Plant-Level Responses to Antidumping Duties:

Evidence from U.S. Manufacturers[1]

Justin Pierce[2]
Georgetown University

October, 2008
Preliminary Version

Abstract: This paper describes the effects of a temporary increase in tariffs on the performance and behavior of U.S. manufacturing establishments (plants). Using antidumping duties as an example of temporary protection, I compare the responses of protected manufacturers to those predicted by new models of trade with heterogeneous firms. I find that apparent increases in revenue-based productivity associated with temporary protection are actually increases in prices and mark-ups in disguise. Antidumping duties lower productivity among the set of protected plants reporting output data in units of quantity. Moreover, antidumping duties allow for the continued operation of low-productivity plants that likely would have otherwise ceased production. As a result, temporary protection slows the process of output rationalization, with less productive plants increasing their share of total output. Importantly, plants producing products that were turned down for protection are no less likely to exit than plants receiving protection. Rather, they adjust by dropping the unprotected product and producing other, potentially higher-productivity products.

JEL Classification: F13, F14, D24

Keywords: Heterogeneous Firms; Antidumping; Temporary Protection

---

[1] I thank Rod Ludema, Anna Maria Mayda and Brad Jensen for continuous support and many useful suggestions. Any remaining errors are my own. The research in this paper was conducted while the author was a Special Sworn Status researcher at the U.S. Census Bureau, Center for Economic Studies. Results and conclusions presented are those of the author and do not reflect the views of the Census Bureau. This paper has been screened to ensure that no confidential data are revealed.

[2] Department of Economics, Georgetown University, 3700 O ST NW, Washington, DC 20057; jrp45@georgetown.edu.

Section 1: Introduction

How do manufacturers respond to the imposition of antidumping duties? Answering this question can provide important new insights into the heterogeneous responses of plants and firms to changes in tariff protection. The novelty of these insights comes in part from the unique characteristics of antidumping duties—they are temporary and they involve tariff increases, rather than decreases. While the effect of permanent tariff reductions on plant and firm-level behavior has been thoroughly studied, it seems natural to question whether tariff increases have the opposite effect of tariff reductions, and whether the temporary nature of antidumping duties matters.

Examining instances of temporary tariff increases allows for closer examination of some unresolved questions in the heterogeneous firm literature, including the effect of changes in tariff rates on plant-level productivity. Moreover, it can provide a new perspective on some of the best-known results from the heterogeneous firm literature. It is well-understood, for example, that tariff reductions speed the exit of the least-productive firms and plants. But do temporary tariff increases lead existing firms to postpone exiting, or induce additional firms to enter? Furthermore, does protection lead to product-switching, with current producers delaying the dropping of unproductive products, or non-producers adding that product? If temporary tariff increases do affect these aspects of plant behavior, what are the implications for output rationalization and aggregate productivity?

I examine these questions using the imposition of antidumping duties as an example of temporary protection. The empirical analysis largely follows a difference in difference framework, in which a treatment group of plants producing products that receive antidumping protection is compared to a control group of plants producing products that applied for, but did not receive protection. This framework nets out macro-level shocks affecting all manufacturers, while providing a useful comparison between plants receiving antidumping protection and highly similar control groups.

There are several benefits to using antidumping duties to study the effect of tariff increases on firms and plants. First, instances of tariff increases are rare, especially in the United States. Second, antidumping duties are particularly useful for examining unresolved questions in the theoretical literature regarding the effect of temporary protection on productivity growth. And third, the increased use of antidumping duties around the world—even as nominal tariff rates have fallen—has made antidumping duties one of the primary forms of trade protection, and hence worthy of study. Providing the first description of the reactions of manufacturers to antidumping duties using U.S. micro-data is another important contribution of this paper.

Using output data measured in units of quantity, I find that apparent increases in revenue-based productivity are actually increases in prices and mark-ups in disguise. Antidumping duties actually lower productivity among the set of plants reporting output data in units of quantity. Moreover, antidumping duties allow for continued production by plants that would have otherwise ceased production. As a result, the level of output rationalization—the process by which high-productivity plants expand their market shares at the expense of low-productivity plants—falls among producers of protected products until it is lower than producers of unprotected products. The combination of these effects indicates that antidumping duties decrease productivity at both the plant and aggregate level. Lastly, and importantly, I find that plants in the control group respond to being turned down for protection not by exiting, but rather by switching products—dropping the unprotected product for other, potentially higher-productivity products.

The paper is organized as follows. Section 2 describes the questions that will be examined and places them within the context of the theoretical and empirical heterogeneous firm literatures. Section 3 describes the data. Section 4 provides a brief discussion of the antidumping investigation process in the United States, as well as a description of the products typically involved in antidumping investigations. Section 5 describes the empirical strategy used to answer each question and reports the results. Section 6 concludes.

Section 2: Research Questions and Background

Do Temporary Tariffs Increase or Decrease Plant-Level Productivity?

Despite extensive research attention, the question of whether temporary tariffs increase or decrease plant-level productivity remains unresolved in the theoretical and empirical heterogeneous firm literatures. On one hand, the best-known theoretical models predict that any increase in tariffs should decrease mean plant-level productivity. In Melitz (2003), an increase in tariffs—or a failure to decrease tariffs—allows for the continued operation of low-productivity plants that would have otherwise exited, resulting in a decrease in mean firm-level productivity. In addition, Goh (2000) and Ethier (1982), describe channels for within-plant productivity growth during trade liberalization, arising from decreased opportunity costs for new investment and increased input variety, respectively.

These theoretical results are supported by a large empirical literature, which has shown that nominal tariff rates and productivity are negatively correlated. Pavcnik (2002) and Fernandes (2007) (for developing countries) and Bernard, Jensen and Schott (2006) (for the U.S.) show that manufacturers exhibited productivity gains as a result of reductions in tariff rates. In addition, Amiti and Konings (2007) and Pierce (2008) show that increases in input variety spurred by trade liberalization increased the productivity of manufacturers in Indonesia and Colombia, respectively. The extent of evidence showing productivity gains due to tariff reductions suggests that tariff increases would lead to productivity losses.

In contrast, there is theoretical and empirical evidence that tariff protection—particularly temporary protection—can increase the incentive to invest in new technology and increase plant-level productivity. Matsuyama (1990) was among the first to show that temporary protection can speed up the time of technology adoption, while noting that the government's threat to remove protection if the domestic firm fails to invest is not credible. Similarly, Miyagiwa and Ohno (1995, 1999) showed that temporary protection can induce investment in a fixed cost technology by increasing the market share of domestic firms. These theoretical models are supported by empirical results in Konings and Vandenbussche (2008) showing that revenue-based productivity increased among E.U. manufacturers receiving temporary antidumping protection.

This paper contributes to this literature in a number of important ways. Most importantly, it shows that while antidumping duties appear to increase revenue-based productivity, they are actually associated with a decline in physical productivity. Using a census of U.S. manufacturers that reports output in terms of quantity, as well as revenue, I find that the gains in revenue-based productivity were driven by increases in price and price-cost mark-ups, rather than true productivity growth. This use of quantity-based output data is particularly important when considering the case of antidumping duties, since increases in prices and markups would likely be taking place at the same time as changes in true productivity growth. My results are in line with recent findings by Foster, Haltiwanger and Syverson (2008) who underscore the critical differences between revenue and physical productivity.

In addition, this paper provides the first micro-level evidence of the effects of antidumping duties in the United States. These results are likely of interest due to the importance of the U.S. in international trade, as well as its history as a frequent user of antidumping protection. Moreover, because successful antidumping investigations in the United States almost always result in ad-valorem tariffs—rather than price undertakings or suspension agreements—I am able to study the effect of variation in the antidumping duty rate on productivity, in addition to a binary measure of protection.

Do Temporary Tariffs Discourage Product-Dropping and Encourage Product-Adding?

Recent research by Bernard, Redding and Schott (BRS) (2006, 2008) identifies a previously unexplored extensive margin for productivity growth and output reallocation, through product-switching by multi-product firms. BRS (2006) provide a model of firms with productivity composed of firm-level and firm-product-level components. Trade liberalization yields productivity growth by inducing firms to drop marginally productive products and by forcing the least productive firms to exit. Trade liberalization, therefore, leads to a reduction in the total number of products produced by each firm, but with an increase in output of products that are still produced.

