Forthcoming, *Review of World Economics/*
*Weltwirtschaftliches Archiv*, December 2006

# NEW EVIDENCE ON THE GAINS FROM TRADE

Robert C. Feenstra

*University of California, Davis*
*and National Bureau of Economic Research*

**ABSTRACT**

We summarize the evidence on the gains from trade in monopolistic competition models, arising from three sources: (i) price reductions due to increasing returns to scale; (ii) increased product variety available to consumers; (iii) self-selection of firms with only the most efficient firms surviving after trade liberalization. There is little direct evidence to support the first source of gains from trade, though some indirect evidence from the European Union. The second and third sources of gains from trade find strong empirical support from studies from various countries, relying on new models and new empirical methods.

JEL no. F10, F12, F15

*Keywords:* tariffs, elasticity of substitution, monopolistic competition, increasing returns to scale

*Remark:* Lecture delivered at the Kiel Institute for World Economics on the occasion of being awarded the Bernhard Harms Prize on September 15, 2006.

*Address for correspondence:* Robert Feenstra, Department of Economics, University of California, Davis, CA 95616 USA; E-mail, rcfeenstra@ucdavis.edu; Fax, (530)752-9382.

## 1. Introduction

A modern-day revolution occurred in international trade theory during the early 1980s with the introduction of economies of scale and monopolistic competition. Since it has now been about 25 years since the first articles by Helpman (1981), Krugman (1979, 1980, 1981) and Lancaster (1980), it is appropriate to look back and seen what the impact of these models has been. On the theoretical side their impact has been very great indeed, with the static models of the 1980s giving rise to the dynamic models of endogenous growth in the 1990s, and in the current decade leading to models with heterogeneous firms, which are being used in both international trade and international macroeconomics. But my interest this afternoon is not so much in the theoretical contributions of these models, but rather, in the empirical work that has developed recently around this class of models.

After all, the models of economies of scale and monopolistic competition were conceived with a very practical application in mind, namely, the gains that would result from large-scale tariff reductions. Whether from multilateral tariff reductions under the WTO, or bilateral tariff reductions under regional trade agreements, these models predicted gains from trade over and above the gains from specialization in conventional models. A favorite example of mine is the gains to Canada from free trade with the United States. There was a literature in Canada dating back to the 1960s that predicted substantial gains from free trade with the US as Canadian firms would then expand their scale of operation and lower their costs. That literature pre-dates the formal models of monopolistic competition, but is very much in the same spirit. A set of simulations exercises performed by Harris (1984a,b) in the mid-1980s were very influential in convincing Canadian policy makers to proceed with the free trade agreement with the US in

1989. With more than 15 years of data now behind us, it is appropriate to look back and compare the outcome of the Canada-US free trade agreement with those simulations.

Of course, North America is not the only testing ground for the models of monopolistic competition and increasing returns to scale: the unification of the European market over the past two decades is perhaps even more important. Simulations done in the late 1980s by Smith and Venables (1988, 1991) predicted large gains to the 1992 Single Market reforms in Europe, allowing for greater unification of the market. Of course, those simulations did not anticipate the subsequent enhancement of labor mobility and also the adoption of a common currency in 2002. In view of all these reforms, it is highly appropriate to see if the simulation results of Smith and Venables have been borne out in the actual experience within Europe.

To organize the discussion, let me deal with three sources of gains from trade predicted by the models of monopolistic competition: first, a fall in prices after tariff reductions due to greater competition between firms; second, an increase in the variety of products available to consumers, allowing for further gains; and third, self-selection of firms with only the more efficient firms surviving after trade liberalization. Not all of these predictions come from the original models of the early 1980's, but have been added in later research, as I shall describe.

## 2. Price reductions

Let us begin with the price reductions that are supposed to follow from trade liberalization. In the original models of Helpman, Krugman and Lancaster, these price reductions are closely related to increasing returns to scale. As tariffs are reduced between two countries, some firms exit the market and the remaining firms expand their outputs and lower their average costs through economies of scale. The reduction in average costs also leads to a reduction in prices in the zero-profit equilibrium.

So the very first empirical question we can ask is whether there is evidence of this expansion of firm scale, and reduction in costs, following the tariff reductions in North America or other regions. The assumption that economies of scale would apply to Canadian firms was made by Harris (1984a,b) in his simulation models of the Canada-US free trade agreement. He used engineering estimates of the extent to which costs would fall as scale expanded in various industries, and based on those estimates, he predicted that firm output would expand by 40 to 70 percent with a rise in labor productivity of 20 to 30 percent (Harris, 1984a, p. 1028). The operation of these scale economies is the principle source of gains from trade within his model.

But surprisingly, when we turn to the actual evidence from Canada, there is very little indication that firms expanded their scale of operations. Work by Head and Ries (1999, 2001) find no systematic indication that Canadian firms grew more in the industries with the greatest tariff reductions, and that negative finding is confirmed by more recent work by Daniel Trefler (2004), which I will refer to again later. Furthermore, when we look at episodes of tariff liberalization in developing countries such as Chile and Mexico, as done by Tybout *et al* (1991, 1995), there is again little indication that a fall in tariffs leads to an expansion in firm scale. So one cornerstone of monopolistic competition models, that tariff reductions lead to an expansion of scale and a fall in prices, does not find empirical support from episodes of trade liberalization.

But there is second route by which free trade areas can lead to a reduction in prices, aside from economies of scale. This second route is stressed in the work by Smith and Venables (1988, 1991), in simulation models they developed to analyze the 1992 Single Market Programme in Europe. Those reforms were not about tariff reductions, since tariffs within the EEC were already zero; rather, the reforms dealt with the elimination of rules and regulations governing the flow of goods between European countries. As these non-tariff barriers were eliminated, Smith and Venables expected that firms would be forced to equalize their selling prices across market. In other words, rather than treating Europe as a collection of segmented markets, where firms could choose their prices in each country separately, Europe would instead become a unified market where firms could not price-discriminate. As price-discrimination is eliminated, then the average prices are expected to fall, providing benefits to consumers.

Given the nearly 15 years since the Single Market reforms of 1992, and the much shorter period since the adoption of the Euro in 2002, we can ask whether the prediction of unified and lower prices within Europe has been realized. The results are mixed. In early work, Smith and his colleagues (Allen *et al,* 1998) were quite optimistic about the pro-competitive effects of the Single Market Programme. But since that work was based on data ending in 1994 the time period was too brief to allow for strong conclusions. A very recent paper by Badinger (2006) uses sectoral data from 1981 to 1999 and finds solid evidence of markup reductions in manufacturing and construction, but not in services. The service industry that we are all perhaps most familiar with is restaurants, where it is widely believed that prices increased following the adoption of the Euro. But a new paper by Hobijn, Ravenna and Tombalotti (2006) argues that this increase can be understood as making up for unusually small price changes prior to the adoption of the Euro, and in fact, the real puzzle is why such price increases were not more widespread. So I conclude that there is some evidence in favor of falling markups in Europe, but not in all sectors.

