

Offshoring and the Shortening of the Quality Ladder: Evidence from Danish Apparel*

Valrie Smeets
Aarhus University[†]

Sharon Traiberman
Princeton University[‡]

Frederic Warzynski
Aarhus University[§]

April 24, 2016

Abstract

An important margin of adjustment for firms responding to trade shocks is in their output quality. Access to either higher or lower quality inputs from abroad may induce higher or lower output quality depending on prices and the demand structure. This paper estimates a structural demand model for domestic Danish apparel firms from 1997 to 2010 in order to analyze how firms responded to China's entry to the WTO and the dismantling of the Multi-Fibre Agreement. We find that dispersion in the distribution of quality decreases in response to cheaper inputs from abroad. We refer to this as a shortening of the quality ladder. We build a simple model that illustrates the mechanisms driving this change. In our model, firms differ in two dimensions: their productivity at producing output of a given quality as well as their ability to increase the quality of their products. Firms of low quality firms tend to increase their quality while high quality firms tend to decrease theirs. We also find that import competition spurs entry of high quality firms and exit of low quality and low productivity producers.

1 Introduction

The gains and losses from trade arise through many channels. For example, breaking away from assumptions of perfect competition and homogeneity have allowed economists to study the variety gains and pro-competitive effects of trade liberalization. Recent work has also highlighted cross-country differences in product quality and how access to new markets can effect the quality of consumer goods. This research has mostly been limited to studying cross-country datasets. However, it is also interesting to ask how the distribution of quality *within* a country changes. For example, if the average quality of a good increases is this because low quality goods left the market

*Sharon thanks Stephen Redding and Jan De Loecker for substantial guidance on this project. A special thanks to Henning Bunzel for help with the data. We are also thankful to Thomas Chaney, Jonathan Eaton, Oleg Itskhoki, Amit Khandelwal, Thierry Mayer, Mark Roberts, John Romalis and participants at various conferences and seminars. Sharon Traiberman received financial support from a Princeton IES Summer Grant, the International Economics Section at Princeton, the Economics Department at Princeton and Aarhus University.

[†]email: vas@asb.dk

[‡]email: straiber@princeton.edu

[§]email: fwa@asb.dk

or because high quality goods became more available? This matters when consumers themselves are heterogeneous in their preferences. Our paper contributes to filling this gap by estimating how the distribution of apparel quality across firms within Denmark changed in response to Chinese accession to the WTO, thus giving firms access to a new set of inputs.

Two recent observations in the literature drive our research question. First, the quality of goods varies substantially across countries.¹ Second, a large fraction of trade growth has been in intermediates, offshoring and supply chain disintegration² Putting these facts side-by-side offers a mechanism for trade to affect the menu of goods available to consumers within a country. In particular, one can ask whether firms' importing of cheaper (and thus potentially lower quality) inputs affects their output quality in an appreciable way. This channel is distinct from the effect of trade on output markets, as it is not about the entry and exit of new goods, but also changes to extant goods.

To determine the impact of these new input markets on firms' domestic output we structurally estimate apparel demand in the Danish economy. We link together and exploit several novel datasets that contain domestic production, trade, and sales of apparel at the firm *and* product level. Following Khandelwal (2010) and Berry (1994), we estimate a logit demand system that allows us to recover a demand shifter from a regression of prices on market shares. We will interpret this shifter as an estimate, or at least proxy, of output quality. Thus, quality can be interpreted as any force that induces to consumers to buy more of a good than would be predicted given the price of that good.³

The major econometric challenge in demand estimation rests in finding an instrument for price, as it is likely correlated with quality. To get around this challenge we use the peculiarities of our environment and the richness of our dataset to construct exogenous cost shifters that act as instruments. In particular, we exploit the fact that apparel firms make their design and sourcing decisions well in advance of their pricing decision. Thus, we use unanticipated shocks to trade costs as instruments for price that are not correlated with the initial quality and sourcing decisions. Our dataset includes firm-product level information on price and quantity at a fine level: our unit of observation are Combined Nomenclature 8 digit goods for each firm, which are a strict refinement of the HS6 classification.

¹See Hallak (2006), Khandelwal (2010), and Hallak and Schott (2011)

²See Yi (2003), Feenstra (2010) and Costinot et al. (2013).

³In this sense, we take the view of Sutton (2012) who writes that quality *“includes not just ‘quality’ in the usual narrow sense (a feature of the product’s physical characteristics), but also a range of characteristics that include, for example: brand advertising... services... and logistics”* (Sutton, 2012, p.17).

With our estimation in hand we turn to analysis of how the distribution of quality has evolved in Denmark in response to trade shocks. The object we focus on is the distance between the highest and lowest quality goods in the market at a point in time. Following Khandelwal (2010) and others, we refer to this as the length of the quality ladder.⁴ At the aggregate level, we find that as trade costs decrease, the quality ladder shrinks. Moreover, we also find a trend break at the time China joins the WTO. In order to understand these trends more concretely we perform a firm-level analysis. We find that changes in the quality ladder are the result of two forces: a tightening of quality of surviving firms; and the exit of relatively low quality producers, as well as entry of relatively high quality producers.

When we examine how surviving firms respond to new sourcing opportunities we find that, on average, sourcing from abroad is associated with a decrease in quality. We also show that imported input quality, proxied by the per capita GDP of income partners, moderates the downgrading channel: firms sourcing from relatively poorer countries experience larger output quality drops than firms sourcing from elsewhere. However, we also find that this average effect masks a great deal of heterogeneity in firms' joint sourcing and quality decisions. In particular, we find that lower quality firms that begin to engage in offshoring tend to upgrade their quality relative to other firms' within the same year, while higher quality firms that increase their offshoring activity tend to downgrade their quality.

In order to understand the forces at play, we build an illustrative model to ground our empirical approach. In our framework, firms endogenously choose their sourcing strategy, their output quality and their price. Firms differ along two dimensions: physical productivity and capability. The first term refers to firms' productivity at making physical output conditional on a choice of quality—a standard feature of many heterogeneous firm models. The second term refers to firms' ability to produce higher quality output. Firms leverage their relative advantages in deciding how much and how high a quality to produce. Thus, our setup is similar in spirit to recent models of quality and quantity choice by firms, such as Hallak and Sivadasan (2013) .

In the model, producing a set output quality requires a commensurate input quality. Trade enters the model by allowing firms to access a menu of inputs in different countries. The relative steepness of this menu with respect to quality determines whether a country's comparative advantage is in quality or quantity. We find that access to new countries can change the quantity and

⁴We use this definition in congruency with similar papers, but in our results section we explore other measures of the spread in quality across goods.

quality tradeoff for firms but that this tradeoff depends on where firms initially are in the quality distribution. For example, if a firm is already producing at a low quality, then cheap inputs abroad may induce little change (or some upgrading) in output quality, with the firm focusing on simply lowering price. However, for a firm producing higher quality output, the ability to lower price and gain market share by using cheaper inputs may induce quality downgrading.

Our work touches on several different strands of the trade literature. First, several recent papers have shown that access to high quality capital or inputs from abroad induces quality upgrading in low and middle income countries. For example, see Fieler et al. (2014) examine how access to new inputs can explain rising skill intensity in Colombia. Our work differs from theirs and others in its focus on a high income country. We also focus on a different tradeoff in this context: the tension between a firm wishing to source cheaper inputs but ceding partial control over the quality of its output in doing so. Moreover, our paper is one of the first to demonstrate that the distribution of quality may contract in response to trade shocks, and we offer a reason why.

In related work, Bloom et al. (2011) and Utar (2014) provide evidence that increased competition from China led to organizational restructuring and increased innovation in the European apparel and textile industry, but do not explicitly focus on product quality. Kugler and Verhoogen (2012) document and model how larger and more efficient firms choose higher quality inputs and produce higher quality output that they sell at a higher price when the scope for differentiation is large enough. Holmes and Stevens (forthcoming) also explore how firms' quantity and quality choices may diverge. They observe that smaller, more focused and higher quality firms were more resistant to the surge of imports from China, despite their size. Our paper differs from these last two in that we look more explicitly at heterogeneity in the sourcing decisions of firms, while also allowing for heterogeneity in input and output quality.

Closer to us, Amiti and Khandelwal (2013) extend Khandelwal's original analysis using product level data from 56 countries to the US and find that lower tariffs are associated with product upgrading for firms close to the world quality frontier, but discourage upgrading for firms distant from the frontier. Roberts et al. (2012) use firm level data on exports by product and destination for Chinese footwear exporters and estimate a firm specific demand component together with a cost and an export market profitability components. They find that both the cost and demand components are related to firms' success and they also document a reallocation of resources towards more productive and higher demand firms following the removal of EU quotas. Piveteau and Smagghue (2015) use similar French data to study the link between product upgrading and import competition.

They find evidence that firms improve the quality of their export products when import competition increases.⁵ However, these papers do not focus on how sourcing decisions in advanced economies are related to product quality. Analyzing this relationship is the main contribution of our paper.

The remainder of the paper proceeds as follows. Section 2 provides a brief discussion on the Danish apparel industry as well as the MFA and also presents. Section 3 describes the various datasets that we use. Section 4 presents an illustrative model of offshoring and quality decisions that guides our empirical analysis. Section 5 details our empirical methodology. Section 6 presents the results of our estimation and a discussion of the results. Finally, Section 7 concludes.

2 The Danish Apparel Industry and the End of the Multi Fibre Arrangement

The Danish Apparel Industry

Historically, Danish industries have been famous for the creative and design aspects of their goods. The Danish Design movement in particular, has had a large and lasting influence on modern furniture and architecture. The Danish apparel industry, concentrated predominantly in the medium and high end segment of the fashion industry, continues to be an important part of Denmark’s creative output. The sector represents more than 25% of the so called creative industries that were recently singled out by the Danish government as a major component for future growth.⁶

In Denmark, the majority of apparel firms can be divided into two groups—“Branded Manufacturers” and “Branded Marketers.”⁷ The latter group focus solely on design, distribution and marketing of products, and have grown in recent years with the rise of fast fashion. The former group, in contrast, engage more explicitly with physical input choices. For example, they may produce domestically, and in Denmark there is a small, persistent set of apparel manufacturers.⁸ Even if not producing domestically, they often engage in outward processing whereby they purchase raw materials and send them directly to assembly plants. They may also do some final packaging and assembly locally. Largely, the difference between these two groups is in the definition of the firm’s boundary: whether or not the firm purchases its intermediates directly. Nevertheless, the

⁵See also Martin and Mejean (2014) who use a different empirical approach to study the same question. They also find evidence of a positive relationship between upgrading and import competition through a reallocation of market share from low quality firms to high quality firms.

⁶See of Business and Growth (2013).

⁷See Gereffi (1999) for more on these concepts and other details of the apparel industry

⁸Utar (2014) calculates that there are 13,000 workers in the combined textile and apparel industries.

distinction is important as we only observe branded manufacturers in our data. This is because in Denmark firms who own their own inputs, even if they ultimately do production abroad, are in the production registry.

As our analysis ultimately rests on a single index of quality, it is helpful to understand how the industry normally thinks of this concept. Quality in apparel is normally broken down into two components – the physical quality of the good (e.g., open-end spinning versus ring-spun cotton, thread count, non-bleeding dyes) and the “fashionability” of the item.⁹ Branded manufacturers exercise a great deal of control over both of these as they often engage directly with the sourcing partners of assemblers. In our analysis, we identify a demand shifter that combines the firm’s quality along with individual tastes and perceptions. Because of this, we cannot tease apart different aspects of quality. This is an issue insofar as the fashionability of a good is a relative concept, and also can substitute with physical quality.¹⁰ But while we cannot separate out these concepts, it also highlights the idea that quality is partially a feature of a particular firm and its capabilities at making likable designs, and partially an outcome of a sourcing decision. Our results ultimately corroborate the idea that firms exercise some control over our estimated demand shifters, even if imperfectly.