BRS (2006) provides a useful framework for examining how multi-product firms react to changes in trade policy. There are, however, important differences between the framework in BRS (2006) and the temporary antidumping protection examined in this paper. First, BRS (2006) is based explicitly on a multilateral trade liberalization occurring as two countries move from a closed economy to an open-economy equilibrium. In antidumping duty proceedings, changes in trade policy are unilateral and are targeted against imports from a particular country. Second, BRS (2006) focuses on trade liberalization for all products. Antidumping duty investigations, on the other hand, involve a single product or a set of closely related products.

By examining the effect of antidumping investigations on plants' decisions of whether to add or drop products, I am able to provide some empirical evidence in this new literature. Importantly, I find that product-switching is the primary method used by plants to adjust their production as a result of antidumping investigations. Specifically, I find that plants are less likely to drop protected products than products that applied for, but did not receive protection. While this result is perhaps not surprising, its importance is magnified when combined with the results regarding plant exit discussed below: plants that apply for but do not receive protection are no less likely to exit than those that apply for and do receive protection. In other words, failure to receive protection does not appear to increase plant-level exit. Rather, plants that do not receive protection respond by switching production to other, potentially more productive products. Finally, I find no evidence that plants that did not produce a protected product, begin production once temporary protection is imposed.

Do Temporary Tariffs Discourage Exit and Encourage Entry?

There is widespread agreement that trade liberalization leads to the exit of the least productive plants and firms. In the theoretical literature, exit of low-productivity firms during trade liberalization is a key result of Melitz (2003), Bernard, Eaton, Jensen and Kortum (2003) and BRS (2006). These theoretical results are also supported by robust empirical evidence. Aw, Chung and Roberts (2000), Pavcnik (2002) and BRS (2006) have all found that decreases in trade costs bring about exit of low productivity firms or plants in Korea and Taiwan, Chile and the U.S., respectively.

While it seems plausible that a temporary increase in tariffs would decrease the amount of exit that might otherwise occur, this question has never been examined in the empirical literature. It is an important question due to its implications for aggregate productivity. If the exit of low-productivity plants is an important contributor to aggregate productivity growth, the slowing of that growth could be a major cost of temporary protection. The imposition of antidumping duties provides a useful setting for examining the effect of temporary protection on exit, since observation of the control group lets us see the extent of exiting that would have occurred without protection. Conversely, it would be interesting to know whether there is increased entry into protected product-groups, or whether they view the temporary protection from antidumping duties as being too short-lived to make protection worthwhile.

I am able to examine these issues by comparing the exit and entry decisions of plants in the treatment and control groups described above. As alluded-to above, I find that plants that apply for, but do not receive protection are no less likely to exit than those that receive protection. Rather, plants react to a failed request for protection by dropping unprotected products. This is a surprising result, given that most of the theoretical and empirical literatures have emphasized the effect of trade liberalization on exit, rather than product-switching. These results suggest that plants adjust quickly to changes in trade policy by altering their product mix, in ways that prevent them from being forced to exit. I find no evidence of increased entry into protected product-groups, perhaps because potential entrants recognize the temporary nature of antidumping protection.

Do Temporary Tariffs Decrease Output Rationalization and Aggregate Productivity?

The past ten years have yielded an immense body of research examining the effect of trade policy on output rationalization among heterogeneous firms. Melitz (2003) is, of course, well-known for showing that trade liberalization leads to expansion of output, and possibly exporting at the most productive firms with output contraction and potentially exit for low-productivity firms. Similarly, Bernard, Redding and Schott (2006) show that trade liberalization leads to reallocation of output toward comparative advantage/high-productivity industries in a model with heterogeneous firm productivity, heterogeneous industry factor intensity and countries with varying factor abundance.

The empirical evidence on the effect of trade on the allocation of output within industries both predates and builds upon these theoretical frameworks. Pavcnik (2002) and Fernandes (2007) each show that trade liberalization was associated with output rationalization in Chile and Colombia, respectively. This output rationalization was a major contributor to overall productivity gains in both countries, and accounted for a majority of Chile's aggregate productivity growth in Pavcnik (2002). In addition, Bernard and Jensen (1999, 1999b) and Bernard, Jensen and Schott (2006), have shown that the reallocation of output toward high-productivity exporters and exit by low-productivity firms have been a significant contributor to U.S. aggregate productivity growth.

The theoretical and empirical literature described above suggests that output rationalization would be lower among protected plants than unprotected plants due to effects on two margins. First, on an intensive margin, there would be less reallocation of output from existing low-productivity plants to existing high-productivity plants, than there would be without protection. Second, along an extensive margin, temporary protection would likely delay exit and encourage entry by low-productivity plants. A similar effect may also be seen along an additional extensive margin if temporary protection causes plants to delay dropping or encourages adding of low-productivity products.

The results above, which show that antidumping duties allow for continued production by low-productivity plants suggest that temporary protection likely slows the process of output rationalization. I examine this question by decomposing aggregate productivity into mean plant-level productivity and a term that measures the degree to which higher-productivity plants produce a larger share of output, as in Olley and Pakes (1996). The decomposition shows that while the degree of output rationalization is significantly higher in the treatment group prior to protection, the control group steadily increases its level of output rationalization as the period of protection sets in. By the end of the period of analysis, the control group of unprotected plants has overtaken the treatment group, to exhibit a higher level of output rationalization.

Section 3: The Data

This analysis uses plant-level and plant-product-level[3] data from the U.S. Census Bureau's (Census) Longitudinal Research Database (LRD) for the years 1987 to 1997. Total factor productivity is calculated using data from the Census of Manufactures (CMF). The CMF contains plant-level data on output (shipments and sometimes quantity), as well as input data including the number of production and non-production employees, raw material usage, investment, depreciation and book value of capital. The CMF is conducted every five years, in years ending in two and seven (e.g. 1987, 1992, 1997) and all U.S. manufacturers, regardless of size, are required by law to respond.[4]

A very important benefit of the CMF is its collection of plant-level output data measured in units of physical quantity. The availability of output data measured in units of quantity allows for the calculation of physical productivity—in addition to the standard revenue productivity—as

---

[3] Plant-product-level data refers to output shipment for every product produced at every plant. These shipment data are measured in revenue for all plants and in units of physical quantity for a subset of plants.

[4] The CMF collects a limited set of data from small manufacturers, referred to in the data as "administrative records." Since input usage data may be imputed for administrative records, they have been excluded from the analysis. This exclusion of administrative records is standard in research employing the LRD. See, e.g. Bernard, Redding and Schott (2008).

well as well as average unit prices and price-cost mark-ups. The ability to examine physical productivity, prices and mark-ups is extremely important when studying antidumping duties, since changes in physical productivity are likely accompanied by increases in prices and mark-ups. The quantity-based output data are the same as those employed in Foster, Haltiwanger and Syverson (2008).

It is important to define a number of terms that will be used throughout this paper. The term plant refers to a manufacturing establishment, which is a production facility located at a single physical location. Products and industries are 5-digit and 4-digit categories of the Standard Industrial Classification (SIC), respectively.[5] A product group is the set of plants producing a particular product.

The use of plant-level data is an important innovation of this paper and provides many advantages over more aggregated data, even including firm-level data. Many firms involved in petitioning for antidumping protection are large multi-product manufacturers. In fact, some firms participated as petitioners in multiple antidumping investigations involving multiple products. Research using firm-level data, therefore, is unable to accurately assign these firms to either the treatment group that received protection or the control group that did not. Individual plants on the other hand, tend to produce a much narrower set of products than firms as a whole. The use of plant-level data, therefore allows for much more accurate matching between the products named in contingent protection investigations and the facilities that actually produce those products.

[5] The 1987 SIC contains 459 four-digit industries and 1,848 five-digit products.

The benefits of plant-level matching can be seen clearly by examining the experience of a specific firm, the integrated steel-maker United States Steel. According to its 2006 annual report, U.S. Steel operated 24 plants in the United States, producing multiple products including flat-rolled sheets, tin mill, strip mill plate, galvanized sheets and tubular products.[6] Moreover, several of U.S. Steel's products have been subject to antidumping or countervailing duty protection over the years including Corrosion-Resistant Carbon Steel Flat Products (1993), Cut to Length Plate (1979, 2003), Seamless Pipe (1995), Oil-Country Tubular Goods (1995), Hot-Rolled Steel Products (2001) and Welded Large-Diameter Line Pipe (2001, 2002).[7] In the case of U.S. Steel, firm-level data are not sufficient for defining when or in what way U.S. Steel should be included in the treatment group. With plant-level data, however, I am able to assign protection to the plants producing the specific products covered by antidumping duties applied in specific years.

