the affiliate's production process.

---

[13] E.g., MNCs may have market power by virtue of fixed investments in establishing brand names or in inventing a proprietary production process. Licensing out foreign operations might dissipate brand equity, or risk revelation of proprietary knowledge. Unfortunately, we cannot directly measure investments in intangible inputs in our data.

5. The affiliate and the parent both face isoelastic demand functions in both the U.S. and Canadian final-goods markets.
6. The parent and affiliate both face labor force adjustment costs.
7. The MNC maximizes the expected present value of profits in U.S. dollars, converting Canadian earnings to U.S. dollars using the nominal exchange rate.
8. The expected rate of profit is equalized across firms.
9. Parameters of technology and of demand are heterogeneous both across firms and within firms over time. We allow for time trends in these parameters to capture "technical change" and shifts in demand. (We discuss our stochastic specification in Section IID.)

Assumption 1 enables us to generate bilateral arms-length trade in finished goods, which is a pervasive feature of the data. The part of Assumption 3 that goods produced by the parent and affiliate are unrelated (e.g., U.S. demand facing the parent is unaffected by imports from the affiliate) is necessary to identify the demand elasticities given only information on revenues and costs, as we show below. This assumption can be motivated by heterogeneity of parents' output (e.g., while our affiliates are almost always single industry (see Section IV), 81 percent of parents are in two or more industries, 53 percent in four or more, and 30 percent in eight or more).[14] Furthermore, we argue below that this assumption is unlikely to influence our main results.

Assumption 2 is motivated by the fact that the annual data on intra-firm intermediate flows are not broken down by stage of processing, and the fact that the data do not contain information on the composition of output. Thus, it is not feasible to estimate multistage or multioutput production functions. The assumption that the affiliate produces a single good that can either be sold as a final good or used as an intermediate by the parent (and the symmetric assumption for the parent) enables our model to generate the pattern (often observed in the data) that affiliates sell intermediates to the parent, and simultaneously sell final goods to third parties in both the United States and Canada (and vice versa for parents).

Assumption 6 is motivated by the fact that the labor input is the only real quantity we observe directly. Thus, it makes sense to put adjustment costs on labor rather than capital. The main role of adjustment costs is to help explain short-run persistence in firm behavior, and we do not think our main results regarding long-run tariff effects are sensitive to this assumption.

Assumption 8 is motivated by the fact that the capital stock data are rather imprecise (i.e., property plant and equipment (PPE) at historical cost). This is, of course, a very general problem not limited to the BEA data. As we discuss below, the assumption of an equalized (expected) profit rate across firms—which can be rationalized on the basis of capital market equilibrium with mobile capital (and ignoring risk adjustments)—will enable us to dispense with the capital stock data entirely and to construct payments to capital as a residual using the other available cost and revenue data.

B. *Basic Structure of the Model*

We present our model graphically in Figure 2. This illustrates the most general case, in which a single MNC exhibits all four potential trade flows (arms-length imports and exports and bilateral intra-firm trade). Instances where an MNC uses a subset of the four flows are special cases. Q_d and Q_f denote total output of the parent and affiliate, respectively. N_d denotes the part of affiliate output shipped to the parent as an intermediate. Similarly, N_f denotes intermediates shipped from parent to affiliate. I (imports) denotes the quantity of goods sold at arms length by the Canadian affiliate to consumers in the United States, and E (exports) denotes arms-length exports from the U.S. parent to consumers in Canada. Thus, $S_d \equiv (Q_d - N_f - E)$ is the quantity of its output the parent sells in the United States, and $S_f \equiv (Q_f - N_d - I)$ is the quantity of 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. Since we do not observe prices and quantities separately in the data, we will work with the six MNC firm-level trade and domestic sales flows, which are $P_d^2 I$, $P_f^1 E$, $P_d^1 N_f$, $P_f^2 N_d$, $P_d^1 S_d$, and $P_f^2 S_f$.

---

[14] The assumption is also suggested by optimal behavior of a multi-product firm. An MNC should obviously design its product line so final goods imported from affiliates are not in close competition with ones made by the parent.

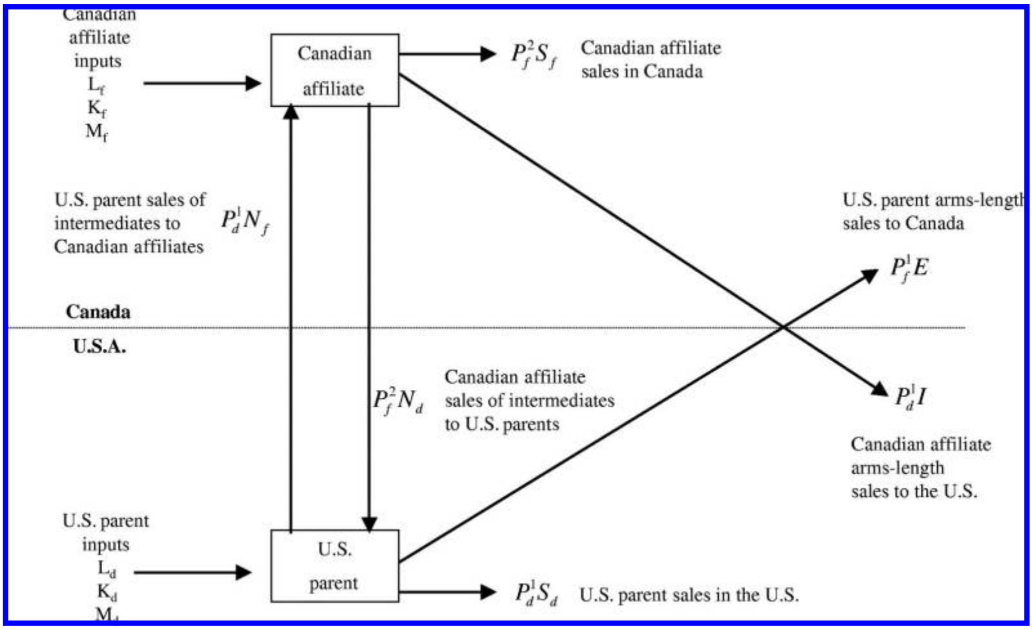

FIGURE 2. STRUCTURAL MODEL OF MNC PRODUCTION AND TRADE FLOWS

The MNC's domestic and Canadian production functions are Cobb-Douglas, given by

$$Q_d = H_d K_d^{\alpha^{Kd}} L_d^{\alpha^{Ld}} N_d^{\alpha^{Nd}} M_d^{\alpha^{Md}}. \tag{1}$$

$$Q_f = H_f K_f^{\alpha^{Kf}} L_f^{\alpha^{Lf}} N_f^{\alpha^{Nf}} M_f^{\alpha^{Mf}}. \tag{2}$$

Note that there are four factor inputs: capital (K), labor (L), intermediate goods (N), and materials (M). We assume that the share parameters α sum to one for both the parent and affiliate (CRTS). We allow the constants H_d and H_f to follow time trends in order to capture TFP growth.[15]

For the domestically (U.S.) produced good (good 1), the MNC faces the following isoelastic demand functions in the United States and Canada:

$$P_d^1 = P_{0d}^1 S_d^{-g_1}, \qquad P_f^1 = P_{0f}^1 E^{-g_1}, \tag{3}$$

$$0 < g_1 < 1.$$

---

[15] Interestingly, estimation of an aggregate MNC production technology can lead to misleading conclusions regarding returns to scale. To see this, note that the production function (1) can be expressed directly in terms of capital, labor, and materials inputs of the parent and affiliate, by substituting out for N_d. If the decision rules for intermediates are $N_d = \lambda_{Nd} Q_f$ and $N_f = \lambda_{Nf} Q_d$ (where, obviously, optimal λ_{Nd} and λ_{Nf} depend on prices), then

$$Q_d = H_d K_d^{\alpha^{Kd}} L_d^{\alpha^{Ld}} M_d^{\alpha^{Md}} (\lambda_{Nd} H_f K_f^{\alpha^{Kf}} L_f^{\alpha^{Lf}} M_f^{\alpha^{Mf}} N_f^{\alpha^{Nf}})^{\alpha^{Nd}}.$$

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

$$Q_d = \{[H_d H_f^{\alpha^{Nd}} \lambda_{Nd}^{\alpha^{Nd}} \lambda_{Nf}^{\alpha^{Nd}\alpha^{Nf}}](K_d^{\alpha^{Kd}} L_d^{\alpha^{Ld}} M_d^{\alpha^{Md}})$$
$$\times (K_f^{\alpha^{Kf}} L_f^{\alpha^{Lf}} M_f^{\alpha^{Mf}})^{\alpha^{Nd}}\}^{1/(1-\alpha^{Nf}\alpha^{Nd})}.$$

Note that λ_{Nd} and λ_{Nf} are combined with the TFP terms H_d and H_f. Thus, the "aggregate" technology for the MNC as a whole exhibits CRTS only if λ_{Nd} and λ_{Nf} are fixed, as capital, labor, and material inputs vary. However, the aggregate technology may exhibit increasing or decreasing RTS if λ_{Nd} and/or λ_{Nf} vary along with capital, labor, and materials inputs, as input prices, output prices, and/or tariffs vary. We would expect tariff reductions to lead to both increases in the size of MNCs (increases in K, L, M inputs) and increases in the λ. Thus, it may appear that MNCs have increasing RTS if an aggregate production function, rather than separate parent and affiliate production functions, is estimated, and if intra-firm flows of intermediates are ignored. Another problem arises if one adopts a standard approach of estimating the $ln(Q_d)$ equation. Then, the stochastic terms will subsume changes in the λ_{Nd} and λ_{Nf} terms. Since these depend on input prices, input prices will not be valid instruments for the input quantities.

Similarly, for the good produced in Canada (good 2), the MNC faces the demand functions

$$(4) \qquad P_f^2 = P_{0f}^2 S_f^{-g_2}, \qquad P_d^2 = P_{0d}^2 I^{-g_2}, \qquad 0 < g_2 < 1.$$

Recall that S_d denotes the quantity of the U.S.-produced good sold in the United States, and S_f denotes the quantity of affiliate sales in Canada. The g_1 and g_2 are the (negative) inverses of the price elasticities of demand for the domestic and foreign produced good, respectively.

Note that the price elasticity of demand is a property of the good, and not the country. For instance, the price elasticity for the U.S.-produced good is $-1/g_1$ in both the U.S. and Canadian markets. But the demand function intercepts, P_{0d}^1 and P_{0f}^1, differ.[16] We also assume the two goods are unrelated. These assumptions are critical in order to identify the demand elasticities using only information on nominal quantities. Given CRTS Cobb-Douglas, when we sum the first-order conditions (FOCs) for all inputs we get two equations in the two unknowns, g_1 and g_2. Complicating the demand functions by letting the demand elasticity for a good differ by country, or by letting demand for good 1 depend on the price of good 2, would leave us with two equations in more than two unknowns, and the demand function parameters would no longer be identified.

Next, we assume the MNC faces labor force adjustment costs. These are often assumed quadratic, e.g.: $AC_{dt} = \delta_d[L_{dt} - L_{d,t-1}]^2$ where $\delta_d > 0$. But the following generalization leads to a substantial improvement in fit and accommodates many reasonable adjustment cost processes:[17]

$$(5) \quad AC_{dt} = \delta_d((L_{dt} - L_{d,t-1})^2)^{\mu}/L_{d,t-1}^{\Delta}, \qquad \text{where } \delta_d > 0,\ \mu > 0,\ \Delta \geq 0.$$

A similar adjustment cost function is specified for the affiliate, which will be allowed to have a different δ parameter (δ_f). The curvature parameters μ and Δ are assumed to be common.

We can write the MNC's period-specific profits (suppressing the time subscripts) as

$$\begin{aligned}(6) \quad \Pi = {} & P_d^1(Q_d - N_f - E) - P_d^1 N_f(T_f + C_f) \\ & + P_f^1 E(1 - T_f - C_f) \\ & + P_f^2(Q_f - N_d - I) \\ & - P_f^2 N_d(T_d + C_d) \\ & + P_d^2 I(1 - T_d - C_d) \\ & - w_d L_d - w_f L_f - \phi_d M_d - \phi_f M_f \\ & - \gamma_d K_d - \gamma_f K_f - AC_d(L_d, L_d^{(-1)}) \\ & - AC_f(L_f, L_f^{(-1)}).\end{aligned}$$

Here, T_f and C_f are the ad valorem Canadian tariff and transportation costs the MNC faces when shipping products from the United States to Canada (and similarly for T_d and C_d). Transportation costs are measured ad valorem because industry specific transport cost measures are available only in that form. Such an assumption is quite common in the trade literature (see, e.g., Keith Head and John Reis, 2001; Hanson et al., 2002).[18]

The exchange rate enters (6) implicitly because Canadian affiliate costs and revenues are converted into U.S. dollars using the nominal exchange rate. Thus, the MNC cares about U.S. dollar profits (and hence U.S. dollar output and input prices). w_d and w_f are the domestic and foreign real wage rates respectively; ϕ_d and ϕ_f are the domestic and foreign materials prices; γ is the price of capital, which we assume is equal for the parent and the affiliate ($\gamma_d = \gamma_f$).

The MNC's problem is to maximize the expected present value of profits in real U.S. dollars, $E \sum_{\tau=1}^{\infty} \beta^{\tau} \Pi_{t+\tau}$, by choice of eight control variables $\{L_{dt}, M_{dt}, K_{dt}, N_{dt}, L_{ft}, M_{ft}, K_{ft}, N_{ft}\}$.

---

[16] With the isoelastic functional form, gross revenue from exports is $P_d^2 E = P_{0d}^2 E^{(1-g_1)}$. Since $0 < g_1 < 1$, as E goes to zero, revenue goes to zero, despite the fact that P_d^2 approaches infinity. Thus, the model can rationalize a decision not to export, although we do not model this decision structurally. The same is true with imports.

[17] For example, if $\mu = 1$ and $\Delta = 0$, we get $\delta_d[L_{dt} - L_{d,t-1}]^2$, while if $\mu = 1/2$ and $\Delta = 1$, we get $\delta_d \cdot |(L_{dt} - L_{d,t-1})/L_{d,t-1}|$.

[18] Given that a large component of transportation cost is insurance, along with the fact that higher-value items tend to be shipped using more expensive means, the ad valorem specification does not seem completely unreasonable.

The solution to this problem generates shadow prices on intermediates shipped from the parent to affiliate, and vice versa, as we will see when we examine the FOCs in Section IIC. These shadow prices are distinct from the "transfer" prices that the MNC attaches to goods shipped intra-firm for tariff and tax reporting purposes.

In Figure 2 and in our equation for Π, we assume that the "transfer price" on N_d, the intermediate shipped from the Canadian affiliate to the parent, is equal to P_f^2, the price of the same good if sold to final customers in Canada (and vice versa for N_f).[19] We do this because U.S. Internal Revenue Service (IRS) code section 482 and Canadian Income Tax Code section 69 impose an "arms-length" standard on transfer prices, meaning they should be set in the same way as prices charged to unaffiliated buyers. We assume MNCs follow this legal standard, and ignore "transfer price manipulation." The corporate tax structures of the United States and Canada do create incentives to manipulate transfer prices to shift profits to Canada, as corporate tax rates are lower in Canada for manufacturing firms. But Lorraine Eden (1998) reviews empirical work on transfer price manipulation and shows that the evidence for such behavior by MNCs in the U.S.-Canada context is quite weak.[20]

Finally, recall Assumption 8: the expected rate of profit is equalized across firms. We adopt this because we lack reliable measures of payments to capital, $\delta_d K_d$ and $\delta_f K_f$, for parents and affiliates. By assuming a particular profit rate, we can back out the payments to capital from data on total revenues and payments to the other factors.[21] Thus, we treat the profit rate, denoted R, as an unknown parameter to be estimated. We discuss the intuition for its identification in Section IIID. In the next section we discuss some features of MNC behavior in this model.

C. *Solution of the Firm's Problem and Derivation of the FOCs*

We can express the FOCs more compactly if we first define

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

$$= \left(\frac{\mathrm{P_d^1S_d - P_d^1N_f(T_f + C_f)}}{\mathrm{P_d^1S_d}}\right),$$

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

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

and express the adjustment cost term in the FOC for domestic labor (L_d) as $E(FD)$, where

---

[19] Suppose we had assumed that the price on N_d, the intermediate good shipped from the affiliate to the parent, was P_d^2, the price on imports from the affiliate to final consumers in the United States, rather than $\mathrm{P_f^2}$. This would make it impossible to solve our model for firms with $N_d > 0$ but $I = 0$, because then P_d^2 is undefined.

[20] MNCs may eschew transfer price manipulation because they can use other means, such as licensing fees for intangible inputs, to shift profits. According to Eden (1998), such methods may be more common since they are more difficult to detect. In light of this, we view our "arms-length" pricing assumption as reasonable.

[21] The procedure works as follows. Denote domestic revenue by RD, domestic costs by CD^*, and domestic costs excluding capital costs by $CD1$. These quantities are given by

$$RD = P_d^1 S_d + P_d^1 N_f + (1 - T_f - C_f)(P_f^1 E),$$

$$CD^* = w_d L_d + \gamma_d K_d + \phi_d M_d$$

$$+ P_f^2 N_d(1 + T_d + C_d) + AC_d,$$

$$CD1 \equiv \mathrm{CD} - \gamma_\mathrm{d} \mathrm{K_d}.$$

Now, let R_K denote the fraction of operating profit that is pure profit, leaving $(1 - R_K)$ as the fraction that is the payment to capital. This gives $\Pi_d = R_K \cdot [RD - CD1]$, and thus

$$\gamma_d K_d = (1 - R_K) \cdot [RD - CD1].$$

Therefore, the rate of profit for domestic operations is $R = \Pi_d / \gamma_d K_d = R_K/(1 - R_K)$. 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

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

$$
\begin{aligned}
FD = {} & \partial \mu((L_{dt} - L_{d,t-1})^2)^{\mu-1}(L_{dt} - L_{d,t-1}) \\
& \div L^{\Delta}_{d,t-1} - [\beta\mu((L_{dt+1} - L_{dt})^2)^{\mu-1} \\
& \times (L_{dt+1} - L_{dt})/L^{\Delta}_{dt}] - [\beta\Delta((L_{dt+1} \\
& - L_{dt})^2)^{\mu}(L_{dt+1} - L_{dt})/L^{\Delta+1}_{dt}].
\end{aligned}
$$

The first-order conditions for parent factor inputs and parent exports to Canada are then

$$
\begin{aligned}
L_d &: \alpha^{Ld}(1 - g_1 A)\left(\frac{P^1_d Q_d}{L_d}\right) \\
& \quad - w_d - E(FD) = 0, \\
K_d &: \alpha^{Kd}(1 - g_1 A)\left(\frac{P^1_d Q_d}{K_d}\right) - \gamma_d = 0, \\
M_d &: \alpha^{Md}(1 - g_1 A)\left(\frac{P^1_d Q_d}{M_d}\right) - \phi_d = 0, \\
N_d &: \alpha^{Nd}(1 - g_1 A)\left(\frac{P^1_d Q_d}{N_d}\right) \\
& \quad + g_2 P^2_f B - (1 + T_d + C_d)P^2_f = 0, \\
E &: (1 - g_1)P^1_f(1 - T_f - C_f) \\
& \quad - (1 - g_1 A)P^1_d = 0.
\end{aligned}
$$

For the affiliate, the first-order conditions for L_f, K_f, M_f, N_f, and I are similar.