Furthermore, there is some *indirect* evidence that we can bring to bear on the question of segmented versus unified markets. I am thinking here of the somewhat controversial work on the effects of currency unions on the volume of trade. Many economists believe that adopting a fixed exchange rate or a common currency across industrial countries should have only a small impact on trade, since even under flexible rates, firms should be able to hedge away their risks using forward markets. But that belief is not supported by the empirical evidence. Due to the work of Rose (2000), we know that common currencies actually have a large and positive impact on trade between countries: countries in a currency union can have *three times* more trade with other than expected, or an increase of 200 percent. That result just seem too big for economists to believe, so a fair amount of work has gone into explaining it away.

Perhaps the most comprehensive paper along these lines is the recent survey by Baldwin (2006a,b), which focuses on the Euro's effect on trade. In his own delightful phrase, Baldwin says that he was asked "to write a short paper on the subject but I didn't have time for that so I wrote a long paper instead." Baldwin systematically goes through the estimate of the Euro on trade, and shows how the specification of the regression equations can heavily influence the results. The preferred estimate seems to be in the range of 8 to 15 percent, which is of course much less than the 200 percent estimate of Rose. The implication of this finding is that we should not expect the Euro to greatly increase trade between the countries using it, and as a corollary, we should not expect the introduction of the Euro to substantially reduce prices.

I do not want to debate with Baldwin or Rose as to the exact magnitude of currency unions on trade, and am willing to accept any result around their preferred estimates. But I still believe that the positive impact of currency unions is larger than many economists expected. And I would like to suggest that such a result is consistent with the simulation models performed by Smith and Venables nearly two decades ago. Just like the elimination of rules and regulations between countries will limit the ability of firms to price-discriminate, and thereby increase trade, so too the elimination of currency fluctuations between countries will limit the ability of firms to charge separate prices across those national markets (Friberg, 2001). With the introduction of the Euro or other currencies unions, these models of segmented versus unified market would predict a significant increase in trade. So we should not be surprised in the first place to discover that currency unions enhance trade, and that result is highly consistent with the models of firm competition used by Smith and Venables. By stressing the ability of firms to price-discriminate or not in homogeneous products, these models are different from the monopolistic competition framework of Helpman, Krugman and Lancaster. But despite that difference, the models give us a framework to understand why prices will fall in unified markets, and Europe is the ideal testing ground for these ideas. So we can hope that further empirical research will add to the results that we have already, and demonstrate the gains from trade due to a unified market in Europe.

**3. Product Variety**

Let me turn now to the second reason that trade liberalization generates gains for consumers in the monopolistic competition model, and that is the expansion in the variety of goods available through trade. Besides economies of scale, another feature of these models is the idea that each firm produce products that are somewhat different from other firms. Whether we are talking about automobiles, or consumer electronics, or food products, or nearly any other industry, it is very plausible that firms will differentiate their products and that trade allows consumer to purchase more varieties. So this second cornerstone of the monopolistic competition framework seems plausible at face value.

However, if we go back to the simulation model for Canada by Harris, he was reluctant to build in the assumption of product differentiation. The reason, I believe, is that Harris realized that the calculated gains from trade would be very sensitive to the extent of differentiation across products, i.e. on the elasticity of substitution. If the elasticity of substitution of high, then products are easily substitute for each other, and consumers do not gain much from having new varieties available. A high elasticity of substitution may describe T-shirts for example, where we probably don't care too much whether our T-shirts comes from China, India, Vietnam, or wherever.

But as soon as we leave such basic items of clothing, and go to higher-fashion items, then the product sold by one company is probably quite different from the product sold by another. Consumers gains by having more choices available when they shop for fashion items, or electronics, or nearly anything else. So the elasticity of substitution between these items is lower, which indicates that consumers benefit more from having greater variety available.

Harris was reluctant to build product differentiation into his simulation model for Canada because he did not know what value to use for the elasticity of substitution in each industry. For technical reasons I will discuss in a moment, the empirical estimates for the elasticity that were available in the 1980s were quite poor. Often the estimated elasticities were too low, which would result in exaggerated estimates of the consumer benefits from product variety. Harris realized this potential for bias in his simulation results, so whereas he always made use of economies of scale, he added product differentiation only a secondary feature to the simulation models.

But now two decades later, we do have the statistical technique we need to estimate the elasticity of substitution between varieties of each and every product. This statistical technique comes from Feenstra (1994), which deals with the empirical methods needed to analyze the gains from trade due to expanding product variety. I applied that statistical technique to just half a dozen products, obtain estimates of the elasticity of substitution for each. Now more than ten years later our computing power has increase by several orders of magnitude, and the same technique been applied to over 30,000 products in recent work by Broda and Weinstein (2006). They obtain estimates of over 30,000 elasticities of substitution, all of which feed into the calculation of the gains from trade. While their work has focused on the gains for the United States, it could readily be applied to Canada, Europe, or any other set of countries. Before I report the results obtained by Broda and Weinstein, I would like to digress and explain the background to my 1994 paper.

***First Digression: Measuring the gains from new goods***

In 1989 I was invited to visit the Institute for Advanced Studies at the Hebrew University of Jerusalem, with a group organized by Elhanan Helpman. At that time the static models of monopolistic competition and trade had been published, and work was beginning on the dynamic models of endogenous growth. Joining the group at the Hebrew University were Paul Krugman, Gene Grossman, William Ethier and James Markusen, in addition to Elhanan, Assaf Razin and others, all of whom had contributed in important ways to the development of the monopolistic competition models in trade. The dynamic models emphasized that as new product varieties become available to firms as inputs, then the firms experience a productivity gain, and by adding this up over all firms we can generate ongoing growth in the economy. While there were many theoretical features of these dynamic models that needed to be worked on, my own interest was in the potential empirical applications.

If we take seriously the idea that new inputs lead to productivity gains for firms, just like new goods lead to welfare gains for consumers, then the key empirical question is now to measure these gains. This is a question that international trade economists are not especially familiar with, but it had been asked in a different branch of economics, namely, within the theory of index numbers. For any price index, such as the consumer price index, the appearance of a new good such as a cell phone raises a problem: if we have the price of the cell phone today, but do not have a price last year or last decade, then how can we measure the price decline for that good? The answer given many years ago by Hicks (1940) was that the relevant price of a product before it is available is the "reservation price" for consumers, namely, a price so high that their demand would be zero. Once the product appears on the market then it will have a lower price, determined by supply and demand. Then the fall in the price from its "reservation" level to the actual price can be used to measure the consumer gains from the appearance of that new good.

