

IMPORT COMPETITION AND QUALITY UPGRADING

Mary Amiti and Amit K. Khandelwal*

Abstract—The production of high-quality goods is often viewed as a precondition for export success and economic development. We provide the first evidence that countries' import tariffs affect the rate at which they upgrade product quality. Our analysis uses highly disaggregated data covering exports from 56 countries across 10,000 products to the United States using a novel approach to measure quality. As predicted by distance-to-the-frontier models, we find that lower tariffs are associated with quality upgrading for products close to the world quality frontier, whereas lower tariffs discourage quality upgrading for products distant from the frontier.

I. Introduction

SUBSTANTIAL evidence suggests that firms in high-income countries produce and export higher-quality goods than those in less-developed nations (Schott, 2004; Hummels & Klenow, 2005; Hallak, 2006). These studies inform our understanding of the determinants of quality within a cross-section, but they do not examine how countries transition from the production of low-quality to high-quality products. This question is important, since the production of high-quality products is often viewed as a precondition for export success and, ultimately, economic development (Kremer, 1993; Grossman & Helpman, 1991). In this paper, we provide the first evidence that import tariffs in the importing country affect the rate at which they upgrade the quality of their products.

We analyze the effect of import competition on quality upgrading using highly disaggregated data on exports to the United States from 56 countries in 10,000 products. We infer the quality of these products using a novel approach that Khandelwal (2010) developed. In contrast to previous approaches that use unit values as proxies for quality, we estimate quality using both prices and quantity information, where, conditional on price, higher quality is assigned to products with higher market shares. This approach has the advantage of accounting for differences in quality-adjusted manufacturing costs, such as wages, that could explain variation in prices. We match these quality estimates to detailed import tariff data for the countries' home markets. The tariffs provide a measure of the import competition in these countries.

To guide our empirical analysis, we draw on the theoretical literature in industrial organization that analyzes competition

and innovation.¹ Although both competition and innovation are much broader concepts than the focus of our analysis, quality upgrading is one important element of innovation, and trade liberalization is one policy tool that stimulates competition in product markets. For example, according to widespread evidence, tariff reductions lead to procompetitive pressures in the liberalizing countries that result in both resource reallocation and lower markups.² We draw on models by Aghion and Howitt (2005), Aghion, Bloom, Blundell, Griffith, and Howitt (henceforth ABBGH, 2005) and Aghion, Blundell, Griffith, Howitt, and Prantil (henceforth ABGHP, 2009) that allow the relationship between competition and innovation to depend on the distance of the product to the world technology frontier.³ These models highlight two forces. First, for firms far from the technology frontier, an increase in competition reduces incentives to innovate because ex post rents from innovation are eroded by new entrants, as in Schumpeter's appropriability argument. As firms approach the frontier, however, competition can increase incentives to innovate because it reduces firms' preinnovation rents by more than it reduces their postinnovation rents. We refer to this force as the escape competition effect. These models therefore nest predictions of both positive and negative effects of competition on the innovative activity of firms.⁴

The empirical strategy we adopt examines the relationship between the change in quality of products over five-year intervals and the five-year lagged tariff levels across many countries between 1990 and 2005. The high level of disaggregation of both the tariff and the quality measures allows us to isolate the effects of import tariffs from other factors that could potentially affect quality. We include product-year fixed effects to account for product-specific productivity shocks or changes in consumer demand and include country-year fixed effects to account for changes in countries' relative

¹Recent work by Verhoogen (2008) and Bustos (2011) provides theoretical and empirical evidence that access to destination markets increases the incentive for firms to improve the quality of their products. Since our context focuses on the impact of competition within a country's home market, we turn to models by Aghion and coauthors.

²For instance, see Levinsohn (1993), Pavcnik (2002), Konings and Vandenbussche (2005), and the comprehensive survey article by Tybout (2003).

³Hausmann and Rodrik (2003) and Acemoglu, Aghion, and Zilibotti (2006) also show that policies that initially facilitate growth could in fact inhibit growth at later stages of economic development.

⁴Hart (1983) argues, for instance, that competition will reduce managerial slack, while Schumpeter (1943) suggests that the appropriability effect would reduce incentives to innovate. These models suggest behavioral changes within firms. Heterogeneous firm models based on Melitz (2003) also suggest that average export quality could rise in response to trade liberalization because less productive firms are driven out of the market. While selection may be one mechanism, we note that it cannot explain the nonmonotonic results that we find, which are consistent with the models by Aghion and coauthors.

Received for publication May 4, 2010. Revision accepted for publication October 28, 2011.

* Amiti: Federal Reserve Bank of New York and CEPR; Khandelwal: Columbia Business School and NBER.

This paper was previously circulated under the title "Competition and Quality Upgrading." We thank two anonymous referees, the editor, Nina Pavcnik, David Weinstein, and participants at seminars at the Federal Reserve Bank of New York, Yale University, Geneva Graduate Institute, and the European Research Workshop in Madrid for many useful suggestions. We thank Spencer Amdur for excellent research assistance. All errors are our own. The views expressed in this paper are our own and do not necessarily represent those of the Federal Reserve Bank of New York.

endowments of skilled workers or changes in countries' institutional structures. To assess the nonmonotonic relationship hypothesized in recent theories, we allow the relationship between tariffs and quality upgrading to depend on the proximity of the product to the world quality frontier.

Our analysis generates a number of new findings. First, we demonstrate the importance of including controls for country-specific variables that could be correlated with import tariffs. In a linear specification between quality change and tariffs, we show that omitting country-year fixed effects flips the sign on the tariff coefficient and therefore can lead to incorrect conclusions about the association between these two variables. With the controls, we find a negative relationship between tariffs and quality upgrading, which suggests that lower tariffs result in faster quality upgrading.⁵ Second, we provide strong evidence in support of a nonmonotonic relationship. We show that low import tariffs promote quality upgrading of products that are initially close to the world technology frontier, whereas lower tariffs discourage quality upgrading of products that are distant from the world frontier. Third, we find that this nonmonotonic relationship holds across countries that span a wide income distribution, provided that the country has reached a minimum level of institutional quality.⁶ This result is intuitive, given that in countries with multidimensional sources of market frictions, changes in import tariffs are likely to have limited effects on the competitive pressures that domestic firms face. Thus, the results suggest that a minimum institutional quality may be needed for the mechanisms of the model to operate. The results are robust to specifications that deal with potential endogeneity concerns regarding the tariffs, including using worldwide export growth as a proxy for productivity shocks and using the end of the Multifiber Arrangement (MFA) in 2005 as a natural experiment that resulted in shocks to certain textile and clothing products in which China had binding quotas.

Our paper builds on and contributes to a number of literatures. The international trade literature on quality has established that improved market access can promote quality upgrading. Verhoogen (2008) shows that exporters upgrade quality in response to a currency depreciation in Mexico, while Bustos (2011) finds that a regional free trade agreement led Argentine firms to invest in new technologies. Both papers have identified an important channel that affects quality upgrading, but neither focuses on the competition mechanism at the core of our analysis.⁷ The international trade literature on tariffs and productivity has shown a positive relationship between tariffs and total factor productivity (see Pavcnik, 2002; Amity & Konings, 2007; and Topalova & Khandelwal,

2011, for firm-level studies within countries, and Romalis, 2007, for a cross-country analysis). In our paper, we focus on quality upgrading, which may be viewed as a specific component of TFP and for the first time allow the potential nonmonotonic relationship to depend on the proximity to the world quality frontier using tariffs as the competition measure. Our results highlight that the distance from the world frontier is important for understanding a country's performance following trade liberalization.

Our paper is also related to studies in industrial organization that examine the relationship between competition, such as foreign entry or regulation, on related but distinct outcomes, such as research and development, patents, or total factor productivity. These studies have found support for the hypothesized nonmonotonic relationship using either aggregate cross-country analysis (Vandenbussche, Aghion, & Meghir, 2004; Acemoglu et al., 2006) or industry studies within one country (ABBGH, 2005; ABGHP, 2009). Our paper differs from these studies in a number of ways. First, the various competition measures (for example, delicensing, FDI policy, Lerner indices) and outcome measures (for example, patents, citations, TFP) used across these studies make it difficult to draw general conclusions. Here, we use quality and tariff measures that are comparable over time and space, which allows us to make comparisons across countries. Moreover, because of the high level of disaggregation, we can maintain the advantages of industry-level studies, which can better isolate how changes in economic environments affect outcomes while maintaining a broad set of countries that span the income distribution.

The remainder of the paper is as follows. In section II, we describe the distance-to-frontier models that serve as the basis for our empirical specification. In section III, we outline our empirical strategy and the methodology used to infer product quality. In section IV, we present the results, and in section V, we conclude.

II. Theoretical Background

Aghion and Howitt (2005) provide an overview of a number of models that analyze how competition affects innovation, which we draw on for our empirical analysis. The key insight behind these models is that the relationship hinges critically on the incumbent's position relative to the world technology frontier. The basic intuition behind this class of models can be illustrated using one of the simpler models in ABGHP (2009). The model assumes there are two firms capable of producing an innovation for each intermediate input under Bertrand competition. In each period and in each of the intermediate sectors, there is a potential entrant who can pay an entry cost to enter. If entry occurs, it takes place at the frontier; the entrant captures the entire market and becomes the new leading firm. Because the entrant observes the postinnovation technology, it would not pay the entry cost if the incumbent was at the frontier after innovation since Bertrand competition would drive profits to zero.

⁵ That is, without the additional controls, one would conclude support for Shumpeterian models, while with the controls, the results are supportive of the procompetitive models of Hart (1983) or Melitz (2003), for example.

⁶ As discussed in more detailed below, we rely on the World Bank's "Doing Business" (2005) report to infer a country's business climate.

⁷ A related paper by Teshima (2008) finds that lower tariffs increase research and development spending that reduces costs in Mexico.

The incumbent laggard never invests in innovation because at best it would catch up to its rival and earn zero profits. In this model, the cost of entry is an exogenous parameter that determines the extent of competition that incumbents face.