The MFA and Trade Patterns

Starting in 1974, the majority of apparel and textile trade was governed by a series of quotas outlined in the Multi-Fibre Arrangement (MFA), and later the Agreement on Textiles and Clothing (ATC). These quotas were set up to allow developed countries to adjust slowly to competitive pressure from abroad. There was an understanding that the quotas would ultimately be lifted. To that end, the MFA was phased out in several stages beginning in 1995 and ending in 2005. In 2005, the majority of apparel quotas were phased out with some countries maintaining the option to reinstitute short-term restrictions on trade if they saw fit.¹¹

In this section we briefly detail how the phase out of the MFA made itself felt in the Danish apparel industry. It is worth noting that as China was not in the WTO until late 2001, quotas

⁹We are thankful to Avinash Vora for walking us through the daily goings-on of an Indian textile and clothing factory. Also, to Line Lyngholm at Bestseller for helping us understand the Danish apparel industry. For an attempt at formally modeling the distinctions noted above, see Paul R. Liegey (1993).

¹⁰This substitution is at the heart of the strategy of many “fast fashion” retailers, such as Zara and H&M.

¹¹Some additional quota restrictions were placed on China for a short period, leading to extended negotiations between China and the EU. See “EU and China in ‘bra wars’ deal,” published in *The Guardian* newspaper on September 5, 2005. Also, see Brambilla et al. (2010) for a history of the MFA as well as how the end of the MFA affected numerous countries.

on Chinese apparel and textile production remained in effect even after the first two rounds of restriction easing had passed. However, in January 2002 many restrictions on Chinese textile and apparel trade were lifted. Thus, China's accession to the WTO provided a large, new sourcing opportunity for Danish firms. While China's entry into the WTO is the largest one-time shock to the Danish fashion industry, the phase out of the MFA/ATC also led to more trade with other Asian and Eastern European countries, such as Turkey and Poland.

Turning to general patterns, from 1997 to 2010, imports of apparel in Denmark grew by 26.5% in real terms. This rise is driven disproportionately by two factors: (1) increased quantities (as prices actually fell in this period) and (2) more re-exporting activity. Figure 1 makes this clear by plotting the nominal value of *net* imports and the weight of net imports. The nominal value is falling. This reflects high re-export activity: total imports have increased while the imports that remain in Denmark has decreased. On the other hand, the weight is increasing. This suggests that the actual quantity of apparel entering Denmark has increased. These two patterns are driven by changes in the structure of the apparel industry in Denmark at this time. In particular, importing by Branded Marketers (those that only engage in importing and re-exporting) grew by 43% while importing activity actually decreased for Branded Manufacturers (who nevertheless maintain around 40% of domestic market share—albeit decreasing over time).

The removal of trade quotas not only increased trade volumes, it also drastically changed the composition of imports. In particular, there is a rapid increase in sourcing from Asian countries and in particular China. Figure 2 documents the rise of China in Danish apparel, and figure 3 breaks this down by type of firm. While movement to China was steadily growing, starting with its entry to the WTO, trade with China began to rise rapidly and constitutes 45% of Danish apparel imports by the end of the sample. For manufacturers, China is responsible for 16% of imports. In this group we also see a substantial decline in importing from Eastern Europe, as work is supplanted by China and Turkey.

While we exploit and study these changes in sourcing patterns, we have no direct evidence that countries like China offer different quality inputs at different prices than Eastern European countries. That is, we cannot observe, directly, input quality. However, others have studied the effect of the end of the quotas on the exports of poor countries. For example, Amiti and Khandelwal (2013) find that China upgraded the quality of many of its products after the fall of the MFA, but this upgrading was heterogeneous and highly dependent on firms' initial quality. Moreover, as documented in Brambilla et al. (2010), Chinese apparel product quality may have risen but

decreased relative to the rest of the world. In a similar vein, Manova and Zhang (2012) find evidence that Chinese exporters “use higher quality inputs to produce higher quality goods.” This latter point suggests that apparel firms may actually make demands of their sourcing partners. The overall picture suggest that the dismantling of the MFA led China to engage in some quality upgrading, but nevertheless they focus on production of lower quality goods.

3 Data

In order to carry out our analysis we rely on several datasets maintained by Statistics Denmark. All of our datasets can be linked through firm identifiers and ultimately allow us to construct a panel of Danish apparel firms running from 1997 to 2011.

First, we use a dataset that contains annual longitudinal production data for all manufacturing firms employing at least 10 individuals or who meet a minimum revenue threshold. The data contains price and quantity data on sales at the Combined Nomenclature 8 (CN8) digit product level for each firm.¹² If firms own their materials (i.e., if they source them and route them to producers on their own), they will appear in this dataset even if production is not physically taking place in Denmark. Hence we observe both those firms that produce domestically and engage actively in sourcing materials. We do *not* observe domestic sales and price data for firms with their headquarters in Denmark, but who do not either manage their own inputs or produce in Denmark. That is to say, we do not observe the firms we described in section 2 as Branded Marketers.

To identify our firms, we focus on those firms producing at least one type of apparel product and who make at least 90% of their revenues in the apparel industry.¹³ Table 1 shows the most common products made by our sample of firms. While basic these products can still incorporate a large design component.

As the data contains quantity and sales data, we are able to construct unit values. It is well known that unit values can be a noisy proxy for prices, and we cut data that are likely to be outliers and erroneous. In particular, we remove observations based on the following criteria:

1. We remove the top and bottom 5% of unit values (many of which are in the pennies or the

¹²The Combined Nomenclature is the EU’s classification system for recording customs transactions. The first 6 digits are the same as HS6 classifications and the last 2 digits are defined by the EU’s documentation. In the case of apparel, the last 2 digits distinguish weight and material used in construction of apparel.

¹³The distribution is heavily bimodal—with most firms making more than 90% or less than 5% of their profits in apparel. Upon examination, most of the firms with a small revenue share of apparel are producing niche products (e.g., industrial wear). Thus, while we cannot know for sure, we think it reasonable to assume they are not competing with the majority of apparel firms, who focus on household consumers.

tens of thousands of DKK).

2. We remove products without at least 45000 DKK in sales. This is because sales are rounded to the nearest 1000, so this avoids large swings due to rounding errors.
3. If a variety (firm-product pair) exits and reappears, we keep only the longest contiguous period. This is because sometimes exit and re-entry is met with large price swings, which we believe may be due to aggregation bias. That is, the product may have changed even within a narrow code.
4. We remove firm-product-years where the unit value jumps above twice the the within-firm-product median price.

We can link our production data at the firm-product level to firm level data on capital, labor and materials use. More importantly, we link our data to customs data. This dataset contains all import and export transactions at the product level (also CN8) for each firm in a given year. From here we can see how many textiles and apparel products each of our firms imports and exports. The dataset is comprehensive and contains information on the trade partner, quantity, price and unit of the good. This last variable matters because customs data and production data are not often identically scaled (e.g., counts in 1000s versus 10s). Thus, this dataset tells us the sourcing strategy of firms.

As a final point, we bring in several other data sources. Data on quotas comes from the EU's SIGL database. This database includes product-level data on quota utilization, quota fill rates and license volume for the entire length of our panel. For data on exchange rates we used data published by the IMF's International Financial Statistics and the Federal Reserve's FRED Database.

4 Illustrative Model

Before turning to the empirical analysis, we present a model that illustrates the main forces that we test in the data. As we neither calibrate the model nor run any counterfactuals, we do not specify the full general equilibrium model. Instead we focus on a firm in partial equilibrium making a decision about its quality and production. In our setting, production requires two decisions: at what quality to produce output, and how much output to ultimately sell. Similarly, firms differ in two dimensions: first, in their ability to produce output at a given quality (what we call productivity); and second, in the ease with which they can increase the quality of their good (what

we call “capability”¹⁴). We show how this simple setup can generate a shortening of the quality ladder in response to changes in the cost of imports. Our approach is similar in spirit to recent models with firms both differing in two-dimensions and making quality choices. For example, our model is closely related to that of Hallak and Sivadasan (2013). To this existing literature we build in, explicitly, an extensive sourcing decision. As a final point, in most of the subsequent discussion we focus only on the importing decisions of a firm, and not on the rest of the trade environment. This is for clarity, as our focus is on sourcing, however we discuss import competition and exports in the context of our model at the end of this section.

Consumer Preferences and Demand

In our empirical section, we will estimate a model based on the random utility framework common in the IO literature. In particular, we assume that workers face a choice over a finite and discrete set of goods and each choose one. While this set up is very useful empirically, working with it theoretically can be difficult.¹⁵ However, in order to maintain some connection to the empirical section that follows we posit that firms face a demand curve that is “approximately” derived from a standard logit model of discrete choice. In particular, we assume that firms face the following demand curve:

$$x(v, p) = A \exp\{\log v - \sigma p\}$$

where A is an aggregate demand shifter, v is quality, p is the price and σ is the semi-elasticity of price. This functional form is similar to the logit demand curve which, indexing the firm momentarily by j , is given by:

$$x_j(v, p) = \frac{1}{A_0 + \sum_{f \in \mathcal{F}} \exp\{\log v_f - \sigma p_f\}} \exp\{\log v - \sigma p\}$$

where A_0 is the value of some outside option and \mathcal{F} indexes all firms including j . Given the demand curves above, the setup is reminiscent of monopolistic competition whereby no firm internalizes their impact on the aggregate shifter. In fact, if we had written $\log(p)$ above then this would be a standard CES demand curve and we could make the limiting argument more formally.¹⁶ The

¹⁴This term is borrowed from Sutton (2012) who develops many models of firms making quality choices.

¹⁵For example, Caplin and Nalebuff (1991) show that even in a simple price setting game where consumers are characterized by a discrete choice, there may be no interior equilibrium where firms’ actions solve the first order conditions of profit maximization. In most of the IO literature (e.g., Berry et al. (1995)) one assumes that such an equilibrium is the one observed.

¹⁶It is further true that our model predictions are robust to a CES specification of preferences, although the particular closed form solutions are not identical.

benefit of this approximation is that we can avoid modeling the strategic interactions of firms, which keeps the analysis tractable. When we derive the solution to the firm’s problem we will draw out more of the connections to the traditional logit demand system.

Quality Production and the Cost Function

Firms maximize profits given the consumers’ demand curve choosing *both* the quality v and price p of their output. Firms in our economy are characterized by two objects: their “capability,” ω and their “productivity,” λ . The first term will describe how well firms can produce quality, while the latter term describes how well firms can scale production, given a choice of quality. We assume these are drawn from some joint distribution F over $\Omega \times \Lambda$. Moreover, we assume that ω is bounded below by a strictly positive number. In the remainder of this subsection we will describe how these two terms enter the firms’ cost and production function.

First, physical output is created by combining a homogeneous factor and a differentiated factor in a Leontief fashion. Thus, production is given by,

$$x(v) = \min\{m_1, m_2(\psi\omega)\}$$

where $x(v)$ is output and m_1 and m_2 are two factors of production. ψ describes input quality. The term ω captures the capability of the firm—for higher values of ω , the firm can purchase a lower quality input and still achieve the same output quality. Thus ω may reflect the firm’s ability to design or brand itself. The mapping from input quality to output quality is given by a simple decreasing returns to scale quality production function:

$$v = (\psi\omega)^\alpha$$

We think of the homogeneous factor as a machine and the differentiated factor as either materials, the skill of workers or both.¹⁷ We assume that the homogeneous good is paid a factor price a and that the schedule for the differentiated good is linear in input quality so that $\tilde{c}(\psi) = b\psi$. With

¹⁷We have chosen this production structure for analytic convenience, however our results extend to a setting in which there is constant marginal cost of production and fixed costs of design. What will ultimately matter for our results is that the marginal cost of production is increasing in quality *and* that the slope of the marginal cost curve with respect to quality is increasing in trade costs. In such a setting, reductions in trade costs will induce firms *already* sourcing from abroad to upgrade their quality and any firms *induced* to source will downgrade their quality. Our functional form assumptions in this section let us demonstrate these forces in a closed form. Fixed costs of design or communication costs (as we use in our model) are necessary to create an extensive margin of sourcing. A generalized version with costs of design, which also features more standard CES preferences, is available upon request.

these prices, the Leontief structure leads to the following unit cost of production a quality ψ :

$$c(\psi) = a + \tilde{c}(\psi)$$

We introduce trade by allowing a and $\tilde{c}(\psi)$ to vary across countries. When firms import, they face a communication cost that makes it difficult to transmit their capability. For an importer the quality production function is given by,

$$v = (\psi\omega^z)^\alpha$$

where z reflects the communication cost. We focus on two countries and normalize $z_H = 1$, in order to treat this as the only quality production function. An exponent of $z < 1$ means that sourcing from abroad makes *vertical* differentiation more difficult, and guarantees that firms will either produce at home or abroad. With multiplicative costs, it would turn out that all firms would make the same decisions about sourcing. Thus, this firm guarantees that not all firms make the same decision. We include the exponent for the sake of generality but at an extreme it has a very intuitive flavor: If $z_F = 0$ then this implies that once sourcing from abroad the firm cannot exploit its own design capabilities. We discuss the role of this exponent in the next subsection, after demonstrating its role in the firms' optimal sourcing decision.