In addition, I am able to greatly refine the identification of plants that did and did not receive contingent protection through the use of plant-product-level data contained in the LRD. These data report—for each plant—the full list of products manufactured at each plant, as well as the value—and sometimes quantity—of shipments attributable to each product. The availability of this plant-product level-data represents a significant improvement over the "major industry" codes generally used to identify plants and firms in micro-level datasets.

---

[6] United States Steel at 32. http://www.uss.com/corp/investors/annual_reports/2006-annual-report.pdf

[7] United States Steel at 15-16. http://www.uss.com/corp/investors/annual_reports/2006-annual-report.pdf

The list of products involved in antidumping investigations in the United States is from version 3.0 of Chad Bown's Global Antidumping Database. Products subject to antidumping investigations are identified using the Harmonized Tariff System (HTS) and products may be defined from the 4-digit level to the 10-digit level.[8] In addition to a description of the products involved in each investigation, the antidumping database provides the dates and outcomes of each phase of the investigation—e.g. preliminary and final injury and dumping determinations—along with the final remedy. These remedies can include ad-valorem antidumping duties or suspension agreements. This analysis in this paper considers the effects of antidumping investigations that were completed during the period from 1988 to 1996. This setup ensures that I am able to observe plant-level outcomes before and after protection for every product group.

Because products in antidumping investigations are classified under the HTS, while products in the LRD are classified under the SIC, it is necessary to concord the two product classification systems.[9] The matching of HTS codes to SIC codes takes place through a set of SIC Base Codes (SICBase) developed by Census. SICBase codes are a bridge that connects the HTS—where products are defined solely based on their physical characteristics—to the SIC, where products are also classified based on their method of production. For this reason, Census assigns a single SICBase to each HTS10. This SICBase may contain a single SIC5 if the HTS10 is a subset of a single SIC5, or multiple SIC5s if the HTS10 fits several SIC5 categories. Using a three step process, I am then able to assign to determine which plants produce products that were involved in antidumping investigations:

---

[8] Although the HTS was not effective until 1989, investigations in Bown's Global Antidumping Database that ended in 1988 were assigned HTS numbers, ex-post.

[9] The procedure for concording HTS and SIC codes is described in detail in Schott (2008).

- Step 1: SICBase codes are assigned to the HTS10 codes contained in the AD dataset (referred to here as BOWN_AD[10] for brevity) using an HTS10-SICBase concordance (HTS_SICBase) published by the Census Bureau.
- Step 2: SICBase codes are assigned to each SIC5 in the plant-product-level data in the LRD using an SIC5-SICBase concordance known as the Principle Differences file (PD).[11]
- Step 3: The BOWN_AD antidumping dataset is merged to the LRD using the SICBase codes.

The process is described visually in Figure 1 below.

Figure 1: HTS-SIC Concordance Process

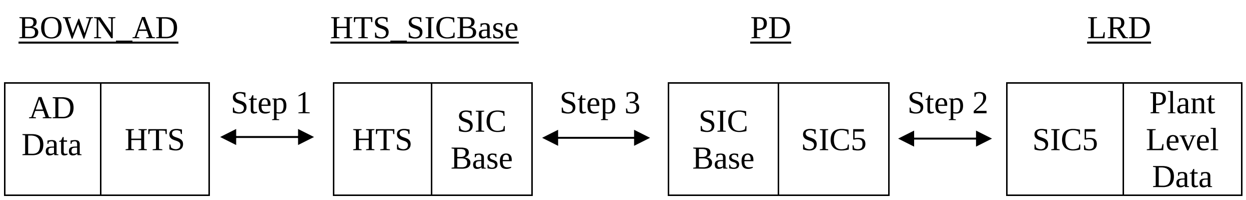

Section 4: Antidumping Duties in the United States

Under GATT Article VI and the WTO's Antidumping Agreement, WTO members are permitted to impose discriminatory tariffs on goods sold by foreign producers at prices that are deemed to be less than fair value (LTFV), if these sales result in material injury to the domestic industry. In the United States, sales are considered to be made at LTFV—i.e. dumped—when a

[10] Most products in the BOWN_AD antidumping file are at the 10-digit HTS level. There are, however, a number of products that are identified at more disaggregated levels, including the 8 and 6-digit HTS levels. In instances where a product is identified as treated (control) at a level that is more aggregated than HTS10, I treat all 10-digit HTS within that category as being treated (control).

[11] The 1992 principle differences file, which is used for the analysis in this paper can be found online at http://www.census.gov/epcd/www/intronet.html.

foreign firm sells a good in the United States at a price that is below that offered on comparable sales in its home market, or below its average total cost (ATC).[12]

Antidumping investigations in the United States are initiated by individual firms, trade associations or sometimes labor unions, which are referred to in AD proceedings as petitioners. The foreign firms selling allegedly dumped merchandise are referred to as respondents. Petitioners apply for AD protection by submitting a petition to the Import Administration of the Department of Commerce (DOC) and the International Trade Commission (ITC). The DOC determines whether sales made by foreign firms in the U.S. are being made at LTFV. The ITC determines whether the U.S. industry has been injured as a result of the dumping.

If the DOC finds that sales have been made at LTFV and the ITC concludes that these sales have injured U.S. producers, an ad-valorem tariff is placed on imports of goods from the respondents' home countries. This ad-valorem tariff, which is known as an antidumping duty is equal to the percentage difference between the U.S. price and the home-market price or ATC. I refer to the magnitude of the antidumping duty as the antidumping rate. Because the antidumping duty is applied to all dumped goods, it benefits the petitioners, as well-as non-participating producers of the investigated product.

[12] There are additional subtleties to the LTFV determination. For market economies, the preferred price comparison is between sales by the foreign producer in the U.S. and its home market. If there are insufficient sales in the foreign producer's home market, U.S. prices are compared to sales in a third country. If there are insufficient sales in the third country, U.S. prices are compared to the "constructed value (CV)" of the foreign producer's merchandise, which is gathered from the firm's cost accounting system and is essentially ATC. Sales made by firms in non-market economies are always compared to the "normal value (NV)" the firm's merchandise, which is again essentially the firm's ATC.

Table 1 reports the types of products involved in antidumping investigations from 1988 to 1996, showing the number of antidumping duty investigations by 2-digit HTS Chapter. The most frequent seekers of antidumping duties were producers of iron and steel (Chapter 72) and articles of iron and steel (Chapter 73). Other active applicants for antidumping protection included producers of machinery and appliances (Chapters 84 and 85), inorganic and organic chemicals (Chapters 28 and 29) and transportation vehicles and parts (Chapter 87). As these examples indicate, antidumping duties are primarily used to protect relatively homogenous manufactured goods.

Table 2 shows the number of antidumping investigations completed, by outcome for the years 1980 to 2005. The number of antidumping investigations tends to increase during and immediately following periods of recession, and we see that this phenomenon did, in fact, occur following the recession of 1990-1991, when the number of new investigations spiked in 1991 and 1992. Aside from this countercyclical trend in new investigations, the period from 1988 to 1996 was typical in terms of the number of investigations initiated.

Section 5: Empirical Strategy and Results

Pre-Estimation Definitions

A. Definition of Treatment and Control Groups

To borrow terms from the program analysis literature, I conduct this analysis by comparing the behavior of plants in a treatment group receiving antidumping protection to plants in a control group that do not. The treatment group consists of plants producing products that applied for and received antidumping protection. Each plant in the treatment group is assigned a date of treatment and an ad-valorem duty rate,[13] which comes from the results of the antidumping investigation associated with the product it produces. If a plant produces more than one product that receives protection, the treatment date and duty come are those associated with the treated product that accounts for the highest share of its output.

Comparing the behavior of these treated plants to a control group—rather than simply examining changes in treated plants over time—allows me to net out changes in plant-level variables over time that are independent of the treatment. Using the difference in difference framework described in more detail below I am able to net out macro-level shocks that affect all manufacturers, while identifying the effect of antidumping duties.