This idea of Hicks has been applied to new products by Hausman (1997, 1999), who analyzes the appearance of cellular telephones or a new breakfast cereal. The empirical method that Hausman uses requires that we estimate a reservation price for each new product. But we run into difficulty when we try to apply the same idea to the appearance of new products varieties from many countries due to trade liberalization. If we assume that each supplying country is providing a different variety from each other country, then we potentially have hundreds if not thousands of new product varieties through trade, and it is impractical to estimate the reservation price of each. So while the method recommended by Hicks is absolutely correct in theory, it is not that useful in practice when there are very many new varieties.

One way to resolve this difficulty is to use an assumption that is common in the monopolistic competition models, namely, that the elasticity of substitution of each product is *constant.* That assumption implies that the reservation price for each new product variety is actually infinite: no matter how high the price is, there is always some consumer who is willing to buy at least a very small amount of the product. So now we can ask: what happens if we take the price index formula that applies to the constant-elasticity case, and allow one price within that formula to approach infinity, indicating that the product is not available in that year? Do we still get a reasonable limiting value for the price index?

That is the question I posed during my visit in Jerusalem. I did not think of the question alone, but it came through a discussion with James Markusen, a colleague whom I rode the bus with each morning on our way to the Hebrew University. Markusen and I actually have a paper (Feenstra and Markusen, 1994) where we thank the Egged bus company for providing us with time, if not seats, each morning. It was on one of those morning bus rides that I realized that no one had yet tried to take the limiting value of a price index formula as one price approaches infinity. So that is what I did as soon as I arrived at the office, and by lunch, I realized that the limiting value was completely sensible. In the constant-elasticity case, the decline in price from infinity to its observed level causes a reduction in the price index that depends on just two pieces of information: first, the share of spending on the new good when it becomes available; and second, the elasticity of substitution between that good and its substitutes. If the share of spending on a new good is higher, or the elasticity is lower, then the appearance of this new good has a greater impact on pulling down the overall price index, which is an indication of greater gains to consumers from its appearance.

To obtain this result, it is convenient to work with the formula for the exact price index corresponding to a CES function due to Diewert (1976).[1] The CES unit-cost function is given by:

$$c(p_t, I_t) = \left[\sum_{i \in I_t} b_i p_{it}^{1-\sigma}\right]^{1/(1-\sigma)}, \ \sigma > 1, \ b_i > 0, \qquad (1)$$

where $\sigma$ is the elasticity of substitution, $p_{it}$ are the prices of some inputs, and $i \in I_t$ is the set of inputs (let us call them goods) available in year t. We are interested in determining how much unit-costs are reduced when the set of product varieties expands. To this end, let us first consider

[1] Actually, Diewert (1976, p. 130) shows that equation (2) holds for a broader class of unit-cost functions, called the “quadratic mean of order r” functions, which include the CES as a special case.

the case where $I_{t-1} = I_t$ so there is *no change* in the set of goods. Let us also assume that the observed inputs purchased are cost minimizing, that is, $x_{it} = y_t(\partial c / \partial p_{it})$, where $y_t$ is the firm's output. In that case, the ratio of unit-costs can be measured by the price index due to Diewert (1976, p. 131):

$$\frac{c(p_t, I_t)}{c(p_{t-1}, I_t)} = P_r(p_{t-1}, p_t, x_{t-1}, x_t, I_t) \equiv \left\{ \frac{\sum_{i \in I_t} (p_{it} / p_{it-1})^{r/2} s_{it-1}(I_t)}{\sum_{i \in I_t} (p_{it-1} / p_{it})^{r/2} s_{it}(I_t)} \right\}^{1/r}, \quad (2)$$

where $r = 1 - \sigma$, and $s_{it}(I_t) \equiv p_{it} x_{it} \Big/ \sum_{i \in I_t} p_{it} x_{it}$ are the expenditure shares. These can be obtained by differentiating the unit-cost function to obtain $x_{it}$, so that:

$$s_{it}(I_t) = \left( \frac{b_i p_{it}^{1-\sigma}}{\sum_{i \in I_t} b_i p_{it}^{1-\sigma}} \right) = \left( \frac{b_i p_{it}^{r}}{\sum_{i \in I_t} b_i p_{it}^{r}} \right). \quad (3)$$

Now suppose that the first good is not available in period t-1, so that we let $p_{1t-1} \to \infty$. Notice that for $\sigma > 1$, as we assume, then $r < 0$. The price term $p_{1t-1}^{-r/2}$ which appears in the numerator of (2) approaches infinity. But a quick calculation from (3) shows that $s_{1t-1} \to 0$ as $p_{1t-1} \to \infty$, and also that $p_{1t-1}^{-r/2} s_{1t-1} \to 0$. In other words, the share $s_{1t-1}$ approaches zero faster than $p_{1t-1}^{-r/2}$ approaches infinity (since $s_{1t-1}$ depends on $p_{1t-1}^{r}$). In the denominator of (2) we also have that $p_{1t-1}^{r/2} s_{1t} \to 0$, since $r < 0$. So the terms in (2) involving the infinite price all cancel out, and we are left with the summations over all goods *except* good 1 in the numerator and denominator. It follows that the shares $s_{it}(I_t)$ appearing in the denominator sum to *less than unity* over the goods $i \in I_t$, $i \neq 1$.

To adjust for this, define the set of goods available in both periods as $I \equiv \{i \in I_t,\ i \neq 1\}$, and then the period t shares defined over just these goods are:

$$s_{it}(I) = \left( \frac{p_{it} x_{it}}{\sum_{i \in I} p_{it} x_{it}} \right) = s_{it}(I_t) \left( \frac{\sum_{i \in I_t} p_{it} x_{it}}{\sum_{i \in I} p_{it} x_{it}} \right). \tag{4}$$

These re-defined shares sum to unity over the set of goods I. We can re-define the period t-1 shares $s_{it-1}(I)$ in a similar fashion, but since good 1 is not purchased in that period, then $s_{it-1}(I)$ $= s_{it-1}(I_t)$. Using (4), the limiting value of the price index in (2) becomes:

$$\begin{aligned} \lim_{p_{1t-1} \to \infty} & \left\{ \frac{\sum_{i \in I_t} (p_{it} / p_{it-1})^{r/2} s_{it-1}(I_t)}{\sum_{i \in I_t} (p_{it-1} / p_{it})^{r/2} s_{it}(I_t)} \right\}^{1/r} \\ &= \left\{ \frac{\sum_{i \in I} (p_{it} / p_{it-1})^{r/2} s_{it-1}(I)}{\sum_{i \in I} (p_{it-1} / p_{it})^{r/2} s_{it}(I)} \right\}^{1/r} \left( \frac{\sum_{i \in I_t} p_{it} x_{it}}{\sum_{i \in I} p_{it} x_{it}} \right)^{1/r}. \end{aligned} \tag{5}$$

We conclude from (5) that as one price approaches infinity, the price index has a well-behaved limit, and the extra term introduced is the last term on the right of (5). This term depends on the share of expenditure on good 1 when it becomes available in period t, as well as on the elasticity of substitution σ (or on $r = 1 - \sigma$). As the share of expenditure on good 1 is higher, or σ is lower, then the last term on the right of (5) will pull down the price index more, indicating greater cost savings due to the appearance of good 1.