We introduce productivity in the standard, unit-cost augmenting way. In particular, a firm's final unit cost is given by $c(\psi)/\lambda$. Combining the above, the firm's profit function *conditional on whether or not they import* is given by,

$$\pi(\psi, p) = A \exp \{ \alpha \log(\psi\omega^z) - \sigma p \} (p - c(\psi)/\lambda) - a_H f$$

where f is a fixed cost or production, paid in units of the homogeneous good at home. From the first order conditions one can derive the following expression for optimal quality given the firm's sourcing strategy:

$$\psi^* = (\omega^{z_s} \lambda) \frac{\alpha}{\sigma b_s} \tag{1}$$

Similarly, one can derive the following expression for optimal price:

$$p^* = c(\psi^*)/\lambda + \frac{1}{\sigma} \tag{2}$$

This states that quality choices are increasing both in capability and productivity and that the price is decreasing in productivity. The optimal pricing equation demonstrates the connection between the typical logit demand system and our model. In particular, recall that for the standard logit model the optimal pricing rule would be given by,

$$p = c/\lambda + \frac{1}{\sigma(1-s)}$$

where s is the market share of the firm and c is marginal cost. Thus in the logit setting, markups are heterogeneous and larger firms charge higher markups. Our setting is the limit case where $s \rightarrow 0$. The benefit of our simplification is that we can solve not only the pricing decision, but the quality decision of the firm, in closed form.

Plugging the optimal quantities back into the profit function leads to the following profit conditional on a sourcing strategy:

$$\pi = \kappa \exp \{ -\alpha \log (b_s / (\lambda \omega^{z_s})) - \sigma a_s / \lambda \} - f$$

where κ is a constant that is independent of productivity and sourcing strategy. The profit function has an intuitive interpretation as it says that profits are determined by a weighted sum of the two components of costs—the unit cost, a_s of the homogeneous good moderated by price sensitivity, and the cost of the differentiated factor, b_s/ω^{z_s} , moderated by tastes for quality.

Trade and the Distribution of Quality

To discuss trade more concretely we focus on two countries, home and foreign. Furthermore, we assume that $a_H > a_F$ but $b_F > b_H$. These restrictions imply that the home country has an absolute advantage in high quality input production, but the foreign country has an absolute advantage in low quality input production. With this in mind, a firm will source from abroad if

$$\frac{\exp \{ \alpha [z \log \omega + \log \lambda - \log b_F] - \sigma a_F / \lambda \}}{\exp \{ \alpha [\log \omega + \log \lambda - \log b_H] - \sigma a_H / \lambda \}} > 1$$

This inequality highlights the role of the exponent in our communication cost. If we had a multiplicative communication cost, for example, it would be the case that the ω terms would cancel and all firms would behave symmetrically—either producing abroad or at home depending on param-

eters.¹⁸ By taking logarithms and rearranging, we arrive at the following cutoff for sourcing from abroad:

$$\omega \leq \exp \left\{ \frac{\frac{\sigma}{\lambda}(a_H - a_F) - \alpha \log \left(\frac{b_F}{b_H} \right)}{\alpha(1 - z)} \right\} \quad (3)$$

This cutoff is a function of both productivity and capability—so that sourcing patterns depend on both the joint distribution of these two terms. In particular, the cutoff value is *decreasing* in productivity. This suggests that conditional on quality, size may be a predictor of sourcing activity. Our general finding that productivity and capability jointly determine firm size and quality echoes the same point made recently by Holmes and Stevens (2014), Hallak and Sivadasan (2013) and Ecklel, Iacovone, Javorcik and Neary (2015): large firms may not be the highest quality firms and vice versa.

In order to discuss how changes in import costs map into the choices of firms, we break up the discussion into looking at the cost of the homogeneous factor, a_s and then at the cost of the differentiated factor, b_s . Notice that (3) shows that lowering trade costs will lower the cutoff rule for firms, so we also break up our analysis by looking separately at firms who were already sourcing abroad and firms who switch their sourcing strategy.

Turning to the homogeneous factor, if a_F decreases then firms that were already sourcing from abroad will *not* change their quality choice. This can be seen in (1), which shows that optimal quality choice only depends on the cost of the differentiated factor. On the other hand, it will induce switching by some middle-quality firms, as it moves the threshold for sourcing decisions. Switchers will *decrease* their quality. This can be seen by dividing the optimal quality decision of firms holding fixed their capability but not their sourcing strategy:

$$\frac{\psi^{*H}}{\psi^{*F}} = \omega^{1-z} \frac{b^F}{b^H}$$

which will be greater than 1 so long as b^F is sufficiently high. Finally, there will be no response of high quality firms.

Next looking at the differentiated factor, if b_F decreases then firms that were already sourcing will *upgrade* their quality. This can be seen in (1), where quality choice, conditional on sourcing, is a decreasing function of b . On the other hand, switchers will once again downgrade their quality

¹⁸More generally, if the cost of sourcing abroad is given by $a + b/\tau(\omega)$ where τ captures communication costs, we need that $\tau(\omega)/\omega$ is *strictly* decreasing. That is, the elasticity of τ with respect to ω would have to be negative, so that it is more difficult to communicate higher capability than lower capability.

(as long as it is still the case that b_F remains sufficiently larger than b_H).

Finally, we briefly discuss the aggregate demand shifter and entry and exit. Thus far the model has been primarily concerned with import costs and quality choice. However, it is an empirical fact that firms may enter and exit in response to changes in import costs. Moreover, changes in import costs are often concurrent with changes in export costs as well as import competition. This is true of our dataset. And, since we only have data on branded manufacturers, we will treat the combined demand of branded marketers (which includes wholesalers and retailers that do not appear in our production data) as an outside good. Thus, lowering import costs, which acts as a positive shock to all firms, can lead to a *negative* demand shock to manufacturing firms facing stiffer import competition. On the other hand, removing apparel quotas also lowered export costs for firms. In the model above, A is a common demand shock facing all firms, and so it can be a useful proxy for thinking about net changes in import competition and export costs. In particular, we think of increased import competition as a decrease in A and new export opportunities as an increase in A . From equations (1) and (2), one can see that this aggregate shifter does not affect optimal prices or quality choices. However, given our assumption that fixed costs of production are strictly positive for all firms, decreases in A will result in the net exit of low profitability firms—which can be either low capability, low productivity or both, depending on the joint distribution of these two firm attributes.

We summarize the above discussion in the following concrete set of observations that guide our empirical analysis:

Observation 1 (Quality). *Conditional on productivity, if b_F decreases but $b_H < b_F$ then,*

1. Firms that were already offshoring will increase their output quality.
2. Some firms will begin to offshore and *decrease* their quality.
3. Firms of sufficiently high quality will not respond.

Conditional on productivity, if a_F decreases then

1. Firms that were already sourcing from abroad will not respond.
2. Some firms will begin to offshore and their quality will decrease.
3. Firms of sufficiently high quality will not respond.

Observation 2 (Entry/Exit). *If A remains constant and a_s and/or b_s decrease then low quality firms enter. However, if A decreases then exit patterns depend on the joint distribution of λ and ω .*

Because the model is highly stylized, these observations are particularly sharp. In the data there are obviously more than two countries and all firms engage in some sourcing activity. Hence one should think of this model as largely descriptive. To operationalize our ideas in the data we will focus on looking at how the distribution of quality changes and how firms of different qualities respond to new importing opportunities. In particular, we modify our predictions about heterogeneity to suggest that lower quality ought to increase their quality relative to other firms, while middle quality firms and high quality firms ought to decrease their quality or have a more muted response.

It is important to highlight the crucial role of vertical differentiation in this context. In this model, vertical differentiation is not simply a productivity shifter or demand shifter as in the standard set up—rather two sources of heterogeneity separate physical productivity and quality capability. More importantly, this quality capability is affected by the firms’ sourcing decisions. Without these separate aspects of the firm, we would predict a strictly monotonic relationship between size, quality and productivity.

5 Econometric Model

This section outlines our econometric model, including consumer demand, timing assumptions and decision making by firms, as well as the details of mapping our model to data and instrumenting strategy. The strategy follows closely the recent work of Khandelwal (2010) and Amiti and Khandelwal (2013) in that we use the discrete choice framework common in IO to estimate consumer demand.

5.1 Consumer Demand

In this section we will derive a logit style demand curve for each product, similar to the demand curve in the section 4. However, things will differ now in two dimensions: we will no longer rely on an approximation to the demand curve and instead assume firms do internalize their impact on aggregate shifters; we will also depart slightly from the logit framework and instead use a nested logit framework. The nested logit setup is almost identical to the standard framework except that goods are grouped together into nests, and substitution patterns are different between and across

nesses. This modification allows for more realistic substitution patterns as one can imagine that not all apparel is equally substitutable (for example, men’s and women’s clothing). In section 5.3 we define the full nesting structure. In the remainder of this subsection, we derive the demand curve for each firm-product pair, which leads to our estimating equation.

In order to derive a demand curve for a product, we model consumers as making a discrete choice over goods in each period. To make things precise, assume that consumer i has indirect utility for good (j, t) given by,

$$V_{ijt} = \nu_j t - \alpha p_{jt} + \epsilon_{ijt}$$

where δ_{jt} is a common taste for product j at time t , p is price and ϵ is a consumer specific taste shock for product jt . We assume that ϵ is distributed as generalized extreme value (GEV), which allows for the shock to be correlated across some goods and not others. Consumer i picks good j iff $V_{ijt} \geq V_{ikt} \forall (k)$ at time t . Berry (1994) shows that if a large number of consumers have preferences as above, then the market share for product j at time t is given by,

$$s_{jt} = \underbrace{\frac{e^{(\delta_{jt} - \alpha p_{jt}) / (1 - \sigma)}}{\sum_{k \in \mathcal{J}_{g_j}} e^{(\delta_k - \alpha p_k) / (1 - \sigma)}}}_{\text{Within Group Share}} \underbrace{\frac{\left(\sum_{k \in \mathcal{J}_{g_j}} e^{(\delta_k - \alpha p_k) / (1 - \sigma)} \right)^{1 - \sigma}}{\sum_g \left[\left(\sum_{k \in \mathcal{J}_g} e^{(\delta_k - \alpha p_k) / (1 - \sigma)} \right)^{1 - \sigma} \right]}}_{\text{Group Share in Total}}$$

where σ is a parameter that governs nest substitution, g indexes nests (groups of goods with correlated shocks) and \mathcal{J}_g is the set of products in nest g . Given this decomposition, one can actually interpret the nested logit structure as one in which a consumer first decides on a type of good he or she wishes to purchase (e.g., “men’s formal wear”) and then chooses a variety from within the nest (e.g., “men’s suit from firm j ”). The first term in the above decomposition is the probability of buying the variety conditional on choosing from within some nest, while the second term is the probability of choosing a particular nest.