In the framework being examined in this paper, a natural concern is that any estimated treatment effects could be affected by a selection bias, because the set of plants that apply for antidumping protection are almost certainly different from those that do not. For example, antidumping applicants produce goods that are subject to import competition, perceive themselves as being injured by imports and operate in industries capable of cooperating to file a case. To minimize this selection bias, the plants in the control group should be as similar as possible to those in the treatment group, differing only in the sense that they do not receive antidumping protection.

I use an approach employed in Konings and Vandenbussche (2008) to define an appropriate control group. That is, I define the control group to be plants producing products that applied for antidumping duties, but were denied protection by the government. Defining the control group in this way should greatly reduce the selection bias associated with plants

[13] Of the 160 antidumping investigations initiated between 1988 and 1996, 5 ended with suspension agreements. For these cases, no ad-valorem antidumping duty rate was available.

receiving antidumping protection. In particular, it will all but eliminate a self-selection bias arising from the fact that firms and plants that apply for antidumping protection are different from those that do not. As with treated plants, control plants produce products characterized by high import competition, perceive themselves as injured by imports and are able to organize the industry to file an AD petition. Moreover, as shown in Table 3, control plants are concentrated in the same sectors that successfully apply for protection—especially primary and fabricated metals, and industrial and electronic equipment.[14]

The treatment and control groups are highly similar, with the exception of their antidumping treatment status. In addition to operating in related major industry categories, they have similar plant-level characteristics. As described in Table 4, plants in the treatment and control groups are comparable in terms of their total value of shipments, number of employees and capital to labor ratios. Importantly, they also display nearly identical mean levels of total factor productivity and labor productivity in the pre-treatment year of 1987.

B. Productivity Measures

I calculate productivity in two ways. The first is a simple, single-factor labor productivity, defined as the real total value of sales (RTVS) per employee:

$$(1)\ LP_{pt} = \frac{RTVS_{pt}}{TE_{pt}}$$

[14] Observations where the treatment and control groups overlap have been dropped from the analysis. The overlapping of treatment and control groups could overlap for two reasons. First, a single SIC5 product could receive protection from one antidumping investigation but be denied protection in another. This is possible if the HTS10 products defined in two different antidumping investigations both map into the same SIC5. In addition, a single plant could produce two products, where one product receives protection and the other is turned down for protection.

where TE_{pt} is the total number of employees at plant p at time t. The second is the superlative TFP index from Caves et al. (1982). As described in Aw, Chen and Roberts (2001), this TFP expression measures the performance of each plant, relative to a hypothetical plant producing the mean level of output with the mean level of inputs, within an industry, in the base period, 1987. The TFP index therefore incorporates a plant's deviation of output and inputs from the industry mean in any given year, but also from the mean in the base period. This calculation yields a TFP measure that is comparable across plants and years:

$$(2)\ \ln TFP_{pt}^{i} = (\ln Y_{pt}^{i} - \ln \overline{Y}_{t}^{i}) + \sum_{s=2}^{t} (\ln \overline{Y}_{t}^{i} - \ln \overline{Y}_{t-1}^{i})$$

$$-[\sum_{m} \frac{1}{2}(S_{mpt}^{i} + \overline{S}_{mt}^{i})(\ln X_{mpt}^{i} - \ln \overline{X}_{mt}^{i})$$

$$+\sum_{s=2}^{t} \ \sum_{m} \frac{1}{2}(\overline{S}_{mt}^{i} + \overline{S}_{mt-1}^{i})(\ln \overline{X}_{mt}^{i} - \ln \overline{X}_{mt-1}^{i})]$$

I construct the TFP index expressed in Equation (2) for each plant *p* in year *t* using the set of inputs *m*={Capital, Raw Materials, Production Workers, Non-Production Workers}. The superscript *i* indicates that mean variables are calculated at the SIC4 industry level. X_{mpt}^{i} is the expenditure of plant p in time t on input m and S_{mpt}^{i} is the share of input m in total revenue. I calculate average input usage and shares at the industry-year-level. Therefore, $\overline{S}_{mt}^{i}$, $\ln \overline{Y}_{t}^{i}$ and $\overline{X}_{mt}^{i}$ are the arithmetic means of industry-level input shares, revenue and input expenditure.

Semi-parametric estimators, including those developed by Olley and Pakes (1996) and Levinsohn and Petrin (2003) have been used extensively in recent papers studying changes in

TFP.[15] As has been established in this literature, these methods can be useful for correcting the simultaneity bias that arises when plants with high TFP consume more inputs and the selection bias associated with only observing surviving plants. These methods are not well-suited for use with the economic census data employed in this paper, due to their use of lagged input values in the TFP calculation. While it could be useful to calculate TFP using one of these semi-parametric methods if annual data were available, I will note that Van Biesebroeck (2004) finds that TFP measures derived from various methods tend to be highly correlated.

C. Deflation

Whenever productivity is calculated using either revenue or value-added data—i.e. data that contain both a price and a quantity component—it is important to separate changes in prices and mark-ups from changes in true productivity. This separation becomes critically important when mark-ups and productivity are likely to move in the same direction, as is the case in the situation examined in this paper. While some have suggested that antidumping duties can increase productivity, through their influence on technology adoption decisions, they almost certainly lead to higher mark-ups as well. Without an adjustment to account for changes in mark-ups, an increase in prices resulting from AD protection would show up as an increase in observed total factor productivity. This means that the results are biased toward finding a positive correlation between AD protection and revenue TFP.

The availability of output data measured in units of quantity, described above, puts me in a unique position to separate changes in prices and mark-ups from changes in physical productivity. In instances in which quantity data are available, physical quantities can be used as

[15] See, for example, Pavcnik (2002), Fernandes (2008) and Konings and Vandenbussche (2008).

a measure of plant-level output and incorporated into the calculation of physical productivity, without deflation.

When calculating revenue productivity, I control for changes in mark-ups—to the extent possible—by deflating revenue using industry-level price indexes, applied to the set of products produced at each plant. This technique results in a plant-level deflator that is constructed by weighting the industry-level deflators according the share of a plant's output that is assigned to that industry. Industry-level output deflators, as well as industry-level deflators for cost of materials and capital are from the NBER-CES Manufacturing database reported in Bartelsman, Becker and Gray (2000).

There are at least two ways in which these plant-level deflators are insufficient for completely separating changes in mark-ups from changes in true productivity. First, since they are based on average price indexes, they do not allow for heterogeneity in pricing across plants. In this sense, plants that charge high prices—due to high local market power, for example—would be misinterpreted as high-productivity plants. Second, because the price indexes are calculated at the industry, rather than the product level, they will not fully reflect increases in product-level prices. This higher level of aggregation means that revenue-based productivity measures will overstate productivity growth in situations where mark-ups are increasing, as is likely the case in the natural experiment considered in this paper.

D. Effective Antidumping Duty Rates

A single antidumping investigation can be filed against imports from multiple countries and if the case ends in with a determination by the DOC and ITC to offer protection, each country may be assigned a different ad-valorem antidumping duty. Naturally, imports from certain countries account for larger shares of U.S. imports of a good than others. In order to account the true importance of an antidumping duty on U.S. trade, therefore, I calculate an effective antidumping duty rate for each product that is assigned an ad-valorem antidumping duty. The effective antidumping rate is calculated as follows:

$$Rate_{gt} = \sum_{c} SHARE_{c,g,t-1} * AVD_{c,g,t}$$

where $SHARE_{c,g,t-1}$ is country c's share of U.S. imports of product g in time t-1 and $AVD_{c,g,t}$ is the ad-valorem duty applied to imports of product g from country c in time t. A country's share is calculated based on imports in time t-1, rather than time t, because antidumping duties often lead to significant reductions in imports from pre-protection levels. Using a pre-protection share, therefore, provides a more accurate representation of a country's importance to U.S. trade.

Research Questions

A. Do Temporary Tariffs Increase or Decrease Plant-Level Productivity?

As discussed above, some have argued that temporary protection can increase within-plant productivity by increasing the incentive to invest in new technology. On the other hand, temporary protection is also likely to lead to higher prices. Because an increase in revenue-based productivity that occurs at the time of protection could be caused by either of these phenomena, however, it can be difficult to determine what is driving gains in revenue productivity. Using output data that is measured in units of quantity, I am able to separate these two effects by calculating both revenue and physical productivity measures. Moreover, I am able to directly measure the effects of AD on plant-level prices and mark-ups. I find that apparent growth in productivity associated with AD protection appears to be driven mainly by higher prices and mark-ups, rather than increases in true productivity.