ABGHP show that increasing competition reduces innovation if the firm is far from the frontier because of a “discouragement” or “appropriability” effect. That is, firms far from the frontier know they cannot survive increased competition even if they successfully innovate. As a result, any policies that promote competition, such as lowering the entry cost, will discourage these lagging firms from spending resources on innovation. In contrast, policies that increase the threat of competition stimulate innovative activity for firms at the technology frontier. The reason is that successful innovation enables the incumbent leader to escape from the the threat of entry. ABGHP refer to this as the escape-entry effect. That is, leading incumbent firms can avoid losses that would result from increased competition by successful innovation. Thus, the implications of the model are that the relationship between innovation and competition will depend, in a non-monotonic way, on firms’ proximity to the world technology frontier:

$$\text{innovation} = f(\text{competition}, \text{proximity to the frontier}). \quad (1)$$

The assumptions can vary across these models. For example, the model in AHHV (2001) does not feature entry but instead focuses on the incentives to innovate among existing firms as product market competition (for example, through antitrust policy) changes. Nevertheless, the core insight that generates the nonmonotonic hypotheses is a distinct feature of this class of models that we use to guide our empirical specification.

III. Empirical Specification

We explore whether the nonmonotonic relationship hypothesized between competition and innovation holds for import competition and quality upgrading. We begin by describing both measures used in the analysis.

A. Import Competition

To measure a country’s import competition, we collect disaggregated import tariffs for each country in our sample. The tariff data are obtained from the World Bank’s World Integrated Trade Solution (WITS) and are specified at the Harmonized System (HS) six-digit level and over time. That is, we measure the import competition faced by an HS six-digit industry in South Korea by South Korea’s tariffs on imports in that industry. The advantage of using tariffs as our measure for competition within a country is that they are readily available at a disaggregated level and comparable across

countries and time.⁸ Importantly, there is widespread evidence that tariff reductions result in procompetitive pressures in the liberalizing countries, which result in both a reallocation of resources toward more competitive firms and exit of inefficient firms, and a reduction in markups.

B. Quality

The innovation in the theories described in section II could comprise a broad range of activities, including patenting, product creation, new processes, and investments in research and development. In this paper, we focus on one element of innovative activity, product quality, using an approach developed in Khandelwal (2010) using export data to the United States.⁹ The advantage of our approach is that we are able to extract measures of quality that are comparable across countries and over time. We rely on a country’s exports to the United States rather than its production to infer quality because the trade data are available at a highly disaggregated level, which is important for our analysis, and because these data are comparable across countries and time. We are likely to capture the highest-quality products within a country given the evidence that exporting firms tend to be more productive, employ higher-skilled workers, obtain more International Organization Standard (ISO) certifications, and produce higher-unit-value products relative to nonexporters (see Bernard et al., 2007; Verhoogen, 2008; Kugler & Verhoogen, 2012). There is also evidence that higher-unit-value goods are exported to higher-income countries (Hallak, 2006; Bastos & Silva, 2009; Manova & Zhang, 2009), and so exports to the United States are likely to be among the highest-quality products.

C. Methodology for Measuring Quality

We estimate a product’s quality from both export prices and market share information, following Khandelwal (2010). This is in contrast to the literature in international trade that often uses prices or unit values (value divided by quantity) as a proxy for quality (Schott, 2004; Hallak, 2006).¹⁰ The obvious advantage of using unit values is that they are easily calculated in the trade data. However, if products possess both

⁸ Nontariff barriers also affect import competition, but these data are notoriously difficult to obtain across all countries and over time. Some countries, particularly China and Mexico, rely on export processing zones where firms receive special exemptions, such as duty-free imports of intermediate inputs. The presence of such export processing zones is likely to lead to a downward bias on the tariff coefficient. We demonstrate that the magnitudes are larger for OECD countries, where export zones are less prevalent, and we include robustness checks that control for tariffs on intermediate inputs.

⁹ These different forms of innovation are often isomorphic. For example, in the class of models that Burstein and Melitz (2011) studied, innovation to improve productivity is equivalent to innovation that improves product quality.

¹⁰ Hummels and Klenow (2005) and Hallak and Schott (2011) also extract quality measures from using price and quantity information in trade data, but they cannot produce quality estimates at the level of disaggregation required for this study. Boorstein and Feenstra (1991) provide a related approach to infer the quality of a basket of goods by comparing the changes in quantity-weighted unit value index to an ideal price index, but this approach assumes that the quality of any individual variety is constant over time.

vertical (for example, comfort) and horizontal (for example, style) attributes, unit values may be inappropriate proxies for quality. For example, consider women’s trousers, defined at the HS 10-digit level (HS 6204624020), exported to the United States in 2005 by India and Venezuela. The unit values (inclusive of transportation and tariff costs) associated with these imports were \$140 and \$163, respectively. Under the price-equals-quality assumption, Venezuelan trousers would be assigned higher quality. However, the income per capita of Venezuela exceeds India’s by tenfold, and so it is possible that the differences in unit values also reflect in part the wage differential. Our measure of quality also takes into account differences in market shares; thus, for two products with identical unit values, the product with a higher market share is assigned a higher quality (how much higher quality depends, as seen below, on the slope of the demand schedule). Indeed, India exported over 1 million units more than Venezuela, and after accounting for these differences in market shares, the methodology we describe assigns a higher quality to Indian trousers despite lower prices.

To estimate quality, we use a nested logit demand framework, based on Berry (1994). In this framework, quality is defined as the vertical component of the model and has a structural definition as the mean valuation that U.S. consumers attach to an imported product. The intuition behind this approach is that higher quality is assigned to products that have higher market shares, conditional on prices. We closely follow the setup in Khandelwal (2010) and summarize the estimation procedure here.

To understand the nested logit structure, first we need to describe how the data are classified. We define products as the HS ten-digit codes, the most disaggregated U.S. trade data classification. A U.S. import from a country within a product is called a variety. All products can be mapped into a coarser five-digit SITC (revision 3) classification code, which we refer to as the industry. For example, an industry may be men’s knit shirts, and within this industry, shirts are classified into products that can vary by fabric material (cotton, wool, and so on). Chinese cotton and Japanese wool shirts are examples of varieties. We use the HS ten-digit products as the nests for our application. As shown below, the nested logit allows more plausible substitution patterns than the logit by allowing differences in the correlation among consumer preferences for varieties within a nest than for varieties across nests.

We derive the structural equation for a single SITC industry, comprising many varieties, and then estimate this equation separately for each industry (thus, we suppress industry subscripts). Consumer n has preferences for HS product h imported from country c (variety ch) at time t . The consumer purchases the one variety that provides her with the highest indirect utility, given by

$$V_{ncht} = \lambda_{1,ch} + \lambda_{2,t} + \lambda_{3,cht} - \alpha p_{cht} + \sum_{h=1}^H \mu_{nht} d_{ch} + (1 - \sigma)\epsilon_{ncht}. \tag{2}$$

The λ terms represent the variety’s valuation that is common across consumers (notice that these terms are not subscripted by n). The first term, $\lambda_{1,ch}$, is the time-invariant valuation that the consumer attaches to variety ch . The second term, $\lambda_{2,t}$, controls for secular time trends common across all varieties. The $\lambda_{3,cht}$ term is a variety-time deviation from the fixed effect that consumers observe but that we as econometricians do not. Consequently, this last component of quality is potentially correlated with the variety’s unit value inclusive of transportation and tariff costs, p_{cht} .

The horizontal component of the model is captured by the random component, $\sum_{h=1}^H \mu_{nht} d_{ch} + \epsilon_{ncht}$. The term ϵ_{ncht} is assumed to be distributed type 1 extreme value and explains why a low-quality variety that is expensive is ever purchased. The former term interacts the common valuation that consumer n places on all varieties within product h , μ_{nht} , with a dummy variable d_{ch} that takes a value of 1 if country c ’s export lies in product h . This term generates the nesting framework because it allows consumer n ’s preferences to be more correlated for varieties within product h than for varieties across products.¹¹ For instance, a consumer who prefers Japanese wool shirts is more likely to prefer other wool shirts than cotton shirts. The nested logit is designed to capture this preference structure.

An outside-variety completes the demand system. The outside option allows consumers to choose a domestically produced variety instead of any imported variety. The consumer chooses this outside option if the utility derived from the outside variety exceeds that from purchasing any inside option. The utility of the outside variety is given by

$$u_{n0t} = \lambda_{1,0} + \lambda_{2,t} + \lambda_{3,0t} - \alpha p_{0t} + \mu_{n0t} + (1 - \sigma)\epsilon_{n0t}. \tag{3}$$

The mean utility of the outside variety is normalized to 0; this normalization anchors the valuations of the inside varieties. In this context, one can think of the outside variety as the domestic substitute for imports, and we therefore set the outside variety market share to 1 minus the industry’s import penetration.¹² Note that the choice of the outside variety proxy affects the absolute growth rate of import qualities but not the relative growth rate because our analysis includes year fixed effects that are common to all varieties. Once the outside variety market share s_{0t} is known, we can compute the industry size: $MKT_t = \frac{\sum_{ch \neq 0} q_{cht}}{1 - s_{0t}}$, where q_{cht} denotes the

¹¹ As discussed in Berry (1994), Cardell (1997) has shown that the distribution of $\sum_{h=1}^H \mu_{nht} d_{ch}$ is the unique distribution such that if ϵ is distributed extreme value, then the sum is also distributed type 1 extreme value. The degree of within-nest correlation is controlled by $\sigma \in (0, 1]$ and is assumed to be identical across all products. As σ approaches 1, the correlation in consumer tastes for varieties within a nest approaches 1; as σ tends to 0, the nested logit converges to the standard logit model.

¹² We obtain import penetration data from Bernard, Jensen, and Schott (2006) for 1990 to 1996. We update the series through 2005 by combining U.S. export, import, and production data from the U.S. Census. Since the United States switched its industry classification from SIC to NAICS in 1997, we use the concordance developed by Pierce and Schott (2012) to link the two series together.

import quantity of variety ch . The market shares for imported varieties are then calculated as $s_{cht} = q_{cht}/MKT_t$.