The FOCs for L_d, K_d, and M_d are familiar, except for the A term that multiplies g_1. Note that A is one minus the ratio of tariff and transport costs on intermediates to domestic sales. With no intra-firm flows, $A = 1$, and the g_1 terms capture the incentive of a firm with market power to hold down output in order to raise the sales price. But, with intra-firm flows, $A < 1$, and the incentive to reduce output is mitigated, since a higher price also raises the firm's tariff costs.

The FOC for N_d, the part of affiliate output used as an intermediate input by the parent, equates the marginal revenue product from increasing the input of N_d in domestic production to the effective cost of importing N_d. This effective cost can be written $(1 - g_2 B)P^2_f + (T_d + C_d)P^2_f$. The first term is lost revenue from failing to sell N_d in Canada. The term $g_2 B P^2_f$ arises because the affiliate has market power. The second term is the tariff and transport cost. Note that the shadow price is a completely separate quantity from the transfer price, P^2_f, since obviously it is not optimal for the affiliate to charge the parent the same price it charges third parties.

Similarly, the FOC for E equates marginal revenue from exports of the U.S.-produced good to Canada with that from domestic (U.S.) sales. If Canadian tariff and transport costs are zero, then $(C_f + T_f) = 0$, so $A = 1$, and the FOC reduces to $P^1_f = P^1_d$. Thus, the MNC exports good 1 up to the point where U.S. and Canadian prices are equalized. Tariffs and transport costs create a wedge between U.S. and Canadian prices for final goods. For instance, if the affiliate does not use the U.S.-produced good as an intermediate, then $A = 1$, and the FOC for exports reduces to $P^1_d/P^1_f = (1 - T_f - C_f) < 1$. Thus, the U.S.-produced good costs more in Canada.[22] If, however, the affiliate uses good 1 as an intermediate, then $A < 1$. The parent then has an incentive to hold down P^1_d because the tariff and transport cost of shipping intermediates to Canada is increasing in P^1_d. This effect is stronger the less elastic is demand (i.e., the larger is g_1) and the higher are tariffs.[23]

D. *Stochastic Specification*

Our model contains eight parameters (R, β, δ_d, δ_f, μ, Δ, H_d, and H_f) which we assume are common across firms. It also contains eight technology parameters (α^{Kd}, α^{Ld}, α^{Nd}, α^{Md}, α^{Kf}, α^{Lf}, α^{Nf}, and α^{Mf}) and six demand function parameters (g_1, P^1_{0d}, P^1_{0f}, g_2, P^2_{0d}, and P^2_{0f}) which we allow to be heterogeneous both across firms

---

[22] Without our assumption that price elasticity of demand for good 1 is equal in both countries, the price ratio would also depend on the difference in demand elasticities. But with equal elasticities, the incentive to raise the price by reducing sales is equally strong in both countries, so price is equalized except for the tariff and transport cost wedge.

[23] If the parent were to lower P^1_d too much relative to P^1_f, it could create an arbitrage opportunity whereby third parties could buy the good in the United States and resell it in Canada. We could invoke a segmented-markets (or no-resale) condition as in Brander (1981) to rule out this possibility. Such an assumption is common in trade models with differentiated products. Alternatively, we could assume transport costs are higher for third parties.

and within firms over time. Given CRTS, two Cobb-Douglas share parameters (α) are determined by the other six. Thus, we have 12 heterogeneous parameters that vary independently. We adopt a random effects specification, where each parameter has a mean, a firm-specific component, and a firm/time-specific component.

A key feature of our specification is that we allow for technical change by letting the means of the Cobb-Douglas share parameters have time trends. If the share parameters are stable over time (i.e., the time trends are insignificant or quantitatively small), it implies that changes in tariffs, input prices, exchange rates, and demand conditions, along with TFP, can explain all the changes in MNC behavior over time, holding technology fixed. But, if the time trends for the share parameters are significant, it implies that changes in tariffs and the other forcing variables cannot, by themselves, explain all the changes in MNC behavior during our analysis period.

In preliminary analysis, we found that parents (affiliates) that do not use intermediates from affiliates (parents) have stable Cobb-Douglas share parameters. Given the large relative changes in factor input prices from 1984 to 1996 (see Figure 3), this is striking evidence for a simple Cobb-Douglas specification. But we also found that parents and affiliates that use intermediates have important time trends in the share parameters. Thus, we allow the time trends and base values of the technology and demand parameters to differ for firms that use intermediates.

Production Function Parameters.—Letting the Cobb-Douglas share parameters be stochastic, while also imposing they are positive and sum to one (CRTS), is challenging. To achieve this, we use a logistic transform, treating the share parameters as analogous to choice probabilities in a multinomial logit (MNL) model. For instance, for the domestic labor share, we have, suppressing firm and time subscripts,

$$(7)\quad \alpha^{Ld} = \frac{\alpha_R^{Ld}}{1 + \alpha_R^{Ld} + \alpha_R^{Md} + \alpha_R^{Kd}} = \frac{\exp\{x\bar{\alpha}^{Ld} + \varepsilon^{Ld}\}}{1 + \exp\{x\bar{\alpha}^{Ld} + \varepsilon^{Ld}\} + \exp\{x\bar{\alpha}^{Md} + \varepsilon^{Md}\} + \exp\{x\bar{\alpha}^{Kd} + \varepsilon^{Kd}\}},$$

where the vector x_{it} includes all firm characteristics that shift the share parameters, and $\bar{\alpha}^{Ld}$ is a vector of parameters. The expressions for α^{Kd} and α^{Md} are similar. Note the expression for α^{Nd} is:

$$(8)\quad \alpha^{Nd} = \frac{1}{1 + \exp\{x\bar{\alpha}^{Ld} + \varepsilon^{Ld}\} + \exp\{x\bar{\alpha}^{Md} + \varepsilon^{Md}\} + \exp\{x\bar{\alpha}^{Kd} + \varepsilon^{Kd}\}}.$$

So α^{Nd} plays the role of the "base alternative" in a multinomial logit model. This specification insures that, for any values of x_{it} and ε_{it}, the share parameters will be positive and sum to one.

Note that in (7) the quantities α_R^{Ld}, α_R^{Md}, and α_R^{Kd} are simply latent variables that map into the share parameters. A natural specification for these latent variables is log normality (although we generalize this below). Then, for example, the stochastic term α_R^{Ld} would be given by: $\ln \alpha_R^{Ld} = x\bar{\alpha}^{Ld} + \varepsilon^{Ld}$, where $\varepsilon^{Ld} \sim N(0, \sigma_{Ld}^2)$, and similar equations could be specified for α_R^{Md}, α_R^{Kd}. This insures that $\alpha_R^{Ld} > 0$ for all ε^{Ld}, which guarantees that the labor share is positive.

In our empirical application, we allow x_{it} to include just an intercept and a time trend t ($t = 0$ in 1983). But, as we noted above, we allow different intercepts and time trends for parents (affiliates) that do and do not use intermediate inputs from affiliates (parents).

If the parent is not structured to use intermediates from the affiliate, then $\alpha^{Nd} = 0$, and the three share parameters, α^{Ld}, α^{Md}, and α^{Kd}, must sum to one. This constraint is imposed just as above, except now α^{Kd} plays the role of the base alternative. A similar construct is used for affiliates that do not use intermediates from the parent. Because the scale of the coefficients in an MNL with three alternatives differs from that in a model with four, we introduce a parameter to scale down the error terms in the three-alternative case. Thus, for the α_R^{Ld} equation, we have

$$(9)\quad ln\ \alpha_R^{Ld} = \alpha_0^{Ld} + \alpha_{shift}^{Ld} I[N_d > 0] + \alpha_{Time}^{Ld} \cdot t \cdot I[N_d > 0] + \alpha_{Time}^{Ld} \cdot t \cdot I[N_d = 0] \cdot SC_d + \varepsilon^{Ld} \times \{I[N_d > 0] + SC_d \cdot I[N_d = 0]\}.$$

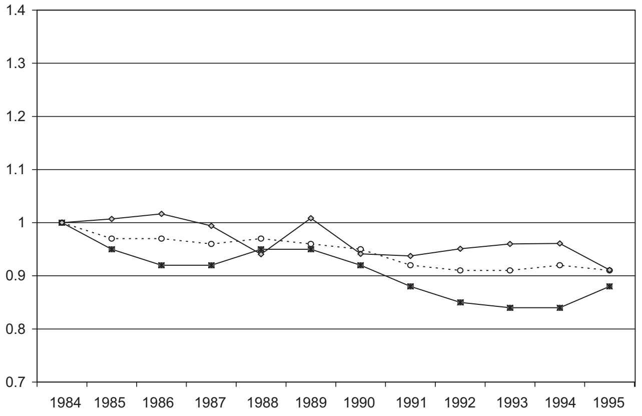

A. U.S. prices

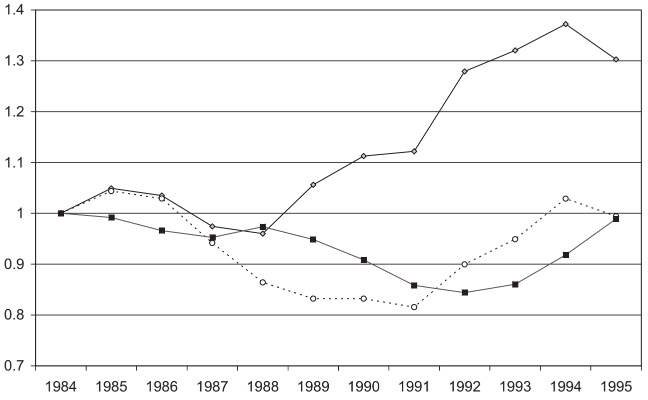

B. Canadian prices (Canadian $)

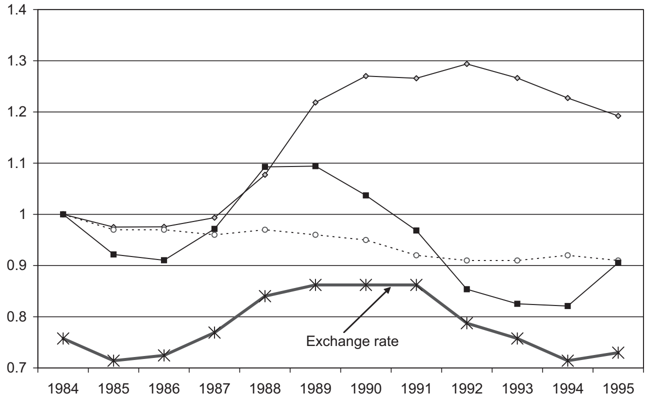

C. Canadian prices (US $)

FIGURE 3. U.S. AND CANADIAN REAL PRICES OF CAPITAL, MATERIALS, AND LABOR

Notes: Real prices are nominal wages or price indices divided by the GDP deflator. Nominal wages were obtained from the BEA sample used in this research. In panel B, Canadian affiliate wages were converted to Canadian dollars using the nominal exchange rate (Canadian dollars/U.S. dollars). All prices in panel B are deflated using the Canadian GDP deflator. Prices in panel A and panel C are deflated using the U.S. GDP deflator. The exchange rate in the panel C is expressed in U.S. dollars/Canadian dollars.

Similar equations hold for α_R^{Md} and α_R^{Kd}, except that we simply have $\alpha_R^{Kd} = 1$ in the $N_d = 0$ case. The same scaling parameter, SC_d, applies in the α_R^{Md} equation in the $N_d = 0$ case. In the equations for the affiliate parameters α_R^{Lf}, α_R^{Mf}, and α_R^{Kf}, we allow for a different scaling parameter, SC_f.

In preliminary work, the log normality assumption was severely rejected for most of the stochastic terms. So we turned to a Box-Cox transformation. Then, equation (9) becomes

$$(10)\quad \frac{(\alpha_R^{Ld})^{bc(1)} - 1}{bc(1)} = x\bar{\alpha}^{Ld} + \varepsilon^{Ld}\{I[N_d > 0] + SC_d \cdot I[N_d = 0]\},$$

where $bc(1)$ is the Box-Cox parameter.[24] Expressions like (10) also hold for α_R^{Md} and α_R^{Kd}, and for the affiliate parameters. The Box-Cox parameters in these equations are $bc(2)$ through $bc(6)$.[25]

Turning to the correlations of the ε, we specify that

$$(11)\quad (\varepsilon^{Ld}\quad \varepsilon^{Md}\quad \varepsilon^{Nd})' \sim N(0, \boldsymbol{\Sigma}^d),$$

where $\boldsymbol{\Sigma}^d$ is unrestricted. Similarly, for affiliates, $\boldsymbol{\Sigma}^f$ is unrestricted. But, in order to conserve on parameters, we do not allow for covariances between the parent and affiliate share parameters.[26]

Finally, consider the TFP parameters H_d and H_f in equations (1) and (2). Since we do not observe output prices and quantities separately, we cannot identify the scale of the H (either absolutely or for the affiliate relative to the parent). We can, however, identify technical progress. Thus, we normalize $H_d = H_f = 1$ at $t = 0$ (1983) and let each have a time trend:

$$(12)\quad H_{jt} = (1 + h_j)^t \qquad for\ j = d, f.$$

We could not reject a specification with equal time trends, so we set $h_d = h_f = h$.

Demand Function Parameters.—Now, we turn to the stochastic specification for the demand function parameters. For the inverse price elasticity of demand, or market power, parameter g_1, we specify

$$(13)\quad \frac{g_1^{bc(7)} - 1}{bc(7)} = g_{10} + g_{1,time} \cdot t + g_{1,shift} \cdot I[Nd > 0] + \varepsilon^{g_1}.$$

A similar equation holds for g_2, and the Box-Cox parameter in that equation is $bc(8)$.

For the demand function intercepts for good 1 in the domestic market, we specify[27]

$$(14)\quad \frac{(P_{0d}^1)^{bc(9)} - 1}{bc(9)} = P_{0d,0}^1 + P_{0d,time}^1 \cdot t + P_{0d,shift}^1 \cdot I[Nd > 0] + \varepsilon^{P_{0d}^1}.$$

Similar equations hold for P_{0f}^2, P_{0f}^1, and P_{0d}^2, and the Box-Cox parameters in these equations are denoted by $bc(10)$, $bc(11)$, and $bc(12)$, respectively.

Preliminary work showed that cross correlations among the three groups of parameters (technology, price elasticities, and demand function intercepts), were not important. Allowing for such correlations leads to a severe proliferation of parameters. Thus, we assume ε^{g_1} and ε^{g_2} are independent of other stochastic

---

[24] Strictly speaking, the Box-Cox transform does not guarantee positive share parameters. But, given our estimates of the Box-Cox parameters and the variances of the stochastic terms, negative outcomes would be extreme outliers.

[25] Note that we only have an equation for α_R^{Kd} in the case of $N_d > 0$, because if $N_d = 0$ we normalize $\alpha_R^{Kd} = 1$. Similarly, if $N_f = 0$ we normalize $\alpha_R^{Kf} = 1$. Thus, for illustration, the equation for α_R^{Kd} is just

$$\frac{(\alpha_R^{Kd})^{bc(3)} - 1}{bc(3)} = a_0^{Kd} + a_{time}^{Kd} \cdot t + \varepsilon^{Kd}.$$

[26] Interpretation of the $\boldsymbol{\Sigma}^d$ and $\boldsymbol{\Sigma}^f$ terms is rather subtle. The logistic transformation already incorporates the negative correlation among the share parameters that is generated by the CRTS assumption. If $\boldsymbol{\Sigma}^d = \mathbf{I}$, we have an "IIA" setup, where if one domestic share parameter increases, the other share parameters decrease proportionately. The correlations in $\boldsymbol{\Sigma}^d$ and $\boldsymbol{\Sigma}^f$ allow firms to depart from this IIA situation. For example, if $\boldsymbol{\Sigma}_{12}^d$ is very large, then we get a pattern where firms with large domestic labor shares also have large domestic materials shares.

[27] Technically, we should impose that g_1 and g_2 are positive and less than one, and that the P_o terms are positive. Equations (12) and (13) do not impose these constraints. But, given our estimates, violations would be extreme outliers.

terms. We let the $(\varepsilon^{P^1_{0d}}, \varepsilon^{P^2_{0f}}, \varepsilon^{P^1_{0f}}, \varepsilon^{P^2_{0d}})$ vector be correlated within itself with covariance matrix $\boldsymbol{\Sigma}^P$, but it is independent of the other stochastic terms.

Labor Force Adjustment Cost Parameters.—Recall that labor force adjustment costs are given by equation (5). The parameters δ_d and δ_f are allowed to vary across firms as follows:

$$(15)\quad \delta_{dt} = exp\{\delta_{0d} + \delta_{dw} w_{dt} + \delta_{d,time} \cdot t + \delta_{dN} \cdot I[N_{dt} > 0]\},$$

$$\delta_{ft} = exp\{\delta_{0f} + \delta_{fw} w_{ft} + \delta_{f,time} \cdot t + \delta_{fN} \cdot I[N_{ft} > 0]\}.$$

As with the other structural parameters, we allow for the possibility that adjustment costs vary over time, and we allow for the possibility that the adjustment costs differ between firms that do and do not have intra-firm flows. We also allow the δ to be functions of the wage rate, on the premise that search and severance costs for high-wage/highly skilled labor are higher.

Serial Correlation.—We model serial correlation of the errors for each firm using a random-effects structure. For example, for the stochastic part of the parent labor share parameter ε^{Ld}, we have:

$$\varepsilon^{Ld}(it) = \mu^{Ld}(i) + \nu^{Ld}(it) \qquad \text{for } t = 1, \ldots, T_i$$

and similarly for the other eleven parameters. Let $\boldsymbol{\mu}_i \sim N(0, \boldsymbol{\Sigma}_\mu)$ denote the 12×1 vector of random effects for firm i, and let $\boldsymbol{\nu}_{it} \sim N(0, \boldsymbol{\Sigma}_\nu)$ denote the 12×1 vector of firm/time-specific error components. Then, $\mathbf{V}_t \equiv Var(\varepsilon_{it}) = \boldsymbol{\Sigma}_\mu + \boldsymbol{\Sigma}_\nu$ and $\mathbf{C}_{t-j,t} \equiv Cov(\varepsilon_{it}, \varepsilon_{i,t-j}) = \boldsymbol{\Sigma}_\mu$. Note that $\boldsymbol{\Sigma}_\mu$ and $\boldsymbol{\Sigma}_\nu$ each contain 78 unique elements, but these are restricted as described earlier. For instance, the nonzero elements of $\boldsymbol{\Sigma}_\mu + \boldsymbol{\Sigma}_\nu$ consist entirely of elements of $\boldsymbol{\Sigma}^d$, $\boldsymbol{\Sigma}^f$, and $\boldsymbol{\Sigma}^P$ along with σ^2_{g1} and σ^2_{g2}. Other cross correlations are set to zero. Defining $\boldsymbol{\varepsilon}_i = (\varepsilon_{i1}, \ldots, \varepsilon_{iT_i})$, we have

$$(16)\quad Var(\boldsymbol{\varepsilon}_i) = \begin{pmatrix} \mathbf{V}_1 & & & \\ \mathbf{C}_{12} & \mathbf{V}_2 & & \\ \vdots & & \ddots & \\ \mathbf{C}_{1T_i} & \cdots & \cdots & \mathbf{V}_{T_i} \end{pmatrix}.$$

So far, we have considered the general case where a firm has all four potential trade flows. If a firm has $N_{dt} = 0$ (or $N_{ft} = 0$) at time t, then there is no value for $\varepsilon^{Kd}(it)$ (or $\varepsilon^{Kf}(it)$), while if $E_t = 0$, there is no value for $P^1_{0f}(it)$, and if $I_t = 0$, there is no value for $P^2_{0d}(it)$. Also, some firms are not observed in all years. In such cases, $Var(\boldsymbol{\varepsilon}_i)$ is collapsed by removing the irrelevant rows and columns.