Empirical Strategy

I examine the effect of temporary protection on plant-level productivity, prices and mark-ups using a difference in difference approach. As discussed above, the treatment group is composed of plants producing products that receive AD protection. The control group is composed of plants producing products that applied for, but did not receive AD protection. The goal of the difference in difference methodology is to isolate the effect of the treatment—AD protection—by eliminating time-invariant differences between the treatment and control group, as well as time-specific effects common to both treatment and control. The difference in difference estimator, therefore, measures not simply the change in the dependent variable that occurs following AD protection, but rather measures the difference between the changes in the treatment group and the changes in the control group.

Let T be the set of plants producing products that receive AD protection and let C be the set of plants producing products that apply for, but do not receive protection. Further, define I_g to be the date that the AD investigation is initiated for product g. I measure the difference in difference effect by estimating Equation (3):

(3) $Prod_{pgt} = \alpha + \beta_1 Treatment_{pgt} * Post_{pgt} + \gamma_t + \delta_g + \varepsilon_{pt}$, where

$Treatment_{pgt} = 1 \ \forall \ p \in T$ and $Treatment_{pgt} = 0 \ \forall \ p \in C$

$Post_{pgt} = 1 \ \forall \ t > I_g$, 0 otherwise

Here, $Prod_{pgt}$ is productivity—measured in TFP or Labor Productivity—at plant p, which produces product g at time t. Year fixed effects capture any macro-level shocks affecting plants in T and C equally. Similarly, product fixed effects, δ_g, capture time-invariant differences between products. Note that Equation (3) contains product-level fixed effects, rather than a more general *Treatment* dummy used in the most basic difference in difference expressions. This specification captures time-invariant differences between producers of different products *within T and C*. This is likely important when dealing with a diverse set of manufacturers from different sectors and industries. Finally, the coefficient β_1 on the interaction term is the coefficient of interest and measures the difference in difference effect of AD protection on the plant-level outcomes discussed below.

Equation (3) defines protection with a binary variable—any plant that receives any AD protection is considered to be equally protected. It seems reasonable to expect, however, that plants' reactions to protection would depend not only on this simple binary classification, but also on the level of protection they receive. That is, plants producing products that receive high ad-valorem AD rates—such as the 376.67 percent AD rate on garlic from China—may respond differently than those producing products that receive low AD rates, such as the 2.98 percent AD rate on Collated Roofing Nails from Taiwan. As these two examples indicate, the variation in AD rates among cases that receive protection is large: the mean is 64 percent and the standard deviation is 60 percent.

I measure the effects of differing AD rates by augmenting Equation (3) with an additional interaction term:

$$(4)\ Prod_{pgt} = \alpha + \beta_1 Treatment_{pgt} * Post_{pgt} + \beta_2 Rate_{pgt} * Post_{pgt} + \gamma_t + \delta_g + \varepsilon_{pt}$$

Here, $Rate_{pgt}$ is the ad-valorem effective antidumping duty rate on product *g*, which is produced by plant *p* at time *t*. By interacting $Rate_{pgt}$ with the $Post_{pgt}$ dummy, I am able to separate the effect of varying rates of protection from the mean response of all plants receiving AD protection.

Equations (3) and (4) provide within-product estimates of the effect of antidumping duties on plants. It is important to note, these results do not reflect the within-plant effect of antidumping duties. Because equations (3) and (4) are estimated an unbalanced panel, coefficient estimates could reflect changes in mean plant-level productivity due to entry in exit. In order to estimate the within-plant effect of antidumping duties, I re-estimate equations (3) and (4) with plant fixed effects for the balanced subsample of plants producing in all three census years.[16] These estimates provide both a useful robustness check for the within product-group estimates.

Lastly, I will employ the difference in difference framework in equations (3) and (4) to examine the effect of antidumping duties on plant-level prices, as well as mark-ups over average variable cost and average total cost. Prices and mark-ups are defined as follows:

$$P_{pgt} = \frac{TVS_{pgt}}{Q_{pgt}}$$

where *TVS* is a plant's total value of shipments and *Q* is the total quantity of units. Plant-level mark-ups over average total cost are defined as:

$$PATC_{pgt} = \frac{P_{pgt}}{ATC_{pgt}}$$

where $ATC_{pgt} = \frac{Wages_{pgt} + CM_{pgt} + RTAE_{pgt}}{Q_{pgt}}$

Here, *Wages* are the wages paid to production workers, *CM* is the cost of materials and *RTAE* is real total value of assets, or capital. Similarly, plant-level mark-ups over average variable cost are defined as:

[16] Estimates with plant fixed effects can only be estimated on the balanced subsample. For plants with two or less years of data, plant and year fixed effects would completely explain variation in plant-level variables, preventing the identification of the antidumping treatment effect.

$$PAVC_{pgt} = \frac{P_{pgt}}{AVC_{pgt}}$$

where $AVC_{pgt} = \frac{Wages_{pgt} + CM_{pgt}}{Q_{pgt}}$

Results

Revenue Productivity

I find that antidumping protection is associated with an increase in revenue productivity, as shown in Table 5. Columns 1 and 2 of Table 5 report the results for equation (3) with TFP and labor productivity and columns 3 and 4 report results for the same specification, with state fixed effects. Estimating equation (4), with the interaction term for the effective duty rate included, I find similar results. As shown in columns 5 through 8 of Table 5, I continue to find a positive and significant relationship between antidumping protection and revenue productivity. While I do not find a significant relationship between the effective duty rate and revenue productivity separate from the binary protection effect, the inclusion of the rate term does allow me to estimate the coefficient on the binary protection variable more precisely.

In Table 6, I report the equivalent results when estimating equations 3 and 4 with plant fixed effects on the balanced subsample. As discussed above, these results are important because they show the effects of antidumping duties within a particular plant. In other words, they are not affected by changes in the composition of plants in the product-group caused by entry and exit. The results are consistent with those obtained with product fixed effects, with protected plants exhibiting increases in plant-level productivity.

Physical Productivity

As reported in Tables 7 and 8, the effects of antidumping duties on physical productivity are starkly different than those reported for revenue productivity. Table 7 shows that antidumping duties are associated with a decrease in physical productivity, both with and without consideration of the effective duty rate. This effect persists when estimating the difference in difference specification with plant fixed effects on the balanced subsample, as shown in Table 8.

A word of warning in terms of interpreting these results is necessary here. It would be inappropriate based on these results to claim that antidumping duties, *in general*, decrease plant-level physical productivity. It is true that antidumping duties were associated with a decline in productivity among the set of plants reporting output data in units of quantity. However, this group is not necessarily representative of the full set of plants subject to antidumping protection. First, as can be seen in Tables 7 and 8, the plants reporting quantity-based data are not necessarily clustered in the sectors most commonly associated with antidumping protection. Second, as shown in Table 8, plants reporting output in units of quantity show no effect on revenue productivity from antidumping protection, unlike the increase experienced in general. Nonetheless the fact that plants in this sub-sample experienced a zero effect of antidumping protection on revenue productivity and a large and highly significant decrease in physical productivity suggests that mark-ups are affecting results based on revenue productivity.

Prices and Mark-Ups

The disparity between results showing the effect of antidumping protection on revenue versus physical productivity suggest that increases in prices and mark-ups are likely playing a role in the apparent increase in revenue productivity. I estimate the effect of antidumping protection on the measures of prices and mark-ups over average total cost and average variable cost described above. As reported in Table 9, plants increase prices when they receive antidumping protection. Moreover, these pricing changes are sensitive to the effective duty rate a plant experiences—the higher the effective duty rate, the higher the prices charged by the plant.

Table 10 reports the effects of antidumping protection on mark-ups over average total cost and average variable cost. I find that antidumping duties increase mark-ups over average total cost across all specifications. The effect of antidumping duties on mark-ups is positive over all specifications, but I am only able to estimate the coefficient precisely when estimating the specification with plant fixed effects. This is not necessarily surprising since the presence of entering and exiting plants in the unbalanced panel used when estimating with product fixed effects will understate the effect of protection on mark-ups if entering and exiting plants charge lower prices.[17]

B. Do Temporary Tariffs Discourage Product-Dropping?

BRS (2006) shows that reductions in trade barriers can increase firm or plant-level productivity by inducing firms to drop their least productive products, while expanding output of their most productive products. It is interesting, then, to consider the effect of temporary protection on plants' product-adding and dropping decisions.