The equation above may appear complicated, but Berry (1994) also demonstrates the following transformation of the data that allows for estimation of model parameters in a linear setting:

$$\log s_{jt} - \log s_{0t} = \delta_{jt} - \alpha p_{jt} - \delta_{0t} + \sigma \log s_{jt/g}$$

where s_{0t} is the market share of some outside good and $s_{jt/g}$ is the within group share of product jt . The equation above is not identified as there are more parameters than data points. To this

end, we follow the literature and treat quality as a residual in an estimation of α via regression. More concretely, we split the quality parameter into a time fixed effect, a firm-product “average quality” fixed effect, and a residual:

$$\delta_{jt} = \delta_j^1 + \delta_t^2 + \delta_{jt}^3$$

Plugging in yields the following estimating equation:

$$\log s_{jt} - \log s_{0t} = \delta_j^1 + \lambda_t - \alpha p_{jt} + \sigma \log s_{j/g} + \delta_{jt}^3 \quad (4)$$

where $\lambda_t = \delta_t^2 - \delta_{t0}^2$ is growth in quality *relative* to the outside good—which we highlight here as it will be important later on. In general, δ_{jt}^3 is correlated both with price and the nest share. This motivates an IV approach, and we pick up the discussion of constructing instruments in subsection 5.4.

5.2 Firms’ Decisions and Timing

Since our estimation only relies on consumer choices, we do not need to fully specify firms’ production. However, our instrumenting strategy relies crucially on the timing of firms’ decisions. To that end, we outline the firms’ decision making in this subsection. For the remainder of this section, we also drop all time subscripts as we will assume that firms solve the same problem anew in every period.

We assume that each period can be broken into three stages. In the first stage, firms decide on their quality and production plan given expectations about costs and demand. In the second stage, a vector of costs shocks is realized and the firm produces. Finally, in the third stage, they set prices and compete. The timing here is typical in quality models and is similar to that found in Sutton (2001, 2012). We think this timing assumption is particularly reasonable in the apparel industry, given its institutional features. For example, Frederick and Staritz (2012) walk through the production process for a typical apparel firm and argue that the design and planning process happens before production takes place while marketing and selling logistics occur after. We also assume that firms can solve their problem statically in each period—so that there are no adjustment costs for price or quality.

To formalize the above, we work backwards. And so, in the final stage of the period, after all cost shocks are realized, firms set their prices and compete. Suppose as in Berry et al. (1995)

(henceforth BLP) that there are F firms active on the market producing differentiated products. Each firm produces a subset Γ_f of the J products available on the market. Consider first the *short run* profit function of firm f :

$$\begin{aligned}\Pi_f &= \sum_{j \in \Gamma_f} (p_j - mc_j) q_j \\ &= \sum_{j \in \Gamma_f} (p_j - mc_j) M s_j(p, \delta; \vartheta)\end{aligned}$$

where q_j is the quantity of good j produced by the firm, p_j is the price of the product, mc_j is the marginal cost, M is the size of the market and s is market share, that depends on the price vector, as well as the unobserved quality of the good, δ_j .

Maximizing profits with respect to price, we get the following FOC:

$$s_j(p, \delta; \vartheta) + \sum_{j \in \Gamma_f} (p_j - mc_j) \frac{\partial s_j(p, \delta)}{\partial p_j} = 0$$

Given the pricing strategy, in the first stage the firm's expected profit is given by,

$$E \left(\sum_{j \in \Gamma_f} (p_j(\boldsymbol{\delta}, \epsilon) - mc(\boldsymbol{\delta}, \epsilon)) M s_j(p, \boldsymbol{\delta}; \vartheta) \right)$$

where bold denotes the whole vector of quality across firms. The key to our instrumenting strategy is the presence of the cost shocks, ϵ . We assume that firms set prices *after* these have been observed but choose quality based on expected profits. This assumption will allow us to exploit the orthogonality between cost shocks and unobserved quality in estimating the demand model's parameters. This strategy parallels the proxy method of estimating production functions where one uses assumptions about timing of investment and hiring decisions relative to realization of productivity innovations to identify certain parameters.

5.3 Nest Structure, Trade and Market Size

In this subsection we give some detail on how apparel products are defined in the Combined Nomenclature, and how we use this to construct nests. Our nests are based on combining CN 4 digits codes. Our nests were chosen so that they approximately group clothing by its purpose (formal

wear, undergarments, etc.) and by the gender of the intended consumer (male or female). Within each nest, we observe products at the 8 digit level. These are highly disaggregated and normally include more detail on the type of garment, the material and sometimes other characteristics or weight. For example, some products are “Men’s suits, of wool or fine animal hair, knitted or crocheted” and “Women’s knee-length stockings, measuring per single yarn less than 67 decitex, of synthetic fiber.” In total there are 16 nests listed in Table 2.

We will define a *variety* as a CN8 code at a particular firm. Thus, if 2 firms both make men’s wool suits, then they are counted as two separate varieties. This structure leaves us with around 3,000 varieties in the sample. For each of these varieties, we know sales to the nearest 1000 DKK and the quantity (usually measured as a count of goods sold).

Before discussing our instrumenting strategy, some discussion of the outside good is in order. In our setting, because we use time-fixed effects, the choice of outside good does not effect our estimates of any parameters or elasticities. However, the outside good will largely determine the shape of the time-fixed effects which determine aggregate changes in quality over time. For the outside good, we use the total quantity of imports into Denmark.¹⁹ This means that after quotas fall and imports into Denmark dramatically increase, the outside good grows and this influences our quality estimates. We will discuss interpreting this more in the results section. But, in effect, this fact implies that we can identify changes in the spread of quality, but cannot easily identify changes in the overall average quality of goods across time.

5.4 Instrumenting Strategy

Our regression equation runs into the standard endogeneity present in demand estimation: price and the nest share will be correlated with a firm’s quality. To circumvent this issue, we employ an IV strategy. However, finding exogenous instruments in our situation is still difficult as quality and price are both decision variables. Thus, many cost shifters that could be used to instrument for price—e.g., wages—also reflect input quality. However, as outlined in the timing section we use *unanticipated* cost shocks as instruments. This strategy is the same as that followed by Foster et al. (2008) who used structural estimates of innovations to firm’s productivity as instruments for demand shifters. However, as estimating production functions in a differentiated market is difficult,

¹⁹At this time, our estimation ignores the distinction between domestic sales and foreign sales by firms. This is because without more data on processing trade, we cannot disentangle between imports for final sale and re-exports. Unfortunately, the extent to which this might bias our results relies on covariance between domestic market shares and the share of sales going to exports. In follow up work we hope to more carefully distinguish these two sources of sales.

we construct directly observable cost shocks.

Denmark’s size and location within the EU leads to an economy where the vast majority of firms engage in some trade. Our instrumenting strategy relies on the idea that trade carries the risk of unanticipated cost shocks to the firm. In particular, we will use forecast errors on exchange rates as instruments. The source of variation arises from cross-sectional heterogeneity in import mixes across firms.

More explicitly, we model log exchange rates as an AR(1) process:

$$\epsilon_{ct} = \mu_c + \rho_c \epsilon_{ct-1} + \sigma_c z_{ct}$$

where c indexes countries, $z_{ct} \sim \mathcal{N}(0,1)$, σ_c is the error variance and (μ_c, ρ_c) govern the AR process.²⁰ After estimation we use forecast errors, $\hat{\eta}_{ct}$, defined as, $\epsilon_{ct} - \widehat{E}(\epsilon_{ct})$, to construct our instrument as follows:

$$\zeta_{ft}^1 = \sum_c \hat{\eta}_{ct} s_{ft,c}^{imps}$$

where f indexes firms and $s_{ft,c}^{imps}$ is the share of firm f ’s apparel and textile imports that are from country c .²¹ In our sample, 97% of firms import apparel or textiles from outside the Eurozone. For the remaining 3% we set the value of their shock to zero, as they experience no exchange rate uncertainty. Notice that this instrument is measured annually at the *firm* level, while the demand equation is at the product level. Hence, we cluster all errors at the firm-year level. Even with this conservative clustering strategy, we show in the results section that this is a powerful instrument for price. We construct an analogous measure using exports and include this as a control. This is an important control because the extent to which exchange rates are cost shocks are implicitly hedged by the fact they are positive shocks to export profitability. This mechanism is highlighted by Amiti et al. (2014).

To instrument for the nest share parameters, we use sales weighted averages of the cost shocks across a firm’s *competitors* within a nest. This is similar in spirit to the approach used by BLP, who use own product characteristics as instruments for price and average characteristics of firms’

²⁰We experimented with different specifications of this process but the best fit is typically a simple random walk with $\rho = 1$.

²¹For those countries that are part of the European Exchange Rate Mechanism (ERM 2), we set the shock equal to 0. This is because the Danish Kroner is pegged to the Euro (and varies less than 1% around the peg).

competing products as instruments for the nest share. The instruments are constructed as follows:

$$\zeta_{ft}^2 = \sum_{f' \neq f} s_{f't}^{sales} \zeta_{f't}^1$$

In the estimation, we then interact ζ_{ft}^2 with a dummy for each nest. With or without these interaction terms, the nest instrument is strong, with a very low first stage p value. However, we found that fitting a linear function given the large differences in nest sizes leads to weak identification in the second stage. Allowing for the interactions yields strong instruments and precision in the second stage. When we present results we present p -values for both the second stage *and* the first stage.

As a final point on our instruments, since sourcing strategies are endogenously determined alongside quality, one may think our instrument invalid. However, even if there is a systematic relationship between quality and the exchange rate risk posed by different countries, this does not mean that *unanticipated* exchange rate errors and unobservable quality are correlated. That is, it should always be the case that $E(\delta_{it} \zeta_{ft}^1) = 0$ given our timing assumption and given that forecast errors are mean 0.

To conclude this section, we briefly discuss the clustering strategy and particular choice of estimation method. Our instruments are firm level while the unit of observation is a product. It is also plausible that unobservable quality decisions may be autocorrelated for a particular product. To address both of these concerns, we employ a two-way clustering strategy. In particular, we allow for arbitrary correlation of demand across products within a firm each period, and across time within each product.

6 Results

In this section we analyze the results of our estimation. In the first subsection, we comment on the parameter estimates and compare them to other estimates in the literature. We find our measures of price elasticity are well within the literature’s bounds, giving us confidence in our measures of quality. Then we turn to analyzing aggregate trends in the shape of the distribution—and demonstrate a break in aggregate trends that occurs when China enters the WTO. In particular, we find that the dispersion in quality across firms, measured as the distance between the lowest and highest quality firms, decreases rapidly after China enters the WTO. With this in mind, we turn

to a firm-level analysis, highlighting the heterogeneous response of firms that induce this reduction in the dispersion of quality.

6.1 Parameter and Quality Estimates Overview

In this subsection we review our parameter estimates. We also analyze how our quality estimates correlate with various dimensions of the firms—such as size, the price of goods, and input prices. Turning to the parameter estimates, for the sake of comparison, we run an OLS regression, a logit model (That is, no nesting) and the full nested logit. Table 3 contains our estimates. As expected, the OLS estimates are biased upwards, and the price coefficient is not significantly different from zero. In the logit specification, the price coefficient is larger and of the correct sign but imprecisely measured. In the complete model all coefficients are significant and of the expected sign.

Our estimated price coefficient of $-.0077$ falls comfortably in the range of parameters estimated by Khandelwal (2010). Khandelwal used HS 2 headings to define a market (whereas we combine headings 61 and 62). He also defined a nest an HS 6 digit code and the unit of analysis was a country-good pair. Despite the differences in aggregation, we believe it is a useful benchmark. He found a median coefficient on price of $-.001$ and an IQR for of $.070$, which places our coefficient close to this range.