III. Estimation

A. *Overview*

In this section we discuss our estimation procedure. Although our model is quite simple, the fact that we allow many parameters to be heterogeneous across firms and time complicates estimation. We base estimation on the FOCs of the MNC's optimization problem. Estimation based on FOCs is often implemented using generalized method of moments (GMM), but this is impractical here because multiple stochastic terms enter the FOCs nonlinearly. Hence, we cannot manipulate the FOCs to obtain a single additive error term from which to construct moments. We instead use the simulated maximum likelihood (SML) approach to estimation based on FOCs developed in Keane (2003).

B. *Derivation of Estimable FOCs*

Since prices and quantities are not separately observed, we cannot take the FOCs from Section IIC directly to the data. We must first manipulate them to obtain estimable equations that contain only observed quantities and unknown model parameters. First, if we multiply each first-order condition by the associated control variable, we obtain

$$(17)\quad L_d : \alpha^{Ld}(1 - g_1 A)(P^1_d Q_d) - w_d L_d - \delta_d E(FD) L_d = 0,$$

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

$$M_d : \alpha^{Md}(1 - g_1 A)(P^1_d Q_d)$$

$$- \phi M_d = 0,$$

$$N_d : \alpha^{Nd}(1 - g_1 A)(P_d^1 Q_d)$$

$$+ g_2(P_f^2 N_d)B$$

$$- (1 + T_d + C_d)(P_f^2 N_d) = 0,$$

$$E : (1 - g_1)(P_f^1 E)(1 - T_f - C_f)$$

$$- (1 - g_1 A)(P_d^1 E) = 0.$$

In the FOC for E, the quantity $P_d^1 E$ is not observable.[28] We can, however, exploit the fact that

$$(P_d^1 E) = \left(\frac{P_d^1}{P_f^1}\right)(P_f^1 E) = \left(\frac{P_{0d}^1}{P_{0f}^1}\right)\left(\frac{P_d^1 S_d}{P_f^1 E}\right)^{-g_1}$$

$$\cdot (P_f^1 E)$$

to express the FOC for E in terms of observable quantities and the demand function intercepts (P_{0d}^1/P_{0f}^1), which we treat as unknown parameters, as follows:

$$(18) \quad E\colon (1 - g_1)(P_f^1 E)(1 - T_f - C_f)$$

$$- (1 - g_1 A)\left(\frac{P_{0d}^1}{P_{0f}^1}\right)\left(\frac{P_d^1 S_d}{P_f^1 E}\right)^{-g_1}$$

$$\cdot (P_f^1 E) = 0.$$

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 $P_d^1 Q_d = P_d^1 S_d + P_d^1 N_f + (P_d^1/P_f^1)(P_f^1 E)$ but, again, $P_d^1 E$ is not observed. We therefore repeat the same type of substitution to obtain, for domestic labor:

$$(19) \quad L_d : \alpha^{Ld}(1 - g_1 A)[P_d^1 S_d + P_d^1 N_f$$

$$+ \left(\frac{P_{0d}^1}{P_{0f}^1}\right)\left(\frac{P_d^1 S_d}{P_f^1 E}\right)^{-g_1} (P_f^1 E)$$

$$- w_d L_d - \delta_d E(FD)L_d = 0.$$

The FOCs for K_d, M_d, and N_d, and for the affiliate, are obtained similarly.

Finally, we deal with the unobserved expectation term $E(FD)L_d$ in (19) by invoking a rational expectations assumption:

$$(20) \qquad E_t(FD_{it})L_{dit} = FD_{it}L_{dit} - \eta_{it}^d,$$

where η_{it}^d is a forecast error assumed orthogonal to all information available at time t.

Substituting (20) into (19), we see that the FOC will contain *five* firm-specific stochastic terms (α_{it}^{Ld}, g_{1it}, P_{0d}^1, P_{0f}^1, and η_{it}^d) that enter nonlinearly. Hence, it is not possible to express the equation as a moment condition with a single additive error term.[29, 30]

C. *The SML Based on FOCs Approach, and the Stochastic Process for Forecast Errors*

The SML approach to estimation based on FOCs developed in Keane (2003) can deal with the fact that multiple stochastic terms enter the FOCs in (18) and (19) in a highly nonlinear way. But the SML approach requires that we specify a distribution for the forecast errors η_{it}^d and η_{it}^f.

Without loss of generality, we rewrite (20), the equation for forecast errors, as follows:

$$(21) \quad E(FD_{it})L_{dit} = FD_{it}L_{dit} + \sigma_{dt}\eta_{dt}^*,$$

where η_{dt}^* is standard normal. We tried a generalization where η_{dt}^* is normal subject to a Box-Cox transform. But the estimated Box-Cox parameter was extremely close to one, implying that the distributions of forecast errors are well described by normality. We expect that σ_{dt}, the standard deviation of the labor adjustment cost forecast error, will be increasing in L_{dit}, so we write:

$$(22) \qquad \sigma_{dt} = exp\{\tau_{d0} + \tau_1 L_{dt}\},$$

$$\sigma_{ft} = exp\{\tau_{f0} + \tau_1 L_{ft}\}.$$

---

[28] 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.

[29] One may implement simulated nonlinear GMM by integrating α_{it}^{Ld}, g_{1it}, P_{0d}^1, P_{0f}^1 out of (19), basing moments on η_{it}^d alone. But this requires distributional assumptions on α_{it}^{Ld}, g_{1it}, P_{0d}^1, P_{0f}^1, losing the key appeal of method of moments relative to ML.

[30] Even if we could linearize (19), finding valid instruments is difficult. Usual candidates like input prices would be correlated with firm-specific technology parameters if technology changes over time in response to price changes.

Regarding serial correlation, we assume that the forecast errors are independent over time, as implied by rational expectations. But we allow parent and affiliate forecast errors to be correlated within a period, as must be the case if their production processes are integrated, or if they face common shocks. Thus, we let $(\eta^d \eta^f)' \sim N(0, \mathbf{\Sigma}_\eta)$. We also let $\mathbf{\Sigma}_\eta = \mathbf{CC}'$, where $\mathbf{C}$ is the lower triangular Cholesky decomposition $\left(\begin{smallmatrix} \mathbf{C}_{11} & \mathbf{0} \\ \mathbf{C}_{12} & \mathbf{C}_{22} \end{smallmatrix}\right)$. Finally, let $\boldsymbol{\tau} = (\tau_{d0}, \tau_{f0}, \tau_1)$.

The alternative to specifying a distribution on forecast errors is to adopt a full information ML approach. This would 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 means completely specifying the stochastic processes for all these forcing variables, which would go well beyond the scope of the present investigation.

The SML based on FOCs approach is a compromise in which we specify parametric distributions for the demand and taste shocks, as in a full ML approach. But, rather than specify stochastic processes for all the forcing variables (e.g., tariffs, wages, etc.), we simply substitute realizations of the $t + 1$ labor demand terms for their expectations, as in a typical GMM approach.

Per Krusell et al. (2000) confronted a similar problem in estimating a production function with quality of skilled and unskilled labor as two latent stochastic inputs. The FOCs of their model also contained an unmeasured expectation of next period's price of capital. Thus, three stochastic terms entered nonlinearly, so the FOCs could not be written in terms of a single additive error. They used a simulated pseudo-ML procedure, assuming normality for forecast errors. Keane (2003) extends this idea by showing how to do SML, relax normality (via the Box-Cox transformation), and test the distributional assumptions.[31]

In terms of what we can do with our model once it is estimated, SML based on FOCs represents a compromise between the full solution ML approach and the GMM approach. Since we 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.[32] 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.[33]

D. *Identification of Parameters*

The parameter vector in our model, which we denote by $\boldsymbol{\theta}$, includes values for the common (or nonstochastic) parameters of the model, which are R, β, δ_d, δ_f, μ, Δ, $\boldsymbol{\tau}$, and h, and also for the parameters of the joint distribution of the 12 firm-specific stochastic terms (see Section IID). The simulation of the likelihood contribution for a firm involves the following steps. First, take a draw from the joint distribution of the forecast errors, η^d and η^f. Second, use the ten FOCs for L_d, K_d, M_d, N_d, L_f, K_f, M_f, N_f, E, and I (see equations (18) and (19)), and the production functions (1) and (2) as a system of 12 equations to solve for the 12 stochastic terms that rationalize the firm's behavior in each time period. Third, calculate the joint density of the stochastic terms, using the multivariate normal distribution with covariance matrix given by (16). Fourth, multiply by the Jacobian to obtain the data density. Repeat this process at independent draws for η^d and η^f, and average the data densities obtained in each case, to obtain a simulation consistent estimate of the likelihood contribution for the firm (see Keane, 2003, for details).

Solving the system of 12 nonlinear equations for the 12 stochastic terms in the model is cumbersome. Keane (2003) provides details. A

---

[31] Keane (2003) also develops new simulation algorithms to implement these procedures. These are recursive importance sampling algorithms that extend the idea of the GHK algorithm to the discrete/continuous and continuous cases.

[32] Note that, even if GMM estimation were feasible for our model, it would not be adequate for this purpose. The usual argument for GMM over ML is that one avoids making distributional assumptions on the stochastic terms, and thereby obtains more robust estimates of model parameters. We must estimate the *distributions* of the firm-specific parameters so we can simulate the response of the population of firms to changes in tariffs and other variables.

[33] It is worth emphasizing that the key difficulty in estimation arises not from dynamics, but rather because *multiple* stochastic terms enter the FOCs. This problem would be present in a static model without labor adjustment costs.

basic understanding of the process, however, provides intuition for how the model parameters are identified, so we give a brief summary here.

The market power parameters, g_1 and g_2, are identified by simple markup relationships (i.e., Lerner conditions), which we can construct using data on sales revenues and costs. These relationships are modified slightly to account for labor force adjustment costs, and also for the complication that the MNC has an incentive to hold down prices of final goods it ships intra-firm as intermediates (in order to avoid tariff costs). The U.S./Canadian price ratios for goods 1 and 2, denoted PR_1 and PR_2, are determined by the tariff and transport cost wedge, again modified by the incentive to hold down prices of intra-firm intermediates. Since the strength of this incentive depends only on g_1 and g_2, we can solve for PR_1 and PR_2 once g_1 and g_2 are obtained. Also, given g_1 and g_2, the Cobb-Douglas share parameters are identified by cost shares of modified revenues.

Finally, given g_1 and PR_1, we can infer the *ratio* of the U.S.-to-Canadian demand function intercepts for good 1 (i.e., P^1_{0d}/P^1_{0f}) by observing the ratio of U.S.-to-Canadian sales for good 1. Similarly, by comparing affiliate sales in Canada versus imports to the United States, we can infer the *ratio* of domestic-to-foreign demand function intercepts for good 2 (i.e., P^2_{0f}/P^2_{0d}).

Thus, without separate data on prices and quantities (except wages and employment, which we need only to identify the labor force adjustment cost function), we can identify the market power parameters, the Cobb-Douglas share parameters, and the ratios of the demand function intercepts for goods 1 and 2. To identify the *levels* of the demand intercepts, we need capital and materials price indices. Then, we can construct real capital and materials inputs and use the production functions (1)–(2) as additional equations to determine quantities of output.[34]

IV. The BEA Data

A. *Construction of the Panel*

Our data are from the Benchmark and Annual Surveys of U.S. Direct Investment Abroad administered by the BEA. These surveys contain the most comprehensive information available on the activities of the population of U.S.-based MNCs and their foreign affiliates. For this study, we use the confidential disaggregate BEA data on U.S. MNCs with one or more Canadian affiliates. For the 1983–1996 period, this subset of the BEA data contains 24,313 affiliate-year observations.

We made several alterations to the original BEA data to construct our panel dataset. First, we use only data on manufacturing affiliates, since many nonmanufacturing industries produce nontradeables. Furthermore, in our model we assume that the intermediates shipped intra-firm are goods destined for further processing. Such an assumption would obviously be invalid if we included, for instance, wholesale and retail trade affiliates. Limiting the dataset to manufacturing affiliates reduced the number of observations from 24,313 to 12,241.

Second, we must assign each affiliate to an industry, in order to match it with the appropriate tariff and transport cost data. The large majority of affiliates are not diversified, so the appropriate assignment is clear.[35] There were, however, cases that appeared to be spurious changes in industry classification, or where affiliates consistently had fewer than 80 percent of sales in a single industry. We dropped such cases, leading to a loss of 1,677 affiliate-year observations.

Third, while the Benchmark Surveys, conducted in 1977, 1982, 1989, and 1994, include all U.S. MNC affiliates, smaller affiliates may be exempt from filing the Annual Surveys. If an affiliate reports data in a Benchmark Survey but is exempt from the Annual Surveys, the BEA carries it forward by estimating data. As a result, most of the data for smaller affiliates in

---

[34] Keane (2003) also gives an intuitive explanation of how the time trends in TFP (h) and in the demand function intercepts (the P_0) are separately identified, and how the profit rate R is identified. Briefly, to the extent that growth is more than proportionately slower for firms with more market power, it implies that growth is induced by TFP rather than growth in demand. Regarding the profit rate, the model implies a relation $g/(1-g) = \alpha^K \cdot R$ between market power and capital share. If the profit rate is low (high), there is a strong (weak) tendency for firms with more market power also to have larger capital shares, so that profits accrue to a larger (smaller) stock of capital.

[35] On average, 91 percent of Canadian manufacturing affiliate sales were in only one industry. The median affiliate sells all output in one industry. By contrast, roughly 65 percent of U.S. parent sales were in one industry.

non-Benchmark years is estimated rather than reported.

Ideally, we would prefer to remove all the estimated data. Since we use leads and lags of employment, however, to construct labor force adjustment costs, dropping one estimated data point could, in some instances, cost us two additional reported observations on either side of it. Thus, we decided to drop most of the estimated data, but to keep estimated observations if: (a) they were bracketed at $t - 1$ and $t + 1$ by valid reported observations; and (b) the affiliate also had, at some point, at least three consecutive valid reported observations. This process led to the elimination of 4,247 affiliate-year (estimated) observations, leaving 6,358 affiliate-year observations.[36]

Next, data on same-industry affiliates of the same parent were combined, leaving 5,583 affiliate-year observations. Finally, after removing observations with missing data, we were left with 5,175 firm-year observations on 551 parents and 716 affiliates.

Our model abstracts from the possibility that an MNC might have multiple Canadian affiliates. We care about how tariffs and other factors affect the allocation of MNC activities between the United States and Canada, not the organizational form of an MNC's Canadian operations. Therefore, if a parent had multiple affiliates, we merged them into one "composite" affiliate in two steps. First, if the parent had multiple affiliates in one industry, we merged them into a composite affiliate by adding up the inputs and outputs of the individual affiliates. Second, if the parent had affiliates in multiple industries, we took only the largest, based on total sales.[37]

Our model does not allow for the possibility that an affiliate would have no domestic sales in Canada (this is a very rare event in the data). We therefore deleted eight cases of parent-affiliate pairs where the affiliate had zero Canadian sales. After this step, our dataset contained 3,385 unique affiliate-year observations on 543 unique parents and affiliates.

B. *Construction of Variables*

The BEA data contain U.S. parents' domestic sales, affiliates' domestic sales (in Canada), the value of intermediates shipped intra-firm (in both directions), and affiliates' arms-length sales to the United States. But, except in benchmark years, they do not contain U.S. parents' arms-length sales to Canada. To construct this variable, we used data from Compustat on total parent sales to Canada, and netted out the value of intra-firm shipments.[38]

We also need measures of capital, labor, and materials inputs. The BEA data contain information on employment and the wage bill for both parents and affiliates. We use the ratio of these variables as our measures of wage rates, which are assumed to be specific to each parent and affiliate.[39] This is the only instance where we observe separate data on price and quantity.

To construct measures of materials input, we use information on cost of goods sold (CGS). For affiliates, the BEA data contain direct information on CGS. We calculate affiliate materials input by subtracting the wage bill, imports from U.S. parents, and current depreciation from CGS. U.S. parents do not report CGS, so we needed to obtain this item from Compustat.[40]

As discussed in Section IIB, we constructed the payments to capital inside our estimation algorithm, since payments to capital are not recorded in the BEA data. The data do include PPE at historical cost, which is often used to construct payments to capital (by assuming a required rate of return and a depreciation rate). Given the

---

[36] There are 41 observations that would have been removed by both the second and third screens.

[37] This size comparison is done after the merging of same-industry affiliates in step 1.

[38] If Compustat data were not available, we assumed that a parent's ratio of arms-length to intra-firm exports was the same for Canada as it is worldwide. Then, we multiplied the parent's intra-firm exports to Canada by the worldwide ratio (which is contained in the BEA data). But if the U.S. parent had zero sales to the affiliate, this formula could not be used. In these cases, we multiplied a U.S. parent's total (worldwide) arms-length sales by the ratio of Canadian to worldwide arms-length sales for *all* U.S. parents (obtained from the Benchmark Survey *published* data).

[39] That is, each firm requires a particular type of labor (e.g., a particular skill level), with its own specific wage rate.

[40] Since CUSIP numbers are not regularly reported by U.S. parents in the BEA data, we generally used the name of the parent firm to match up the BEA and Compustat data. For parents with no Compustat data, we used the average value of the CGS to sales ratio for their industry to calculate CGS.

well-known limitations of historical PPE data, however, we preferred to implement our new procedure, described in Section IIB, of estimating a profit rate and backing out payments to capital as a residual. Interestingly, our measures turned out to be highly correlated with PPE.

All data on nominal quantities (i.e., sales, trade flows, costs of labor, and materials inputs) were put in real terms using the U.S. GDP deflator (results were little changed by using the consumer price index (CPI) or producer price index (PPI) instead). The BEA instructs firms to report Canadian values in current U.S. dollars, so we did not need to make any exchange rate adjustment. We implicitly assume that firms use the nominal U.S.-Canada exchange rate to do the conversion.

Our use of the GDP deflator amounts to an assumption that the MNC's objective is to maximize the present value of profits in GDP-deflated U.S. dollar terms. Thus, when a firm considers how its current-period labor input will affect future labor force adjustment costs, it evaluates future costs in GDP-deflated U.S. dollar terms. It is crucial to note that this assumption has no implication for the within-period *producer* prices of capital, labor, and materials. This is because we estimate firm-specific demand functions, implying firm-specific output prices. The ratio of nominal input prices to these firm-specific output prices determines producer prices.