Empirical Strategy

The effect of antidumping duties on the probability of dropping products is investigated using a difference in difference specification similar to that employed to study changes in plant-level productivity. By comparing changes in dropping that occur in protected plants to changes among unprotected plants, I am able to estimate the effect of antidumping duties on product-switching.

[17] Foster, Haltiwanger and Syverson (2008) have found that entrants charge lower prices than continuing plants.

An important difference between this product-switching analysis and the plant-level productivity regressions describe above, however, is that the product-switching data are defined at the plant-product-level. This means that I have dropped the restriction that each plant be assigned to a particular treatment or control product. In doing so, I am able to consider the full set of products that applied for protection. Using these plant-product-level data, I estimate the following two regressions:

$$(5)\ \ Drop^{i}_{pgt} = \alpha + \beta_1 Treatment^{i}_{pgt} * Post^{i}_{pgt} + \beta_2^{\prime} X^{i}_{pgt} + \gamma_t + \delta_g + \varepsilon^{i}_{pgt}$$

$$(6)\ \ Drop^{i}_{pgt} = \alpha + \beta_1 Treatment^{i}_{pgt} * Post^{i}_{pgt} + \beta_2 Post^{i}_{pgt} * Rate^{i}_{pgt} + \beta_3^{\prime} X^{i}_{pgt} + \gamma_t + \delta_g + \varepsilon^{i}_{pgt}$$

Drop is a binary variable that equals 1 if product g is produced by plant p at time t, but not time t+5. Add is a binary variable that equals 1 if product g is produced by plant p at time t, but was not produced in time t-5. X is a matrix of plant-product-level variables found to be determinants of product-dropping in Bernard, Redding and Schott (2008), namely product-level sales and the number of years a product has been produced (tenure). The superscript i denotes that the data are at the plant-product level. To be clear, the variable *Drop* only takes into account product switching by continuing plants. Exiting plants are not considered product-droppers.

Results

I do find that plants are less likely to drop protected products as reported in Table 11. Moreover, this product-switching behavior is sensitive to the value of the effective duty rate applied to a product. I find that plants are less likely to drop a product, the higher the effective rate of protection they receive. In the product-dropping regression, the results are robust to the inclusion of plant-level sales and product tenure, which are both negative and significant, as predicted by theory. These results make clear that more plants produce a given protected product than would be the case if the product was unprotected.

C. Do Temporary Tariffs Discourage Plant-Level Exit?

Trade protection can slow aggregate productivity growth by preventing the exit of low-productivity plants that would otherwise cease to operate. I find strong evidence of this effect among plants protected by antidumping duties. Comparing the treatment and control groups defined above, I find no increased probability of exit among plants in the control group. This result is important and suggests that unprotected plants are dynamic and able to change their product mix in order to survive.

Empirical Strategy

I begin by examining the probability of exit at the plant-level. I define a plant as exiting in year t if it appears in the CMF in year *t*, but not in year *t+5*. To be clear, a plant that halts production of the investigated product between year *t* and year *t+5*, but continues to operate, is not counted as an exit. A binary exit variable is defined in this way for the years 1987 and 1992. The exit variable is missing in 1997, because I am unable to determine whether plants operating in 1997 will exit in 2002.

I estimate the relationship between antidumping protection and the probability of exit using a difference in difference framework, identical to the specification used to study changes in plant-level productivity above. I employ a linear probability model, to allow for the inclusion of fixed effects and clustering of standard errors.

(7) $Exit_{pgt} = \alpha + \beta_1 Treatment_{pgt} * Post_{pgt} + \gamma_t + \delta_g + \varepsilon_{pt}$

The binary dependent variable, *Exit* was described above. The coefficient β_1 is the primary parameter of interest and estimates the effect of receiving antidumping protection on the probability of exit. As in Equation (3), year and product-group fixed effects are included. Estimates with robust standard errors and clustering at the product-group level are reported in Table 8.

Next, I expand Equation (7) to include plant-level variables that have been found to be important determinants of exit in the large empirical literature on the effects of changes in trade costs on exit. Using determinants of exit from Bernard, Jensen and Schott (2006), I estimate Equation (8):

(8) $Exit_{pgt} = \alpha + \beta_1 Treatment_{pgt} * Post_{pgt} + \beta' X_{pgt} + \gamma_t + \delta_g + \varepsilon_{pt}$

where X is a matrix of plant-level variables including log of total employment, plant age, log of capital-labor ratio, log of average wage and indicators for whether the plant is a multi-product plant, or a part of a multi-unit firm.

Results

Results are reported in Table 12. I find that antidumping protection has no effect on the probability of exit. This result holds regardless of whether the additional plant-level covariates or the interaction with the effective duty rate are included. The results suggest that plants that do not receive protection respond not by exiting, but rather by changing their product mix, dropping the unprotected product and switching to production of other products.

D. Do Temporary Tariffs Decrease Output Rationalization and Aggregate Productivity?

A number of theoretical models including Melitz (2003) and Bernard, Redding and Schott (2006) predict that tariff increases allow for the continued operation of firms that might otherwise stop production. If this is the case with antidumping duties, we would expect the level of output rationalization to increase more in the control group, than in the treatment group.

Empirical Strategy

I examine the process of output rationalization using a decomposition from Olley and Pakes (1996), which has also been used in Pavcnik (2002) and Fernandes (2007). This procedure decomposes growth in aggregate productivity into two components, shown below:

$$(12)\ W_{gt} = \sum_{p} s_{pgt} TFP_{pgt} = TFP_{tg}^{mean} + (s_{pgt} - \bar{s}_g)(TFP_{pgt} - TFP_{gt}^{mean})$$

The first term of the final expression represents mean plant-level productivity at time t. The second term is a covariance term representing the degree to which greater output is produced by higher-productivity plants.

The covariance term measuring the degree of output rationalization will be the primary variable of interest. Ideally, I would simply examine the effects of antidumping duties on aggregate productivity, W_{gt} directly. A number of data problems would make this comparison unreliable, however. First, as mentioned above, the use of revenue-based aggregate productivity measures would overstate productivity gains among protected product-groups, since I have shown that protected plants respond to temporary protection by increasing prices. Moreover, quantity-based productivity measures are not useful in settings where analysis is taking place at the product-group level or higher, since quantity data is only available for producers of a limited set of products.

The use of revenue-based productivity measures is less problematic for analyzing output rationalization. Assuming that prices increase uniformly among all producers of a given product once it receives protection, the covariance term will still accurately reflect the degree of output rationalization within a product group. After calculating aggregate productivity, mean plant-level productivity and the output rationalization term at the product-group-level, I report their output-weighted means by year, treatment group and a dummy variable indicating whether the antidumping investigation for product g has already taken place. The results of the decomposition described are reported in Table 13.

In addition, I repeat the exit exercise described above to examine the more general question of whether antidumping protection affects the probability of a plant stopping production. I consider a plant to have stopped production if it either exits completely or stops producing the investigated product. If antidumping protection allows for continued production by low-productivity plants that would have otherwise stopped producing, it will have a negative effect on the level of output rationalization. After examining the effect of antidumping protection on the probability of stopping production, I will then examine whether plants that do stop producing are of lower-productivity than continuing plants.

Results

I find that antidumping protection decreased the process of output rationalization, compared to the control group, particularly for TFP. As can be seen in Table 13, the level of output rationalization was higher in the protected group in 1987, before the initiation of any of the antidumping investigations considered in this analysis. But output rationalization increases in the unprotected group, as the least-productive plants exit and plants drop the products that applied for but did not receive protection. By 1997, the unprotected group has overtaken the protected group in terms of output rationalization. The evolution of output rationalization in terms of labor productivity is less dramatic, although the unprotected group always exhibits output rationalization that is higher than the protected group, and growing.

Finally, in Table 14, I report the results of my examination of the effect of antidumping protection on the probability of a plant stopping production of an investigated product. I do find that antidumping protection decreases the probability of production-stopping. Moreover, I find that plants that stop production are lower-productivity than continuing plants. These two results suggest that antidumping protection allows for the continued operation of low-productivity plants that would have otherwise ceased production, leading to a lower level of output rationalization.