Before turning to the main analysis, we explore the plausibility of our results by seeing how our quality estimates compare to observable firm characteristics. We focus on three features of the data: the price elasticities implied by the model, the size-price premium, and wages. The first is useful as there are other studies on what reasonable elasticities of substitution ought to be. The second feature of the data, the size-price premium, is a stylized fact explored by other authors. We show that their findings are corroborated by our estimates, lending credibility to our interpretation of our demand shifters as quality. Finally, we use the mean wages as way to check whether higher output quality is associated with higher input quality. This implicitly assumes that higher wages reflect higher input quality. While all of these checks are imperfect, combined they give credibility to the idea that our estimates of quality reflect meaningful differences in firms outputs and inputs.

To this end, first we calculate the price elasticities implied by the nested logit model according to the following formula:

$$\left| \frac{d \log s_j}{d \log p_j} \right| = \varepsilon = \alpha p_j \left[\frac{1}{1 - \sigma} - s_j - \frac{\sigma}{1 - \sigma} s_{j/g} \right]$$

where α is the price coefficient and σ the substitution parameter. Figure 4 contains the density of elasticities implied by our estimates. The mean elasticity is 1.90, while the median is 1.66. These are in line estimates of price elasticities in the IO literature using similar demand techniques. There is substantial heterogeneity within nests, and table 4 contains summary statistics by nest for the 5 largest nests. In this table we see that cross-nest heterogeneity in elasticities can be very high. Indeed, the elasticities may seem implausibly high for some nests (e.g., women’s coats). This likely reflects the model’s rigidity in having a single price parameter. However, we think this is a reasonable cost to pay for the simple and straightforward estimation. Crucially, there is a tight correlation between the implied price elasticities and quality (which is not imposed by the model). This lines up with intuition that purchasers of higher quality goods tend to be more price inelastic. Nevertheless, formally capturing this kind of heterogeneity in our model is difficult without data on the identity of consumers.

Next, we explore how our estimates of quality compare to prices and firm size. Table 5 below summarizes the correlation between price, quality, elasticity and size. As one can see, quality and price are highly correlated. To see the relationship further, figure 5 plots this relationship with nest and year means removed. While substantial residual heterogeneity remains, a glance at the red curve (a lowess fit) shows the strong positive relationship between price and quality. Turning to size, measured by employment, the price-size correlation is tighter than the quality-size correlation. Kugler and Verhoogen (2012) suggests that quality may account for the correlation between size and price. This is because larger firms may be able to access higher quality inputs and thus produce higher quality outputs. We assess the validity of our quality estimates by testing this relationship in our data. To that end, we run the following regression:

$$\log P_{jft} = \alpha_j + \alpha_t + \beta \log Emp_{ft} + \gamma \delta_{jft} + \epsilon_{jft}$$

where δ_{jft} is our estimate of quality, and α_j and α_t are product and time fixed effects. We also run a regression with firm-product pair fixed effects to see if changes within a firm-product align with Kugler and Verhoogen’s observation. The results of this set of regressions are in table 6. These regressions confirm that on average, larger firms supply a higher quality and charge a higher price. Nevertheless, there appears to be a residual relationship—suggesting that other forces (e.g., market power) are likely still important.

Finally, we turn to a comparison of our quality estimates and wages in the firm—which we

think may proxy for input quality. Recall that our concept of quality is multifaceted and captures the whole range of firm decisions that increase the attractiveness of a good—for example, design, marketing, and materials. However, we still assume that the firm has *control* over these decisions and that quality does not merely reflect taste shocks. To be clear, we are not estimating a quality production function here; rather we are demonstrating that a relationship between input and output quality exists.

In order to perform this exercise, we first aggregate quality to a firm-level measure. This is necessary because we do not observe product-specific inputs, but is difficult because we do not want to conflate product composition across firms with the quality of a firm. Part of this difficulty arises because of the difficulty in defining a “good” precisely, and the desire to separate certain goods in our analysis. For example, we may want to treat women’s pantyhose as separate from women’s bras as one may have different market shares than the other because of tastes not related to quality. On the other hand, we may *not* want to separate men’s coats of fur from men’s coats of wool, as these sorts of difference may reflect actual differences in the physical quality of a good. The compromise we adopt is removing fixed effects at the 5 digit level from goods.²² Let $\tilde{\delta}_{jft}$ be the quality of product j at firm f with fixed effects removed. Then define

$$\tilde{\delta}_{ft} = \sum_{j \in \mathcal{J}_f} \frac{s_{fjt}}{S_{ft}} \tilde{\delta}_{jft}$$

to be the firm-level quality where s_{fjt} represents the sales of product j by firm f at time t , and S_{ft} is the total sales of firm f at time t .²³ With firm level quality defined, we examine the relationship between firm’s average wage ($\log \bar{w}_{ft}$) and average quality through the following regression:

$$\log \bar{w}_{ft} = \beta_t + \beta_1 \tilde{\delta}_{ft} + \beta_2 \mathbf{X}_{ft} + \varepsilon_{ft}$$

where β_1 is our object of interest and \mathbf{X} collects control variables such as size (to proxy for productivity). The results of these regressions are in Table 7. As can be seen from the first two columns, firm quality and average wages are strongly positively correlated. When we include a measure of the skill share in the firms (defined as the ratio of workers with college or post-college credentials over all workers), the coefficient on quality becomes insignificant. This suggests that much of the “qual-

²²In the apparel industry, the 5 digit level *almost* exactly corresponds to the 6 digit level. We did not use HS 6 fixed effects because for a handful of products these 6 digit codes are not consistent over time. Nevertheless, this is almost like removing an HS6 fixed effect.

²³Our results are robust to quantity based weighting.

ity premium” may be a composition effect—higher quality firms seem to be paying more because they are employing better workers. This is reinforced in the final two columns of the table. We regress separately wages for high-skilled (those with 4+ years of education) and low-skilled (other) workers. We find that among high-skilled workers there is a quality premium, but not among low skilled workers. The collective evidence demonstrates that our quality measures reflect a feature of products that is both meaningful to consumers (in that they are willing to pay a higher price) and to firms (in that they pay higher input prices). Having established the validity of our quality measure, we turn to an analysis of quality and trade in the next two subsections.

6.2 Quality Evolution and Quality Ladders

Now we turn to looking at how the distribution of quality shifts over time. Keep in mind that we cannot credibly identify growth in the level of quality over time without strong assumptions on the behavior of the outside good. However, our model makes strong predictions about how the *shape* of the quality ladder evolves in response to new sourcing opportunities. In particular, it suggests that dispersion in quality across firms should decrease. Moreover, we know that import competition should have implications for entry and exit of firms. To understand the dynamics at play, we perform two exercises in this section. First, we decompose changes in aggregate quality in order to highlight the problems with looking at means over time as well examine entry and exit patterns. In the second part of this section, we turn to looking at the length of the quality ladder, which we define as the distance between lowest and highest quality goods in a market. With this definition in hand, we examine how this length changes in response to Chinese entry to the WTO.

In order to decompose how quality evolves over time, we first define an aggregate measure of quality by using a sales weighted mean across firm-product pairs:

$$Qual_t = \sum_j \delta_{jt} \frac{sales_{jt}}{\sum_i sales_{it}}$$

We can decompose the above measure in a way similar to that used in the productivity literature. In particular, we break changes in aggregate quality into an average growth term, as well as within-firm-product, across-firm-product and a covariance term as follows:

$$\Delta Qual_t = \underbrace{\delta_t - \delta_{t-1}}_{\text{Time Trend}} + \underbrace{\frac{N_{Ent.}}{N_t} \delta_{it}^{Ent.} - \frac{N_{Exit}}{N_{t-1}} \delta_{it-1}^{Exit}}_{\text{Entry/Exit}} + \underbrace{\Delta \delta_{it}^{Stay}}_{\text{Within Growth}} + \underbrace{[Cov(\delta_{jt}, s_{jt}) - Cov(\delta_{jt-1}, s_{jt-1})]}_{\text{Reallocation}}$$

where s_{jt} is the sales share of product j at time t . The first term measures the secular change in quality, contaminated by changes to the outside good. The second captures the effect of entry and exit. The third captures the changes to firms that are present in both periods (Notice that within this effect are two smaller effects: the actual idiosyncratic changes to quality of surviving firms as well as the shifts in weight that these firms receive in the aggregate calculation). The last term captures the covariance between market share and quality. Figure 7 plots the evolution of our aggregate measure of quality and its erraticism suggests care needs to be taken in interpretation.

First, we explain the problem with analyzing the time trend. Figure 8 plots the time fixed effects and shows that they drive the majority of aggregate changes in quality. Ideally this term would reflect secular growth in quality and we could look at quality changes relative to this trend to determine if goods are downgrading or upgrading in absolute terms. The issue is that our method can only identify $\delta_{it} - \delta_{0t}$ where δ_{0t} is the outside good's quality. So if δ_{0t} is changing then we cannot estimate the trend of δ_{it} . In particular, a positive supply shock to the outside good (which the MFA would be) will look like a negative demand shock to the inside good. However, there is still much to be learned from the other terms, which we turn to now.

Now looking at entry and exit, figure 9 plot their respective contributions to changes in aggregate quality. There is a sharp upward trend in this graph and the sign reverses when China enters the WTO. This means that after the WTO shock, new entrants tend to produce higher quality goods relative to incumbents. Our model predicted that entry and exit patterns ought to depend on the joint distribution of productivity and quality. While we cannot directly observe productivity, the findings on exit and entry are certainly plausible to the extent that for marginal firms, quality and profitability are positively correlated. However, we cannot directly corroborate this. Regardless of its connection to the model, the analysis of the time fixed effects and entry/exit patterns both suggest that the end of the MFA not only opened up new sourcing opportunities, but also led to substantial import competition.

To conclude the discussion of the decomposition, we now turn to the covariance terms. Figure 10 plots the evolution of this term. This term is flat for some time before beginning a downward trend in 2003. Two forces may drive this decrease in the covariance. First, it could be that there is less dispersion in quality. To see why, notice that in the extreme if the distribution of quality was degenerate, then this covariance would trivially be zero. On the other hand, this could be driven by changes in other drivers of consumer behavior, such as price. Thus, we avoid drawing strong conclusions from this term; instead, in the remainder of this section, we define our quality ladder

concept and demonstrate directly how the dispersion in quality has declined.

To formally explore our predictions, we first define the quality ladder concretely. To this end, first define good j 's position in the quality ladder at time t as:

$$l_{jt} = (\delta_{jt}) - \frac{1}{n_t} \sum_{i=1}^n (\delta_{it})$$

Thus, it is the product's quality relative to the mean quality of all other products present in that year. Clearly this definition removes the aggregate time component from each year, thus we cannot explore changes to average quality over time. This also means that our definition of the quality ladder is insensitive to the choice of the outside option. With this definition, the quality ladder is *cardinal*: the magnitude in difference between positions is a measure of the quality difference between products. This is clear from the estimation strategy since a larger magnitude of quality, holding price fixed, maps one-for-one into higher market share. Thus, the changing shape of the quality ladder gives insight into aggregate changes.

To understand how the quality ladder looks, figure 6 plots the density of the ladder measure at the beginning and the end of our sample. There is less dispersion over time—so that the quality ladder tightens. To make this clearer, figure 11 plots the evolution of ladder length over time, where length is measured by subtracting the max and min quality.²⁴ To analyze the possible impact of China's accession to the WTO and the MFA, we run a structural break test for 2001, the results of which are in table 8. The trend coefficient nearly doubles in magnitude after 2001. With such a short panel, it is hard to determine statistical significance, so in the next section we use firm level variation to explore how firms respond to new import opportunities. Nevertheless, this evidence is suggestive that the end of the MFA played an important role in aggregate trends.

The tightening of the quality ladder is consistent with two forces at play: import competition driving out lower quality firms and offshoring opportunities inducing compression of the ladder as some firms upgrade and other firms downgrades. We have already established that the first force—entry and exit—is important. In the next section, we will use firm-level variation to show evidence of heterogeneity in firms' response to new offshoring opportunities.

²⁴We also used standard deviations and the distance between 95th and 5th percentiles as measures of length. The general trends are robust to these differences, but not all regressions remain statistically significant. In table 8 we report results for different measures of length.