That is the calculation I did one morning at the Hebrew University, and I recall going into Helpman's office and saying that it would take me five years to work out all the implications of this formula. I was only partially correct: it would take me about five years to publish the results in my 1994 article, but in fact, many of the implications of this formula are still being worked out today. One of the reasons it has taken so long for this formula to be useful is that it still requires knowledge of the elasticity of substitution for each product.[2] Needing that information puts us in the same quandary that Harris found himself in with the simulation models of Canada: he did not know what elasticity values to use, because the existing estimates tended to be too low, which would lead to exaggerated estimates of the gains from trade. To explain why the methods of the 1980s resulted in poor estimates for the elasticity, I need to digress even further.

***Second digression: Measuring the elasticity of substitution***

There is an old and well-known problem in econometrics called the "identification problem." To explain this, think about the market for any particular good and the price charged and quantity sold over a series of years. Depending on the product we pick, the data on the price and quantity can look quite different. For personal computers for example, the price has been declining and the quantity increasing over time. For housing in the United States, the price has recently been increasing along with the quantity sold, though some people believe this real-estate bubble will soon burst. The pattern of price and quantity movements depends a great deal on whether the supply curve for the product has been shifting, due to a cost reduction, for example, or the demand curve has been shifting, as with the demand for housing in some US cities, or both. The "identification problem" states that unless we know quite a bit about these shocks to supply and demand, then it is really impossible to estimate the elasticities of the supply and demand curves themselves.

---

[2] There are two terms on the right of (5), both of which require knowledge of the elasticity of substitution, or r. But the first term on the right is simply the price index $P_r(p_{t-1}, p_t, x_{t-1}, x_t, I)$ defined over the set of products I. From (2), that price index equals the ratio of CES unit-costs, defined over the set of products I in both periods. But instead we can use the price index formula due to Sato (1976) and Vartia (1976), which is also exact for a CES unit-cost function but does not require knowledge of the elasticity of substitution. In other words, the first term on the right of (5) can be replaced by the Sato-Vartia formula, as in Feenstra (1994). However, the second term always requires knowledge of the elasticity of substitution, and is quite sensitive to that value.

A wonderful discussion of the identification problem is contained in a 1981 paper by Leamer, entitled “Is it a Demand Curve or is it a Supply Curve?” Leamer poses the identification problem in the following way. Suppose that we have collected the data on prices and quantities for a particular good over time, but that we do not have any additional information on the shocks to supply or demand. Using just the price and quantity data, and assuming normally distributed errors on the supply and demand curves, we can still ask what the maximum likelihood estimates of the supply and demand elasticities are. The answer is that the maximum likelihood estimates are not unique: the estimates can be anywhere along a hyperbolic curve that Leamer describes. The fact that the estimates are not unique is just another way of saying that we cannot identify the supply and demand elasticities without further information.

Existing econometric practice during the 1980s and 1990s was to search for additional information that would help to measure the shocks to supply or demand. Such information is called “instrumental variables” or “instruments” for short. In the housing market, for example, we can use interest rates in the economy as an instrument to help explain demand. For computers, we might use speed of the latest chips available as a instrument for the cost of production, and therefore an instrument for supply.  If we have enough instrumental variables, then the identification problem can be overcome and we can estimate the supply and demand elasticities. But if our instruments are poor, then the elasticities estimates we obtain will still be mismeasured and biased towards zero.  That was the case for estimates of the elasticities of substitution for product varieties from different countries: we did not have good instruments to measure the supply shocks from each country, so the elasticities obtained were too small.

There is a fascinating paragraph in Leamer’s article where he describes a historical debate between Leontief (1929) and Frisch (1933), concerning the identification problem. Let us go back to the idea that without any instruments, our estimates of the supply and demand elasticities are not unique, and lie along a hyperbolic curve. Then Leontief made the suggestion that we split the sample in half. The first half of the sample could give us estimates of the supply and demand elasticities over one hyperbolic curve, and the second half of the sample could give us estimates along a second curve. It would seem, then, that we could take the intersection of these two curve to obtain unique estimates of the supply and demand elasticities, and thereby overcome the identification problem.

Leamer (1981, p. 321) reports that: "This procedure brought down upon [Leontief] the wrath of Frisch's 1933 book, which is devoted almost completely to debunking the method." The reason the idea does not appear to work is that, by just splitting a sample in half, there is not reason to expect the two curves obtain to be different from each other. If the first half of the sample is drawn from the same statistical population as the second half, then the maximum likelihood estimates of the supply and demand elasticities would lie along the same hyperbolic curve in both cases. If the two curves are just the same, then their intersection is still the same curve, and we have not made any headway at all!

But, Leontief may have been right after all, if we just add another feature to the data. Rather than having the price and quantity of just one good over time, suppose that we have the price and quantity for that good exported from multiple countries over time. So in addition to the time dimension of the dataset we also have a country dimension, making it a panel. We continue to assume that the elasticity of substitution between the goods from each country is constant over time, and also the same across countries. In other words, the variety supplied by one country is different from that supplied by any other country, but a German variety is just as different from a French variety as it is from an American variety. This assumption of a constant elasticity of substitution over time and across countries is a simplification, of course, but it allows us to make great progress on the identification problem.