As we discussed in Section IIID, all the parameters of our model except the *levels* of the demand function intercepts are identified without needing measures of real capital and materials inputs. But to identify these intercepts in levels, we need to measure real inputs. To put materials costs in real terms, we use, for U.S. parents, the Bureau of Labor Statistics (BLS) PPI for intermediate supplies, materials, and components for manufacturing, and, for Canadian affiliates, the PPI for manufacturing intermediates, obtained from Statistics Canada. For the cost of capital, we use the price index for gross private domestic (nonresidential) investment in producer durables from the BLS. The cost of capital is assumed to be the same in the United States and Canada.[41]

Key variables of interest in our model are tariffs and transport costs. We measure U.S. and Canadian tariffs on an ad valorem basis for each of our 50 manufacturing industries. That is, the tariff on imported goods in industry j in year t is measured as the ratio of duties paid to the value of the imports. A 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.[42] Such ad valorem freight-rate measures are commonly used in empirical work (see, e.g., Head and Reis, 2001; Hanson et al., 2005). And, while our tariff measures are more aggregate than the level at which tariffs are actually imposed, they are more disaggregated than measures often used in empirical work (see Harry Grubert and John Mutti, 1991).

As shown in Figure 1, U.S. and Canadian tariffs fell substantially over the 1983–1996 period. Canadian tariffs fell from an average of nearly 6 percent to 1.75 percent, and U.S. tariffs fell from 4 percent to less than 1 percent. There is also considerable cross-industry variation in tariffs. U.S. tariffs are highest in tobacco (average 13 percent) and lowest in motor vehicles and pulp and paper (average less than 0.2 percent). Canadian tariffs are highest in tobacco and apparel (both averaging over 17 percent), and lowest in agricultural chemicals, autos, and farm machinery (all averaging approximately 1 percent).

As the affiliates in our dataset are predominantly single-industry, it is straightforward to assign them appropriate U.S. tariff and transport cost data. For diversified U.S. parents, we construct sales-weighted average Canadian tariff and transport cost measures across the (up to) eight industries in which they report sales.

---

[41] Large differences in prices of physical capital between the United States and Canada would create arbitrage opportunities.

[42] U.S. tariff and transportation cost data were obtained from the U.S. Census Bureau. Canadian tariff data were obtained from Statistics Canada. Canadian tariffs were reported at the three-digit SIC code level, which were converted into U.S. SIC codes, then BEA ISI codes. Transportation cost data for imports into Canada were not available from Statistics Canada, so we simply assumed that these costs were the same as for imports into the United States.

TABLE 1—DESCRIPTIVE STATISTICS
(*Thousands of U.S. dollars*)

	Analysis dataset $n = 2335$	Complete dataset $n = 3855$
	Mean (Standard deviation)	Mean (Standard deviation)
Canadian affiliate flows		
Sales in Canada $P_f^2 S_f$	167,469	147,382
	(603,795)	(561,257)
Sales of intermediates to U.S. parents $P_f^2 N_d$	108,167	88,164
	(740,065)	(667,864)
Percent positive observations	75.2%	71.8%
Arms-length sales to U.S. $P_d^2 I$	12,393	10,660
	(41,440)	(38,389)
Percent positive observations	41.9%	39.1%
U.S. parent flows		
Sales in the U.S. $P_d^1 S_d$	2,768,442	2,534,900
	(7,168,396)	(6,668,197)
Sales of intermediates to Canadian affiliates $P_d^1 N_f$	96,772	78,218
	(673,554)	(598,080)
Percent positive observations	82.4%	81.2%
Arms-length sales to Canada $P_f^1 E$	29,316	25,290
	(136,145)	(119,988)
Percent positive observations	86.1%	84.9%
Canadian affiliate inputs		
Employee compensation	38,803	33,082
	(123,593)	(110,112)
Employment	1,279	1,110
	(3,724)	(3,315)
Materials cost	93,551	83,345
	(373,064)	(353,366)
U.S. parent inputs		
Employee compensation	674,557	605,531
	(1,881,555)	(1,711,614)
Employment	20,781	19,111
	(47,196)	(43,198)
Materials cost	1,240,438	1,132,683
	(3,546,697)	(3,293,377)

Notes: All variables are in thousands. All dollar variables are in 1984 U.S. dollars. The "analysis" dataset is smaller than the "complete" dataset because observations from 1983 and 1996 are dropped, as these years are used only to construct leads and lags. Estimation also requires three consecutive observations to create each data point, and some observations are lost due to missing observations.

V. Empirical Results

A. *Descriptive Statistics*

Table 1 gives descriptive statistics for the firms in our dataset. To estimate our structural model, we need one-period lags and leads of the labor force data, in order to construct current and future labor force adjustment costs. Thus, we can use only those firms with at least three consecutive observations. Also, while our panel extends from 1983 to 1996, we use the 1983 and 1996 data only to model adjustment costs. These additional screens take us from the "complete" dataset (543 firms, 3,885 affiliate-year observations) to the "analysis" dataset (446 firms and 2,335 affiliate-year observations). Table 1 reports summary statistics on both datasets.

In the analysis dataset, parents' intra-firm sales to the Canadian affiliates make up 34 percent of affiliates' total sales (which average about $288 million in 1984 U.S. dollars). Affiliates' intra-firm sales to parents are 38 percent of affiliate total sales. These figures imply a high degree of integration of the production processes of parents and affiliates. Using the

TABLE 2—INPUT SHARES BY TIME PERIOD
(*Percentages*)

	Input shares (mean)				Input shares (median)			
	Labor	Capital	Materials	Intra-firm intermediates	Labor	Capital	Materials	Intra-firm intermediates
Parents, $N_d > 0$								
1984–1986	28.7	30.7	39.4	1.2	28.0	31.5	38.5	0.2
1987–1989	26.7	31.9	39.7	1.7	25.6	32.3	40.6	0.3
1990–1992	26.3	32.5	39.2	2.0	25.7	32.6	38.5	0.5
1993–1995	25.2	32.4	39.6	2.8	23.7	32.6	39.5	0.9
Parents, $N_d = 0$								
1984–1986	27.8	33.1	39.1		26.5	33.3	38.8	
1987–1989	25.9	31.6	42.5		24.9	32.0	42.4	
1990–1992	26.0	31.9	42.1		25.6	32.5	42.2	
1993–1995	25.7	32.8	41.5		25.0	33.1	42.3	
Affiliates, $N_f = 0$								
1984–1986	21.3	28.1	35.5	15.0	20.8	27.5	36.5	9.8
1987–1989	20.5	26.2	38.1	15.3	19.8	26.3	38.4	10.5
1990–1992	22.0	22.4	40.0	15.5	21.8	22.7	38.5	9.7
1993–1995	20.5	21.2	39.5	18.9	19.7	21.2	38.4	13.6
Affiliates, $N_f = 0$								
1984–1986	25.6	31.4	43.0		23.3	30.8	44.6	
1987–1989	24.8	31.4	43.7		22.9	31.7	45.2	
1990–1992	26.2	28.7	45.1		26.2	30.0	46.3	
1993–1995	24.9	28.2	46.9		24.1	29.2	47.8	

Notes: Parents, $N_d > 0$, are parents that use imports of intermediates from affiliates. Similarly, affiliates, $N_f > 0$, are affiliates that use imports of intermediates from parents.

Benchmark survey data from 1989 and 1994, we verified that approximately 93 percent of the goods shipped intra-firm from U.S. parents to Canadian manufacturing affiliates were intermediates destined for further processing.

Table 1 also reports the fraction of firms that utilize each of the four trade flows. In the analysis dataset, affiliate arms-length sales to the United States are positive for 42 percent of firm-year observations, while affiliate intra-firm sales to parents are positive 75 percent of the time. For parents, the fraction with arms-length sales to Canada is 86 percent, and the fraction with intra-firm sales to the affiliate is 82 percent. These percentages are quite stable over time, despite tariff reductions.

Table 2 provides descriptive statistics on the cost shares for labor, materials, and capital. It reports mean and median cost shares for the firms in the analysis dataset. Statistics are presented for three-year intervals, because there is some noise in the means and medians induced by the rather frequent entry and exit of small firms from the dataset.

Our model with Cobb-Douglas technology implies that cost shares should be stable over time, except for minor fluctuations induced by labor force adjustment costs, unless there is technical change.[43] Table 2 provides evidence that cost shares were indeed quite stable for those parents and affiliates that did not use intermediates. Given the large movements in relative factor prices apparent in Figure 3, this suggests that a Cobb-Douglas specification is reasonable. For instance, between 1984 and 1986 and 1990 and 1992, the price of labor in Canada rose roughly 20 percent relative to the prices of materials and capital. Yet, for affiliates that did not use intermediate inputs from parents ($N_f = 0$), the mean labor share rose only 0.6 percentage points (about 2.4 percent) over this period. Thus, unit elastic demand does not seem like too bad an assumption. Input cost shares also seem stable for parents that do not use intermediates shipped from the affiliate ($N_d = 0$).

---

[43] In Table 2, we assume zero profits and measure the capital share as the residual of revenues minus other factor costs. This is not exactly consistent with our model, in which firms have market power and positive profits. Thus, in Table 2, the capital share could fluctuate due to changes in the price elasticity of demand facing firms.

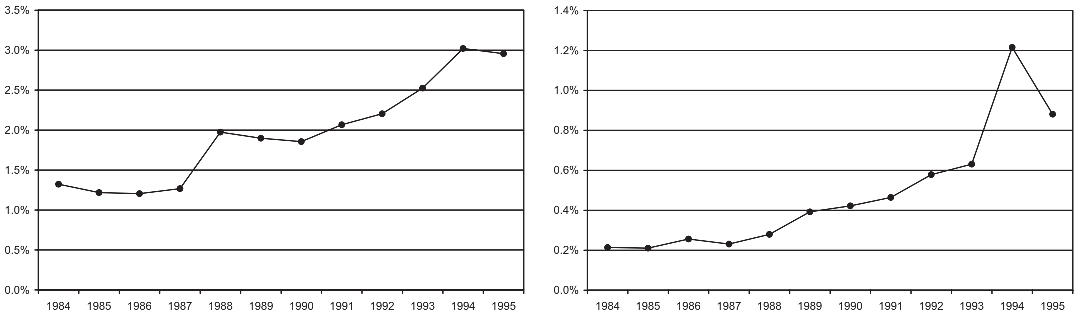

A. Parent intermediate input shares (mean)

B. Parent intermediate input share (median)

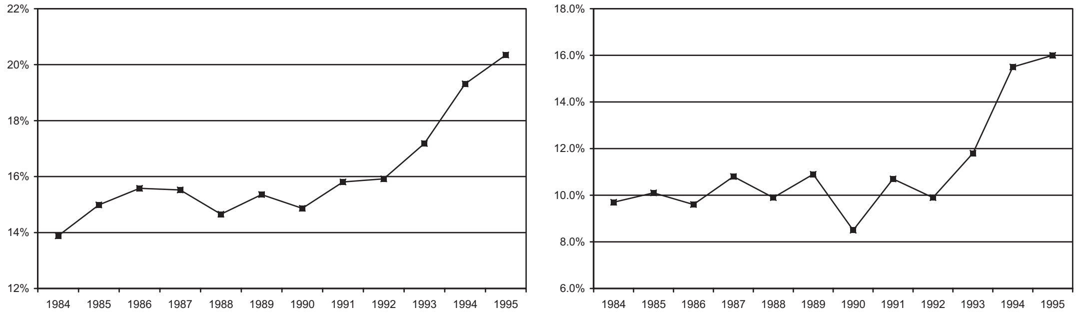

C. Affiliate intermediate input share (mean)

D. Affiliate intermediate input share (median)

FIGURE 4. AVERAGE AND MEDIAN INTERMEDIATE REAL INPUT SHARES FOR U.S. PARENTS AND CANADIAN AFFILIATES

Note: Intermediate shares for parents and affiliates are calculated conditional on the shares being greater than zero.

But cost shares were much less stable for affiliates that use intermediates from parents. For such affiliates, the mean (median) capital share dropped from 28.1 percent (27.5 percent) in 1984–1986 to 21.2 percent (21.2 percent) in 1993–1995. At the same time, the mean (median) cost share for intermediates shipped from the parent rose substantially, from 15.0 percent (9.8 percent) in 1984–1986 to 18.9 percent (13.6 percent) in 1993–1995. These figures understate the growth in intermediate shares, because there was rapid growth in both the first and last three years of the analysis period. Taking a three-year average dampens this out. Figure 4 makes the rapid growth in affiliates' intermediate input shares more obvious.

Figure 4 also shows that the share of intrafirm intermediates in parents' costs increased by roughly 130 percent over the analysis period. At the same time, parents that use intermediates shipped from affiliates ($N_d > 0$) had downward-trending labor shares. Prima facie, it appears quite implausible that the drop in the mean U.S. tariff rate from about 3.8 percent in 1984 to 1.2 percent in 1995 could explain this huge increase in intermediates shipped from affiliates. This tariff decline decreased the cost of intermediates to parents by only about 2.5 percent. Thus, enormous demand elasticities would be needed to rationalize the observed increase in affiliate-to-parent flows of intermediates on this basis. A similar story holds for parent-to-affiliate flows of intermediates.

B. *Parameter Estimates for the Structural Model*

Table 3 reports estimates of our structural model of MNCs' marginal production and trade decisions. The first panel reports parameters that determine the labor share in the parent's Cobb-Douglas production function. These map into the share parameter itself through the transformation described in equations (7) to (10). The first term is the intercept (α_0^{Ld}). The second and third terms are time trends, for parents that do and do not use intermediate inputs from the affiliate, respectively. The fourth term, α_{shift}^{Ld}, is

TABLE 3—TECHNOLOGY, ADJUSTMENT COST, MARKET POWER, AND DEMAND PARAMETER ESTIMATES

	Parameter name	Symbol	Estimate standard error
U.S. parent technology			
Parent labor	Intercept	α_0^{Ld}	−0.2288 (0.0925)**
	Time trend ($[N_d > 0]$)	$\alpha_{time}^{Ld} \cdot t \cdot I[N_d > 0]$	−0.1135 (0.0089)***
	Time trend ($[N_d = 0]$)	$\alpha_{time}^{Ld} \cdot t \cdot I[N_d = 0]$	−0.0149 (0.0199)
	Intercept shift	$\alpha_{shift}^{Ld} \cdot I[N_d > 0]$	5.4032 (0.1096)***
	Box Cox parameter	$bc(1)$	−0.0182 (0.0035)***
Parent materials	Intercept	α_0^{Md}	0.2311 (0.0834)***
	Time trend ($[N_d > 0]$)	$\alpha_{time}^{Md} \cdot t \cdot I[N_d > 0]$	−0.0922 (0.0083)***
	Time trend ($[N_d = 0]$)	$\alpha_{time}^{Md} \cdot t \cdot I[N_d = 0]$	−0.0044 (0.0168)
	Intercept shift	$\alpha_{shift}^{Md} \cdot I[N_d > 0]$	4.8860 (0.1142)***
	Box Cox parameter	$bc(2)$	−0.0398 (0.0038)***
Parent capital	Intercept	α_0^{Kd}	4.5094 (0.1125)***
	Time trend ($[N_d > 0]$)	$\alpha_{time}^{Kd} \cdot t \cdot I[N_d > 0]$	−0.0764 (0.0058)***
	Box Cox parameter	$bc(3)$	−0.0635 (0.0045)***
Canadian affiliate technology			
Affiliate labor	Intercept	α_0^{Lf}	−0.2373 (0.1172)**
	Time trend ($[N_d > 0]$)	$\alpha_{time}^{Lf} \cdot t \cdot I[N_f > 0]$	−0.0448 (0.0058)***
	Time trend ($[N_d = 0]$)	$\alpha_{time}^{Lf} \cdot t \cdot I[N_f = 0]$	0.0167 (0.0134)
	Intercept shift	$\alpha_{shift}^{Lf} I[N_f > 0]$	1.5053 (0.0964)***
	Box Cox parameter	$bc(4)$	−0.0771 (0.0091)***
Affiliate materials	Intercept	α_0^{Mf}	0.4295 (0.1920)**
	Time trend ($[N_d > 0]$)	$\alpha_{time}^{Mf} \cdot t \cdot I[N_f > 0]$	−0.0301 (0.0103)***
	Time trend ($[N_d = 0]$)	$\alpha_{time}^{Mf} \cdot t \cdot I[N_f = 0]$	0.0259 (0.0186)
	Intercept shift	$\alpha_{shift}^{Mf} \cdot I[N_f > 0]$	1.3688 (0.1694)***
	Box Cox parameter	$bc(5)$	−0.0277 (0.0052)***
Affiliate capital	Intercept	α_0^{Kf}	1.2639 (0.2013)***
	Time trend ($[N_d > 0]$)	$\alpha_{time}^{Kf} \cdot t \cdot I[N_f > 0]$	−0.0808 (0.0084)***
	Box Cox parameter	$bc(6)$	−0.0260 (0.0062)***
Parent and affiliate scaling parameters for N_d, $N_f = 0$	U.S. parent scaling parameter	$SC_d \cdot I[N_d = 0]$	0.4065 (0.0117)***
	Affiliate scaling parameter	$SC_f \cdot I[N_f = 0]$	0.6334 (0.0141)***
Common parameters	Profit rate	R_K	0.1652 (0.0033)***
	TFP growth rate	h	0.0453 (0.0080)***
	Discount factor	β	0.9500 —
U.S. parent labor adjustment costs	Intercept	δ_{0d}	−5.1078 (0.0995)***
	Domestic wage	$\delta_d w_{dt}$	1.0194 (0.0267)***
	Time trend	$\delta_{d,time} \cdot t$	0.0240 (0.0038)***
	Intercept shift	$\delta_d \cdot I[N_{dt} > 0]$	−0.0151 (0.0287)
CA affiliate labor adjustment costs	Intercept	δ_{0f}	−5.4817 (0.0506)***
	Foreign wage	$\delta_f w_{ft}$	0.9667 (0.0073)***
	Time trend	$\delta_{f,time} \cdot t$	0.0884 (0.0049)***
	Intercept shift	$\delta_f \cdot I[N_{ft} > 0]$	0.1563 (0.0500)***
Common adjustment cost parameters	numerator exponent	μ	0.7161 (0.0015)***
	denominator exponent	Δ	0.2216 (0.0047)***
Forecast Errors—Standard deviation and correlation	parent intercept	τ_{d0}	0.3186 (0.1409)**
	affiliate intercept	τ_{f0}	−0.2072 (0.1247)*
	labor force size	τ_1	1.1368 (0.0143)***
	Correlation	$\mathrm{CORR}(\tau_d, \tau_f)$	0.2974 (0.0808)***
Inverse price elasticity of demand for domestically produced good	Intercept	$g_{1,0}$	−1.3610 (0.0960)***
	Time trend	$g_{1,time} \cdot t$	0.0001 (0.0005)
	Intercept shift	$g_{1,shift} \cdot I[N_d > 0]$	0.0055 (0.0061)
	Box Cox parameter	$bc(7)$	0.6084 (0.0674)***
Inverse price elasticity of demand for foreign-produced good	Intercept	$g_{2,0}$	−1.0878 (0.0435)***
	Time trend	$g_{2,time} \cdot t$	−0.0021 (0.0004)***
	Intercept shift	$g_{2,shift} \cdot I[N_f > 0]$	−0.0002 (0.0037)
	Box Cox parameter	$bc(8)$	0.8424 (0.0433)***
Demand function parameters			
Domestic demand for domestically produced good	Intercept	$P_{0d,0}^1$	2.6898 (0.0908)***
	Time trend	$P_{0d,time}^1 \cdot t$	−0.0526 (0.0085)***
	Intercept shift	$P_{0d\ shift}^1 \cdot I[N_d > 0]$	0.0647 (0.0346)*
	Box Cox parameter	$bc(9)$	−0.0049 (0.0157)

TABLE 3—*Continued.*

	Parameter name	Symbol	Estimate standard error
Domestic demand for foreign-produced good	Intercept	$P^2_{0d,0}$	2.5330 (0.1875)***
	Time trend	$P^2_{0d,time} \cdot t$	−0.0740 (0.0123)***
	Intercept shift	$P^2_{0d\,shift} \cdot I[N_f > 0]$	0.3912 (0.1018)***
	Box-Cox parameter	$bc(12)$	0.1013 (0.0235)***
Foreign demand for domestically produced good	Intercept	$P^1_{0f,0}$	2.4520 (0.0816)***
	Time trend	$P^1_{0f,time} \cdot t$	−0.0539 (0.0083)***
	Intercept shift	$P^1_{0f\,shift} \cdot I[N_d > 0]$	0.0583 (0.0269)**
	Box-Cox parameter	$bc(11)$	−0.0135 (0.0136)
Foreign demand for foreign-produced good	Intercept	$P^2_{0f,0}$	2.8451 (0.2123)***
	Time trend	$P^2_{0f,time} \cdot t$	−0.0888 (0.0128)***
	Intercept shift	$P^2_{0f\,shift} \cdot I[N_f > 0]$	0.4714 (0.1129)***
	Box-Cox parameter	$bc(10)$	0.1568 (0.0195)***

*** Significantly different from 0 at the 1-percent level.
** Significantly different from 0 at the 5-percent level.
* Significantly different from 0 at the 10-percent level.

a shift parameter that allows the labor share to differ for the subset of firms with positive intra-firm flows (i.e., it multiplies $I[N_d > 0]$). Finally, the fifth term is the Box-Cox parameter from equation (10) that captures departures of the stochastic term in the labor share equation from log normality ($bc(1)$).