6.3 Firms' Quality and Sourcing Decisions

In this final subsection we use firm-level variation, as well as the fall of quotas with the end of the MFA, to analyze the relationship between firms' sourcing and quality decisions. In our model we had a sharp distinction between sourcing destinations. In reality, there is a mixture of home production and foreign production. In this section we refer to this mixture of domestic and foreign production as offshoring and we also define a measure of this activity. With this measure, the key predictions of our model we explore are (1) the negative relationship between offshoring activity and quality, (2) heterogeneity in firms' quality upgrading and downgrading decisions in response to new offshoring opportunities and (3) the relationship between a firms' quality and whether or not they should respond to new offshoring opportunities (in particular, the extent to which quality predicts purchases of Chinese textile and apparel goods).

In order to discuss the sourcing decisions of firms we focus on imports of finished or nearly finished apparel goods. This corresponds to what Hummels et al. (2014) call "narrow offshoring." The idea is to reflect the part of the supply chain that would most likely be performed by manufacturing plants in apparel that performed all production domestically. For example, it is highly unlikely that apparel firms would manufacture their own textiles from raw cotton, but could presumably make shirts at home or abroad. We focus on this concept because it represents the imports that are closest to the product that firms ultimately sell. Following Hummels and coauthors, as well as Autor et al. (2013), we define offshoring to be imports of apparel divided by the number of domestic employees. Thus, while we use the term offshoring here for convenience, it is shorthand simply for apparel imports per head at each apparel firm. With this measure we hope to capture the idea that the less employees per import volume one has at home, the less these employees are directly involved in production—whether it be actual manufacturing work or even quality control, distribution or other activities.²⁵ Figure 12 plots the time series of average offshoring by firms. While there is some volatility in the measure, there is a broad upward trend in offshoring over the sample period. Figure 13 focuses on offshoring to China. Overall, there is a .8 log point increase in overall offshoring activity in our sample and a 1 log point increase in offshoring to China. Thus, off-

²⁵Discussions with people in the industry as well as the literature on fashion suggests that outsourcing abroad is normally done at arms length. After production, goods are often shipped back to their home country for quality control, inspection and distribution. This is especially true of branded manufacturers in Europe. This suggests a vital role for domestic employees at these firms, even if not directly involved in production. Moreover, by lowering overhead without a concurrent decrease in output, firms may be spending less time on quality control, managing their sourcing decisions or design. We test and verify this empirically by showing the relationship between our measures of quality and offshoring.

shoring to China doubled and, in general, imports per head increased substantially for the apparel industry.

Before moving on, we address two potential concerns with our measure of offshoring. First, more productive firms may require less workers to generate the same amount of final output, suggesting a positive bias between offshoring and productivity. To deal with this fact, we include export values and the value of intermediates as controls for productivity and size. Many papers in the trade literature suggest that exports and intermediates should be highly correlated with productivity (for example, see Melitz (2003) and Halpern et al. (2005)).

Second, the value of imports contains price information and will be correlated with quality. To address this source of bias we do two things: first, we scale by the wage bill at a firm instead of the headcount of employees, and all of our results are robust to this; second, we use the import share weighted per capita GDP of import partners as a proxy for input quality. We think this is a good proxy as recent work in the trade literature, such as Manova and Zhang (2012) and Khandelwal (2010), suggest that per capita GDP is highly correlated with the quality of a country’s exports.

With our definition of offshoring in hand, we turn to testing our model predictions. First we explore the relationship between offshoring and quality by running regressions of the form:

$$l_{jt} = \alpha_g + \beta \times \text{Offshoring}_{ft} + X_{ft}\gamma' + \epsilon_{jt}$$

where l_{jt} is the relative ladder position of good j at time t , α_g is a product fixed effect, and X is a set of controls and year fixed effects. Since our measure of offshoring is constant within a firm, we follow the same clustering strategy as in our structural estimation. Table 9 displays the results of our regressions for several choices of controls. In the first column, the coefficient on offshoring is positive, which may partially reflect productivity as an omitted variable. After we control for productivity, using exports and intermediates as proxies, the coefficient on offshoring is negative. In column (3) we allow for the effect of offshoring to depend on input quality, which makes the negative relationship stronger. The relationship outlined in this section suggests that, on average, firms that intensively source production from abroad will produce lower quality goods than other firms.

While the preceding paragraph discussed the relationship between offshoring and quality in a cross-section of firms, we now turn to the heterogeneous response of firms to *new* offshoring opportunities. To assess how cheaper inputs induce changes in firms’ relative position in the quality

ladder we run regressions of the form:

$$\Delta l_{jt} = \beta_1 \times \Delta \text{Offshoring}_{ft} + \beta_2 \times \Delta \text{Offshoring}_{ft} \times l_{t-1} + X_{ft} \gamma' + \epsilon_{jt} =$$

where once again X is a set of controls that includes year fixed effects. We also include lagged offshoring as a control variable. We find this is important as it appears that there may be diminishing returns to offshoring, or that firms find it more difficult to increase offshoring beyond a base amount. Our objects of interest are β_1 and β_2 . In particular, β_2 allows movement along the quality ladder to one's initial position. Our model predicts that $\beta_1 > 0$ and $\beta_2 < 0$.²⁶ Table 10 contains the results of this regression. In the next two paragraphs we discuss β_1 and β_2 in turn.

Notice that the coefficient on growth in offshoring, β_1 , is positive in all specifications. Thus, firms experiencing an increase in offshoring activity also experience an increase the quality of their output. This fact, along with the fact that offshoring firms tend to have lower quality output, suggests precisely that increasing offshoring is associated with a tightening of the quality ladder. In particular, lower quality firms are more likely to engage in offshoring and also upgrade their quality relative to other firms.

We explore heterogeneity in offshoring and quality, measured by β_2 , in column 4. We find that, controlling for productivity and input quality, the response of quality to offshoring is negative for middle and higher quality firms while remaining strongly positive for lower quality firms. This need not be the case and our model demonstrates how the marginal cost of increasing quality in different countries determines firms' response to new sourcing opportunities. As a caveat, remember that these regressions are about relative movement over time, not absolute changes. So, it could be that all firms upgrade their output quality but some firms engage in more or less than others.

In the last part of this section, we turn our focus to China. China is a useful test case both because of its rapid global rise but also because we can exploit that China's access to the WTO led to lower quotas under the MFA. However, these quotas were not applied to all products—hence, we can use quota fill rates as a measure of whether the China shock actually mattered to a particular product. The first thing we do is rerun our first set of regressions with a focus on offshoring to China:

$$l_{jt} = \alpha_g + \beta \times \text{ChinaOffshoring}_{ft} + X_{ft} \gamma' + \epsilon_{jt}$$

²⁶Technically, it suggests a more complex relationship with difference between middle and high skilled firms. We attempted to include a quadratic term, and while an F test found all coefficients jointly significant they were not well identified separately.

The first two columns of table 11 present the results of these regressions. Once we control for confounders, we find a significant negative relationship between offshoring intensity to China and output quality. Next we turn to our predictions regarding firms' heterogeneous responses along the quality ladder. Here, a regression that ignores the possibility of heterogeneity finds that, even with a full set of controls present, there is no relationship between increasing offshoring and changes in quality. However, when the response is allowed to vary along the quality ladder the coefficients become significant and mirror previous findings.

Next, we use China's entry to the WTO and the dismantling of the MFA to assess how firms respond to an exogenous change in offshoring opportunity. While typical models predict a positive relationship between import intensity and productivity, our model suggests a negative correlation between entry into offshoring to low quality countries and initial quality. To explore this effect quantitatively, we use the sharp drop in quotas that China experienced upon accession to the WTO, first in 2001, then again in 2005 as the MFA came to a close. Our source of variation comes from the fact that these quotas did not affect all goods, as the quotas were only occasionally binding. We use a continuous measure (quota fill rates) but one may think of this as a treatment on certain products and not others. Our measure of quota fill rates refers to the quantity of apparel imports from China over the total quota allotted to China. Note that this is set at the EU level.²⁷ With this discussion in mind, we run the following probit regression:

$$Y_{\text{Source from CN}}^* = \beta_0 + \beta_1 \times \text{Quality}_{jt} + \beta_2 \times \text{Quota Fill Rate}_{jt} + \beta_3 \times \text{Quality}_{jt} \times \text{Quota Fill Rate}_{jt} + X'_{jt}\gamma + \varepsilon_{jt}$$

where Y^* is the latent propensity to offshore, quality refers to our quality ladder measure and the quota fill rate was described above. Table 12 contains the regression results for various choices of controls. In all specifications, the coefficient on the ladder position is negative, and is significant when one controls for productivity. Interestingly our final specification shows that size is also an important determinant of offshoring activity. This suggests that ultimately both productivity and capability determine one's offshoring choices, and understanding which firms and how firms respond to trading opportunities depends on the quality of their outputs.

²⁷All quota information comes from the EU's SIGL database.

7 Conclusion

In this paper, we developed a model of offshoring and quality decisions by firms. In particular, we showed that when firms can outsource production, whether they choose to downgrade or upgrade their product quality will hinge on whether they are a low-capability or high-capability firm. We use detailed information on Danish apparel, including products level data on production, imports and exports to estimate a demand model and recover unobserved product quality. Our demand estimates are found to be in line with the previous literature. For example, we find that our estimated quality differences between firms explain the size-price relationship documented by Kugler and Verhoogen (2012). Moreover, we use our demand estimates and our trade data to demonstrate a heterogeneous response of firms. We find that firms producing at an initially medium or high quality downgraded their quality relative to low quality firms, who upgraded theirs.

To make more headway on the role of trade we use the dismantling of apparel quotas as a source of exogenous variation to the Danish apparel industry's access to foreign input markets. We use two facts: first that China's entry to the WTO immediately lowered quotas on Chinese goods; second these quotas fell unevenly across different products. We find that firms' product quality is strongly affected by the change in the competitive environment and offshoring opportunities. In particular, we find strong evidence that the distribution of quality not only tightens, but patterns of entry into new import markets are driven by quality. The importance of this dimension in firms' decision making holds even when we control for size and other proxies of productivity.

This tightening of the quality ladder, and the heterogeneous response causing it, complicate previous analyses of firms, especially those in vertically differentiated industries where quality plays a major role. In particular, our model and our results demonstrate that two separate attributes of a firm determine how they respond to trade liberalization: their productivity as well as their capability, the ability to produce high quality goods. How different industries respond to shocks will depend, ultimately, on the joint distribution of these two firm attributes. This suggests that the response of other variables, such as prices and wages, to a trade shock may depend on whether the industry affected by a shock is more vertically or horizontally differentiated.

Our work leaves open several avenues for future research. In particular, while we have documented how offshoring impacts the quality ladder, we have said nothing about the actual welfare changes induced by such a cost shock. While exact welfare calculations are sensitive to the demand system considered, as well as how one defines consumers' outside options, it would still be interest

to determine the extent to which the cost savings induced by offshoring are passed through into quality-adjusted prices. This would require modeling more carefully the quality and physical production function of firms—which requires product-level input data or a model of how inputs are allocated across products. Future work should also bring in wholesale firms and foreign competitors at a more granular level. This would allow researchers to analyze how offshoring opportunities change the entire industrial structure in vertically differentiated markets, and not just the menu of quality offered by domestic firms. Finally, our estimation procedure relies on income-independent tastes for price and quality. However, if one also had information on the consumers at various firms, future work could attempt to estimate richer demand systems that allow for more precise estimates of quality and richer interactions between price, income heterogeneity and quality.