Section 6: Conclusion

This paper describes the effects of a temporary increase in tariffs on the performance of U.S. manufacturers. By considering the case of antidumping duties, I compare the responses of protected manufacturers to a similar control group of unprotected plants. I find that these apparent gains in revenue productivity are driven predominantly by increases in prices and mark-ups, rather than true productivity growth. Moreover, I find that antidumping protection allows for continued production by plants that might have otherwise ceased production of the investigated product. As a result, temporary protection slows the process of output rationalization, with less productive plants increasing their share of total output. Lastly, the results suggest the importance of thinking about plants and firms in a multi-product setting, since unprotected plants respond to a failure to receive protection primarily by switching products, rather than exiting.

References

Amiti, Mary and Konings, Jozef (2007), "Trade Liberalization, Intermediate Inputs and Productivity: Evidence from Indonesia," *American Economic Review*, 97, pp. 1611-1638.

Aw, B.Y., Chung, S. and Roberts, M. (2000), "Productivity and Turnover in the Export Market: Micro-Level Evidence from Korea and Taiwan (China)," *World Bank Economic Review*, 14, pp. 65-90

Bartelsman, Eric J., Becker, R. A. and Gray, Wayne B. (2000) "NBER-CES Manufacturing Industry Database," NBER Technical Working Paper 205.

Bernard, Andrew and Jensen, J. Bradford (1999), "Exceptional Exporter Performance: Cause, Effect, or Both?," *Journal of International Economics*, 47, pp. 1-25.

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

Bernard, Andrew, Jensen, J. Bradford and Schott, Peter (2006), "Trade Costs, Firms and Productivity," *Journal of Monetary Economics*, 53, pp. 917-937.

Bernard, Andrew, Redding, Steven and Schott, Peter (2006), "Multi-Product Firms and Trade Liberalization," mimeo.

Bernard, Andrew, Redding, Steven and Schott, Peter (2008), "Multi-Product Firms and Product-Switching," mimeo.

Caves, Douglas W., Laurits R. Christensen and W. Erwin Diewert (1982), "The Economic Theory of Index Numbers and the Measurement of Input, Output and Productivity," *Econometrica*, 50(6), pp. 1393-1414.

Crowley, Meredith (2006), "Do Antidumping Duties and Safeguard Tariffs Open or Close Technology Gaps?," *Journal of International Economics*, 68, pp. 469-484.

Fernandes, Ana (2007), "Trade Policy, Trade Volumes and Plant-Level Productivity in Colombian Manufacturing Industries," *Journal of International Economics*, 71(1), pp. 52-71.

Feenstra, Robert, John Romalis and Peter Schott (2002), "U.S. Imports, Exports and Tariff Data," NBER Working Paper # 9387.

Goh, A.-T. (2000), "Opportunity Cost, Trade Policies and the Efficiency of Firms," *Journal of Development Economics*, 62, pp. 363–383.

Konings, Jozef and Vandenbussche, Hylke (2008), "Heterogeneous Responses of Firms to Trade Protection," *Journal of International Economics*, forthcoming.

Levinsohn, James and Amil Petrin (2003), "Estimating Production Functions Using Intermediates to Control for Unobservables," *Review of Economic Studies*, 70, pp. 317-341.

Matsuyama, Kiminori (1990), "Perfect Equilibria in a Trade Liberalization Game," *American Economic Review*, 80(3), pp. 480-492.

Melitz, Marc J. (2003), "The Impact of Trade on Aggregate Industry Productivity and Intra-Industry Reallocations." *Econometrica*, 71(6), pp. 1695-725.

Miyagiwa, Kaz and Ohno, Yuka (1995), "Closing the Technology Gap Under Protection," *American Economic Review*, 85(4), pp. 755-770.

Miyagiwa, Kaz and Ohno, Yuka (1999), "Credibility of Protection and Incentives to Innovate," *International Economic Review*, 40(1), pp. 143-163.

Pavcnik, Nina (2002), "Trade Liberalization, Exit and Productivity Improvements: Evidence from Chilean Plants." *Review of Economic Studies, 69*(1), pp. 245-76.

Pierce, Justin (2008), "Trade Liberalization, Productivity and Input Variety: Evidence from Colombian Plants," Georgetown University, mimeo, 2008.

Serti, Francesco and Tomasi, Chiara (2007), "Self Selection and Post-Entry Effects of Exports. Evidence from Italian Firms," mimeo.

United States International Trade Commission (2006), "Import Injury Investigations: Case Statistics: FY 1980-2005."

Van Biesebroeck, Johannes (2004), "Robustness of Productivity Estimates," National Bureau of Economic Research Working Paper 10303.

Table 1: Antidumping Investigations by HTS Chapter, 1988-1996

HS2	Description	Investigations
73	Articles of Iron and Steel	27
72	Iron and Steel	20
84	Machinery	16
28	Inorganic Chemicals	14
85	Eletrical Machinery	13
29	Organic Chemicals	12
87	Transportation Vehicles and Parts	11
90	Precision Instruments and Apparatus	8
39	Plastics and articles thereof	6
25	Plastering, Lime and Cement	5
81	Other Base Metals	5
30	Pharmaceutical Products	4
40	Rubber and articles thereof	4
56	Certain Textiles	4
83	Misc. Articles of Base Metal	4
Other		45
Total		198

Table 2: All Antidumping Cases, by Outcome and SIC2

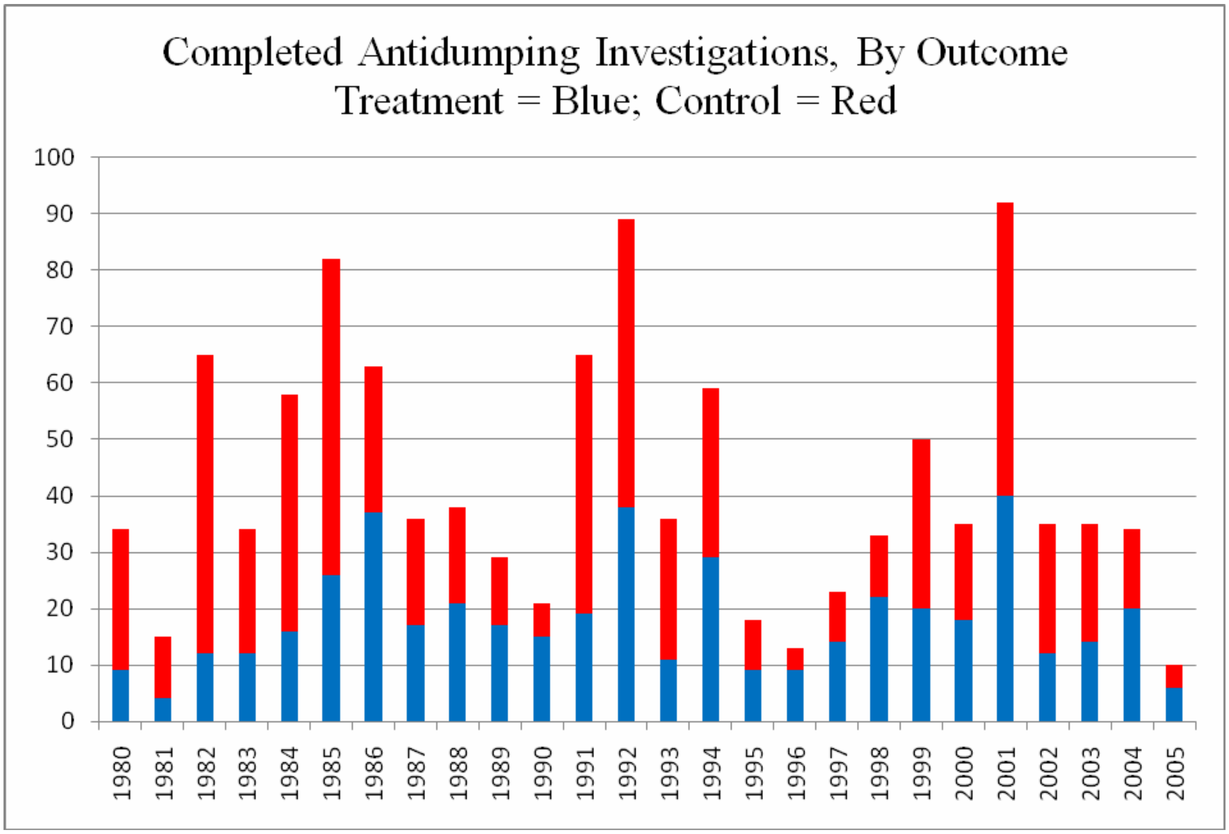

Table 3: Plant-Level Observations, by SIC2

Total Observations				Observations With Quantity			
SIC2	Control	Treatment	Total	SIC2	Control	Treatment	Total
20	163	1,462	1,625	20		1,363	
22	1,061	891	1,952	22	782	593	1,375
23	8,283	1,725	10,008	23	1,047	713	1,760
26	2,602	0	2,602	26	1,613	0	1,613
28	815	3,566	4,381	28	273	1,077	1,350
30	13,681	2,996	16,677	30	424	192	616
32	2,081	582	2,663	32	430	485	915
33	468	3,266	3,734	33	16	2,044	2,060
34	13,244	4,318	17,562	34	892	351	1,243
35	3,884	16,066	19,950	35	157	437	594
36	650	7,540	8,190	36	42	55	97
37	2,869	889	3,758	37	723		
38	75	3,071	3,146	38	25		
39	0	413	413	39	0	62	62