The consumer chooses variety ch if $V_{ncht} > V_{nc'h't}$. Under the distributional assumptions for the random component of consumer utility, Berry (1994) shows that the demand curve from the preferences in equation (2) is

$$\ln(s_{cht}) - \ln(s_{0t}) = \lambda_{1,ch} + \lambda_{2,t} - \alpha p_{cht} + \sigma \ln(vs_{cht}) + \lambda_{3,cht}, \quad (4)$$

where vs_{cht} is variety ch 's share within product h at time t (the nest share).¹³

The trade data do not record detailed characteristics of varieties, so we exploit the panel dimension of the data by specifying a time-invariant component of quality ($\lambda_{1,ch}$) with variety fixed effects and the common quality component ($\lambda_{2,t}$) with year fixed effects. The third component of quality, $\lambda_{3,cht}$, is not observed and plays the role of the estimation error.

Since $\lambda_{3,cht}$ and the nest share are potentially correlated with the variety's price, instrumental variables are required to identify the parameters. We instrument the price with the variety's transportation costs, which are obviously correlated with prices but may also be correlated with quality if firms ship higher-quality goods in order to lower their per unit transport costs. This practice potentially raises concerns that trade costs may be correlated with a variety's quality (Hummels & Skiba, 2004). However, the exclusion restriction remains valid as long as transport costs do not affect the period-specific deviations ($\lambda_{3,cht}$) from a variety's average quality ($\lambda_{1,ch}$). In other words, if an Australian firm chooses on average to export a high-quality variety to the United States because of distance, the instruments remain valid as long as shocks to transportation costs do not affect the deviations from the firm's average quality choice. Indeed, the Washington apples phenomenon discussed in Hummels and Skiba (2004) identifies the impact of distance on prices using cross-country variation in distance rather than shocks to transport costs over time. We also include exchange rates and the interaction of distance to the United States with oil prices as additional instruments, which vary at the country-year level. Finally, vs_{cht} is also endogenous, and so we instrument this term with the number of varieties within product h and the number of varieties exported by country c . The identification assumption is that the entry and exit of other varieties will be correlated with ch 's share within the nest but uncorrelated

¹³ If one adopts a logit demand system, the nest share disappears from equation (4). To understand why this nest share term is important for inferring quality, consider the following example. Imagine there are two shirts—Japanese wool and Italian cotton—that are identical in every dimension (including prices) and evenly split the market. We would infer their qualities also to be equal. Now suppose an identical Chinese cotton shirt enters, and the market shares for the cotton shirts are 25% each and the wool shirt captures the remaining 50%. Without the nesting structure, we would infer that the quality of the Italian cotton shirt has fallen in half (since its market share has fallen while its price remains the same), even though its underlying attributes have not changed. The nested logit takes into account the correlated preferences within nests. So although the market share for the Italian cotton shirt falls, its nest share (vs_{cht}) also falls, and so its inferred quality would remain unchanged.

with quality shocks. This would occur in a model of monopolistic competition where all varieties are atomistic, or in an oligopoly model where entry and exit decisions occur in the first stage and Nash prices and qualities are chosen in the second stage of the game.¹⁴

A second issue that arises in estimating equation (4), first noted by Feenstra (1994) and also by Hallak and Schott (2011), is the problem of unobserved, or "hidden," varieties. To understand how hidden varieties could confound the measurement of quality, suppose that the reason India exported far more women's trousers than Venezuela was simply that India exported more unobserved twelve-digit HS varieties (for instance, more colors). If the Venezuelan and Indian varieties were identically priced with equal market share, then when aggregating to the observed ten-digit HS level, we would assign a larger market share to the Indian varieties. From equation (4), India's estimated quality would be biased upward simply due to the hidden varieties. Drawing on standard models (Krugman, 1980) that predict that the number of varieties produced is increasing in a country's population, we use (the log of) population as an additional covariate in equation (4).¹⁵

The demand curve that controls for the hidden-varieties problem is given by

$$\ln(s_{cht}) - \ln(s_{0t}) = \lambda_{1,ch} + \lambda_{2,t} - \alpha p_{cht} + \sigma \ln(vs_{cht}) + \gamma \ln pop_{ct} + \lambda_{3,cht}, \quad (5)$$

where pop_{ct} is the population in country c at time t . The estimated parameters and the residual of the regression define the quality of variety ch at time t as

$$\lambda_{cht} \equiv \hat{\lambda}_{1,ch} + \hat{\lambda}_{2,t} + \hat{\lambda}_{3,cht}. \quad (6)$$

From equation (5), we see that the quality of an imported variety is defined relative to its market share after controlling for exporter size and price. More generally, our notion of quality is an attribute that allows a variety's price to rise without losing market share. One potential concern about this interpretation is that many factors unrelated to quality could affect market shares and therefore confound our measure of quality. However, it is important to note that this set of factors is made much smaller by conditioning on prices. For example, a

¹⁴ We note that the validity of using a count of varieties to instrument for vs_{cht} relies on weaker assumptions than those typically made in the discrete choice literature, which typically instruments vs_{cht} with the average characteristics of varieties (Berry, Levinsohn, & Pakes, 1995). This practice assumes that the firms' quality choices are fixed (or chosen before prices). Here, we only need that the number of varieties is uncorrelated with the deviation from average quality, $\lambda_{3,cht}$. This will be the case in a model where entry and exit occur prior to the firms' quality choice. For example, this occurs in models where firms choose quality in the final stage of a multistage game of location, price, and quality decisions (Vogel, 2008).

¹⁵ The specification constrains the coefficient on population to be the same for all countries within an industry. If the true elasticity of variety growth with respect to population differs across countries, this will be absorbed by our measure of quality. In principle, one could allow this coefficient to be country specific if there are enough varieties exported from a country within each industry.

variety may have a large market share if the exporting country is geographically close to the United States. However, since the price includes transportation costs, the quality estimate is not capturing purely gravity effects such as distance. We note that defining quality to be inclusive of a residual is analogous to the productivity literature that interprets total factor productivity as the residual from conditioning output on observable inputs. A regression of λ_{cht} on year fixed effects yields an adjusted R^2 of 0.01, and adding the variety fixed effects increases the adjusted R^2 to 0.77. This implies that the residual ($\hat{\lambda}_{3,cht}$) accounts for 23% of the variation in the quality estimate.

Proximity to frontier. We estimate equation (5) separately for each five-digit SITC (revision 3) industry and use the estimated parameters of the regressions to define the qualities according to equation (6).¹⁶ We construct the frontier measures by first taking a monotonic transformation of the quality measures to ensure that all qualities are nonnegative: $\lambda_{cht}^F = \exp[\lambda_{cht}]$.

We define a variety's proximity to the frontier as the ratio of its (transformed) quality to the highest quality within each HS ten-digit product:

$$PF_{cht} = \frac{\lambda_{cht}^F}{\max_{c \in ht} (\lambda_{cht}^F)}, \quad (7)$$

where the max operator chooses the maximum λ_{cht}^F within a product-year and $PF_{cht} \in (0, 1]$. For varieties close to the frontier, this measure is close to 1. For varieties far from the frontier, this measure is close to 0.

IV. Data Description

To estimate quality, we use annual export data to the United States from 1990 to 2005 at the HS ten-digit level for the manufacturing sector (SITC 5–8). Since unit values are notoriously noisy (General Accounting Office, 1995), prior to estimating the demand systems in equation (5), we trim the data along three dimensions: we drop variety-year observations above or below the 1st and 99th percentile of unit values, exclude varieties with annual price increases of more than 200% or price declines of more than 66%, and drop varieties with export quantities of fewer than ten. The quality estimates obtained from equation (6) are also noisy, and so we trim the qualities at the 5th and 95th percentiles. We also drop any observations with five-year quality growth outside the 1st and 99th percentiles. We trim along five-year intervals since our dependent variable will be defined as quality growth over five-year intervals.

We obtain six-digit HS import tariffs for 56 countries for 1990, 1995, and 2000 from the World Bank WITS database.

¹⁶ Thus, the coefficients that we estimate from equation (4) are industry specific. Since we run the estimation separately, quality cannot be compared across industries. We include appropriate fixed effects in our analysis below to account for this.

TABLE 1.—SUMMARY STATISTICS

Variables	OECD	Non-OECD	High DB	Low DB
$\Delta Quality_{cht}$	0.277 (1.116)	0.405 (1.074)	0.287 (1.111)	0.449 (1.063)
PF_{cht-5}	0.508 (0.351)	0.559 (0.332)	0.520 (0.347)	0.557 (0.333)
$Tariff_{c,h6,t-5}$	0.076 (0.061)	0.251 (0.262)	0.091 (0.104)	0.310 (0.280)
Countries	20	36	29	27
Observations	74,053	57,204	94,274	36,983

Table reports summary statistics of changes in quality, lag proximity to frontier and lag tariffs. Mean values are reported with standard deviations in parentheses. DB refers to the World Bank's "Doing Business" report (2005), and we split countries between those above and below the median values. Countries above the median DB in our sample are Australia,* Austria,* Belgium/Luxembourg,* Canada,* Chile, Denmark,* Finland,* France,* Germany,* Hong Kong, Ireland,* Italy,* Japan,* Malaysia, Netherlands,* New Zealand,* Nicaragua, Norway,* Poland, Singapore, South Africa, South Korea, Spain,* Sweden,* Taiwan, Thailand, Tunisia, Turkey,* and United Kingdom*. Countries below the median DB in our sample are Argentina, Bangladesh, Brazil, China, Colombia, Costa Rica, Egypt, El Salvador, Greece,* Guatemala, Honduras, India, Indonesia, Kenya, Mexico, Morocco, Nepal, Pakistan, Paraguay, Peru, Philippines, Portugal,* Saudi Arabia, Sri Lanka, Uruguay, Venezuela, and Vietnam. Asterisks denote OECD countries.