References

- Amiti, Mary and Amit Khandelwal**, “Import Competition and Quality Upgrading,” *The Review of Economics and Statistics*, 2013, *95* (2), 476–490.
- , **Oleg Itskhoki, and Jozef Konings**, “Importers, Exporters, and Exchange Rate Disconnect,” *American Economic Review*, 2014, *104* (7), 1942–78.
- Autor, David H., David Dorn, and Gordon H. Hanson**, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, October 2013, *103* (6), 2121–68.
- Berry, Steven, James Levinsohn, and Ariel Pakes**, “Automobile Prices in Market Equilibrium,” *Econometrica*, 1995, *63* (4), 841–90.
- Berry, Steven T.**, “Estimating Discrete-Choice Models of Product Differentiation,” *RAND Journal of Economics*, Summer 1994, *25* (2), 242–262.
- Bloom, Nicholas, Mirko Draca, and John van Reenen**, “Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity,” NBER Working Papers 16717, National Bureau of Economic Research, Inc 2011.
- Brambilla, Irene, Amit Khandelwal, and Peter Schott**, “China’s Experience under the Multi-Fiber Arrangement (MFA) and the Agreement on Textiles and Clothing (ATC),” in R. Feenstra and S. Wei, eds., *China’s Growing Role in World Trade*, National Bureau of Economic Research, Inc, 2010, pp. 345–387.
- Costinot, Arnaud, Jonathan Vogel, and Su Wang**, “An Elementary Theory of Global Supply Chains,” *Review of Economic Studies*, 2013, *80* (1), 109–144.
- Feenstra, Robert C.**, *Offshoring in the Global Economy: Microeconomic Structure and Macroeconomic Implications*, MIT Press, 2010.

- Fieler, Ana Cecilia, Marcela Eslava, and Daniel Xu**, “Trade, Skills, and Quality Upgrading: A Theory with Evidence from Colombia,” NBER Working Papers 19992, National Bureau of Economic Research, Inc 2014.
- Foster, Lucia, John Haltiwanger, and Chad Syverson**, “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?,” *American Economic Review*, March 2008, 98 (1), 394–425.
- Frederick, S. and C. Staritz**, “Success and Upgrading After the End of the MFA,” in Gladys Lopez-Acevedo and Raymond Robertson, eds., *Sewing Success? Employment, Wages, and Poverty following the End of the Multi-Fibre Arrangement*, The World Bank, 2012, pp. 87–136.
- Gereffi, Gary**, “International trade and industrial upgrading in the apparel commodity chain,” *Journal of International Economics*, 1999, 48 (1), 37–70.
- Hallak, Juan Carlos**, “Product quality and the direction of trade,” *Journal of International Economics*, January 2006, 68 (1), 238–265.
- **and Jagadeesh Sivadasan**, “Product and process productivity: Implications for quality choice and conditional exporter premia,” *Journal of International Economics*, 2013, 91 (1), 53–67.
- **and Peter K. Schott**, “Estimating Cross-Country Differences in Product Quality,” *The Quarterly Journal of Economics*, 2011, 126 (1), 417–474.
- Halpern, Lszl, Mikls Koren, and Adam Szeidl**, “Imports and Productivity,” CEPR Discussion Papers 5139, C.E.P.R. Discussion Papers July 2005.
- Hummels, David, Rasmus Jorgensen, Jakob Munch, and Chong Xiang**, “The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data,” *American Economic Review*, June 2014, 104 (6), 1597–1629.
- Khandelwal, Amit**, “The Long and Short (of) Quality Ladders,” *Review of Economic Studies*, 2010, 77 (4), 1450–1476.
- Kugler, Maurice and Eric Verhoogen**, “Prices, Plant Size, and Product Quality,” *Review of Economic Studies*, 2012, 79 (1), 307–339.
- Liegey, Jr Paul R.**, “Adjusting Apparel Indexes in the Consumer Price Index for Quality Differences,” in M. Foss, M. Manser, and A Young, eds., *Price Measurements and Their Uses*, National Bureau of Economic Research, Inc, 1993, pp. 209–226.
- Manova, Kalina and Zhiwei Zhang**, “Export Prices Across Firms and Destinations,” *The Quarterly Journal of Economics*, 2012, 127 (1), 379–436.
- Martin, Julien and Isabelle Mejean**, “Low-wage country competition and the quality content of high-wage country exports,” *Journal of International Economics*, 2014, 93 (1), 140–152.
- Melitz, Marc J.**, “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, November 2003, 71 (6), 1695–1725.
- of Business, Danish Ministry and Growth**, “Summary: Denmark at Work - Plan for Growth in the Creative Industries and Design,” February 2013.

Piveteau, Paul and Gabriel Smagghue, “Estimating firm product quality using trade data,” 2015.

Roberts, Mark J., Daniel Yi Xu, Xiaoyan Fan, and Shengxing Zhang, “The Role of Firm Factors in Demand, Cost, and Export Market Selection for Chinese Footwear Producers,” NBER Working Papers 17725, National Bureau of Economic Research, Inc January 2012.

Sutton, John, *Technology and Market Structure: Theory and History*, 1 ed., Vol. 1, The MIT Press, 2001.

– , *Competing in Capabilities: The Globalization Process*, Oxford University Press, May 2012.

Utar, Hale, “When the Floodgates Open: ”Northern” Firms’ Response to Removal of Trade Quotas on Chinese Goods,” *American Economic Journal: Applied Economics*, October 2014, 6 (4), 226–50.

Yi, KeiMu, “Can Vertical Specialization Explain the Growth of World Trade?,” *Journal of Political Economy*, 2003, 111 (1), 52–102.

Appendix A: Tables

Table 1: Most Popular Products

Top 5 Products (by # of Producers)	
1997-2002	2002-2010
Cotton tee shirts	Cotton tee shirts
Cotton women’s jerseys	Cotton women’s jerseys
Syn. fiber women’s blouses	Syn. fiber women’s jerseys
Syn. fiber women’s trousers	Syn. fiber tee shirts
Syn. fiber women’s skirts	Cotton women’s blouses

Table reports most popular products by number of producers in Denmark. Products are defined at the Combined Nomenclature 8 digit level.

Table 2: Description of Nests

Men’s	Women’s	Gender Neutral
Coats and jackets	Coats and jackets	Sweaters, jerseys, cardigans
Suits, jackets, blazers, trousers	Suits, jackets, dresses, skirts, trousers	t-shirts
Shirts	Shirts, blouses	Miscellaneous
Underwear, pajamas, gowns	Underwear, lingerie, gowns	Accessories
Sweaters, jerseys, cardigans	Sweaters, Jerseys, Cardigans	
Miscellaneous	Miscellaneous	

Apparel products defined as products in CN categories 61 (Articles of Apparel, Knitted or Crocheted) and 62 (Articles of Apparel, Not Knitted or Crocheted). Nests are based on the first four digits, but ignore the distinction between knitted/crocheted wear and non-knitted/crocheted wear.

Table 3: Demand Estimation for Domestic Apparel

	OLS	IV: Logit	IV: Nested Logit
Dep Var:	$\log(s_{jft}/s_{0t})$	$\log(s_{jft}/s_{0t})$	$\log(s_{jft}/s_{0t})$
pf_{jt}	-.00013 (-1.16)	-.02129* (-1.82)	-.00768* (-1.89)
$\log s_{jgft}$.901*** (95.32)		.321*** (3.43)
Fixed Effects:	Firm-Product, Year	Firm-Product, Year	Firm-Product, Year
Clusters:	Firm	Product, Firm-Year	Product, Firm-Year
	(188)	(1554,953)	(1554,953)
n	7,586	7,586	7,586
1st Stage p-value - Price	-	.0928	.0369
1st Stage p-value - Nest	-	-	.0000
2nd Stage p-value	.0000	.0000	.0000

Point estimates reported with t-statistics in parentheses. ***-1%, **-5%, *-10%. Dependent variable is the market share of firm-product pair f, j at time t relative to the outside good. p refers to the price measured in 2000 DKK. All specifications include firm-product and year fixed effects. Every specification drops 792 firm-products that are only observable in one period. Column (1) reports an OLS regression. Column (2) reports the results from the regression implied by a logit demand using import share weighted exchange rate shocks as an instrument for price. Column (3) reports the results from a regression based on nested logit system. There are 12 nests today, constructed based on CN 4 digit headings of goods in apparel (CN headings 61 and 62). Included instruments are imported share weighted exchange rate shocks for each firm, and, for each firm-product, the average exchange rate shock across competing firms. This latter variable is also interacted with a dummy for each nest.

Table 4: Distribution of Elasticity Estimates

Nest	Mean	Q25	Q50	Q75
Women's Dresses	2.17	1.15	2.06	2.88
Women's Shirts	1.65	.911	1.61	2.16
Men's Suits	2.53	1.39	2.27	3.36
Women's Sweaters	1.45	.690	1.24	1.99
Women's Coats	3.43	2.07	3.33	4.82

Absolute values of elasticities reported. Elasticities are calculated for each firm-product-year triple according to the equation $|d \log s / d \log p| = \alpha p \left[\frac{1}{1-\sigma} - s - \frac{\sigma}{1-\sigma} s_g \right]$ where s is total market share and s_g refers to within-nest market share. Reported means and quantiles are unweighted across all varieties within a nest.

Table 5: Correlation between Price, Size and Quality

	Quality	log(Price)	log(Employment)	Elasticity
Quality	1			
log(Price)	.1408	1		
log(Employment)	.0978	.1785	1	
Elasticity	.2234	.9210	.1749	1

All values are unweighted correlations across firm-product-year triples. Absolute values of elasticities reported. Elasticities are calculated for each firm-product-year triple according to the equation $|d \log s / d \log p| = \alpha p \left[\frac{1}{1-\sigma} - s - \frac{\sigma}{1-\sigma} s_g \right]$ where s is total market share and s_g refers to within-nest market share. Quality is estimated from a nested logit demand system, price is in 2000 DKK and employment is measured in number of employees at the firm.

Table 6: Estimating the Size-Price Correlation

	(KV)	(1)	(2)
Dep Var: $\log P_{jft}$			
log(Employment)	.1130*** (2.92)	.0989*** (2.72)	.0123 (.44)
Quality (δ_{fjt})		.0509*** (12.92)	.1189*** (5.83)
Fixed Effects:	Year, CN8	Year, CN8	Year, Firm-CN8
Cluster:	Firm	Firm	Firm
	177	177	177
R^2	.5709	.5760	.6462
N	8,132	8,132	8,132

Point estimates reported with t-statistics in parentheses. ***-1%, **-5%, *-10%. Dependent variable in all regressions is log price of a firm-product-year triple, with price measured in 2000 DKK. Employment is calculated as the number of employees at the firm-year level and quality is estimated from a nested logit demand system.

Table 7: Relationship Between Wages and Firm Quality

	(1)	(2)	(3)	Non-High Only	High Only
Dep Var: $\log w_{ft}$					
Quality (δ_{fjt})	.0245** (2.20)	.0224* (1.73)	.0155 (1.25)	.0149 (.92)	.028** (2.26)
log(Employment)		.0055 (0.32)	.0053 (0.33)	.0076 (.35)	.0079 (.61)
Skill Share			.3947*** (6.87)		
Cluster:	Firm	Firm	Firm	Firm	Firm
	(160)	(160)	(160)	(157)	(159)
R^2	.1775	.1782	.3324	.0917	.1324
N	811	811	811	804	803

Point estimates reported with t-statistics in parentheses. ***-1%, **-5%, *-10%. Dependent variable is mean wage across workers in 2000 DKK for each firm-year. High-skilled workers are defined as those with at least 4 years of college. Employment and skill shares are based on head counts of employees, while quality is calculated from a nested logit model.

Table 8: Aggregate Changes in Ladder Length over Time

Length Measures			
	$l_{max} - l_{min}$	$l_{p95} - l_{p5}$	σ_l
t	.269** (2.00)	.032 (.57)	.004 (.32)
$t \times \delta_{MFA}$	-.509*** (-1.14)	-.047 (-.79)	-.019 (-1.38)
δ_{MFA}	1017.78*** (.14)	95.58 (.79)	39.74 (1.38)
Constant	-530.20** (1.42)	-58.92 (-.57)	-7.01 (-.26)
N	14	14	14
R^2			

Point estimates reported with t-statistics in parentheses. ***-1%, **-5%, *-10%. The dependent variable, l , refers to the time-demeaned quality of a product-firm-year triple. δ_{MFA} is a dummy for $t \geq 2002$, for when the penultimate round of quota drops occurred and China was in the WTO for the first full year.