The second and third panels report parameters relevant to the parent's material and capital shares, respectively. Note that the capital share equation has fewer parameters. If a parent does not use intermediates from the affiliate, it has only three inputs, so the capital share is just $1 - \alpha^{Ld} - \alpha^{Md}$. Thus, the capital share equation is relevant only for parents that do use intra-firm intermediates, and so it does not include the shift parameter or the extra time trend that are included in the labor and material share equations. The fourth through sixth panels of Table 3 contain exactly the same types of parameters, but for the affiliate.

A fascinating aspect of the results is that the time trends on the share parameters are small and insignificant for parents and affiliates that do not use intermediates that are shipped intra-firm. That is, in Table 3, the terms $\alpha^{Ld}_{Time} \cdot t \cdot I[N_d = 0]$, $\alpha^{Md}_{Time} \cdot t \cdot I[N_d = 0]$, $\alpha^{Lf}_{Time} \cdot t \cdot I[N_f = 0]$, and $\alpha^{Mf}_{Time} \cdot t \cdot I[N_f = 0]$ are all insignificant and quantitatively small. Thus, the behavior of these parents and affiliates is well described by a CRTS Cobb-Douglas technology with fixed-share parameters, consistent with the descriptive statistics on cost shares presented in Table 2.

In contrast, for the subset of MNC parents that do utilize intermediates from affiliates, and affiliates that use intermediates from parents, the time trends for the share parameters are all highly significant. The direction of all the trends is negative, but this result should be interpreted with care. Since the trends feed into logistic transformations like (7) and (8), it is only the share of the input with the largest negative time trend that necessarily falls. The behavior of shares for inputs with smaller negative trends is ambiguous. Clearly, however, the fact that the time trends are negative for labor, capital, and materials means the share of the omitted category (intermediates) must be rising. Thus, conditional on MNCs having had positive intra-firm flows initially, the estimates imply that "technical change" was driving up the share of intermediates.

For parents, the strongest negative time trend is on the labor share, while for affiliates it's on the capital share. This is consistent with the descriptive statistics in Table 2, which showed that the labor share trended downward for parents while the capital share trended downward for affiliates.

For all the technology parameters, the Box-Cox parameters (i.e., $bc(1)$ through $bc(6)$) 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.[44]

---

[44] The next panel of Table 3 contains estimates of the parameters, SC_d and SC_f (see equation (10)), which scale the stochastic terms in the share equations between the three-outcome ($N_d = 0, N_f = 0$) and four-outcome ($N_d > 0, N_f > 0$) cases. As expected, both parameters are less than one, since the scale should be reduced in the three-outcome case.

Table 3 reports estimates of some key parameters of the model that are common across firms. The estimated profit rate, R, is 16.5 percent and the estimated rate of TFP growth, h, is 4.5 percent. We could not reject a specification in which h was common across the parent and affiliate. Both of these estimates seem reasonable. The high rate of TFP growth for U.S. MNCs is consistent with prior work in this area (see Jason G. Cummins, 1998), and the profit rate seems reasonable given prior estimates (see, e.g., Barbara M. Fraumeni and Dale W. Jorgenson, 1980).

Table 3 also contains estimates of the labor force adjustment cost parameters. The estimates for $\delta_d w$ and $\delta_j w$ 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, high-skill workers. The intercept shift for $N_f > 0$ is significant and positive in the δ_f equation, suggesting that labor force adjustment costs are greater for affiliates that ship intermediates to parents.[45] The time trends in the δ_d and δ_f equations are both significant and positive, suggesting that labor force adjustment costs increased over time.

The next section contains estimates of the generalized labor force adjustment cost function. The estimates of μ and Δ imply that adjustment costs are not well described by the common linear-quadratic in levels specification, since they depart substantially from $\mu = 1$ and $\Delta = 0$. The fact that $\mu < 1$ and $\Delta > 0$ implies that the cost of a given absolute change in labor force size is smaller, to the extent that the change represents a smaller fraction of the initial labor force.

Next, Table 3 contains estimates of the parameters that determine the variance of labor adjustment forecast errors (τ). Not surprisingly, the forecast error variances are an increasing function of labor force size. Interestingly, the domestic forecast error variance is higher (for a given labor force size), and domestic and foreign forecast errors are positively correlated.

Table 3 also contains estimates of the parameters that determine the (negative) inverse price elasticities of demand for the U.S. and Canadian-produced good (g_1 and g_2). These market power parameters are estimated with an intercept shift and a differential time trend for firms that use intra-firm flows. Note that the intercept shift is not significant for either the domestic or foreign market power parameter. Also, domestic market power does not vary significantly over time. But market power for Canadian-produced goods trends downward. The Box-Cox parameter for g_1 implies that a transform close to the square root is needed to induce normality of the residuals, while the transform for g_2 is closer to linear. The estimates imply price elasticities of demand in the -15 to -20 range.[46]

Finally, the bottom of Table 3 contains estimates of the four demand function intercept parameters (i.e., U.S. and Canadian demand for the U.S.- and Canadian-produced goods). All four demand function intercepts exhibit significant negative time trends, implying reduced demand at any given price level. The Box-Cox transform parameters for these four equations are all close to zero, implying that the demand shocks are well described by log normality.

Recall that each of the 12 technology and demand function parameters is heterogeneous across firms and over time. To conserve on space, Table 3 does not present the covariance matrix estimates, which are presented in Keane (2003). We highlight some key points: each parameter has a firm/time-specific component consisting of a random effect, μ_i, and a transitory error, ν_{it}. This structure implies equal correlations at all leads and lags, all the way out to $t + 11$. As we would expect, all the firm-specific parameters are highly serially correlated. The technology and demand-function intercept parameters show more persistence (i.e., correlations in the 0.72 to 0.80 range) than do the market power parameters (i.e., correlations of 0.46 and 0.53). The cross correlations within a time period imply that the demand shocks

---

[45] A dummy for $N_d > 0$ was not significant in the affiliate adjustment cost equation, and vice versa.

[46] These demand elasticities imply markups in the 5-7-percent range. As the mean capital share is about 30 percent, these figures are roughly consistent with the 16.5-percent profit rate estimate. It is important to note that these elasticities alone cannot be used to predict responses of a firm's imports or exports to tariff changes, because of other complicating features of the model. For example, as we discussed in Section IIC, if the MNC sells the same good both to third parties and intra-firm, then it has an incentive to produce and sell more of the good than it otherwise would, in order to depress the price and hold down tariff costs. A tariff reduction reduces this incentive, which may mitigate its sales-enhancing effect. Thus, we must rely on simulations to examine the model's predictions for tariff effects on trade.

TABLE 4—SIMULATED RESPONSES TO TARIFF CHANGES

	U.S. parent sales in the U.S. $P_d^1 S_d$	Canadian affiliate sales in Canada $P_f^2 S_f$	U.S. parent intermediates' sales to Canadian affiliates $P_d^1 N_f$	U.S. parent arms-length sales to Canada $P_f^1 E$	Canadian affiliate intermediates' sales to U.S. parents $P_f^2 N_d$	Canadian affiliate arms-length sales to the U.S. $P_d^2 I$	U.S. parent labor L_d	Canadian affiliate labor L_f
1995 baseline simulation steady-state level	3,619,683	213,074	241,630	37,700	325,587	35,256	24,377	1,579
Experiment: fixed tariffs at 1984 level	3,588,695	198,422	229,852	30,116	317,977	24,737	23,996	1,433
Percent difference from base	**−0.9%**	**−6.9%**	**−4.9%**	**−20.1%**	**−2.3%**	**−29.8%**	**−1.6%**	**−9.2%**
Experiment: eliminate tariffs	3,640,632	226,825	253,316	48,014	330,548	41,955	24,602	1,667
Percent difference from base	**0.6%**	**6.5%**	**4.8%**	**27.4%**	**1.5%**	**19.0%**	**0.9%**	**5.6%**
Percent difference from tariffs = 1984	**1.4%**	**14.3%**	**10.2%**	**59.4%**	**4.0%**	**69.6%**	**2.5%**	**16.3%**

Notes: Trade flows are expressed in thousands of 1984 dollars. The table reports means for firms with a positive value for the indicated flow.

for the same good across the two countries are extremely highly correlated (roughly 0.97), while demand shocks for the two different goods within the same country are not so highly correlated (roughly 0.40).[47]

Keane (2003) develops econometric methods to evaluate distributional assumptions in a model like ours. For our model, he shows that normality is rejected at the 1-percent level for only one of the 12 residuals—the affiliate materials share—and at the 5-percent level for only one additional residual—the parent capital share. Recall that a novel feature of our approach is the assumption of normal forecast errors. Keane (2003) shows that simulated posterior distributions (conditional on the data and the model parameters) for the domestic and foreign labor force adjustment cost forecast errors, η^d and η^f, are essentially indistinguishable from normality. Thus, our assumptions about the distributions of the stochastic terms appear (for the most part) to be supported.

VI. Simulations of Tariff Effects on Trade

In Table 4, we use simulations of our model to examine the effect of tariff reductions on MNC-based trade. Specifically, using a steady-state simulation, we compare the baseline levels of intra-firm and arms-length trade in the last year of our data (1995), given actual tariffs in that year, to levels that would have obtained under two different counterfactual scenarios: (a) no tariff reductions (i.e., keeping tariffs at their 1984 levels); and (b) complete tariff elimination.

Since we have not implemented a full-solution algorithm,[48] we cannot simulate the short-run outcomes generated by our dynamic model. We can, however, use a steady-state version of our model (which assumes no labor force adjustment costs) to simulate the long-run response of the population of firms to changes in the policy environment. The simulations are done using 446 vectors of technology and demand parameters drawn from the posterior distribution of the firm-specific parameters of our model. Keane (2003) describes this procedure.

---

[47] For completeness, we note that the composite standard deviations (i.e., combining both the μ_i and ν_{it} components) of the 12 stochastic terms are, for the domestic shares, 1.98, 1.76, 1.50; for the affiliate shares, 1.44, 1.60, 1.57; for the g's, 0.041, 0.025; and for the demand functions, 0.406, 0.804, 0.393, 0.679.

[48] To do so would require us to specify the driving processes for demand shocks, tariffs, wages, prices of materials and capital, and technical change. This is well beyond the scope of the paper.

The first row of Table 4 shows mean levels of domestic sales, intra-firm and arms-length trade flows, and U.S. and Canadian labor that MNCs are predicted to choose (in the long-run steady state) under the actual policy regime that prevailed in 1995. The second row shows levels that would have obtained under the counterfactual that U.S. and Canadian tariffs are held fixed at their 1984 levels, and the third row simulates the complete elimination of tariffs. In each case, the technology and demand parameters are held fixed at their 1995 levels. The simulations assume that firms adjust to tariff changes only on the intensive margin, since results in Feinberg and Keane (2005) suggest that tariffs have no significant effect on the extensive margin.

Our simulations imply that arms-length trade is quite sensitive to tariffs. If tariffs were returned to their 1984 levels, U.S. arms-length exports to Canada would drop 20 percent, while affiliate arms-length sales to the United States would drop 30 percent. Since average tariff reductions over this period were about 4 and 3 percentage points for Canada and the United States, respectively, these figures imply demand elasticities for imports and exports in the ballpark of -10 to -5. Head and Reis (2003) obtain an elasticity of substitution between U.S. and Canadian good varieties of about 8, which seems broadly consistent with our results.

A key result in Table 4 is the difference in sensitivity of the intra-firm and arms-length trade flows to changes in tariff levels. For instance, the intra-firm trade flow from U.S. parents to Canadian affiliates is predicted to drop only about 4.9 percent if tariffs were raised back to their 1984 level, and the intra-firm flow from affiliates to parents is predicted to drop only about 2.3 percent.

Although we highlight the finding that arms-length trade is much more sensitive to tariffs than intra-firm trade, we are not suggesting that intra-firm trade is insensitive. For example, our model predicts that if tariffs were eliminated completely (starting from the 1995 baseline), then U.S. parents would ship 4.8 percent more products to Canadian affiliates, and affiliates would ship 1.5 percent more intermediates back to parents. These are quantitatively important effects, and they imply that tariffs continue to be a factor restraining intra-firm trade in some industries.

Table 4 also reports effects of the tariff experiments on employment and sales. When tariffs are raised to 1984 levels, Canadian affiliate domestic sales and employment fall by 5.6 percent and 8.5 percent, respectively.[49] Not surprisingly, U.S. parent domestic sales and labor change less under the high-tariff regime, with the former dropping by 0.7 percent and the latter dropping by 1.2 percent.

As we saw in Figure 3, tariffs were not the only factor that changed substantially during our analysis period. There were also substantial changes in relative wages, and in prices of capital and materials. And our estimates imply substantial technical change. In Table 5 we decompose changes in key variables of interest into parts due to tariffs versus technology versus wages versus all other factors. The first column of Table 5 reports the actual changes in several variables of interest that occurred from 1984 to 1995. The second column reports the predicted change from 1984 to 1995 in the steady-state level of each variable, given all changes in the environment. The third, fourth, and fifth columns report the predicted changes due to changes in tariffs, technology, and wages, respectively. The last column reports the combined effect of all other factors.

Of course, the predicted change in steady-state levels is not directly comparable to the change in actuals, since the latter include transition dynamics. From a face validity standpoint, however, it is comforting that the predicted changes line up reasonably well with the actual changes. Our model predicts that all factors combined led to an increase of 99 percent in U.S. parent intra-firm sales to affiliates, and a 79-percent increase in U.S. parent arms-length sales to Canada. The actual changes were 73 percent and 75 percent, respectively. And the model predicts that all factors combined led to a 123-percent increase in affiliate intra-firm sales to parents, and a 9-percent increase in affiliates arms-length sales to the United States. The actual changes were 95 percent and 4 percent, respectively.

Interestingly, the model predicts that tariff reductions alone would have increased affiliates' arms-length sales to the United States by 32.5 percent. The model implies that technical change

---

[49] Our finding that tariff reductions increase affiliate employment may appear to contradict Noel Gaston and Trefler (1997), who find (small) negative employment effects. But they examine 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.

TABLE 5—PERCENTAGE CHANGES, 1984–1995: DATA AND MODEL PREDICTIONS

			Model decomposition			
	Data	Model	Tariffs	Technology	Wages	Other
U.S. parent sales						
Domestic ($P_d^1 S_d$)	−19.6	−9.0	0.8	4.0	−4.3	−9.5
Intrafirm ($P_d^1 N_f$)	73.3	99.1	9.7	71.8	14.4	3.2
Arms-length ($P_f^1 E$)	75.3	79.1	36.0	9.0	−1.3	35.4
Total	−17.1	−6.0	1.1	7.3	−3.9	−10.5
Total net	−20.2	−9.8	1.0	3.8	−3.9	−10.7
Canadian affiliate sales						
Domestic ($P_f^2 S_f$)	−17.9	−6.9	6.4	−71.9	18.3	40.3
Intrafirm ($P_f^2 N_d$)	94.5	123.0	5.2	84.1	−2.9	36.6
Arms-length ($P_d^2 I$)	4.0	8.9	32.5	37.9	−20.5	−41.0
Total	18.7	38.3	7.3	7.8	11.2	12.0
Total net	−0.5	14.2	6.3	−23.3	9.9	21.3
MNC sales	−19.0	−8.4	1.3	2.4	−3.0	−9.1
U.S. parent employment	−24.4	−12.6	1.4	12.4	−1.6	−24.8
CA Employment	−14.1	−3.0	9.0	15.5	−7.0	−20.5

Notes: "Data" shows the percentage change in our analysis dataset. "Model" shows the percentage change from 1984 to 1995 in the model simulation. Under "model decomposition," we compare the percentage change predicted by the model under the baseline simulation, with the percentage change the model predicts under the counterfactual that the indicated forcing variable (tariffs, technology, or wages) stayed fixed at the 1984 level. The difference is the percentage change attributable to changes in that particular forcing variable. For example, the average level of ND increased from $145,979,000 in 1984 to $325,586,000 in 1995 in the baseline simulation (see Table 4), a 123-percent increase. If tariffs stayed fixed at the 1984 level, we predict that ND would instead be $317,977,000 in 1995, a 117.8-percent increase. The difference of 5.2 percent is the percentage increase attributable to tariff reductions.

was also driving up affiliate arms-length sales. Yet, overall, affiliate arms-length sales to the United States increase only 9 percent in the model simulation and 4 percent in the data. According to the model, rising Canadian real wages were a key factor holding down exports to the United States. As we saw in Figure 3, the real wage (in U.S. dollar terms) paid by Canadian affiliates increased by 20 percent from 1984 to 1995. Our model predicts that this reduced affiliates' arms-length exports to the United States by 20 percent.

Table 5 reveals the important role our model assigns to technical change in increasing intra-firm trade. Of the 99-percent increase in parents' intra-firm sales to affiliates, the model attributes 72 percentage points to technical change. And of the 123-percent increase in affiliates' intra-firm sales to parents, the model attributes 84 percentage points to technical change. The changes attributed to tariff reductions are only about 10 and 5 percentage points, respectively. While these tariff effects are not trivial, they are an order of magnitude smaller than the impact of technology.