For now we can use the price and quantity exported of the German variety to get one curve of maximum likelihood estimates of the supply and demand elasticities; and then the French data on price and quantity exported to get a second curve; and then the American data to get a third curve, and so on. The elasticity of substitution or demand is the same across countries, and we might assume the same is true for the elasticity of supply (though that assumption can be generalized). Then a point near to the intersection of these multiple curves gives us an estimate of the supply and demand elasticities. Furthermore, in contrast to the proposal of Leontief, there are very good reasons to expect that these hyperbolic curves for each country will differ from each other. It turns out the curve for each country depends on the *variances and covariances* of the supply and demand shocks, which can depend on the variance of exchange rates and other macroeconomic variables. Provided that our panel of countries includes those with differing variances and covariances of shocks, then the method I have described should result in good estimates for the supply and demand elasticities, even though we do not have instrumental variables in the conventional sense.[3]

In fact, the elasticities of substitution obtained from this method are very plausible indeed. In my 1994 article I considered products like mean’s leather athletic shoes, or cotton knit shirts, or various types of steel, and obtained estimates of the elasticity between 3 and 8. These were much higher values than obtained previously, and more in line with what trade economists would expect. I even added gold bullion and silver bullion as additional test case, and obtained estimates of the elasticity of substitution for each of these products of 25 and 40, respectively. Those high estimates are essentially infinite, indicating that there is perfect substitution between

[3] See Feenstra (1994, pp. 162-165) and the unpublished Appendix to that paper.

country sources of gold or silver. To conclude, for the six products that I analyzed in 1994, the new method for estimating the elasticity of substitution worked very well indeed. Furthermore, the method of splitting a sample to obtain identification also works in other contexts, such as modern finance, as shown by Rigobon (2003). He calls this method “identification through heteroskedasticity.” Like me, Rigobon has essentially re-discovered and justified the method proposed more than 75 years ago by Leontief.

***Recent Research***

Let me return, then, to the research being done today in international trade, by Broda and Weinstein (2006). My estimates of the elasticity over six products pales in comparison to their approach of using all 30,000 products available in the Harmonized System of trade data and its earlier version, the Tariff Schedule of the US. Modern computing power allows us to estimate all 30,000 elasticity of substitution in demand along with elasticities of supply. For the United States, they also use the trade data to measure the value of imports from new supplying countries, relative to some base year. By combining the data on imports from new supplying countries with estimates of the elasticity of substitution for each, Broda and Weinstein come away with an estimate of the gains from trade for the US due to the expansion of import varieties, which amount to 2.6 percent of GDP in 2001.

Economists sometimes get very excited over small numbers, but 2.6 percent of GDP is actually a very large number. It indicates the *ongoing annual gains* from having the import varieties available from new supplying countries. I would expect that a number of the same or greater magnitude would apply to the gains to the EU from having new import varieties available, and there could quite possibly be even larger gains from internal-EU trade. Those estimates for Europe have not yet been made, but I hope that they will.

Looking beyond the US and the EU, there are a few other studies measuring the extent to which trade depends on new product varieties. For example, Hummels and Klenow (2005) look at a cross-section of countries in a single year, 1996, and contrast the trade of larger versus smaller countries. As countries grow then so do their exports and imports, and the question is whether this growth is due to importing and exporting a more diverse range of products, what they call the "extensive margin" in trade, or due to trading more of the same products, the "intensive margin" in trade. It turns out that about 2/3 of the differences in the amount of trade between countries is explained by importing and exporting a more diverse range of products, or the extensive margin, while 1/3 of the differences in trade is due to trading more of the same products, or the intensive margin. While Hummels and Klenow do not estimate the gains from trade, their calculations suggest that the gains due to new product varieties must be very important indeed.[4] That conclusion is confirmed by Broda, Greenfield and Weinstein (2006) in another study, dealing with the contribution of new input varieties to productivity growth. They estimate that in a typical country, new import varieties account for 15 percent of productivity growth, but the effect is large in developing countries that rely more on imported materials.

**4. Selection of Firms**

Let me turn now to the third source of gains from trade in the monopolistic competition model, which is the self-selection of firms, with only the more efficient firms surviving after trade liberalization. This prediction did not come from the original models of the early 1980's, because those models used the simplifying assumption that all firms are the same, producing at

---

[4] Funke and Ruhwedel (2001a,b, 2002) have applied the CES measure of product variety to analyze economic growth across the OECD countries and the east Asian countries. They find that a country's export variety relative to the US is a significant determinant of its exports and GDP per-capita. Feenstra *et al* (1999) apply the same measure of product variety to analyze industry productivity growth in South Korea and Taiwan.

the same scale and with the same costs. The number of firms surviving in the market will depend on the level of tariffs, but there is no essential difference between those that drop out and those that remain. So in the models of the 1980s, there was no reason for the efficiency of surviving firms to differ from those that exit the market.

The most recent advance in the monopolistic competition model is to introduce heterogeneous firms that differ in their productivities. That breakthrough has been made in different ways by Eaton and Kortum (2002), and by Melitz (2003). In extending the monopolistic competition model with CES preferences to allow for heterogeneous firms, Melitz has demonstrated a whole range of new and exciting results.

For example, the opening of trade in a sector will bid up the wage and other factor prices, which forces the least efficient firms to exit the market. The more efficient firms will be able to cover the fixed costs of marketing overseas and will therefore begin exporting, while those firms in the middle-range of efficiency will continue producing for just the domestic market. So the overall distribution of firm outputs shifts in favor of those that are most efficient, since they are producing for the domestic market and exporting, while the least efficient firms have exited. As a result, average productivity in the industry rises due to trade. The same outcome occurs when there are reductions in tariffs, transport costs, or just a growth in the size of the export market: all these will improve average productivity calculated over domestic plus export sales.

These predictions have received very strong support from empirical work utilizing firm-level datasets. For the US, Bernard, Eaton, Jensen and Kortum (2003) have shown that only a small fraction of the firms in any industry are exporters. But this small fraction of firms accounts for a very large amount of total sales within the industry. The implication is that these exporting firms are substantially more productive than other firms in the industry. That finding is reinforced using firm-level data for France, by Eaton, Kortum and Kramarz (2002, 2003).

These studies from the US and France focus on the differences between firms in a single year, or over several years, but do not necessarily include a major trade liberalization. A recent study for Canada, by Trefler (2004), does just that. Trefler uses firm-level data during the decades before and after the Canada-US free trade agreement, and is interested in the impact of the agreement on the selection and productivity of firms. He obtains a number of clear results.[5] First, Canadian industries that had relied on tariffs saw their employment fall by 12 percent due to the elimination of tariffs. In manufacturing overall, the trade agreement reduced employment by 5 percent. Second, these job losses were a short-term effect, and over a 10 year period, employment in Canadian manufacturing did not drop. While low-productivity plants shut down, high-productivity Canadian manufacturers expanded into the United States. Third, the trade agreement set off a productivity boom. Formerly sheltered Canadian companies began to compete with, and compare themselves to, more efficient American businesses. Some went under, but others significantly improved operations. In the formerly sheltered industries most affected by the tariff cuts, labor productivity jumped 15 percent, at least half from closing inefficient plants. That correspond to a compound annual growth rate of 1.9 percent.