Table 4: Summary Statistics by Treatment Group, Year

Table 5: The Effect of Antidumping Duties on Revenue TFP – Within Product-Group Estimators

	TFP	LP	TFP	LP	TFP	LP	TFP	LP
Treatment*Post	0.0643	0.0504	0.0647	0.0504	0.0785	0.0588	0.0784	0.0585
	0.0329	0.0211	0.0328	0.0208	0.0335	0.0228	0.0332	0.0223
Post*Rate					-0.0011	-0.0006	-0.001	-0.0006
					0.0022	0.0014	0.0022	0.0013
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	96,662	96,662	96,662	96,662	96,662	96,662	96,662	96,662
R-Squared	0.666	0.309	0.668	0.313	0.666	0.39	0.668	0.313

Robust standard errors are reported with clustering at the product-group level.

Table 6: The Effect of Antidumping Duties on Revenue TFP – Within Plant Estimators

	TFP	LP	TFP	LP
Treatment*Post	0.0487	0.0256	0.0632	0.0312
	0.0104	0.0082	0.0124	0.0098
Post*Rate			-0.0011	-0.0004
			0.0005	0.0004
Year FE	Yes	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes
Observations	27,699	27,699	27,699	27,699
R-Squared	0.909	0.874	0.99	0.874

Robust standard errors are reported with clustering at the plant level.

Table 7: The Effect of Antidumping Duties on Revenue TFP – Within Product-Group Estimators

	TFPQ	LPQ	TFPQ	LPQ	TFP	LP	TFP	LP
Treatment*Post	-0.4669	-0.4515	-0.0495	-0.0561	-0.001	-0.0174	0.0497	0.0142
	0.2198	0.2002	0.1337	0.1288	0.0575	0.0275	0.0493	0.0389
Post*Rate			-0.0214	-0.0203			-0.0026	-0.0016
			0.0051	0.0049			0.0018	0.0016
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12,634	12,634	12,634	12,634	12,634	12,634	12,634	12,634
R-Squared	0.551	0.539	0.553	0.541	0.756	0.469	0.756	0.469

Robust standard errors are reported with clustering at the product-group level.

Table 8: The Effect of Antidumping Duties on Physical TFP – Within Plant Estimators

	TFPQ	LPQ	TFPQ	LPQ	TFP	LP	TFP	LP
Treatment*Post	-0.7006	-0.6338	-0.1539	-0.1456	-0.0177	0.0248	0.025	0.0135
	0.107	0.1059	0.1303	0.1291	0.0268	0.0211	0.0357	0.0281
Post*Rate			-0.0309	-0.0276			-0.0024	0.0006
			0.0066	0.0066			0.0013	0.0011
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,702	3,702	3,702	3,702	3,702	3,702	3,702	3,702
R-Squared	0.816	0.824	0.822	0.828	0.914	0.915	0.914	0.915

Robust standard errors are reported, with clustering at the plant-level.

Table 9: The Effect of Antidumping Duties on Plant-Level Prices

	Price	Price	Price	Price
Treatment*Post	0.388	0.0136	0.5946	0.098
	0.1906	0.1367	0.1082	0.1325
Post*Rate		0.0192		0.0281
		0.0035		0.0065
Year FE	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	No	No
Plant FE	No	No	Yes	Yes
Observations	12,634	12,634	3,702	3,702
R-Squared	0.565	0.567	0.818	0.822

Table 10: The Effect of Antidumping Duties on Mark-Ups

	P/ATC	P/ATC	P/ATC	P/ATC	P/AVC	P/AVC	P/AVC	P/AVC
Treatment*Post	0.0383	0.0391	0.0761	0.055	0.0091	0.0101	0.0416	0.025
	0.0206	0.0241	0.0186	0.0251	0.0145	0.019	0.0147	0.0234
Post*Rate		0		0.0012		-0.0001		0.0009
		0.0006		0.001		0.0005		0.001
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	No	No	Yes	Yes	No	No
Plant FE	No	No	Yes	Yes	No	No	Yes	Yes
Observations	12,634	12,634	3,702	3,702	12,634	12,634	3,702	3,702
R-Squared	0.286	0.286	0.752	0.752	0.152	0.152	0.668	0.668

Table 11: Antidumping Duties and the Probability of Product-Dropping

	Drop	Drop	Drop	Drop	Drop	Drop	Drop	Drop
Treatment*Post	-0.06	-0.0603	-0.0615	-0.0617	-0.0232	-0.0231	-0.0247	-0.0254
	0.0168	0.0167	0.0162	0.0162	0.0238	0.0237	0.0219	0.0218
Post*Rate					-0.0028	-0.0028	-0.0028	-0.0027
					0.0013	0.0013	0.0012	0.0012
Product Shipments			-0.0762	-0.0764			-0.0762	-0.0764
			0.0024	0.0025			0.0024	0.0025
Product Tenure			-0.1188	-0.1182			-0.1189	-0.1183
			0.0123	0.0123			0.0123	0.0123
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	46,742	46,742	46,742	46,742	46,742	46,742	46,742	46,742
R-Squared	0.118	0.119	0.204	0.205	0.118	0.119	0.205	0.205

Table 12: Antidumping Duties and the Probability of Plant Exit

	Exit	Exit	Exit	Exit	Exit	Exit	Exit	Exit
Treatment*Post	-0.0018	-0.0018	0.0169	0.0165	-0.0024	-0.0025	0.0029	0.0027
	0.0119	0.012	0.0137	0.0137	0.0119	0.012	0.015	0.015
Post*Rate			-0.0013	-0.0013			-0.0004	-0.0004
			0.0006	0.0006			0.0007	0.0007
Log TE					-0.092	-0.09	-0.0919	-0.0899
					0.0031	0.0031	0.0031	0.0031
Age					-0.0022	-0.0022	-0.0022	-0.0022
					0.0004	0.0004	0.0004	0.0004
Log RTAE/TE					-0.0168	-0.0155	-0.0168	-0.0155
					0.0025	0.0025	0.0025	0.0025
Log Avg_Wage					-0.0746	-0.0824	-0.0746	-0.0824
					0.0077	0.0076	0.0077	0.0076
Multi-Unit					0.0934	0.0978	0.0934	0.0978
					0.0072	0.0071	0.0072	0.0071
Multi-Product					-0.0181	-0.0159	-0.0181	-0.0159
					0.0043	0.0043	0.0043	0.0043
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	62,285	62,285	62,285	62,285	62,285	62,285	62,285	62,285
R-Squared	0.059	0.065	0.059	0.065	0.116	0.12	0.116	0.12

Table 13: Antidumping Duties and Output Rationalization

Table 14: Antidumping Duties and Production-Stopping

	Stop	Stop	Stop	Stop	TFP	LP	TFPQ	LPQ
Treatment*Post	-0.06	-0.0603	-0.0247	-0.0254				
	0.0168	0.0167	0.0219	0.0218				
Post*Rate			-0.0028	-0.0027				
			0.0012	0.0012				
Stop					-0.0502	-0.1235	-0.177	-0.2199
					0.0089	0.013	0.0696	0.0775
Year FE	Yes	Yes	Yes	Yes				
Product FE	Yes	Yes	Yes	Yes				
Observations	62,285	62,285	62,285	62,285	62,285	62,285	9,212	9,212
R-Squared	0.074	0.077	0.074	0.077	0.682	0.332	0.552	0.547

Robust standard errors are reported, with clustering at the product-group-level.