The sample of countries is limited by the availability of tariff data for those years.¹⁷ There is wide variation in the tariff levels across countries and over time. The average tariff across countries ranges from very low rates of 0% in Hong Kong and Singapore and 5% in Norway to very high rates of 100% in Bangladesh and 63% in Morocco and Pakistan. Between 1990 and 2000, average tariffs for OECD countries fell from 16% to 12%, whereas for non-OECD countries, they fell from 29% to 18%. The world quality frontier for each product in each year is defined from the set of countries for which we have tariff data. The proximity to the frontier for each country's products in each year is then matched to its import tariffs. Table 1 reports summary statistics of the change in quality, proximity to frontier, and tariff levels for OECD and non-OECD countries,¹⁸ as well as statistics for countries classified as having strong and weak business environments (high doing business and low doing business), which we define in section V. The table shows that non-OECD countries have faster quality growth than OECD countries and higher rates of protection. The table also shows that non-OECD countries have slightly higher proximity to frontier measures than OECD countries, but this is related to product composition. Controlling for product-year fixed effects, there is a positive and statistically significant correlation between proximity to the frontier and income per capita.

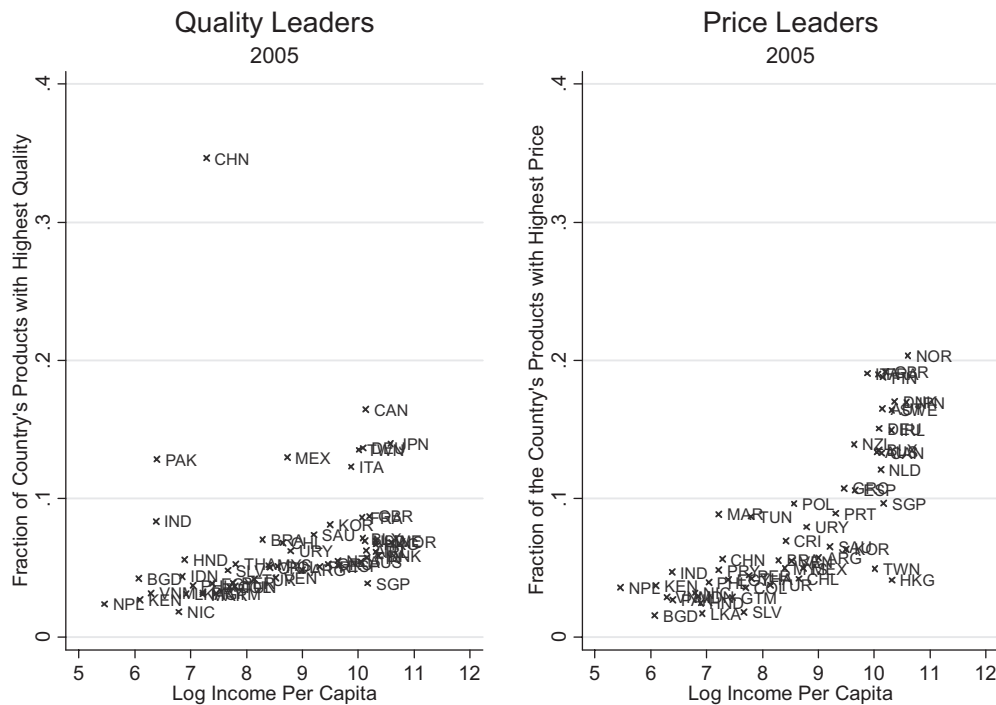
As would be expected, the quality estimates indicate that richer countries export higher-quality varieties within products.¹⁹ Thus, on average, more advanced countries sit atop a product's quality ladder, while developing countries are further from the frontier. The relationship between income and quality in 2005 is depicted in figure 1. The left panel of figure

¹⁷ If tariff data were unavailable for a particular year, we included the data for the preceding year. Note that tariffs are common for all countries within the European Union.

¹⁸ OECD countries include all those that joined the OECD before our sample period in 1990. Countries in our sample that joined the OECD after 1990 are Mexico in 1994 and South Korea and Poland in 1996. Taiwan is a member but cannot vote.

¹⁹ This is shown in Khandelwal (2010) and is consistent with findings in the previous literature.

FIGURE 1.—PRICE AND QUALITY LEADERS



1 plots the proportion of the total number of products a country exports for which it is the quality “leader,” $PF_{cht} = 1$, against its income per capita, showing a positive and statistically significant relationship.²⁰ Similarly, there is a positive relationship between income and the fraction of the highest-priced varieties in the right panel of the figure. Notice in this panel that the positive relationship is steeper than with the quality-based measure.

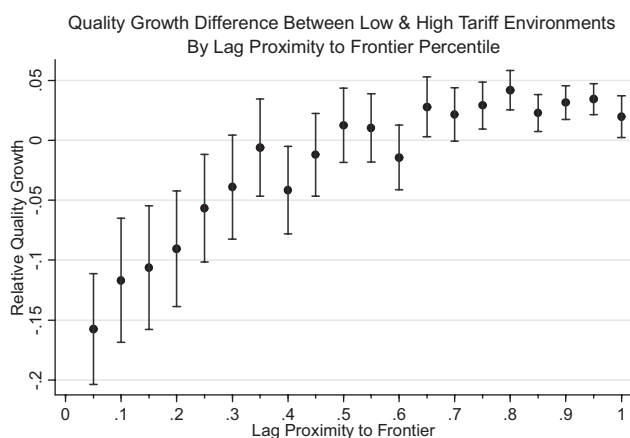
It is clear from the graph that China is an outlier. Although China exported the highest-priced variety in 9% of products in 2005, the quality-based measure indicates that China was the leader in 44% of the total number of products it exported to the United States. There are several reasons for this discrepancy. First, although China exports low-priced varieties, it has exceptionally high market shares (it has the highest quantity in 59% of the products it exports), particularly for labor-intensive products. That is, the procedure above yields high-quality estimates for China because its market shares are larger than the predicted market shares given its price and the estimated elasticity of demand. Thus, the methodology will record higher quality for China in these products. Second, trade statistics are recorded on a total value basis rather than a value-added basis. Given the importance of processing trade for Chinese exports, its value added will vary across sectors. For example, the Apple iPod is “made in China” even though China’s value added accounts for a fraction of the production (Linden, Kraemer, & Dedrick 2007). More generally, Koopman, Wang, and Wei (2008) estimate that China’s value

added in computers may be as low as 5%. If the U.S. Census collected value-added trade data, China’s inferred quality would presumably be much lower. In section VC, we adopt several robustness checks in our analysis, including using unit values as a proxy for quality and excluding China from the analysis and the frontier definition.

Before turning to the estimation, we preview the data to see how the quality distribution changes with tariffs over time without imposing functional forms on the relationship. For varieties close to the quality frontier, we would expect that quality growth would be relatively higher in a low-tariff environment compared to those in a high-tariff environment. In contrast, for varieties distant to the quality frontier, we would expect that quality growth would be relatively lower in a low-tariff environment compared to those in a high-tariff environment. To check for this, we plot the difference in the quality growth for low tariffs versus high tariffs, dividing the data into twenty bins of the proximity-to-frontier measure. High tariffs are defined as those above the mean and low tariffs as those below the mean, and we de-mean the quality estimates to make them comparable across products and time by regressing quality growth on country-year and product-year fixed effects and obtaining the residual. Figure 2 shows an upward-sloping curve, with negative relative growth for varieties distant from the frontier and positive relative growth for varieties close to the frontier. This implies that the quality distribution for low-tariff varieties spreads out over time relative to high-tariff environments. Indeed, the standard deviation of quality growth in low-tariff environments is 0.69 compared to 0.63 in high-tariff environments.

²⁰ The income per capita and population variables are obtained from the the World Development Indicators database.

FIGURE 2.—RELATIVE QUALITY GROWTH IN LOW- VERSUS HIGH-TARIFF ENVIRONMENTS, BY PROXIMITY



Change in quality is demeaned with country-year and product-year fixed effects. Dashed lines denote 5% confidence intervals.

V. Quality Upgrading and Import Competition Results

With the import tariffs and quality measures in hand, we can analyze the effect of competition on quality upgrading, allowing for the discouragement and escape-competition forces. We use the following empirical specification to relate quality growth to import tariffs, proximity to the frontier, and the interaction of the two, which allows for the nonmonotonic relationship highlighted in ABGHP (2009):²¹

$$\Delta \ln \lambda_{cht}^F = \alpha_{ht} + \alpha_{ct} + \beta_1 PF_{cht-5} + \beta_2 \text{tariff}_{cht,t-5} + \beta_3 (PF_{cht-5} \times \text{tariff}_{cht,t-5}) + \varepsilon_{cht}. \quad (8)$$

The dependent variable, $\Delta \ln \lambda_{cht}^F$, is the change in a variety's quality between period t and $t - 5$. All the explanatory variables are in levels for the period $t - 5$. Our specification includes both product-year fixed effects (α_{ht}) and country-year fixed effects (α_{ct}), which are critical to the analysis. The product-year fixed effects deal with two issues. First, because the qualities are estimated separately across industries, the quality estimates are comparable only within the industry or product. Including the product-year fixed effects ensures that the estimation exploits the variation only between comparable qualities. Second, product-year fixed effects control for shocks that are common to all varieties within a product such as demand shocks or worldwide technology shocks that could also influence quality upgrading. The country-year fixed effects sweep out country-level shocks such as technological shocks, changes in relative endowments, and changes in institutions that affect competition. Thus, this specification flexibly controls for different shocks that may be correlated with tariff changes and affect quality growth.

The ABGHP model suggests that $\beta_2 > 0$ and $\beta_3 < 0$. Thus, a fall in tariffs would spur a variety's quality growth in subsequent periods only if the product variety is close to the world quality frontier (PF_{cht-5} close to 1); this is consistent

²¹ ABGHP (2009) specifies a similar estimating equation in their context.