Finally, a fascinating aspect of both the data and the simulations is the radically changing nature of the parent-affiliate relationship. Affiliate sales in Canada have been falling at the same time that shipments to parents have doubled. Feinberg and Keane (2001) pointed out that the growth of affiliate shipments of intermediates to parents was clear evidence that Canadian manufacturing is not being "hollowed out" by free trade, as many FTA opponents had feared.[50] But Canadian manufacturing affiliates are clearly being transformed into production units that are more fully integrated into MNCs' overall production process. In 1984, sales of intermediates to parents were about 29 percent of affiliate total sales. By 1995, this figure had risen to 51 percent! Our model attributes most of

---

[50] Many argued that tariff-jumping FDI led to inefficiently small Canadian plants (see Harry C. Eastman and Stefan Stykolt, 1967; Richard E. Caves, 1982, 1990; John R. Baldwin and Paul Gorecki, 1986). This led to concern that the FTA would "hollow out" Canadian industry, since U.S. MNCs could most efficiently serve both markets from large U.S. plants. Feinberg et al. (1998) and Feinberg and Keane (2001) present firm-level evidence that contradicts this view. Gaston and Trefler (1997) and Trefler (2001) also examine the impact of the FTA on Canadian manufacturing, using industry-level data. Their identification strategy exploits the fact that tariff decreases differed substantially across industries. Trefler (2001) concludes that the FTA reduced employment while raising earnings of production workers, reduced output, and raised labor productivity. He concludes the FTA did not affect the number or scale of Canadian manufacturing plants. Head and Ries (1997) also estimate that the FTA had little effect on scale of Canadian plants.

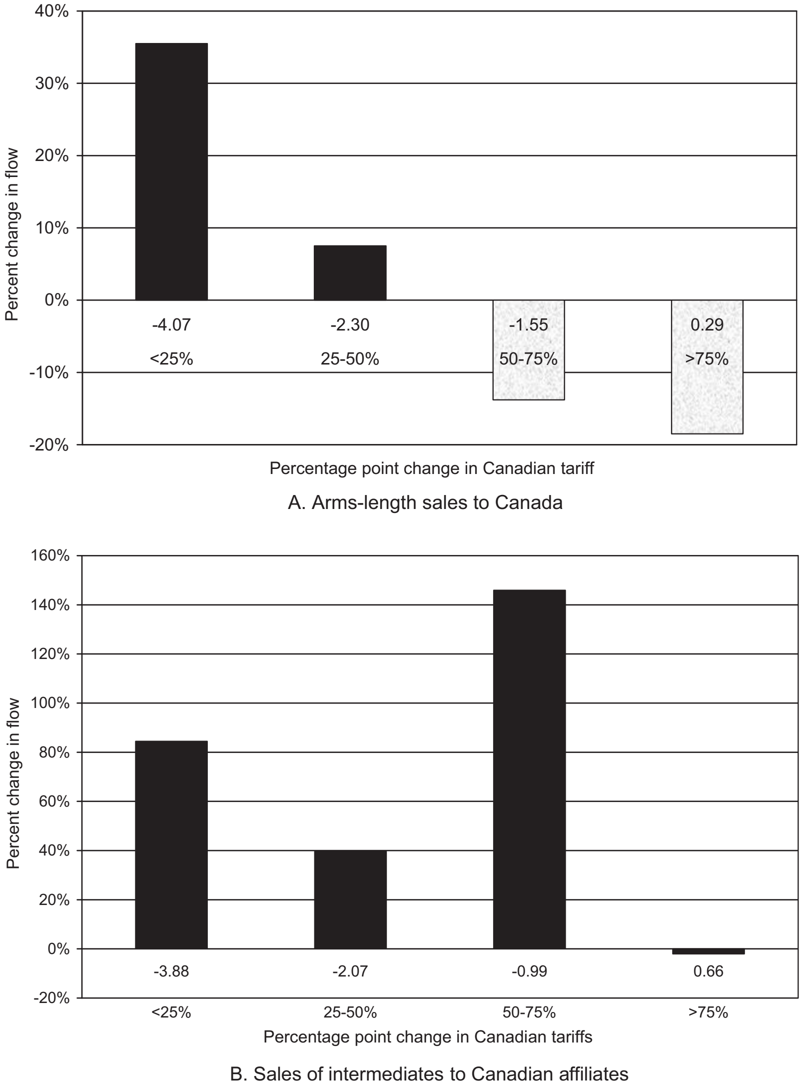

A. Arms-length sales to Canada

B. Sales of intermediates to Canadian affiliates

FIGURE 5. CHANGE IN U.S. PARENT FLOWS TO CANADA BY CHANGES IN CANADIAN TARIFFS, 1989–1995

Notes: Tariff changes are calculated conditional on each flow being greater than zero in both periods. Tariff changes are in percentage points, and the numbers shown are the medians of each quartile.

this dramatic change not to tariff reductions, but to "technology."

VII. Analyzing the Sources of Technical Change: Why Do the Share Parameters Increase?

Next, we explore the sources of technical change driving up the intra-firm intermediate shares. But first we ask if our key finding—that tariffs can explain a substantial increase in arms-length MNC-based trade between the United States and Canada, but little of the growth of intra-firm trade—is an artifact of some special feature of our model, or if the data clearly "speak" to this point. Figures 5 and 6 address this issue. The top panel of Figure 5 plots changes in U.S. parent arms-length exports between 1989 (the year of the FTA) and

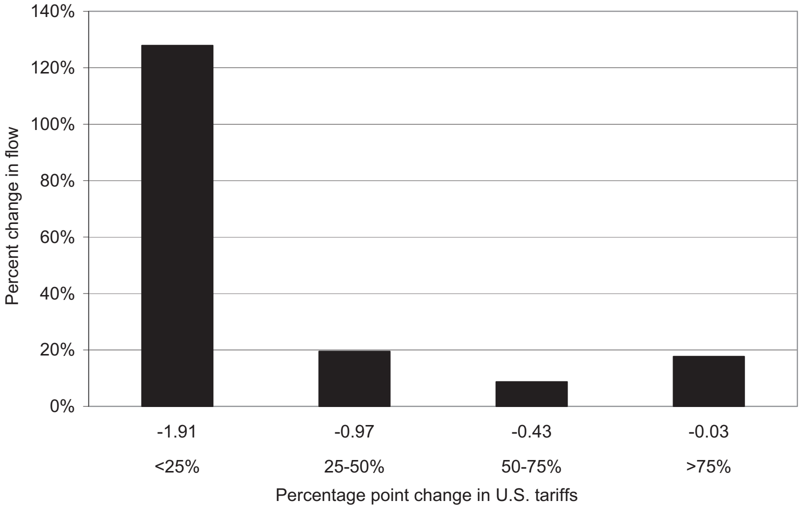

A. Arms-length sales to the United States

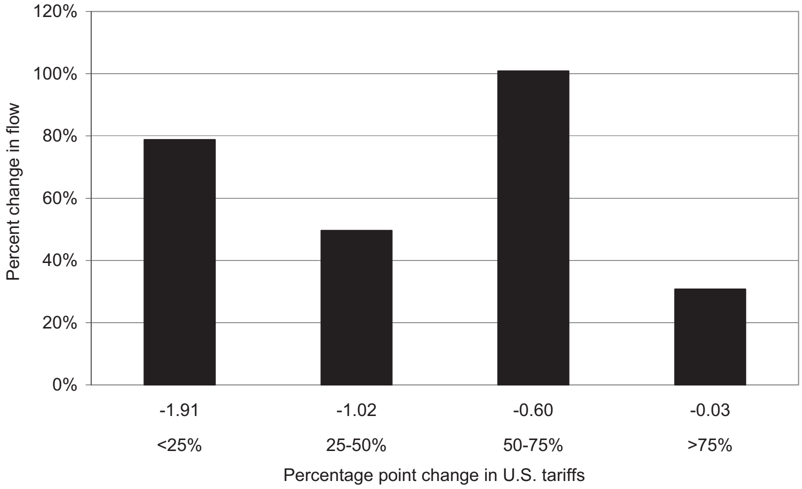

B. Sales of intermediates to U.S. parents

Figure 6. Change in Canadian Affiliate Flows to the United States by Changes in U.S. Tariffs, 1989–1995

Notes: Tariff changes are calculated conditional on each flow being greater than zero in both periods. Tariff changes are in percentage points, and the numbers shown are the medians of each quartile.

1995, for four groups of firms. Firms are grouped into quartiles based on the magnitude of Canadian tariff reductions in their industry over the period.[51] Clearly, exports to Canada increased more in industries where Canadian tariffs fell most. For industries with negligible Canadian tariff reductions, arms-length exports actually fell. In industries with the greatest tariff reductions (4

---

[51] Note, in Figures 5 and 6, that Canadian tariffs actually increase slightly for firms in the fourth quartile. This is because we construct weighted average tariffs for U.S. parent flows to Canada. A change in the industry mix of sales may lead to an increase in our tariff measure, even if industry-level tariffs are decreasing. This is an unavoidable consequence of the need to aggregate sales and tariffs across industries.

percentage points at the median), growth in arms-length exports was roughly 35 percent.

The bottom panel of Figure 5 reports a similar graph, except for U.S. parent intra-firm sales to Canadian affiliates. Here, there is no clear relationship between the magnitude of the tariff reduction and the increase in intra-firm flows. In fact, the largest increase in intra-firm trade (roughly 140 percent!) occurred in the third-quartile industries, where Canadian tariff reductions were quite small. Many of these industries had very low tariffs to begin with.

Figure 6 tells a similar story for Canadian affiliate arms-length and intra-firm shipments to the United States. The former are closely related to tariff changes, while the latter are not. Even in the third- and fourth-quartile industries where U.S. tariffs changed very little (and were small to begin with), intra-firm shipments to parents grew roughly 100 percent and 30 percent, respectively.

Given these figures, it seems obvious that trade liberalization cannot explain much of the growth in intra-firm trade. For instance, to explain the average 100-percent increase in intra-firm trade in intermediates we see among the third-quartile firms in the right panel of Figure 6, for whom the median tariff decline was only 0.6 points, we would need demand elasticities of at least -100. It would be difficult for any plausible model to generate such large price elasticities of demand for intra-firm intermediates,[52] so it does not appear that our Cobb-Douglas specification for technology could be critically driving our results.[53] Given the lack of correlation between tariffs and intra-firm trade, it is not surprising that our model attributes most of the growth of intra-firm trade to "technical change" (i.e., trends in the Cobb-Douglas share parameters).[54]

Nevertheless, we proceed to test the Cobb-Douglas specification by regressing the intermediate share parameters (here denoted by ND and NF) on a generic time trend and the factor input prices. If factor shares are invariant to factor prices, but technical change shifts the share parameters, then factor prices should be insignificant, and the time trend should capture the technical change. We report the results of this regression in Table 6. Not surprisingly, the Cobb-Douglas is literally rejected, as some of the factor price coefficients are statistically significant. Of course, we don't think technology is literally Cobb-Douglas, so this is to be expected. What is surprising is how minor the rejection is. The bottom panel of Table 6 interprets the

---

[52] Hanson et al. (2005), using a very different methodology from ours, exploit cross-industry and cross-country variation in tariffs to estimate a price elasticity of demand for intermediates by U.S. MNC affiliates of about -3. Given this figure, tariffs can explain only a small fraction of the increase in intra-firm trade. In the model of Yi (2003), tariff effects on the costs of intermediates are magnified if goods-in-process must cross the border several times. His model generates demand elasticities for intermediates on the order of -10 (see his Table 3), which is still far too small to explain a near doubling of intra-firm trade given tariff declines of a few points or less. Indeed, our results are actually similar to Yi's, since his calibrated model still leaves most of the increase in trade in the 1980s and 1990s unexplained, and in his conclusion he speculates that technical change may explain the residual growth of trade.

[53] Another assumption in our model is that final goods produced by parents and affiliates are unrelated (i.e., zero cross-price elasticity of demand). Allowing substitutability would lead, *ceteris paribus,* to a lower level of bilateral MNC-based trade in final goods for any given level of tariffs, because the MNC would have an incentive to avoid competing with itself (see Brander, 1981; Baldwin and Ottaviano, 2001). Thus, the amount of bilateral MNC-based trade at any given tariff level will be sensitive to the assumed degree of substitutability of the final goods. Two points are worth noting, however. First, our identification discussions in Sections IIB and IIID show that our model can generate any given level of bilateral trade (while maintaining a zero cross-price elasticity) by choice of the other free parameters. Second, there is no obvious way that allowing substitutability between final goods could possibly make trade in intermediates extremely sensitive to tariffs (and in such a way as to maintain near zero correlation between tariffs and intermediate trade at the industry level). Thus, as we noted in Section IIA, our key result that tariffs cannot explain the growth of intra-firm trade appears quite robust to this assumption as well.

[54] In recent models of intra-firm and intra-industry trade by Yi (2003) and Jonathan Eaton and Samuel S. Kortum (2002), there are many varieties of intermediates, and tariff reductions increase the number of varieties that are traded. If varieties are close substitutes in production, small tariff reductions can have big effects on intra-firm trade (i.e., their models give demand elasticities on the order of -10, which, as we noted above, is still too small to explain most of the growth of intra-firm trade). If this variety mechanism is at work, we cannot see it, because we observe only the total value of intermediates traded intra-firm, not the number of varieties. But, in this case, the total value of intra-firm trade should appear very sensitive to tariffs at the industry level. This is inconsistent with our finding of essentially no relationship between tariffs and intra-firm trade. It is clearly difficult to reconcile a story designed to make intra-firm trade very sensitive to tariffs with data where there is essentially no correlation at the industry level.

TABLE 6—RELATION BETWEEN TARIFFS, FACTOR PRICES, AND INTERMEDIATE INPUT SHARE PARAMETERS

Variable	N^d share for U.S. parents				N^f share for foreign affiliates			
Trend	0.096***	0.094***	0.107***	0.116***	0.514***	0.514***	0.486***	0.496***
	(0.016)	(0.016)	(0.019)	(0.019)	(0.063)	(0.063)	(0.077)	(0.079)
US tariff +			−0.061	−0.022			0.135	0.422*
transport cost			(0.044)	(0.048)			(0.187)	(0.203)
Canadian tariff +			0.051	0.085**			−0.312***	−0.354***
transport cost			(0.034)	(0.036)			(0.0013)	(0.0014)
U.S./Canadian			−0.012	−0.021			1.063	0.962
wage			(0.174)	(0.175)			(0.716)	(0.719)
U.S./Canadian			−0.0407**	−0.0416**			−0.1537*	−0.155*
materials cost			(0.0163)	(0.0163)			(0.0669)	(0.0668)
Industry dummies		YES		YES		YES		YES
GLS-R^2	0.0194	0.1987	0.0135	0.1942	0.0049	0.2262	0.0160	0.2345
N	1756	1756	1756	1756	1924	1924	1924	1924

Means and standard deviations of the variables, and effects of one-standard-deviation changes

	N^d Equation			N^f Equation		
Variable	Mean	Standard deviation	Effect	Mean	Standard deviation	Effect
N^d or N^f/trend	0.0189	0.0394	0.0128	0.1698	0.1668	0.0546
U.S. tariff + transport cost	3.267	2.295	−0.0005	3.325	2.216	0.0094
Canadian tariff + transport cost	4.557	3.275	0.0028	4.736	3.456	−0.0122
U.S./Canadian wage ratio	1.120	0.417	−0.0001	1.116	0.415	0.004
U.S./Canadian materials price	0.955	0.029	−0.0012	0.954	0.029	−0.0045

Note: Tests are two-tailed.
*** Significantly different from 0 at the 1-percent level.
** Significantly different from 0 at the 5-percent level.
* Significantly different from 0 at the 10-percent level.

magnitudes of the coefficients by asking how one-standard-deviation changes in each variable are predicted to shift the share parameter. In all cases, these shifts are quantitatively quite small.

For instance, in the ND share equation, the U.S. tariff plus transport cost is insignificant. The Canadian tariff coefficient is 0.085 and significant in the most general specification (which includes industry dummies and firm-specific random effects). This coefficient, however, is the "wrong" sign—from the perspective of trying to use misspecification to explain away our result—as it implies that Canadian tariff reductions actually *reduce* the ND share. More importantly, the magnitude is small. As we see in the bottom panel of Table 6, a one-standard-deviation drop in the tariff is 3.275 points. The coefficient estimate of 0.085 implies that this would lower the ND share by only −0.278 points.[55] In contrast, the time trend accounts for a 1.276-point increase in the ND share over the 1984–1995 period. Recall that the average ND share increased from 1.4 percent in 1984 to 3.0 percent in 1995, an increase of 1.6 points. Thus, the estimates in Table 6 imply that the time trend "explains" most of that increase, while factor price changes operating through misspecification of the Cobb-Douglas can only account for small changes in the ND share.

Another way to see this is to note that, in Table 6, when we add factor prices to the ND regression, the coefficient on the generic time trend actually increases (slightly). So changes in

---

[55] In a typical industry, the Canadian tariff dropped by about 4.1 points over our analysis period. Thus, the shift in the ND share that arises due to changing Canadian tariffs and misspecification of the Cobb-Douglas is only −0.35 points.

factor prices combined with misspecification of the Cobb-Douglas actually cause us to slightly *understate* the importance of the technical change captured by the time trend. The story is basically the same in the NF equation. There, including factor prices causes the generic trend coefficient to fall very slightly. Again, the tariff variables are significant, but the quantitative magnitude of their coefficients is very small.

Thus, we feel we have convincingly shown that tariff reductions simply cannot plausibly explain very much of the increase in intra-firm trade between the United States and Canada. It seems nearly impossible to construct a model where tariffs have a big effect on intra-firm trade, yet there is almost no correlation between tariffs and intra-firm trade. Our model's attribution of the growth in intra-firm trade to "technical change" is not the result of model misspecification.[56]

Next, we ask, if factor prices can't explain the shifts in the intermediate input share parameters, then what can? If we could find measurable variables that are quantitatively important in the ND and NF share equations, and that significantly reduce the generic time trend coefficients, it might give us important clues about the underlying causes of increased intra-firm trade. Thus, we decided to conduct an exploratory regression analysis, in which we regress the share parameters on a wide range of industry and firm characteristics.

In Table 7, the dependent variable is the ND share. The independent variables include tariff and transport costs, factor prices, and the following industry/firm-specific variables: the R&D intensity of the industry, the Japanese import penetration rate, the average ratio of information technology capital to sales (IT/S) in the industry (used in Kevin Stiroh, 2002), the average inventory-to-sales (I/S) ratio in the industry, the scale of the MNC (as measured by total third-party sales in logs), the growth of the MNC (as measured by the ratio of sales at time t to average sales over the analysis period), and several dummy variables to capture the structure of the MNC, such as dummies for whether the affiliate sells to third parties in the United States, whether the parent exports to third parties in Canada, whether the parent ships intermediates to the affiliate, and the size of the MNC's worldwide affiliate network. In our companion paper (Keane and Feinberg, 2005), we describe in detail the reasoning behind including this set of variables.

Many of the variables in the regression are entered in two ways. First, the initial, 1983, level of the variable is interacted with trend. This allows us to determine whether, for instance, industries that were more R&D intensive in 1983 had greater trend growth in the ND share. Second, variables are also entered in current levels. This allows us to determine if, for instance, MNCs that grew more or made greater IT investments also had greater growth in the ND share.