Table 9: Offshoring and Quality Ladder Position in the Cross-Section

	Dependent Variable: l_{jt}		
	(1)	(2)	(3)
$\log(\text{Offshoring})$.0599*** (2.55)	-.0298 (-1.31)	-.2356*** (-4.44)
$\delta_{ft,input} \times \log(\text{Offshoring})$.0213*** (3.99)
$\log(\text{Intermediates})$.0330*** (2.18)	.0358*** (2.29)
$\log(\text{Exports})$.1351*** (5.39)	.1448*** (5.28)
Fixed Effects: Cluster:	Product, Year Firm-Year 890	Product, Year Firm-Year 785	Product, Year Firm-Year 785
R^2	.3839	.4140	.4159
N	7,905	7,397	6,836

Point estimates reported with t-statistics in parentheses. ***-1%, **-5%, *-10%. Dependent variable is a firm-product-year's position in the quality ladder, measured as the time-demeaned quality of a product-firm-year triple. Offshoring is measured as the value of imports of apparel products (CN headings 61-62) in 2000 DKK divided by the number of workers in a firm. Exports refer specific to the value of apparel exports (CN headings 61-62) measured at 2000 DKK. Intermediates are reported for each firm-year by Stats DK, and deflated to 2000 DKK. $\delta_{ft,input}$ is a proxy for input quality at the firm level measured by the import share weighted GDP per capita of all import partners, where imports again refer to goods in CN headings 61-62.

Table 10: Offshoring and Ladder Movement - Overall and Heterogeneous Effects

	Dependent Variable: Δl_{jt}			
	(1)	(2)	(3)	(4)
$\Delta \log(\text{Offshoring})$.0596*** (2.03)	.0620* (1.92)	.0623* (1.92)	.0885*** (2.69)
$\Delta \log(\text{Offshoring}) \times l_{jt-1}$				-.0827*** (-3.22)
$\delta_{ft,input}$			-.0202 (-.52)	-.0103 (-.27)
$\log(\text{Intermediates})$.0069 (.83)	.0065 (.78)	.0074 (.99)
$\log(\text{Exports})$.0107 (.77)	.0099 (.72)	.0125 (.95)
$\log(\text{Offshoring})_{t-1}$		-.0275 (-1.29)	-.0269 (-1.26)	-.0228 (-1.14)
Fixed Effects:	CN8, Year	CN8, Year	CN8, Year	CN8, Year
Cluster:	Firm-Year	Firm-Year	Firm-Year	Firm-Year
	701	631	631	631
R^2	.0105	.0110	.0112	.0195
N	5,199	4,923	4,923	4,923

Point estimates reported with t-statistics in parentheses. ***-1%, **-5%, *-10%.
 dependent variable is the change in a firm-product-year's position in the quality ladder. This position is measured as the time-demeaned quality of a product-firm-year triple. Offshoring is measured as the value of imports of apparel products (CN headings 61-62) in 2000 DKK divided by the number of workers in a firm. Exports refer specific to the value of apparel exports (CN headings 61-62) measured at 2000 DKK. Intermediates are reported for each firm-year by Stats DK, and deflated to 2000 DKK. $\delta_{ft,input}$ is a proxy for input quality at the firm level measured by the import share weighted GDP per capita of all import partners, where imports again refer to goods in CN headings 61-62. Columns 2-3 only include those firms reporting intermediates to stats DK.

Table 11: Offshoring and Ladder Movement - Overall and Heterogeneous Effects

	Dependent Variable: l_{jt}		Dependent Variable:		Δl_{jt}
	(1)	(2)	(3)	(4)	(5)
$\log(\text{Offshoring}_{CN})$	-.0230 (-1.49)	-.0290* (-1.90)			
$\Delta \log(\text{Offshoring}_{CN})$			-.0044 (-.30)	-.113 (-.70)	.0055 (.32)
$\Delta \log(\text{Offshoring}_{CN}) \times l_{jt-1}$					-.0521*** (-3.57)
$\log(\text{Intermediates})$.0246 (1.16)		.0053 (.30)	.0106 (.71)
$\log(\text{Exports})$.1373*** (3.03)		.212 (1.27)	.0192 (1.16)
$\log(\text{Offshoring}_{CN,t-1})$				-.0141 (-1.58)	-.0141 (-1.56)
Fixed Effects:	CN8, Year	CN8, Year	Year	Year	Year
Cluster:	Firm-Year	Firm-Year	Firm-Year	Firm-Year	Firm-Year
	397	366	287	269	269
R^2	.4563	.4754	.0050	.0059	.0213
N	5,173	5,006	3,159	3,107	3,107

Point estimates reported with t-statistics in parentheses. ***-1%, **-5%, *-10%.

In columns (1) and (2), dependent variable is the change in a firm-product-year's position in the quality ladder. This position is measured as the time-demeaned quality of a product-firm-year triple. In columns (3) - (5), dependent variable is the change in ladder position over time. In all specifications, offshoring is measured as the value of imports of apparel products (CN headings 61-62) in 2000 DKK divided by the number of workers in a firm. Exports refer specific to the value of apparel exports (CN headings 61-62) measured at 2000 DKK. Intermediates are reported for each firm-year by Stats DK, and deflated to 2000 DKK. $\delta_{ft,input}$ is a proxy for input quality at the firm level measured by the import share weighted GDP per capita of all import partners, where imports again refer to goods in CN headings 61-62. Columns 2-3 only include those firms reporting intermediates to stats DK.

Table 12: Probit Regression: Probability of Offshoring to China Projected on Quality


Dep Var: Import from China	(1)	(2)	(3)	(4)
Quality (δ_{jt})	-.0313*	-.0305	-.0233	-.0757***
	(-1.66)	(-1.62)	(-1.22)	(-3.74)
Quota Fill Rate		-.2189***	-.2016***	-.1765**
		(-3.46)	(-3.20)	(-2.36)
Quality \times Quota Fill Rate			-.1161**	-.0976*
			(-2.56)	(-1.84)
log(Employment)				.5736***
				(5.01)
log(Intermediates)				.1115***
				(2.61)
Fixed Effects:	Year	Year	Year	Year
Cluster:	Firm-Year	Firm-Year	Firm-Year	Firm-Year
	941	941	185	173
Pseudo- R^2	.1080	.1090	.1096	.2618
N	8071	8071	8235	8057

Point estimates reported with t-statistics in parentheses. ***-1%, **-5%, *-10%.

All specifications are coefficients of a probit regression measuring the probability of importing apparel goods (CN headings 61-62) from China. Quality is measured as the demand shifter derived from a nested logit demand system. Quota fill rates refer to the fraction of total European imports from China over the quota on European imports from China. This quota is calculated at the product level from the EU's SIGL database. The product codes in the SIGL database strictly contain those in the Combined Nomenclature. Thus, more than one good may be bound under the same quota, but there is no ambiguity in which quota applies to which good. Column (4) includes intermediates and exports as proxies for productivity. Intermediates are reported for each firm-year by Stats DK, and deflated to 2000 DKK, while employment is measured as a headcount of employees at time t at the firm producing good j .


Appendix B: Figures

Figure 1: Time Series of Danish Apparel Import




Apparel goods are defined as those in CN headings 61 and 62. Net imports are measured as the sum of all imports minus the sum of all exports, including re-importing and re-exporting activity. Nominal apparel imports are recorded in 1000s of DKK. Weight data is in kilograms. Danish customs records include both count measures and weight measures of goods. We use weight as the units are common across goods.

Figure 2: Changes in the Danish Apparel Industry



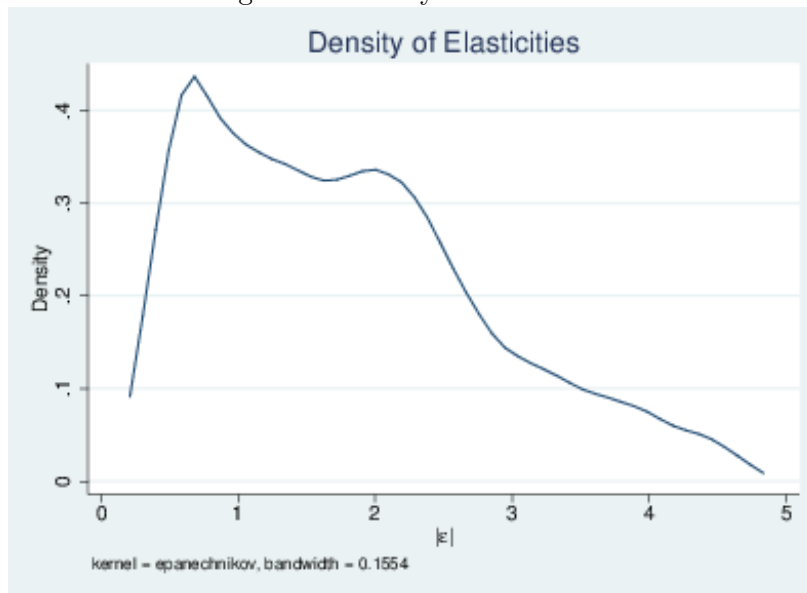
Apparel goods are defined as those in CN headings 61 and 62. We determine the largest import partners by taking the average import share across years.

Figure 3: Growth of Chinese Share in Apparel Imports



Apparel goods are defined as those in CN headings 61 and 62. We plot the 5 largest partners as measured by their average share in imports over the sample period.

Figure 4: Density of Elasticities




Elasticities are calculated with the formula $|d \log s / d \log p| = \alpha p \left[\frac{1}{1-\sigma} - s - \frac{\sigma}{1-\sigma} s g \right]$, using estimates from a nested logit demand system. This density is pooled across all firm-product-years ($N = 8132$).

Figure 5: Price versus Quality




Unit of observation is a firm-product-year triple ($N = 8132$). Price is recorded in 2000 DKK, while quality is derived from a nested logit demand system. We remove nest and year fixed effects.

Figure 6: Evolution of Quality Ladders




Unit of observation is a firm-product-year triple. The quality ladder is the demeaned distribution of residuals from a nested logit demand system for a given year. The above pools over all triples.

Figure 7: Evolution of Quality




Aggregate quality is defined as the unweighted mean quality across all product-firm-year triples in a given year. Quality for each firm is calculated from a nested logit demand system.

Figure 8: Series of Time Fixed Effects



Aggregate quality shock refers to time fixed effects estimated in a nested logit demand system.

Figure 9: Entry and Exit Component of Quality Growth




Quality growth is calculated from an Olley-Pakes decomposition of aggregate quality. The entry exit term is calculated as the difference in the mean quality of entering and exiting firm-products weighted by their respective share in total firms: $N_{entrant}/N_{firms} \overline{Quality}_{entrant} - N_{exits}/N_{firms} \overline{Quality}_{exits}$.

Figure 10: Market Share-Quality Covariance Evolution




The covariance component is calculated from an Olley-Pakes decomposition of aggregate quality. This measures the covariance between the market share of a firm-product (measured in sales) and the firm-product's quality.

Figure 11: Evolution of Ladder Length




Quality for each firm is calculated from a nested logit demand system. The length of the quality ladder is the distance between the maximum and minimum observed quality in a given year.

Figure 12: Evolution of Offshoring Activity



Narrow imports are measured as the value (in 2000 DKK) of apparel (CN headings 61 and 62) imports. Offshoring is defined as the value of imports divided by the number of employees in an apparel firm. In order to report an aggregate number we reported total imports over total employees. That is, indexing firms by f , we report the log of $\sum_f \text{Imports}_{ft} / \sum_f \text{Employees}_{ft}$.

Figure 13: Evolution of Offshoring Activity in China



Narrow imports are measured as the value (in 2000 DKK) of apparel (CN headings 61 and 62) imports. Offshoring is defined as the value of imports divided by the number of employees in an apparel firm. In order to report an aggregate number we reported total imports over total employees. That is, indexing firms by f , we report the log of $\sum_f \text{Imports}_{ft} / \sum_f \text{Employees}_{ft}$.