TABLE 2.—QUALITY UPGRADING, COMPETITION, AND DISTANCE TO FRONTIER

Regressors	(1)	(2)	(3)	(4)
PF_{cht-5}			-0.823***	
$Tariff_{c,h6,t-5}$	0.207***	-0.070**	0.214***	
$PF_{cht-5} \times Tariff_{c,h6,t-5}$	0.014	0.033	0.048	
OECD Indicator interacted with:				
PF_{cht-5}				-0.737***
$Tariff_{c,h6,t-5}$				0.021
$PF_{cht-5} \times Tariff_{c,h6,t-5}$				0.416**
Non-OECD Indicator interacted with:				
PF_{cht-5}				0.165
$Tariff_{c,h6,t-5}$				-0.973***
$PF_{cht-5} \times Tariff_{c,h6,t-5}$				0.220
Product-year FEs	Yes	Yes	Yes	Yes
Country-year FEs	No	Yes	Yes	Yes
R^2	0.54	0.54	0.58	0.58
Observations	131,257	131,257	131,257	131,257

Table reports regression results of change in (log) quality of a variety on the lag HS6-level tariff faced in the origin country, the varieties lag proximity to frontier, and the interaction. Column 1 reports quality growth on tariffs. Column 2 introduces country-year fixed effects. Column 3 reports the baseline specification with the interaction between proximity to frontier and tariffs. Column 4 estimates separate coefficients for the OECD and non-OECD countries (the OECD dummy is not reported). All regressions include product-year fixed effects. Standard errors clustered by exporting country (with EU countries treated as one country because of its common trade policy). Significance: *10, **0.05, and ***0.01.

with the escape competition effect already discussed. In contrast, if a product variety's quality is a long way from the frontier, a fall in tariffs could reduce quality upgrading due to the discouragement effect. That is, products a long way from the frontier need high tariffs to protect rents in order to promote quality upgrading. Note that $\beta_1 < 0$ implies that varieties that are far from the frontier (PF_{cht-5} close to 0) experience faster quality upgrading, implying convergence in quality.

A. Baseline Results

Before estimating equation (8), we first look for a monotonic relationship between competition and quality growth by regressing the growth in a variety's quality on the home market's import tariffs and product-year fixed effects, as in the trade and growth literature (Romalis, 2007). The first column of table 2 shows that a fall in tariffs is associated with slower quality upgrading. However, once we include country-year fixed effects in column 2, to control for factors such as changes in a country's relative factor endowments or technology shocks, the tariff coefficient switches sign and is now negative, indicating that a fall in tariffs is associated with faster quality upgrading. Thus, the table suggests that the relationship between import competition and quality upgrading may be confounded by unobservables. In particular, when country-year fixed effects are excluded, the results support the appropriability argument in Schumpeter (1943), but this finding reverses in favor of arguments related to

Hart (1983) once we control for country-level changes in economic environments. These results highlight the importance of controlling for country-year effects that may be correlated with industry-level competition measures such as tariffs. In all subsequent regressions, we therefore include both country-year and product-year fixed effects.

Next, we examine whether the relationship between quality upgrading and tariffs depends on a variety's proximity to the frontier according to baseline equation (8). Column 3 in table 2 shows a negative coefficient on the lag proximity to the frontier, which implies a faster catch-up for varieties far from the frontier. The positive coefficient on tariffs and the negative coefficient on the interaction of tariffs with the proximity to frontier provide support for the effects highlighted in ABGHP (2009). The negative coefficient on the interaction implies that the varieties close to the world frontier are more likely to upgrade quality in response to tougher competition in the domestic market (the escape-competition effect). And the positive coefficient on the linear tariff variable implies that tariffs are likely to have the opposite effect for varieties distant from the world frontier (the discouragement effect).

In column 4, we examine heterogeneity in the discouragement and escape-competition effects by allowing separate effects for OECD and non-OECD countries. The results hold across both groups, but the magnitudes of the tariff coefficients are much larger for OECD countries. For OECD varieties that are distant from the frontier (PF_{cht-5} close to 0), a 10 percentage point fall in tariffs is associated with a 4.2% fall in quality upgrading. However, for OECD varieties close to the frontier, a fall in tariffs has the opposite effect: a 10 percentage point fall in tariffs is associated with a 5.6% increase in quality upgrading. For non-OECD varieties far from the frontier, a 10 percentage point fall in tariffs is associated with a 1.1% fall in quality upgrading, and for varieties close to the frontier, a 10 percentage point fall in tariffs is associated with a 1.3% growth in quality.

B. Institutions and Quality Upgrading

The results in column 4 of table 2 raise the question as to why there are larger quality responses in OECD countries than in non-OECD countries as tariffs change. For the effects in the theory to be present, market forces need to be able to operate. In particular, the potential for entry and exit of firms is crucial for tariffs to invoke more competition in the home market. However, nontariff barriers, bureaucratic red tape, and other entry regulations are likely to imply heterogeneity in the impact of tariffs on the competitive environment across countries. In countries with more regulation, additional domestic reforms may be needed for lower tariffs to induce further competition in the market.

We test for heterogeneous effects in the tariff-frontier interaction coefficient according to institutional quality in the first column of table 3. To assess the quality of a country's institutions, we rely on a measure of the regulatory environment

TABLE 3.—QUALITY UPGRADING, COMPETITION, AND INSTITUTIONS

Regressors	All Countries (1)	OECD Indicator Interactions (2a)	Non-OECD Indicator Interactions (2b)	HDB Countries Only (3)
Countries above median DB				
PF_{cht-5}	-0.769***	-0.739***	-0.836***	-0.810***
	0.017	0.021	0.027	0.019
$Tariff_{c,h6,t-5}$	0.438***	0.420**	0.413***	0.524***
	0.089	0.167	0.104	0.094
$PF_{cht-5} \times Tariff_{c,h6,t-5}$	-0.790***	-0.991***	-0.622***	-0.907***
	0.111	0.221	0.130	0.119
Countries below median DB				
PF_{cht-5}	-0.992***	-0.652***	-1.037***	
	0.026	0.114	0.027	
$Tariff_{c,h6,t-5}$	0.033	0.590	-0.017	
	0.058	0.702	0.059	
$PF_{cht-5} \times Tariff_{c,h6,t-5}$	-0.121**	-0.740	-0.051	
	0.057	0.985	0.058	
Product-year FEs	Yes		Yes	Yes
Country-year FEs	Yes		Yes	Yes
R^2	0.58		0.58	0.58
Observations	131,257		131,257	94,274

Column 1 reports regression results of change in (log) quality of a variety on the lag HS6-level tariff faced in the origin country, the varieties lag proximity to frontier, and the interaction, with each coefficient interacted with a dummy variable if the country is above (HDB) or below (LDB) the median doing business value. Panel 2 introduces an additional interaction if the country is an OECD country. Column 2a reports the OECD interactions, and column 2b reports the non-OECD interactions; these coefficients are estimated in a single regression. Column 3 reports the baseline specification for just countries above the median doing business values. See the footnote to table 1 for a list of the country classifications. All regressions include product-year and country-year fixed effects. Standard errors clustered by exporting country (with EU countries treated as one country because of the common trade policy). Significance: *.10, **.05, and ***.01.

from the World Bank's Doing Business Survey.²² The index ranges from 0 to 1, with a higher value indicating a better business environment. We separate countries into two groups, with HDB comprising countries with a doing business indicator greater than the median and LDB comprising countries with a doing business indicator lower than the median.²³ We construct dummies for HDB and LDB countries and interact these with all the key variables. Column 1 shows that for countries with weak business environments, the magnitudes and significance on the tariff variables are much lower than for countries with strong business environments, with the coefficient on the linear tariff term insignificant.

Interestingly, the business environment indicator is picking up an effect beyond differences in income per capita. To see this, we allow additional flexibility in the coefficients for strong and weak business environments further broken down by OECD and non-OECD countries in table 3, columns 2a and 2b. The results indicate that even non-OECD countries characterized by strong business environments display both the discouragement and the escape-competition forces (see the "countries above median DB" rows in column 2b). Yet

²² We construct an aggregate doing business index by following the procedure outlined in World Bank (2005). The "Doing Business" database tracks constraints along several dimensions, including the ease of starting a business, enforcing contracts, obtaining credit, and hiring and firing. We compute each country's percentile ranking for each outcome. The aggregate doing business measure takes the (simple) average of a country's percentile rankings across the outcomes. A higher value indicates an environment more conducive to conducting business.

²³ See table 1 for the list of countries classified as above and below the median doing business index.

for countries characterized by weak doing business indicators, the coefficients on the tariff variables are insignificant on both tariff terms for both OECD countries and non-OECD countries, and the coefficient on the tariff term for non-OECD countries becomes negative (see the “countries below median DB” rows in column 2b).²⁴ This result suggests that a minimum institutional “quality,” and not simply high income per capita, is required for the two forces to operate. In particular, the lack of support for the models among weaker business climate countries appears consistent with a variant on the Acemoglu et al. (2006) model that discusses how political economy factors can inhibit the escape-competition effect from operating (see section 5.2 of Acemoglu et al., 2006). Since countries with poorer business climates are unlikely to fit the theory, we restrict the subsequent analysis to the set of countries characterized by a relatively stronger business environment.

In column 3, we therefore reestimate equation (8) with only the sample of countries with business environments above the median. The results indicate that for varieties far from the frontier, a 10 percentage point fall in tariffs is associated with a 5.2% decline in quality growth, while an equivalent tariff decline for varieties close to the frontier is associated with a 3.8% increase in quality growth. To get a sense of the economic significance of these point estimates, we evaluate what a 10 percentage point change in tariffs implies for varieties close to the frontier and for those distant from the frontier and compare these predicted changes to the actual change in quality for these varieties. Thus, for varieties close to the frontier ($PF > 0.9$), the predicted mean change in quality is 3%, whereas the actual mean change in quality for these varieties is 13%. This calculation implies that a 10 percentage point change in tariffs can account for around 20% of the actual change in quality. An analogous calculation for varieties distant from the frontier ($PF < 0.1$) implies that a 10 percentage point change in tariffs can account for around 10% of the actual change in quality.