To summarize, Trefler finds overwhelming evidence that the Canada-US free trade agreement resulted in the self-selection of Canadian firms, with only the more productive firms surviving. Productivity in Canadian manufacturing overall rose 6 percent. This productivity gain translates directly into higher wages or lower prices, and is a gain from trade for consumers. Trefler's estimates of the gains are the highest we have seen from any country study, and

---

[5] These results are drawn from an interview of Trefler by Virginia Postrel, "Economic Scene," *The New York Times*, January 27, 2005, p. C2, posted on the home page for Daniel Trefler at the University of Toronto.

probably higher than would occur in the US, simply because trade is a much higher fraction of GDP for Canada. But we could expect some of the European countries to show productivity gains of the same magnitude as Canada, provided that the labor market institutions are flexible enough to allow for the entry and exit of firms on the same scale as occurred in Canada.

For developing countries, the gains from the self-selection of firms due to trade should also be substantial. I have recently completed a study of 44 countries over two decades, 1980 – 2000 (Feenstra and Kee, 2006). Over this period, export variety to the United States increases by 4.6 percent per year, so it more than doubles over these two decades. Furthermore, that increase in export variety is associated with a 4.5 percent productivity improvement for exporters over the two decades. These gains for the exporters are actually larger than the gains to the United States (of 2.6 percent) from the increase in its import variety over the past several decades.

## 5. Gains from Trade and the Penn World Table

This completes my description the gains from trade, but my discussion raises an important question: if the gains from trade are as large and widespread as I have suggested, then why don't such gains appear more prominently in national and international statistics? Setting aside the obvious answer that the gains are hard to measure, I think there is a deeper reason why the gains from trade do not appear in the statistics that we most commonly publish and cite. The reason, I believe, is that the gains are often hidden within the productivity growth of countries. It follows that the increase in GDP due to trade might be attributed to domestic sources, such as research and development or skill-biased technological change, rather than to international trade. I have already given some examples about how the variety of inputs made available through trade, or the self-selection of firms due to tariff reductions, lead to productivity improvements, and this illustrates how productivity and the gains from trade can be conflated. But this issue is much broader than I have suggested so far. To illustrate this, let us consider how the gains from trade and productivity are intertwined in the most widely used dataset by economists: the Penn World Table.

The Penn World Table (or PWT for short) is a dataset originating at the University of Pennsylvania that provides comparable figures on the real GDP of countries, converted to a common currency. When comparing the GDP of any two countries, it is not appropriate to use the nominal exchange rate to convert the GDP of one country into units of the other, since nominal exchange rates fluctuate a great deal over time due to financial reasons, which would distort that comparison. Instead, what the PWT does is collect data on the actual price of consumer goods in various countries, and then compare those prices in, say, Rupees to the price of the same item in Euros. Averaging those price ratios over many products we get a constructed Rupee-Euro exchange rate, which can be used to convert Indian GDP to Euros, thereby allowing a comparison of the Indian standard of living with that in Europe. Of course, the same calculation can be done to compare the standard of living across European countries, and *should* be done whenever the Euro prices differ across nations, as they do. In fact, the European Commission makes a calculation of this type to determine the allocation of special assistance funds to regions within Europe.

The PWT probably receives more downloads that any other dataset in economics. Its website at the University of Pennsylvania (PWT.econ.upenn.edu) is accessed by over 100,000 unique users per year. The peak months are at the end of the academic year, in March and April, with about 100 hits each day,12 to 13 thousand unique visitors, and downloads of about 100,000 pages each month. This level of activity suggests the importance of the PWT for teaching and research. Given this widespread use, it is worth asking what role international trade statistics play in the calculation of real GDP within the Penn World Table. For example, if a country has particularly strong access to international markets, thereby receiving greater gains from trade, then how would that information be incorporated into the PWT?

The short answer is that the PWT uses practically no data at all on international trade in its calculations. Other than making a simple correction for a country's overall trade balance, the PWT uses no information whatsoever on the prices of exports or imports, or the amount traded of individual goods. This is a very surprising omission. It means that a country with very strong access to international markets – such as Germany – will have a higher real GDP as a result, but that level of GDP will not be attributed to any international cause. A researcher seeing high GDP in that country would most likely attribute it to high productivity, so that the gains from trade are missed entirely.

To make this point more sharply, it is useful to distinguish two concepts of real GDP. The first can be called real GDP on the *expenditure side,* and is meant to measure the standard of living of a typical consumer in that country. The second concept can be called real GDP on the *output side,* and is meant to measure the production possibilities of a country. Even though *nominal* GDP measure from the expenditure-side and output-side are equal, there is no reason for the *real* GDP magnitudes to be the same. Rather, the difference between these real concepts arises from the opportunity for countries to trade on international markets. Countries with better trading opportunities should have real GDP on the expenditure-above that on the output-side, and conversely for countries with poor trading opportunities.

Now we can ask the key question: which concept does the PWT measure – real GDP on the expenditure-side or the output-side? The answer is that PWT tries to measure real GDP on the expenditure-side, representing the standard of living for consumers. This is well-known to

Alan Heston and Robert Summers, who originated the Penn World Table, as well as to experts in the field, but it is not widely appreciated by economists as a whole. Among growth economists, for example, it is common to use real GDP from the Penn World Table to measure the productivity differences across countries and over time. But that is not accurate, because real GDP on the expenditure-side reflects both productivity and the trading opportunities that a country faces. What appears to be high productivity could instead be gains from trade, and what appears to be low productivity could instead be a lack of access to international markets. There is simply no way to disentangle these concepts currently in the PWT.

This problem will not last for much longer, however. A team of researchers including Heston, Summers, myself and colleagues at the University of Groningen (Feenstra, *et al*, 2004) as well as in Australia, have begun work to obtain distinct measures of real GDP on the expenditure-side and on the output-side. The difference between these two concepts can be measured using international data on the prices of exports and imports that each country faces. In our preliminary work, it appears that wealthier countries in Europe and worldwide do enjoy higher prices for their exports than less-wealthy countries, giving them a boost in GDP due to the gains from trade. This preliminary conclusion raises as many questions as it answers, however. Why do the wealthier countries have higher export prices? Could it be due to their proximity to each other, reducing transport costs, or is it due to lower tariffs with their trading partners, or some other cause? Furthermore, can we be sure that the differing prices we are measuring for exports do not simply reflect differences in the quality of goods exported by countries? And if that is the case, how can we make a correction for export and import quality in the calculation of real GDP?

This list of questions sets a research agenda that will take many years to fulfill, but is, I believe, well worth the time and effort. The ability to distinguish the contribution of international trade to a country's well-being, as distinct from domestic productivity, would allow for better dialogue among economists and policy-makers, and ultimately, better economic policies.