A striking result emerges from the regression. With a couple minor exceptions that we will note later, *the only variable that is statistically significant and quantitatively important in predicting growth of shipments of intermediates by affiliates to parents is the industry I/S ratio.*[57]

---

[56] Could transfer price manipulation explain the growth of intra-firm trade? U.S. corporate tax rates are higher than Canadian rates in manufacturing (46 percent versus 36 percent in 1995), creating an incentive to shift profits to Canada. To do this, the MNC should overprice intermediates shipped from affiliate to parent, and vice versa. But, as G. Frank Mathewson and G. David Quirin (1979) note, tariffs create incentives to undervalue all intra-firm shipments, so tariff reductions might lead to higher transfer prices. But this transfer price manipulation story for increased intra-firm trade is implausible because, as we note in Section IIB, MNCs' ability to manipulate transfer prices is heavily constrained by U.S. and Canadian tax regulations. In fact, there is little evidence that MNCs engage in transfer price manipulation in the U.S.-Canada context. More importantly, if tariff-induced transfer price manipulation led to the appearance of greatly increased intra-firm trade, we would see a strong correlation between tariffs and intra-firm trade, which we do not.

Nevertheless, we examined data on real trade quantities to validate our finding of substantial increases in intra-firm trade. The U.S. International Trade Commission (ITC) provides census data on real trade quantities at its Web site (dataweb.usitc.gov). The data are not broken down into arms-length versus intra-firm data. Our data indicate, however, that the vast majority of U.S.-Canada trade in the computer and office equipment industry (SIC 357) is intra-firm trade. This is also the industry with the second highest volume of intra-firm trade (after autos and auto parts). In SIC 3577 (computer peripheral equipment and parts), the quantity of U.S. exports to Canada increased 446 percent from 1989 to 1995, while the quantity of imports from Canada increased 169 percent. Since almost all trade in this industry is intra-firm, this is clear evidence that the quantity of intra-firm trade did increase substantially. The same is true in other subcategories of SIC 357.

[57] This result is extremely robust to many changes in specification. It is little affected if we use the I/S ratio for the

TABLE 7—EXPLAINING THE INCREASE IN N^d SHARE

Variable	Coefficient	Standard error	*T*-statistic
Trend	0.0484	0.0256	1.89
Trend interacted with			
I/S(83)	−0.0213	0.0029	**−7.34**
Log sales(83)	0.0495	0.0108	**4.59**
R&D/S(83)	−0.0001	0.0056	−0.02
IT/S(83)	0.0311	0.0139	**2.24**
Japan import share (83)	−0.0017	0.0038	−0.44
No imports to U.S.	0.0985	0.0305	**3.23**
No exports to Canada	0.1127	0.0494	**2.28**
$NF = 0$	0.0545	0.0590	0.92
Parent industry different	−0.0628	0.0363	−1.73
Industry characteristics			
I/S(*t*)	−0.1786	0.0444	**−4.02**
I/S(*t*) · *I/S*(83)	0.0096	0.0027	**3.53**
R&D/S(*t*)	0.0309	0.0532	0.58
IT/S(*t*)	−0.3376	0.1645	**−2.05**
Japan import share (*t*)	0.0642	0.0350	1.84
Control variables for MNC structure			
NF share (*t*)	0.0015	0.0057	0.27
No. of worldwide affiliates (*t*)	0.0084	0.0042	**1.98**
MNC mean log sales	−0.6395	0.1704	**−3.75**
MNC sales growth (*t*/83)	0.0015	0.0028	0.55
No imports to U.S.	−0.1652	0.2131	−0.78
No exports to Canada	0.6401	0.3657	1.75
$NF = 0$	−0.1905	0.4384	−0.43
Parent industry different	0.2726	0.2956	0.92
Tariffs, transport costs, and factor prices			
U.S. tariff + transport cost	−0.0362	0.0462	−0.78
Canada tariff + transport cost	0.0577	0.0361	1.60
U.S./Canada wage ratio	−0.0174	0.1663	−0.10
U.S./Canada material price ratio	−0.0403	0.0163	**−2.47**

Notes: The regression also includes industry dummies and is estimated with random effects. Variables are de-meaned before being interacted with trend or I/S(*t*) (so the main effects are unaffected by inclusion of the interactions).

Effects of One-Standard-Deviation Changes in Each Variable on the ND Share:

	Coefficient on:		Δ ND from 83 to 96 due to:			Total Δ ND from 83 to 96 if:	
$\dot{I/S}(83)$	Trend	*I/S*(*t*)	Trend	$\Delta I/S = -3.66$	$\Delta I/S = -7.32$	$\Delta I/S = -3.66$	$\Delta I/S = -7.32$
0.00	0.0484	−0.1786	0.629	0.654	1.307	**1.283**	**1.936**
−4.00	0.1336	−0.2170	1.737	0.794	1.588	**2.531**	**3.325**
4.00	−0.0368	−0.1402	−0.478	0.513	1.026	**0.035**	**0.491**

Note: Define $\dot{I/S}(83) = I/S(83) - \overline{I/S}(83)$.

Also, the generic trend coefficient drops from 0.116 to 0.048, suggesting that the factors we have included "explain" roughly 60 percent of the trend growth in intra-firm trade at the aggregate level.[58]

---

firm instead if the industry, and across a wide range of alternative specifications that involve adding or deleting other variables from the equation, or changing the functional form with which I/S is entered. Our results are also little changed if we simply use the ND share of total sales as the dependent variable, rather than the estimated ND share parameters from the production function. These two quantities are very highly correlated.

[58] For all variables that are interacted with trend, we always de-mean the variable before constructing the interaction. Thus, the generic trend coefficient retains its interpretation as the generic trend for a typical firm. This trend captures our "ignorance," since it measures the trend in the ND share that is not explained by any variables in the model.

We originally decided to include the I/S ratio in the regression as a proxy for advances in logistics management and computer-based logistics systems over the past two decades. We suspected that these advances might have enabled MNCs to organize intra-firm flows of intermediates more efficiently, thus reducing the level of work-in-progress inventories—and hence the inventory carrying cost—needed to sustain any given level of intra-firm trade. A reduced I/S ratio implies success in adopting these methods.[59] The result that the I/S ratio is so closely associated with growth in intra-firm trade suggests that this factor may indeed be important.

The quantitative magnitude of the I/S coefficients is difficult to interpret directly, because we enter I/S in the equation in a rather flexible way. It is entered in level form, I/S(*t*), its 1983 level is interacted with trend, I/S(83) · trend, and its 1983 level is interacted with its current level. We did this because it allows not just the absolute change in I/S to matter, but also its change relative to its 1983 level. The bottom panel of Table 7 clarifies the meaning of the estimates. There, we calculate the implied change in the ND share under six different scenarios: the I/S ratio could be average, 4 points above average, or 4 points below average in 1983;[60] and the change in the I/S ratio could be −3.66 points (the average) or −7.32 points (twice the average).

The calculations show that, for a firm that is completely "average" (i.e., average I/S ratio in 1983, and average decline in the I/S ratio from 1983 to 1996), the predicted increase in the ND share is 1.283 points. However, a firm that had a relatively good I/S in 1983 (4 points below average) and had twice the average decline in I/S (7.32 points) is predicted to have a much larger 3.325-point increase in the ND share. And a firm that had a relatively bad I/S ratio in 1983 (4 points above average), and had only the average decline in I/S (3.66 points), and hence no improvement in its relatively bad position, is predicted to have almost no change in the ND share. Thus, the regression model says that improvement in the I/S ratio *relative* to the manufacturing average is a strong predictor of increasing ND share.

Although we hypothesized a relationship between improved logistics and intra-firm trade, we were surprised at the strength of the relationship between trade and I/S, as well as the lack of significance of other factors. To gain a deeper understanding of these results, we undertook case studies of several Canadian affiliates of U.S. MNCs across several industries. We detail our findings in the companion paper (Keane and Feinberg, 2005). Here, we give a brief summary.

In the 1970s and early 1980s, many U.S. MNCs suffered serious market-share losses. This was largely due to a challenge from Japanese manufacturers, who produced higher-quality products in greater variety, yet at lower unit costs. Many U.S. MNCs sent study teams to Japan in the early 1980s to learn how they did it. A key discovery was that many Japanese manufacturers had superior logistics management practices: they could sustain given levels of output with much lower levels of work-in-progress and final goods inventories. The main reason was the development of the just-in-time (JIT) production system by Taichii Ohno at Toyota in the 1950s and 1960s (see Ohno, 1988; Shigeo Shingo, 1989). JIT spread throughout much of Japanese manufacturing in the 1970s. Yasuhiro Monden (1981) and Richard J. Schonberger (1982) provided early expositions of the JIT system in English. General awareness of JIT among U.S. manufacturers developed in the 1983–1985 period. In order to respond to the Japanese manufacturing challenge, many U.S. MNCs began to adopt JIT during the 1980s and 1990s. However, adopting JIT requires a radical reorganization of the firm (i.e., around products rather than functional areas), so successful adoption was gradual.[61]

---

[59] For example, materials requirement planning (MRP) and Enterprise Resource Planning (ERP) systems collect and integrate transactional data within the firm and feed the data into manufacturing scheduling and other basic business functions, such as finance and accounting. These systems help firms reduce inventory by synchronizing upstream and downstream activities both across units of the firm and between a firm and its customers and suppliers.

[60] The mean I/S ratio in 1983 was 16.8 percent.

[61] The JIT system is not about reducing inventories just to reduce carrying costs. The central idea of the system is that inventories provide a buffer that lets a production process keep running despite quality problems (e.g., faulty parts) and despite process inefficiencies (e.g., bottlenecks). Inventories hide problems. By reducing inventories, the

The adoption of JIT techniques by U.S. manufacturers had dramatic effects. In Figure 7, we see that the I/S ratio in U.S. manufacturing hovered in a narrow range from 1958 to 1982. But in 1983 there is a structural break—a clear downward trend begins that has persisted ever since. Even more importantly, there is great heterogeneity across industries in the extent and timing of I/S reduction. Figure 8 presents I/S ratios for several industries over the 1981–1996 period. For example, in computers, the drop in I/S early in the period is dramatic—from 28 percent in 1984 to 16 percent in 1987—but it is then rather flat until another sharp drop in 1996. In contrast, in appliances, there is little change until the early 1990s; but there is a sharp drop from 16 percent in 1990 to 11 percent in 1992. Interestingly, our case studies reveal that this is exactly the period when GE's appliance division adopted JIT. The industries in Figure 8 were not chosen to be representative, but merely as illustrations; indeed, there are also several industries where I/S ratios declined very little.

Based on our case studies of individual firms and the evidence in Figures 7 and 8, we believe the I/S ratio serves largely as a measure of firms' success in adopting the JIT production system. Thus, the results in Table 7 imply that those industries where JIT production techniques have been most successfully adopted are also the industries where U.S. manufacturing parents have most increased their imports of intermediates from Canadian affiliates.

We emphasize that this result is not based on correlating two trending variables (i.e., I/S and intra-firm trade) in the aggregate. It is important that the timing and success of JIT adoption (and hence, inventory reduction) varies considerably across industries and firms, because this creates the leverage to identify the relationship between inventories and intra-firm trade at the industry/firm level. Furthermore, tariffs also exhibit a strong trend over the analysis period at the aggregate level (see Figure 1), but they nevertheless fail to correlate with intra-firm trade at the industry/firm level. Finally, our regression also includes a time trend, so the impact of I/S on intra-firm trade is identified only from industry-specific deviations from the aggregate I/S trend.

Having found empirical support for the hypothesis that advanced logistics practices like JIT increased intra-firm trade in intermediates, we ask if this finding makes sense theoretically. We believe it does. An important cost of transporting intermediates intra-firm, not captured by physical transport costs or tariffs, is the inventory carrying cost—arising from the time goods are in transport and/or sit in stock before they are used in the next stage of production. Inventory costs are higher to the extent that larger buffer stocks of intermediates are needed to insure against faulty shipments shutting down the next stage of production. Under JIT, the required buffer stock of work-in-progress inventory needed to support any given level of intra-firm trade is lower. Thus, *the JIT system lowers the inventory carrying cost of intra-firm trade in intermediates.* At a time when tariff and physical transport costs were already quite low (i.e., 1984), it seems plausible that inventory carrying costs were a substantial part of the cost of intra-firm trade in many industries. Indeed, the only two studies of the subject that we are aware of, the Hewlett-Packard (HP) study by Hau L. Lee et al. (1993) and the Digital Equipment Corporation (DEC) study by Bruce C. Arntzen et al. (1995), both found this to be true. In fact, HP and DEC used these studies to reorganize their worldwide manufacturing and distribution systems so as to reduce inventory carrying costs of intra-firm trade.

Three other aspects of the regression results in Table 7 also seem consistent with the idea that improved logistics management and JIT adoption reduced the costs of intra-firm trade. First, the coefficient on IT/S is actually negative

---

production process becomes more fragile, its problems are revealed, and solutions can be found. This leads to improved quality and greater productive efficiency in the long run. The JIT system also reduces the efficient scale of a plant by about two-thirds, since inventory storage, quality inspection, and rework areas are no longer needed (or at least greatly reduced). A key feature of JIT is its total systemic nature (see Peter E. Drucker, 1990). For instance, quality problems may stem from the design of the product itself, so manufacturing must be considered during the design stage. And keeping extremely low inventories requires holding down variance of demand, which has an impact on marketing and distribution. It also requires quick changeovers between producing different varieties of differentiated products, making it desirable to design varieties with many common components. Such complementary between aspects of design, manufacture, distribution, and marketing in "modern manufacturing" is discussed by Paul R. Milgrom and D. John Roberts (1990).

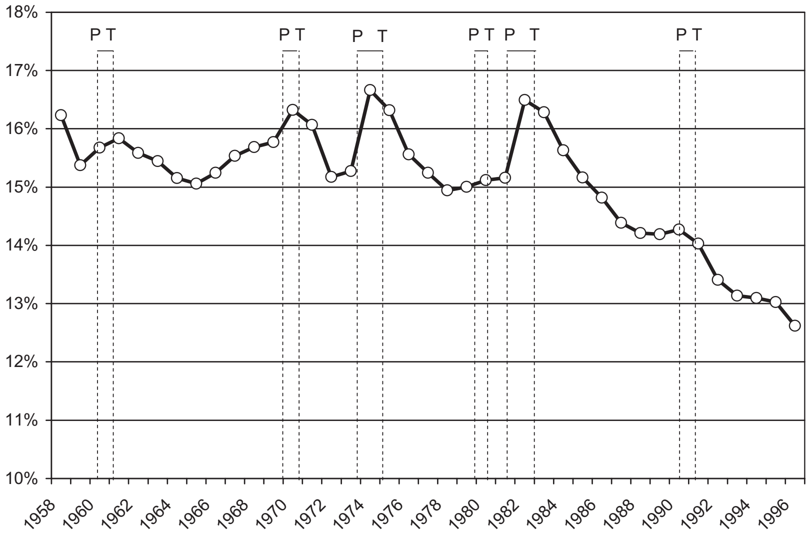

FIGURE 7. INVENTORY/SALES RATIOS IN U.S. MANUFACTURING, 1958–1996

Notes: Data are from the (annual) NBER Center for Economic Studies Manufacturing Industry Database compiled by Eric J. Bartelsman et al. (2000) (see NBER Technical Working Paper #205 for a discussion of the 1985–1991 data). Inventories to sales were defined as inventories divided by the value of shipments. Data were aggregated to three-digit SIC codes and matched with corresponding BEA industries. "P" and "T" denote business cycle peaks and troughs, respectively, which were constructed using the NBER's business cycle dates.

(−0.3376). This surprised us, because our original hypothesis was that IT (i.e., computers running MRP systems, bar codes, etc.) allowed easier communication between divisions of a firm. We thought this would have lowered the cost of coordinating fragmented production processes across locations, reducing the cost of intra-firm trade. But, as we later found, there is strong consensus in the industrial engineering literature that mere adoption of IT, without the substantial changes in management practice and organizational structure needed to implement a JIT system, does not improve logistics management.[62] Similarly, the coefficient on IT/S captures the effect of increasing IT investment, holding I/S *fixed.* Thus, a negative sign on IT/S makes sense: if a firm that adopts IT cannot improve I/S, it implies its logistics management is poorly organized. Such a firm could not successfully organize increased fragmentation of production across locations and increased intra-firm trade. Thus, we argue it was the JIT system, rather than computers per se, that reduced the cost of intra-firm trade.

Second, the interactions of the time trend with firm size in 1983 (as measured by log third-party sales of the entire MNC) and with IT/S in 1983 are both positive. This implies MNCs that were larger and more technologically advanced in 1983 tended to have larger ND share increases over our analysis period. This is consistent with our case studies in Keane and Feinberg (2005), which suggest that larger and more technically advanced MNCs were the first to learn about the JIT system, and were the earliest adopters of advance logistics management practices.

Finally, it is worth emphasizing that tariffs and transport costs are not significant in the regression in Table 7. The significant coefficient

---

[62] It is notable that Toyota originally implemented JIT using the "kanban" system, where physical signals, like cards or empty boxes, are sent back up the supply chain to signal parts requirements, rather than using sophisticated IT technology (see, e.g., Ohno, 1988; Shingo, 1989). Schonberger (1982) argued that computerized MRP logistics systems are highly effective at reducing inventories only if they are used to facilitate JIT implementation by computerizing the manual kanban system. See Keane and Feinberg (2005) for further references and discussion.