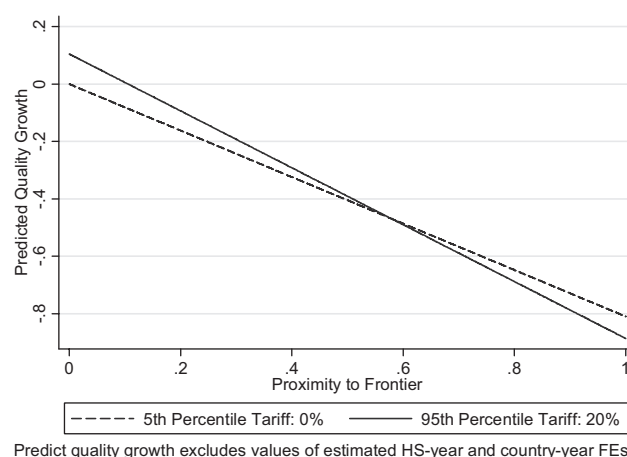
Figure 3 provides an illustration of the relationship between competition and quality upgrading for the set of countries with strong business environments, highlighting the discouragement and escape-competition effects in column 3 of table 3. The figure plots the predicted quality growth,

$$\Delta \ln \hat{\lambda}_{cht}^F = \hat{\beta}_1 PF_{cht-5} + \hat{\beta}_2 \text{tariff}_{cht-5} + \hat{\beta}_3 (PF_{cht-5} \times \text{tariff}_{cht-5}),$$

against the PF_{cht-5} , evaluated at the 5th (dashed line) and 95th tariff percentiles. The downward-sloping lines indicate

²⁴Note that there are only two countries, Greece and Portugal, in the LDB-OECD grouping. The alternative OECD group, with all current members, would add Mexico to that subgroup, which results in significant tariff terms. One explanation for this finding is that the business climate in the *maquiladora* region of Mexico, where the majority of Mexico’s exports to the U.S. originate, are not accurately reflected in the DB measure. More important, the results for all the other three groupings are unaffected by the OECD definition.

FIGURE 3.—PREDICTED QUALITY GROWTH AND PROXIMITY TO FRONTIER



convergence in the data; varieties far from the frontier experience faster quality upgrading than those that are proximate to the frontier. The predicted quality growth line evaluated at the 95th percentile tariff (a 20% tariff) is a clock-wise rotation of the 5th percentile tariff (a 0% tariff), and this reflects the two forces. For varieties far from the frontier, moving from a high tariff to a low tariff is associated with a decrease in the rate of quality upgrading, which is consistent with the Schumpeterian discouragement effect. However, for varieties close to the frontier, moving from a high to a low tariff is associated with a faster rate of quality upgrading, which illustrates the escape-competition effect.

C. Robustness

In the remaining tables, we check the robustness of the results to additional controls. One potential concern is that the proximity variable is measured with error due to randomness or outliers of the highest quality variety. In table 4, we demonstrate that our results are robust to alternative measures of the world frontier. In column 1, we check the sensitivity of the results by excluding varieties at the world frontier (and so exclude observations for which $PF_{cht-5} = 1$). In column 2, we drop the top two varieties and redefine the frontier in equation (7) using the third-highest-quality variety rather than the maximum. In column 3, we redefine the frontier based on the sample of varieties exported by HDB countries. We next redefine the frontier using qualities inferred from the data set after excluding China’s exports. Recall that in section IV, we noted that China’s export quality may be overstated because the U.S. import data record gross values rather than value added. To check that our results are not driven by this, we exclude China’s exports to the United States from the data, reestimate quality using equation (5), and redefine the proximity-to-frontier measures excluding China. In column 4, we report the baseline specification using these (China-excluded) quality measures. Across all four columns, the results remain robust to these alternative measures of the world frontier.

TABLE 4.—ALTERNATIVE PROXY TO FRONTIER MEASURES

Regressors	Exclude PF = 1 (1)	Frontier Defined after Dropping Top Two Qualities (2)	Frontier Defined on HDB Countries (3)	Exclude China from Quality Estimation (4)	Change in Quality Percentile (5)	Unit Values (6)
PF_{cht-5}	-1.025*** 0.024	-0.784*** 0.022	-0.778*** 0.017	-0.884*** 0.019	-0.263*** 0.005	-1.106*** 0.018
$Tariff_{c,h6,t-5}$	0.423*** 0.099	0.529*** 0.123	0.511*** 0.094	0.418*** 0.095	0.187*** 0.028	0.285*** 0.059
$PF_{cht-5} \times Tariff_{c,h6,t-5}$	-0.726*** 0.127	-0.774*** 0.146	-0.841*** 0.115	-0.875*** 0.117	-0.331*** 0.036	-0.496*** 0.115
Product-year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Country-year FEs	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.58	0.59	0.58	0.57	0.15	0.21
Observations	83,552	73,326	94,274	91,205	94,274	91,754

Table reports regression results of change in (log) quality of a variety on the lag HS6-level tariff faced in the origin country, the varieties lag proximity to frontier, and the interaction along with additional controls. Column 1 excludes observations with a proximity to frontier equal to 1. Column 2 removes the top two qualities from each product and redefines the proximity-to-frontier measure (that is, the third-highest quality becomes the frontier). Column 3 redefines the frontier measure using just the sample of HDB countries. Column 4 uses quality measures from estimating equation (5) excluding China and then rerunning the baseline regression (8). Column 5 uses the change in the percentile of the variety's quality, rather than the actual quality measure, as the dependent variable. Column 6 uses unit values as the proxy for quality and so it regresses the change in unit values on tariffs, a unit value proximity to frontier measure, and the interaction. All regressions include country-year and product-year fixed effects and are run on the set of high-DB countries. Standard errors clustered by exporting country (with EU countries treated as one country because of the common tariff). Significance: *.10, **.05, and ***.01.

An alternative way to construct the quality measures is to compute the percentile of a variety's quality within each product-year pair. The advantage of using the quality percentile, as opposed to the actual measure of quality, is that the variety's ranking is easier to compare across products. Rerunning the baseline specification with the change in quality percentile as the dependent variable also has the additional advantage of not being a function of the lag proximity measure on the right-hand side of the regression (we discuss this further below). Column 5 indicates that the results are robust to this alternative measure. The magnitude of the coefficient on the proximity to frontier falls as expected (but is still negative, indicating that it captures convergence), and the signs of the coefficients on tariffs and the interaction remain unchanged from the baseline specification.

In column 6, we re-estimate equation (8) using prices (unit values), the more common proxy for quality in international trade.²⁵ Specifically, we define the proximity to the frontier based on how far a variety's price is from the maximum price and define the dependent variable as the change in log prices.²⁶ Column 6 shows that the magnitudes of the coefficients are similar to our baseline estimates; for varieties close to the unit value frontier, there is a negative relationship between tariffs and subsequent price growth, and for varieties far from the frontier, there is a positive association. This result shows that the discouragement and escape competition effects appear when using prices instead of the alternative measure of quality that Khandelwal (2010) proposed.

In table 5, we address issues surrounding omitted variables and endogeneity. First, a potential concern is that firms are upgrading quality in response to lower tariffs on intermediate inputs rather than tariffs on final goods. As input tariffs fall, firms can access cheaper higher-quality inputs, which can lead to higher-quality outputs. If tariffs on intermediate

inputs and final goods are correlated, this omitted variable could bias our coefficients. In column 1 of table 5, we include input tariffs in the baseline specification and find that it has a positive coefficient, but it is statistically insignificant.²⁷ What is important is that the significance and magnitudes of all the other variables are unaffected by the inclusion of input tariffs.

Second, if products close to the quality frontier reflect a country's comparative advantage varieties, the observed mechanism underlying the quality upgrading may reflect resource reallocation from a country's comparative disadvantage to its comparative advantage sectors. In Bernard et al. (2007), liberalizing trade induces an expansion and productivity gain in countries' comparative advantage sectors because resources reallocate across and within sectors. In that model, part of the reallocation is driven by the bilateral nature of the trade liberalization, which expands firms' market access in the destination countries. Here, we focus on unilateral trade liberalization, which increases import competition. Nevertheless, to examine this alternative channel, we control for the (lag) revealed comparative advantage (RCA) of countries' sectors, which we compute at the HS6 level from the UN Comtrade database. Revealed comparative advantage is defined as the HS6 share of a country's exports divided by the HS6 share in total world exports. As shown in column 2 of table 5, the coefficient on lag RCA is positive and significant. This implies that more rapid quality growth is associated with comparative advantage sectors, as we might expect. However, even after controlling for this additional factor, the baseline results continue to hold.

Third, there may be endogeneity concerns arising from countries possibly liberalizing their industries selectively

²⁵ See Schott (2004), Hallak (2006), Baldwin and Harrigan (2007), and Johnson (2009).

²⁶ We exclude observations that report unit value changes above the 99th and below the 1st percentiles.

²⁷ Interestingly, reestimating column 1 of table 5 on the countries below the median DB yields a significant negative coefficient on input tariffs. This is consistent with Amiti and Konings (2007) and Topalova and Khandelwal (2011), who find stronger productivity effects from lowering input tariffs than by increasing competition with lower tariffs on final goods. Halpern, Koren, and Szeidl (2009) and Goldberg et al. (2010) demonstrate that lower input tariffs improve access to higher-quality foreign inputs in developing countries.

TABLE 5.—ROBUSTNESS CHECKS

Regressors	Input Tariffs (1)	RCA (2)	World Export Growth (3)	MFA (final period) (4)	MFA (all periods) (5)	IV Regression (6)	Current PF (7)
PF_{cht-5}	-0.815*** 0.019	-0.841*** 0.028	-0.806*** 0.019	-1.205*** 0.055	-1.133*** 0.040	-0.404*** 0.033	
$Tariff_{c,h6,t-5}$	0.511*** 0.100	0.604*** 0.121	0.487*** 0.095			0.508*** 0.166	0.177* 0.092
$PF_{cht-5} \times Tariff_{c,h6,t-5}$	-0.907*** 0.121	-0.934*** 0.150	-0.870*** 0.121			-0.769*** 0.210	
$World\ Export\ Growth_{c,h6,t}$			0.047*** 0.007				
$Input\ Tariff_{c,h6,t-5}$	0.055 0.169						
$RCA_{c,h6,t-5}$		0.094*** 0.015					
$PF_{cht-5} \times B_h$				0.241*** 0.089	0.090 0.069		
$PF_{cht-5} \times B_h \times PostWTO_t$					0.241*** 0.089		
PF_{cht}							0.883*** 0.019
$PF_{cht} \times Tariff_{c,h6,t-5}$							-0.160 0.111
Product-year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country-year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.58	0.61	0.59	0.62	0.65	0.07	0.58
Observations	92,612	43,393	90,868	9,423	28,444	39,532	94,274