## 6. Conclusions

To summarize my talk today, the monopolistic competition models developed during the early 1980s promised a very wide range of applications within international trade. That promise has been realized theoretically through the dynamic models of endogenous growth, and also the new models with heterogeneous firms. The empirical applications have been slower to come, but are now are being realized through the use of highly disaggregate trade data combined with new empirical techniques.

There are three sources of gains from trade in the monopolistic competition models: the gains from reduced prices due to competition between firms; the gains from the expansion in product varieties; and the gains due to the self-selection of firms as only the most efficient firms survive. Surprisingly, the first source of gains from trade – due to reduced prices – has the least empirical support to date. There is little direct evidence to support the idea that firms expand their output following trade liberalization and enjoy economies of scale, leading to lower prices. But there is indirect evidence that in environments where trade barriers are really minimized, such at the EU internal market, then prices become more similar across countries and that trade grows. This indirect evidence suggests that firms cannot price-discriminate in unified markets, which can lead to substantial gains.

The second source of gain from trade – due to increase product variety – has received a good deal of support in current research using disaggregate data. The assumption used to analyze these data is that each country provides different product varieties than every other country. New statistical methods allow use to estimate the degree of substitution between product varieties across countries, and with those elasticities of substitution, measure the consumer gains from importing more product varieties. These gains are substantial for the United States, the country that has received the most study, but I would expect them to be even larger for many other countries with a higher ratio of trade to GDP, such as in Europe. The third source of gains from trade – due to the self-selection of firms – has received overwhelming support from recent study of the Canada-US free trade agreement, as well as firm-level datasets for the US, France, and other countries.

Besides the gains from trade, there are many other features of the monopolistic competition model that can remain to be explored, both theoretically and empirically. Let me conclude by mentioning just two areas of current research.

First, the monopolistic competition model with heterogeneous firms can easily be extended to incorporate foreign direct investment. If we suppose that firms face some fixed costs of establishing a plant in a foreign market, then only the most efficient firms will want to do so. So if we line up firms by their productivities, we will find that the least productive firm exit the market; the next set of firms produce for the domestic market but do not export; the next most productive firms sell at home and also export; and the most productive firms sell at home and also abroad through establishing a foreign subsidiary. This ranking of firms is supported by the empirical work of Kimura (2006), Raff and Ryan (2006) and Raff, Ryan and Stähler (2006), for Japanese firms. In the model, we can solve for the share of firms engaged in each activity, and find that industries with a greater variance of productivities will also be those with a greater fraction of foreign investment. That prediction is found to hold for US multinational activity by Helpman, Melitz and Stephen Yeaple (2004).

So the models with heterogeneous firms help us to understand patterns of foreign direct investment. In principle, we should also be able to use the new models to measure the gains from foreign investment, though that has not yet been attempted. The magnitude of goods sold by foreign-owned firms is just as large as the magnitude of international trade flows, so I would conjecture that the gains from foreign investment are just as large as those from international trade in goods. I would hope that future research would allow us to measure these gains.

A second area of very recent research is to introduce multi-product firms into the monopolistic competition model. That is important because in reality, the large trading firms are also multi-product firms. Bernard, Jensen and Schott (2005) have coined the term "most globally engaged" firm, or MGE, to describe these large firms that both import and export as well operate in several countries. These multinationals account for less than one percent of US firms, but, but a remarkable 80 percent of trade and 18 percent of total US employment. The presence of these firms suggests even greater potential for gains from trade, as they exploit economies of scale and economies of scope. On the other hand, the MGE might also be able to strategically raise prices across multiple markets, suggesting that the low prices due to trade liberalization would not be realized. It is important, therefore, that we understand the potential for these multi-product firms to enhance efficiency through economies of scale and scope, but also exercise their pricing decisions over multiple markets simultaneously.

Balancing the beneficial economies against the potentially harmful market power brings us back to the reason that the monopolistic competition model was introduced into international economics 25 years ago, and that was to bring the *trader* back into *international trade.* The perfectly competitive paradigm did not allow us to distinguish the trading firm from the industry as a whole, whereas the monopolistic competition model has a distinct role for the *trading firm*. We have made a good deal of progress in understanding how the actions of trading firms contribute to the gains from trade, but there is still much more to be done. I look forward to continued progress on this research agenda during the next 25 years.

**References**

Allen, C., M. Gasiorek and A. Smith (1998) "The Competition Effects of the Single Market in Europe", *Economic Policy*, 13(27), October, 441 -486.

Badinger, H. (2006) "Has the EU's Single Market Programme fostered competition? Testing for a decrease in markup ratios in EU industries," *Oxford Bulletin of Economics and Statistics*, forthcoming.

Baldwin, R. (2006a) "The Euro's Trade Effects," European Central Bank, Working Paper no. 594, March.

Baldwin, R. (2006b) *In or out: does it make a difference? An evidence based analysis of the trade effects of the Euro,* Center for Economic Policy Research.

Bernard, Andrew B., Jonathan Eaton, J. Bradford Jensen and Samuel Kortum (2003) "Plants and Productivity in International Trade," *American Economic Review*, 93, September, 1268-1290.

Bernard, A. B., J. B. Jensen and P. K. Schott (2005) "Importers, Exporters, and Multinationals: A Portrait of Firms in the U.S. that Trade Goods" NBER working paper no. 11404.

Broda, C. and D. E. Weinstein (2006) "Globalization and the Gains from Variety," *Quarterly Journal of Economics,* May, 121(2), 541-585.

Broda, C., D. Greenfield and D. E. Weinstein (2006) "From Groundnuts to Globalization: A Structural Estimate of Trade and Growth," University of Chicago and Columbia University.

Diewert, W. E. (1976) "Exact and Superlative Index Numbers," *Journal of Econometrics,* 4, 115-145.

Eaton, Jonathan and Samuel Kortum (2002) "Technology, Geography and Trade," *Econometrica,* 70(5), September, 1741-1780.

Eaton, J., S. Kortum and F. Kramarz (2003) "An Anatomy of International Trade: Evidence from French Firms," unpublished, CREST, University of New York and University of Minnesota.

Eaton, J., S. Kortum and F. Kramarz (2004) "Dissecting Trade: Firms, Industries, and Export Destinations," *American Economic Review*, 94(2), 150-154.

Feenstra, R. C. (1994) "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84(1), March, 157-177.

Feenstra, R. C., A. Heston, M. P. Timmer, and H. Deng (2004) "Estimating Real Production and Expenditures Across Countries: A Proposal for Improving the Penn World Tables," NBER Working Paper no. 10866.