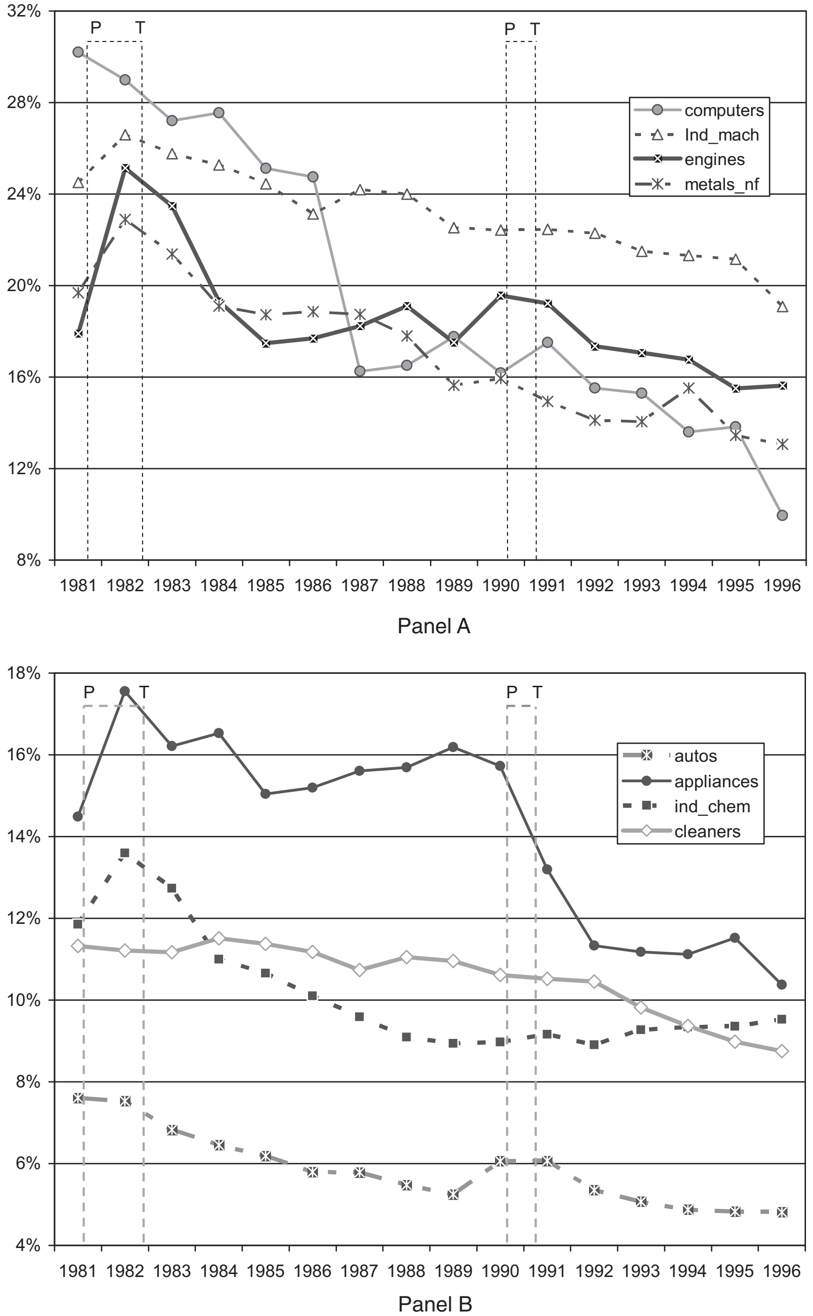

FIGURE 8. INVENTORY/SALES RATIOS IN SELECTED INDUSTRIES, 1981–1996

Notes: Data are from the (annual) NBER Center for Economic Studies Manufacturing Industry Database compiled by Bartelsman et al. (2000) (see NBER Technical Working Paper #205 for a discussion of the 1985–1991 data). Inventories to sales were defined as inventories divided by the value of shipments. Data were aggregated to three-digit SIC codes and matched with corresponding BEA industries. "P" and "T" denote business cycle peaks and troughs, respectively, which were constructed using the NBER's business cycle dates.

on the Canadian tariff plus transport cost that we found in Table 6 is wiped out by inclusion of I/S and the other control variables.

Next, we ran two additional regressions where the dependent variables are (a) the NF share, and (b) a measure of overall intra-firm trade in both directions. That is, we take ND plus NF and divide by total MNC sales to unaffiliated parties to get the intra-firm trade share of total MNC sales. The results of the ND + NF share regression look very similar to those for the ND share—in fact, they imply an even stronger correlation between improvement in the I/S ratio and increases in intra-firm trade—so we will not discuss them in detail.

In a regression with the NF share as the dependent variable, however, we basically find that nothing is significant except the generic time trend. We conjecture this occurs because the NF share may increase for two diametrically opposed reasons. First, the affiliate may become more integrated into the parent's production process, leading to increases in both NF and ND. Second, NF may increase because an affiliate is "hollowed out"—i.e., converted into a low-valued-added "screwdriver" plant, with high-value-added subassemblies all imported from the parent. Suppose these two cases are present in the data with roughly equal frequency. We would then expect NF to be uncorrelated with advances in logistics that enhance the role of affiliates, while ND and ND + NF would both be positively correlated with such advances.

VIII. Conclusions

In this paper, we estimate a structural model of MNC production and trade decisions to gain insight into the factors causing the substantial increase in MNC-based trade between the United States and Canada over the past two decades. Our model implies that bilateral tariff reductions led to a roughly one-third increase in the volume of *arms-length* MNC-based trade between the United States and Canada over the 1984–1995 period. On the other hand, tariff reductions can account for only a 5- to 10-percent increase in *intra-firm trade.* Intra-firm trade nearly doubled over this period, and our model implies that "technical change" accounted for most of the increase.

Our structural model is silent about the underlying source of the technical change driving the dramatic increase in intra-firm trade. Hence, we conduct an exploratory investigation, in which we regress firms' cost shares for intra-firm intermediates on a large set of industry/firm characteristics, in addition to tariffs, transport costs, and factor prices. We obtain the surprising result that *success in reducing inventories has a very strong positive relationship with increased intra-firm trade.* The 1983–1995 period is precisely when many U.S. MNCs and Canadian affiliates began in earnest to adopt advanced logistics management practices, such as the JIT production system pioneered by Toyota in the 1950s and 1960s. The impact of adopting these methods is reflected in the dramatic declines in manufacturing inventories that occurred in many industries during the 1980s and 1990s. Based on this, we argue that improved logistics management in general, and JIT in particular, is a key reason for increased intra-firm trade.

This conclusion is theoretically plausible, because improved logistics enable firms to better organize "convergent" production processes that involve frequent intra-firm transfers of goods (see Troy J. Strader et al., 1999; Michael E. McGrath and Richard W. Hoole, 1992), and reduces inventory-carrying costs of intra-firm trade. In the relatively low-tariff environment that already existed between the United States and Canada in 1984, it is plausible that inventory carrying costs were often a more important component of trade costs than were tariffs. Indeed, the industrial engineering studies by HP (see Lee et al., 1993) and DEC (see Arntzen et al., 1993) both concluded this was the case.[63]

We emphasize that our conclusion that improved logistics management led to increased intra-firm trade is not based simply on correlating two trending variables (i.e., inventories and

---

[63] It is possible that increased intra-firm trade occurred in part because MNCs acquired firms that were previously third-party suppliers. This is consistent with a story of technical change (i.e., improved communications technology) reducing governance costs (see Grossman and Helpman, 2002). But this story is implausible because, as the figures in Table 5 indicate, in terms of both net sales and employment, MNCs shrink slightly over our analysis period.

intra-firm trade) in the aggregate. The timing and success of JIT adoption (and, hence, inventory reduction) varied considerably across industries and firms, creating the leverage to identify the relationship between inventories and intra-firm trade at the industry/firm level.

Our companion paper, Keane and Feinberg (2005), presents case studies that describe how many U.S. MNCs restructured their manufacturing facilities on a worldwide basis in the 1980s and 1990s. Typically, they shut down many plants, and, in those that remained, moved toward JIT production systems, with more common components across products, and global sourcing of those components. This led to closer coordination between plants, greater intra-firm trade, and less local duplication of effort. We show that many Canadian affiliates, facing potential shutdown, took the initiative to shift from final goods to high-value-added intermediate production, while also shifting to JIT in order to become low-cost global suppliers to the MNC. In all cases, it appears that loss of market share, rather than tariff reductions, drove the restructuring.

The work presented here may have implications that go beyond the U.S.-Canada context. The magnitude of the increase in world trade over the past few decades is generally considered an important "mystery" (see Burgeoning and Kehoe, 2001; Yi, 2003). It is hard to explain based on tariffs and transport costs alone, because the growth of trade was so massive while declines in tariffs and transport costs were so modest. The mystery has become particularly severe since the mid-1980s when the growth of world trade accelerated noticeably, even though tariffs were already quite low by the early 1980s.[64]

The finding that improved I/S ratios at the industry level are closely associated with growth of trade provides an important clue about what may be going on. The 1980s were precisely when many manufacturing firms in the United States and Western Europe began in earnest to adopt advanced logistics management systems like the JIT system. Since JIT lowers the inventory carrying-cost component of trading goods intra-firm, it may plausibly account for a decline in the cost of intra-firm trade well beyond that due to declines in tariffs and transport costs. Prior empirical work on the growth of trade has focused on tariffs and physical transport costs (see Yi, 2003, p. 91, for a good review), but has not paid attention to inventory carrying costs. Our work suggests that these may be crucial, especially for explaining the acceleration in trade growth in the 1980s.

Finally, we used our model to perform one more interesting simulation, where we reduce the Canadian wage rate by 1 percent. The model predicts that this would increase Canadian affiliate employment by 4.2 percent, and *increase* U.S. parent employment by 0.08 percent. Thus, while the MNC substitutes Canadian for U.S. labor, the scale effect counteracts this, leading to a slight overall *positive* effect on domestic employment. This illustrates how "offshoring" part of the production process to lower-wage foreign labor can actually increase demand for domestic labor by reducing the parent's production costs.

REFERENCES

Arntzen, Bruce C., Gerald G. Brown, Terry P. Harrison, and Linda L. Trafton. 1995. "Global Supply Chain Management at Digital Equipment Corporation." *Interfaces*, 25(1): 69–93.

Baldwin, John R., and Paul K. Gorecki. 1986. *The Role of Scale in Canada-U.S. Productivity Differences in the Manufacturing Sector, 1970–1979*. Toronto: University of Toronto Press.

Baldwin, Richard E., and Gianmarco I. P. Ottaviano. 2001. "Multiproduct Multinationals and Reciprocal FDI Dumping." *Journal of International Economics*, 54(2): 429–48.

Bergoeing, Raphael, and Timothy J. Kehoe. 2001. "Trade Theory and Trade Facts." Federal Reserve Bank of Minneapolis, Staff Report 284.

---

[64] Yi (2003) is the most successful attempt to explain the growth of trade using tariff reductions in a general equilibrium model, but, as he notes, his model still explains only half of the growth of U.S. exports in the post-1962 period, and "falls short of capturing the nonlinear export surge beginning in the late 1980s" (p. 85). In 1989–1999, U.S. exports grew 80 percent, but his model generates only a 27-percent increase (p. 88). In his conclusion, Yi speculates that one reason for the residual may be "technology induced increases in the ... possibilities for vertical specialization." This is basically another way to state our explanation based on technology-induced increases in intermediate input shares.

Bernard, Andrew B., and J. Bradford Jensen. 2004. "Why Some Firms Export." *Review of Economics and Statistics,* 86(2): 561–69.

Brainard, S. Lael. 1993. "An Empirical Assessment of the Factor Proportions Explanation of Multinational Sales." National Bureau of Economic Research Working Paper 4583.

Brainard, S. Lael. 1997. "An Empirical Assessment of the Proximity-Concentration Trade-Off between Multinational Sales and Trade." *American Economic Review*, 87(4): 520–44.

Brander, James A. 1981. "Intra-Industry Trade in Identical Commodities." *Journal of International Economics*, 11(1): 1–14.

Caves, Richard E. 1982. *Multinational Enterprise and Economic Analysis*. Cambridge: Cambridge University Press.

Caves, Richard E. 1990. *Adjustment to International Competition: Short-Run Relations of Prices, Trade Flows, and Inputs in Canadian Manufacturing Industries*. Ottawa: Economic Council of Canada.

Cummins, Jason G. 1998. "Taxation and the Sources of Growth: Estimates from United States Multinational Corporations." National Bureau of Economic Research Working Paper 6533.

Das, Sanghamitra, Mark J. Roberts, and James R. Tybout. 2001. "Market Entry Costs, Producer Heterogeneity, and Export Dynamics." National Bureau of Economic Research Working Paper 8629.

Davis, Donald R., and David E. Weinstein. 2001. "An Account of Global Factor Trade." *American Economic Review*, 91(5): 1423–53.

Deardorff, Alan V. 2001. "Fragmentation in Simple Trade Models." *North American Journal of Economics and Finance,* 12(2): 121–37.

Dixit, Avinash K., and Gene M. Grossman. 1982. "Trade and Protection with Multistage Production." *Review of Economic Studies*, 49(4): 583–94.

Drucker, Peter E. 1990. "The Emerging Theory of Manufacturing." *Harvard Business Review,* 68(3): 94–101.

Eastman, Harry C., and Stefan S. Stykolt. 1967. *The Tariff and Competition in Canada*. Toronto: Macmillan of Canada.

Eaton, Jonathan, and Samuel S. Kortum. 2002. "Technology, Geography, and Trade." *Econometrica*, 70(5): 1741–79.

Eden, Lorraine. 1998. *Taxing Multinationals: Transfer Pricing and Corporate Income Taxation in North America*. Toronto: University of Toronto Press.

Feinberg, Susan E., and Michael P. Keane. 2001. "U.S.-Canada Trade Liberalization and MNC Production Location." *Review of Economics and Statistics*, 83(1): 118–32.

Feinberg, Susan E., and Michael P. Keane. 2005. "Tariff Effects on MNC Decisions to Engage in Intra-Firm and Arms Length Trade." http://www.business.uts.edu.au/finance/staff/michaelk/extensive_margin_final_4.pdf.

Feinberg, Susan E., Michael P. Keane, and Mario F. Bognanno. 1998. "Trade Liberalization and 'Delocalization': New Evidence from Firm-Level Panel Data." *Canadian Journal of Economics*, 31(4): 749–77.

Fraumeni, Barbara M., and Dale W. Jorgenson. 1980. "Rates of Return by Industrial Sector in the United States, 1948–76." *American Economic Review (Papers and Proceedings)*, 70(2): 326–30.

Gaston, Noel, and Daniel Trefler. 1997. "The Labour Market Consequences of the Canada-U.S. Free Trade Agreement." *Canadian Journal of Economics*, 30(1): 18–41.

Griliches, Zvi, and Jacques Mairesse. 1995. "Production Functions: The Search for Identification." National Bureau of Economic Research Working Paper 5067.

Grossman, Gene M., and Elhanan Helpman. 2002. "Integration versus Outsourcing in Industry Equilibrium." *Quarterly Journal of Economics*, 117(1): 85–120.

Grubert, Harry, and John Mutti. 1991. "Taxes, Tariffs and Transfer Pricing in Multinational Corporate Decision Making." *Review of Economics and Statistics*, 73(2): 285–93.

Hanson, Gordon H., Raymond J. Mataloni, Jr., and Matthew J. Slaughter. 2005. "Vertical Production Networks in Multinational Firms." *Review of Economics and Statistics,* 87(4): 664–78.

Head, Keith, and John Ries. 1997. "Market-Access Effects of Trade Liberalization: Evidence from the Canada-U.S. Free Trade Agreement." In *The Effects of U.S. Trade Protection and Promotion Policies,* ed. Robert C. Feenstra, 323–42. Chicago: University of Chicago Press.

Head, Keith, and John Ries. 2001. "Increasing Returns versus National Product Differentiation as an Explanation for the Pattern of U.S.-

Canada Trade." *American Economic Review*, 91(4): 858–76.

▶**Helpman, Elhanan.** 1984. "A Simple Theory of International Trade with Multinational Corporations." *Journal of Political Economy*, 92(3): 451–71.

▶**Helpman, Elhanan.** 1985. "Multinational Corporations and Trade Structure." *Review of Economic Studies*, 52(3): 443–57.

Helpman, Elhanan, and Paul R. Krugman. 1985. *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*. Cambridge, MA: MIT Press,

▶**Ihrig, Jane.** 2000. "Multinationals' Response to Repatriation Restrictions." *Journal of Economic Dynamics and Control*, *24*(9): 1345–79.

Katayama, Hajime, Shihua Lu, and James R. Tybout. 2003. "Why Plant-Level Productivity Studies Are Often Misleading, and an Alternative Approach to Interference." National Bureau of Economic Research Working Paper 9617.

Keane, Michael P. 2003. "SML Estimation Based on First Order Conditions." http://www.business.uts.edu.au/finance/staff/michaelk/sml-foc.pdf.

Keane, Michael P., and Susan E. Feinberg. Forthcoming. "Advances in Logistics and the Growth of Intra-firm Trade: The Case of Canadian Affiliates of U.S. Multinationals, 1984–1995." *Journal of Industrial Economics*.

▶**Klette, Tor Jacob, and Zvi Griliches.** 1996. "The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous." *Journal of Applied Econometrics,* 11(4): 343–61.

▶**Krusell, Per, Lee E. Ohanian, José-Victor Ríos-Rull, and Giovanni L. Violante.** 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica*, 68(5): 1029–53.

Lee, Hau L., Corey Billington, and Brent Carter. 1993. "Hewlett-Packard Gains Control of Inventory and Service through Design for Localization." *Interfaces*, 23(4): 1–11.

Levinsohn, James A., and Marc J. Melitz. 2002. "Productivity in a Differentiated Products Market Equilibrium." University of California at Santa Cruz Center for International Economics Working Paper 01–23.

▶**Markusen, James R.** 1984. "Multinationals, Multi-Plant Economies, and the Gains from Trade." *Journal of International Economics*, 16(3–4): 205–26.

Markusen, James R., and Keith E. Maskus. 1999. "Multinational Firms: Reconciling Theory and Evidence." National Bureau of Economic Research Working Paper 7163.

▶**Markusen, James R., and Keith E. Maskus.** 2002. "Discriminating among Alternative Theories of the Multinational Enterprise." *Review of International Economics,* 10(4): 694–707.

Mathewson, G. Frank, and G. David Quirin. 1979. *Fiscal Transfer Pricing in Multinational Corporations*. Toronto: University of Toronto Press.

McGrath, Michael E., and Richard W. Hoole. 1992. "Manufacturing's New Economies of Scale." *Harvard Business Review,* 70(3): 94–102.

Milgrom, Paul R., and John D. Roberts. 1990. "The Economics of Modern Manufacturing: Technology, Strategy, and Organization." *American Economic Review*, 80(3): 511–28.

Monden, Yasuhiro. 1981. "What Makes the Toyota Production System Really Tick?" *Industrial Engineering*, 13(1): 36–46.

Ohno, Taiichi. 1988. *Toyota Production System: Beyond Large-Scale Production*. New York: Productivity Press.

Rugman, Alan M. 1985. "Transfer Pricing in the Canadian Petroleum Industry." In *Multinationals and Transfer Pricing,* ed. Alan M. Rugman and Lorraine Eden, 173–92. London: Croon Helm, Ltd.

Rugman, Alan M. 1988. *The Determinants of Intra-Industry Direct Foreign Investment*. New York: St. Martin's Press.

▶**Schonberger, Richard J.** 1982. "Some Observations on Advantages and Implementation Issues of Just-in-Time Production Systems." *Journal of Operations Management*, 3(1): 1–11.

Shingo, Shigeo. 1989. *A Study of the Toyota Production System from an Industrial Engineering Viewpoint*. New York: Productivity Press.

Soboleva, Nadia. 2000. "Foreign Direct Investment: A Dynamic Model of a Firm's Location Choice." Unpublished.

▶**Solow, Robert M.** 1957. "Technical Change and the Aggregate Production Function." *Review of Economics and Statistics*, 39(3): 312–20.

▶**Stiroh, Kevin.** 2002. "Information Technology and the U.S. Productivity Revival: What Do the Industry Data Say?" *American Economic Review,* 92(5): 1559–76.

▶**Strader, Troy J., Fu-Ren Lin, and Michael J. Shaw.** 1999. "Business-to-Business Electronic Commerce and Convergent Assembly Supply Chain Management." *Journal of Information Technology, 14*(4): 361–73.

▶**Trefler, Daniel.** 2004. "The Long and Short of the Canada-U.S. Free Trade Agreement." *American Economic Review*, 94(4): 870–95.

▶**Yi, Kei-Mu.** 2003. "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy*, 111(1): 52–102.

Zeile, William J. 1997. "U.S. Intrafirm Trade in Goods." *Survey of Current Business*, 77(2): 23–38.