Table reports regression results of change in (log) quality of a variety on the lag HS6-level tariff faced in the origin country, the varieties lag proximity to frontier and the interaction along with additional controls. Column 1 includes lag input tariffs as a control. Column 2 includes countries' lag revealed comparative advantage as a control. Column 3 includes a measure of a country's HS6-level world export growth. Column 4 restricts the analysis to 2000–2005 and to products that were subject to quotas under the MFA. These products are taken from Brambilla, Khandelwal, and Schott (2010). The B dummy takes a value of 1 if China's fill rate exceeded 90% in 1991 for that product; fill rate information is also taken from Brambilla et al. (2010). Column 5 includes all periods and interacts $PF \times B$ with a dummy "PostWTO" that takes a value of 1 in the period 2000–2005. We suppress the coefficient on $PF \times PostWTO$ for readability purposes (the coefficient is not statistically significant). Column 6 instruments the lag proximity to frontier with the previous period lag PF measure (the lag-lag proximity to frontier). Using the lag-lag PF measure is why the number of observations falls, and the low R^2 is because the data were de-measured prior to running the regression. Column 7 runs the baseline specification, but replaces the variety's lag proximity to frontier with its current period proximity to frontier. All regressions include country-year and product-year fixed effects and are run on the set of high-DB countries. Standard errors clustered by exporting country (with EU countries treated as one country because of the common tariff). Significance: *.10, **.05, and ***.01.

based on forces that we are unable to observe. For instance, if countries receive productivity shocks that enable them to improve the quality of their products, pressures against liberalizing those markets may subside. To the extent that these shocks are country specific, the country-year fixed effects will control for productivity shocks. Likewise, productivity (or demand) shocks that are common across all varieties within a product will be captured by the product-year fixed effects. However, productivity shocks could be country-industry specific. To address this concern, we include the change in a country's total exports to the world for each HS 6-digit industry by year, taken from Comtrade. The change in industry-level exports for each country is a plausible proxy for productivity shocks: higher productivity shocks are likely to be reflected in higher export growth. In column 3 of table 5, we see that while the change in exports is positively correlated with quality upgrading, its inclusion leaves the key results unchanged. Moreover, the magnitudes are extremely close to the baseline results reported in column 3 of table 3.²⁸

We further address potential endogeneity concerns by exploiting a specific liberalization episode: the ending of textile and clothing (T&C) quotas under the Multifiber Arrangement (MFA) in 2005. A major breakthrough in the Uruguay Round was the agreement by developed countries

to end their stringent quotas on developing countries' T&C exports. While the liberalization episode was anticipated, the quota removal was plausibly exogenous for countries that exported T&C because of the WTO mandate to end the quota regime.²⁹ In 1995, the United States published the HS product schedule to phase out the quotas over ten years. Brambilla et al. (2010) show that China's T&C exports surged following MFA quota removals (it was eligible for quota removals after joining the WTO). The surge was most pronounced in the products in which China was "bound"—products in which China's quota fill rates exceeded 90%. After the quotas were removed, China's exports of bounded products immediately increased by more than 450%. Brambilla et al. (2010) show that with few exceptions (notably Bangladesh and India), virtually all countries' T&C exports to the United States contracted because of China's export explosion.³⁰ Thus, China's exports following the end of the MFA represent a substantial increase in product market competition for T&C HS products, especially in the set of products in which China was bound.

²⁹ Nonetheless, the United States did reimpose quotas on China's exports in a handful of T&C HS codes in 2006. However, the reimposition of the quotas was due to China's export surge in 2005. The lobbying for new quotas therefore precisely reflects the substantial increase in product market competition in 2005, the last year of our sample.

³⁰ This finding underscores the point that the quotas actually ensured market access for many countries' textile and clothing exports.

²⁸ In unreported results (available on request), we also find that the results are robust to including input tariffs, sector-level world growth, and RCA, all in the same specification.

We exploit the MFA episode by restricting our analysis to the T&C products that were covered by the MFA, and we assign an indicator variable, B_h , to equal 1 if China was subject to binding quotas in that product.³¹ As Brambilla et al. (2010) show, product market competition was the most severe in the bound products. We then estimate an analog to the baseline specification in equation (8) on the set of T&C products:

$$\Delta \ln \lambda_{cht}^F = \alpha_{ht} + \alpha_{ct} + \gamma_1 PF_{cht-5} + \gamma_2 (PF_{cht-5} \times B_h) + \varepsilon_{cht}. \quad (9)$$

The specification in equation (9) regresses quality upgrading on product-year and country-year fixed effects, the lag proximity to frontier, and its interaction with B_h .³² We again restrict the analysis to the countries defined as having a relatively stronger business climate. Since China became eligible for quota removals only after it entered the WTO in December 2001, we focus on the period from 2000 to 2005. As before, we should observe $\gamma_1 < 0$; varieties that are far from the frontier experience faster growth due to convergence. The coefficient of interest is γ_2 , which captures the differential quality upgrading in bound and unbound products according to a variety's PF . This coefficient should be positive: varieties that are close to the frontier should experience relatively faster quality upgrading in China's bound products. We report the results in column 4 of table 5. Consistent with the theory and our earlier evidence, we observe that quality upgrading among high PF varieties is faster in products that faced stiffer competition over this period. We also run an additional placebo test by including the earlier periods of our sample in column 5. Since China's exports remained constrained by quotas prior to 2000, product market competition should not have increased in the B products over 1990 to 95 and 1995 to 2000. We check this placebo test by interacting PF and $PF \times B$ with a post-WTO indicator that takes a value of 1 in the period 2000 to 2005 and a value of 0 in the periods 1990 to 1995 and 1995 to 2000. The coefficient on the triple interaction term is statistically significant, implying that, as shown in the previous column, there is a differential change in quality upgrading across China's bound and unbound products after China's WTO entry. Moreover, there is no statistical difference between quality upgrading across products in the periods before China's WTO entry. These findings are entirely consistent with product market competition stiffening in the bound products in the final period of our sample, but having no differential impact in the earlier years.³³ Thus, the predictions of the model and our baseline results are verified using the MFA shock.

³¹ We choose 1991 as the year to determine whether China was binding in the product because this is the earliest year for which the binding quota information is available at the HS level. China was bound in 18% of the T&C products.

³² Note that because B_h is time invariant, we are able to identify only the interaction effect.

³³ Note that we do not report the interaction $PF \times \text{PostWTO}$ in the table for readability purposes. This coefficient is insignificant, which is consistent with product market competition not changing substantially in China's unbound products when the quotas were removed.

A related concern with the baseline specification is that the coefficients on PF_{cht-5} and $(PF_{cht-5} \times \text{tariff}_{cht-5})$ might be downward biased because, all else equal, a high λ_{cht-5} implies a high PF_{cht-5} but a low $\Delta \ln \lambda_{cht}$. Following Acemoglu et al. (2006), we instrument PF_{cht-5} (and the interactions) with its five-year lag value. Column 6 of table 5 shows that the results are robust, with the coefficients on the proximity to frontier and the interaction becoming a little smaller in magnitude (compare with column 3 of table 3).³⁴ Another way to check that the regressions are not simply picking up mean reversion is to use the variety's current period proximity to frontier, PF_{cht} , instead of the lagged measure.³⁵ We should still find the pattern of quality upgrading if the current period measure truly captures a variety's underlying proximity to frontier. Column 7 of table 5 reports the results. The coefficients on the tariff terms have the same signs as the baseline specification. It is not surprising that the coefficient on the interaction is somewhat weaker since the current period measures proxies the lag frontier measure with error (statistically significant at the 15% level), but nevertheless, the result demonstrates that the baseline specifications are unlikely to be driven solely by mean reversion.

Finally, an alternative source of bias could arise from selection. If a country does not export a particular product, we do not observe the quality of that good. Further, the observation is included in the estimation only if the quality is observed in both periods t and $t - 5$. It is difficult to sign the selection bias since it is likely to depend on how tariffs affect entry and exit and where the varieties are located on the PF distribution. For example, suppose lower tariffs result in a country-product pair exiting from our sample. While its quality change would be missing (the current period quality is not observed), one might expect its quality to have fallen (i.e., its quality-adjusted price rises to the point where no U.S. consumer chooses to import the variety). If the variety was already far away from the frontier, then our coefficient on tariffs (β_2) is biased downward. If the exiting variety was close to the frontier, then the interaction coefficient (β_3) is also biased downward, implying that the estimated coefficient is "too negative." On the other hand, if lower tariffs result in more varieties entering the sample, then the selection bias implies that our coefficients are biased upward. This is because while the quality change for entering varieties is missing (because the previous period quality is not observed), entry into our sample can be viewed as a positive change in quality. Thus, for varieties distant from the frontier, our estimate of β_2 is biased upward (that is, it is "too positive"), while for varieties close to the frontier, the baseline estimate of β_3 is not negative enough. This makes it difficult to sign the overall bias on both coefficients due to any selection issues. Nonetheless, to address this selection issue, we implement a two-step Heckman correction. For this estimation, we use

³⁴ The results are also unaffected by the inclusion of the growth in world exports in the instrumental variables estimation (results available on request).

³⁵ We thank an anonymous referee for suggesting this specification test.