Feenstra, R. C. and H. L. Kee (2006) "Export Variety and Country Productivity: Estimating the Monopolistic Competition Model with Endogenous Productivity," University of California, Davis and World Bank Policy Research Group.

Feenstra, R. C. and J. R. Markusen (1994) "Accounting for Growth with New Intermediate Inputs," *International Economic Review*, 35(2), May 1994, 429-447.

Feenstra, R. C., D. Madani, T. H. Yang and C. Y. Liang (1999) "Testing Endogenous Growth in South Korea and Taiwan," *Journal of Development Economics*, 60, 317-341.

Feenstra, R. C., Alan Heston, Marcel P. Timmer, and Haiyan Deng (2004) "Estimating Real Production and Expenditures Across Countries: A Proposal for Improving the Penn World Tables," NBER Working Paper no. 10866.

Feenstra, R. C. and H. L. Kee (2006) "Export Variety and Country Productivity: Estimating the Monopolistic Competition Model with Endogenous Productivity," World Bank Policy Research Group and University of California, Davis.

Friberg, R. (2001) "Two monies, two markets? Variability and the option to segment", *Journal of International Economics,* 55, 317-327.

Frisch, R. (1933) *Pitfalls in the Statistical Construction of Demand and Supply Curves,* Leipzig: Hans Buske, Verlag.

Funke, M. and R. Ruhwedel (2001a) "Product Variety and Economic Growth: Empirical Evidence from the OECD Countries," *IMF Staff Papers,* 48(2), 225-242.

Funke, M. and R. Ruhwedel (2001b) "Export Variety and Export Performance: Evidence from East Asia," *Journal of Asian Economics,* 12, 493-505.

Funke, M. and R. Ruhwedel (2002) "Export Variety and Export Performance: Empirical Evidence for the OECD Countries," *Weltwirtschaftliche Archiv,*138(1), 97-114.

Harris, R. (1984a) "Applied General Equilibrium Analysis of Small Open Economies with Scale Economies and Imperfect Competition," *American Economic Review,* 74(5), December, 1016-32.

Harris, R. (1984b) *Trade, Industrial Policy, and Canadian Manufacturing.* Toronto: Ontario Economic Council.

Hausman, J. A. (1997) "Valuation of New Goods under Perfect and Imperfect Competition," in Timothy F. Bresnahan and Robert J. Gordon, eds., *The Economics of New Goods,* NBER Studies in Income and Wealth, no. 58, Chicago: University of Chicago, 209-237.

Hausman, J. A. (1999) "Cellular Telephone, New Products, and the CPI," *Journal of Business and Economic Statistics*, 17(A2), April, 188-94.

Head, K. C. and J. Ries (1999) "Rationalization effects of tariff reductions," *Journal of International Economics*, 47(2), April, 295-320.

Head, K. C. and J. Ries (2001) "Increasing Returns Versus National Product Differentiation as an Explanation for the Pattern of US-Canada Trade," *American Economic Review*, 91(4), September, 858-876.

Helpman, E. (1981) "International Trade in the Presence of Product Differentiation, Economics of Scale and Monopolistic Competition: A Chamberlin-Heckscher-Ohlin Approach," *Journal of International Economics,* 11, 305-340.

Helpman, E., M. Melitz and S. Yeaple (2004) "Exports vs. FDI with Heterogeneous Firms," *American Economic Review*, March, 300-316.

Hicks, J. R. (1940) "The Valuation of Social Income," *Economica,* 7, 105-24.

Hobijn, B., F. Ravenna and A. Tombalotti (2006) "Menu Costs at Work: Restaurant Prices and the Introduction of the Euro," *Quarterly Journal of Economics,* 121(3), August, 1103-1131.

Hummels, D. and P. Klenow (2005) "The Variety and Quality of a Nation's Trade," *American Economic Review,* 95(3), June, 704-723.

Kimura, F. (2006) "Exports, FDI, and Productivity: Dynamic Evidence from Japanese Firms," *Review of World Economics /Weltwirtschaftliches Archiv*, 142(4), forthcoming.

Krugman, P. R. (1979) "Increasing Returns, Monopolistic Competition and International Trade," *Journal of International Economics,* 9, 469-479.

Krugman, P. R. (1980) "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review,* 70, December, 950-959.

Krugman, P. R. (1981) "Intra-Industry Specialization and the Gains from Trade," *Journal of Political Economy,* 89, 959-973.

Lancaster, K. (1980) "Intra-Industry Trade under Perfect Monopolistic Competition," *Journal of International Economics,* 10, May, 151-175.

Leamer, E. E. (1981) "Is it a Demand Curve, or is it a Supply Curve? Partial Identification Through Inequality Constraints," *Review of Economics and Statistics,* 63(3), August, 319-327.

Leontief, W. (1929) "Ein Versuch Zur Statistichen Analyse von Angebot und Nachfrage," *Weltwirtshaftliches Archiv,* 30(1), 1-53.

Melitz, M. J. (2003) "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica,* 71(6),1695-1725.

Raff, H., and M. J. Ryan (2006) "Firm-Specific Characteristics and the Timing of Foreign Direct Investment Projects," University of Kiel and Western Michigan University.

Raff, H., M. J. Ryan, and F. Stähler (2006) "Firm Productivity and the Foreign Market Entry Decision," University of Kiel, Western Michigan University and University of Otago.

Rigobon, R. (2003) "Identification Through Heteroskedasticity," July, *Review of Economics and Statistics*, 85 (4), 777-792.

Rose, A. K. (2000) "One Money, One Market: Estimating the Effect of Common Currencies on Trade", *Economic Policy* 30, 9-45.

Sato, K. (1976) "The Ideal Log-Change Index Number," *Review of Economics and Statistics* 58, May, 223-228.

Smith, A. and A. J.Venables (1988) "Completing the Internal Market in the European Community: Some Industry Simulations," *European Economic Review,* 32, 1501-1525.

Smith, A. and A. J. Venables (1991) "Economic Integration and Market Access," *European Economic Review, Papers and Proceedings,* 35, 388-395.

Trefler, D. (2004) "The Long and Short of the Canada-U.S. Free Trade Agreement," *American Economic Review,* 94(4), September, 870-895.

Tybout, J. R., J. de Melo, and V. Corbo (1991) "The effects of Trade Reform on Scale and Technical Efficiency," New Evidence from Chile," *Journal of International Economics,* 31(3/4), 231-250.

Tybout, J. R. and M. D. Westbrook (1995) "Trade Liberalization and Dimensions of Efficiency Change in Mexican Manufacturing Industries," *Journal of International Economics,* 39(1/2), 53-78.

Vartia, Y.O. (1976) "Ideal Log-Change Index Numbers", *Scandinavian Journal of Statistics* 3, 121-126.