TABLE 6.—HECKMAN SELECTION CORRECTION

Regressors	First-Stage Probit (Heckman) (1)	Heckman Selection Correction (2)
PF_{cht-5}		-0.818*** 0.019
$Tariff_{c,h6,t-5}$		0.538*** 0.094
$PF_{cht-5} \times Tariff_{c,h6,t-5}$		-0.913*** 0.119
$Log\ Potential\ Freight_{c,h6,t-5}$	-0.097*** 0.002	
$Inverse\ Mills\ Ratio_{c,h6,t}$		-2.844*** 0.239
Product-year FEs	No	Yes
HS2-year FEs	Yes	No
Country-year FEs	Yes	Yes
R^2	0.17	0.58
Observations	729,319	94,274

Table reports a Heckman selection correction for the baseline specification. Column 1 reports the first-stage probit regression of the probability of being observed in the sample on the (log) potential freight measure, country-year fixed effects, and HS2-year fixed effects. Column 2 includes the inverse Mills ratio obtained from column 1 to correct for selection. The specifications are run on the set of high-DB countries, and column 2 includes country-year and product-year fixed effects. Standard errors clustered by exporting country (with EU countries treated as one country because of the common tariff). Significance: *, 10, **, 05, and ***, 01.

freight costs that a country would have to incur if it were to export that product. This variable plausibly affects entry and exit decisions into the U.S. export market but does not affect the change in quality. We calculate this potential freight cost by taking the cost of the closest neighboring country that does export that product. The first stage probit in column 1 of table 6 shows that the coefficient on the freight variable is negative and significant; this suggests that higher potential freight costs reduce the probability of being in the sample. In the second stage, we include the inverse Mills ratio from the first-stage regression, which is significant, implying that the error terms in both regressions are correlated. However, the results in the second step (column 2) indicate that the main coefficients of interest are unchanged from the baseline specifications. Thus, our results remain robust to controlling for potential selection biases.³⁶

VI. Conclusion

In this paper, we show a significant relationship between import tariffs and quality upgrading. The direction of the effect depends importantly on how far the product is from the world quality frontier. For products close to the frontier, low tariffs encourage quality upgrading, whereas for products distant from the frontier, low tariffs have the opposite effect, discouraging quality upgrading. The disaggregated nature of the quality and tariff measures enables us to isolate this relationship while controlling for country-year specific effects such as changes in institutions that could be correlated with industry competition measures or product demand

³⁶ The results are also robust to including tariffs and growth in world exports in the first-stage probit, but the sample is smaller because we do not have tariff information for all censored varieties.

shocks. Our findings support the theories by Aghion and Howitt (2005) and ABGHP (2009) that highlight two forces: the escape-competition effect that induces a firm close to the frontier to invest in quality upgrading in order to survive competition from potential new entrants and the “appropriability” effect that discourages firms distant from the frontier from investing in quality upgrading because they are too far from the frontier to be able to compete with potential new entrants.

Our results show that support for these theories is strongest in countries with good business climates, a finding that is perhaps not surprising given that lower tariffs are unlikely to alter significantly the competitive environments in countries that face many other restrictions on competition. Interestingly, the nonmonotonic relationship between tariffs and quality upgrading holds for both OECD and non-OECD countries with strong business climates. Thus, our results suggest that a minimum institutional quality, and not simply higher income per capita, is required for the two opposing forces in ABGHP (2009) to operate.

These findings have a number of implications for future research. First, the initial distance to the world frontier within industries is important for understanding subsequent industry performance following trade liberalization policies. Second, we have focused exclusively on the effect of import competition on quality upgrading. However, import competition can also affect other forms of innovation, such as investments in technology, research and development, and the introduction of new types of products. These measures of firm performance would also be interesting avenues to explore.

REFERENCES

- Acemoglu, D., P. Aghion, and F. Zilibotti, “Distance to Frontier, Selection, and Economic Growth,” *Journal of European Economic Association* 1:3 (2006), 37–74.
- Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt, “Competition and Innovation: An Inverted U Relationship,” *Quarterly Journal of Economics* 120 (2005), 701–728.
- Aghion, P., R. Blundell, R. Griffith, P. Howitt, and S. Prantl, “The Effects of Entry on Incumbent Innovation and Productivity,” this REVIEW 91 (2009), 20–32.
- Aghion, P., C. Harris, P. Howitt, and J. Vickers, “Competition, Imitation and Growth with Step-by-Step Innovation,” *Review of Economic Studies* 68 (2001), 467–492.
- Aghion, P., and P. Howitt, “Appropriate Growth Policy: A Unifying Framework,” mimeograph, Harvard University (2005).
- Amiti, M., and J. Konings, “Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia,” *American Economic Review* 97 (2007), 1611–1638.
- Baldwin, R., and J. Harrigan, “Zeros, Quality and Space: Trade Theory and Trade Evidence,” NBER working paper 13214 (2007).
- Bastos, P., and J. Silva, “The Quality of a Firm’s Exports: Where You Export to Matters,” mimeograph, University of Nottingham (2009).
- Bernard, A., J. B. Jensen, S. J. Redding, and P. K. Schott, “Firms in International Trade,” *Journal of Economic Perspectives* 21 (2007), 105–130.
- Bernard, A., J. B. Jensen, and P. K. Schott, “Survival of the Best Fit: Exposure to Low-Wage Countries and the (Uneven) Growth of US Manufacturing Plants,” *Journal of International Economics* 68 (2006), 219–237.

- Berry, S., "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics* 25 (1994), 242–262.
- Berry, S., J. Levinsohn, and A. Pakes, "Automobile Prices in Market Equilibrium," *Econometrica* 63 (1995), 841–890.
- Boorstein, R., and R. C. Feenstra, "Quality Upgrading and Its Welfare Cost in U.S. Steel Imports, 1969–74," in Elhanan Helpman and Assaf Razin (eds.), *International Trade and Trade Policy* (Cambridge, MA: MIT Press, 1991).
- Brambilla, I., A. K. Khandelwal, and P. Schott, "China's Experience under the Multifiber Arrangement (MFA) and the Agreement on Textiles and Clothing (ATC)," in R. Feenstra and S.-J. Wei (eds.), *China's Growing Role in World Trade* (Chicago: University of Chicago Press, 2010).
- Burstein, A., and M. Melitz, "Trade Liberalization and Firm Dynamics," NBER working paper 16960 (2011).
- Bustos, P., "Trade Liberalization, Exports and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinean Firms," *American Economic Review* 101 (2011), 304–340.
- Cardell, N. S., "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity," *Econometric Theory* 13 (1997), 185–213.
- Feenstra, R., "New Product Varieties and the Measurement of International Prices," *American Economic Review* 84 (1994), 157–177.
- General Accounting Office, "U.S. Imports: Unit Values Vary Widely for Identically Classified Commodities," GAO report GAO/GGD-95-90 (1995).
- Goldberg, P. K., A. K. Khandelwal, N. Pavcnik, and P. Topalova, "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India," *Quarterly Journal of Economics* 125 (2010), 1727–1767.
- Grossman, G., and E. Helpman, *Innovation and Growth in the Global Economy* (Cambridge MA: MIT Press, 1991).
- Hallak, J. C., "Product Quality and the Direction of Trade," *Journal of International Economics* 68 (2006), 238–265.
- Hallak, J. C., and P. K. Schott, "Estimating Cross-Country Differences in Product Quality," *Quarterly Journal of Economics* 126 (1) (2011), 417–474.
- Halpern, L., M. Koren, and A. Szeidl, "Imports and Productivity," mimeograph, Central European University (2009).
- Hart, O. D., "The Market Mechanism as an Incentive Scheme," *Bell Journal of Economics* 14 (1983), 366–382.
- Hausmann, R., and D. Rodrik, "Economic Development as Self-Discovery," *Journal of Development Economics*, 72 (2003), 603–633.
- Hummels, D., and P. J. Klenow, "The Variety and Quality of a Nation's Exports," *American Economic Review* 95 (2005), 704–723.
- Hummels, D., and A. Skiba, "Shipping the Good Apples Out," *Journal of Political Economy* 112 (2004), 1384–1402.
- Johnson, R., "Trade and Prices with Heterogeneous Firms," mimeograph, University of California at Berkeley (2009).
- Khandelwal, A. K., "The Long and Short (of) Quality Ladders," *Review of Economic Studies* 77 (2010), 1450–1476.
- Konings, J., and H. Vandenbussche, "Antidumping Protection and Markups of Domestic Firms," *Journal of International Economics* 65 (2005), 151–165.
- Koopman, R., Z. Wang, and S. Wei, "How Much of Chinese Exports is Really Made In China? Assessing Domestic Value-Added When Processing Trade is Pervasive," NBER working paper 14109 (2008).
- Kremer, M., "The O-Ring Theory of Economic Development," *Quarterly Journal of Economics* 108 (1993), 551–575.
- Krugman, P., "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review* 70 (1980), 950–959.
- Kugler, M. and E. Verhoogen, "Prices, Plant Size, and Product Quality," *Review of Economic Studies* 79 (2012), 307–339.
- Levinsohn, J., "Testing the Imports-as-Market-Discipline Hypothesis," *Journal of International Economics* 35 (1993), 1–22.
- Linden, G., K. Kraemer, and J. Dedrick, "What Captures Value in a Global Innovation System?" mimeograph, UC Irvine (2007).
- Manova, K., and Z. Zhang, "Export Prices and Heterogeneous Firm Models," NBER working paper 15342 (2009).
- Melitz, M., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica* 71 (2003), 1695–1725.
- Pavcnik, N., "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants," *Review of Economic Studies* 69 (2002), 245–276.
- Pierce, J. R., and P. K. Schott, "Concording U.S. Harmonized System Codes over Time," *Journal of Economic and Social Measurement* 37 (2012), 61–96.
- Romalis, J., "Market Access, Openness and Growth," NBER working paper 13048 (2007).
- Schott, P., "Across-Product versus Within-Product Specialization in International Trade," *Quarterly Journal of Economics* 119 (2004), 647–678.
- Schumpeter, J., *Capitalism, Socialism and Democracy* (London: Allen Unwin, 1943).
- Teshima, K., "Import Competition and Innovation at the Plant Level: Evidence from Mexico," mimeograph, Columbia University (2008).
- Topalova, P., and A. K. Khandelwal, "Trade Liberalization and Firm Productivity: The Case of India," this REVIEW 93(3), 995–1009.
- Tybout, J., "Plant- and Firm-Level Evidence on 'New' Trade Theories," in E. Choi and J. Harrigan (eds.), *Handbook of International Trade* (Oxford: Basil-Blackwell; 2003).
- Vandenbussche, J., P. Aghion, and C. Meghir, "Growth, Distance to Frontier and Composition of Human Capital," mimeograph, Harvard University (2004).
- Verhoogen, E., "Trade, Quality Upgrading and Wage Inequality in the Mexican Manufacturing Sector," *Quarterly Journal of Economics* 123 (2008), 489–530.
- Vogel, J., "Spatial Competition with Heterogeneous Firms," *Journal of Political Economy* 116 (2008), 423–466.
- World Bank, "Doing Business 2005" (2005). www.doingbusiness.org.