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

Valérie Smeets, Aarhus University[†]
Sharon Traiberman, Princeton University[‡]
Frederic Warzynski, Aarhus University[§]

September 15, 2014

Abstract

Recently a small and growing empirical literature has attempted to analyze the role that quality plays in our understanding of trade. In particular, the recent work of Khandelwal (2010) has brought the insights of structural IO models of demand to bear into trade data. Our work builds on this new structural literature; we use similar demand estimation techniques on a panel of 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 explore the implications of offshoring and import competition on the distribution of apparel quality within Denmark, and demonstrate the firm-level mechanisms that induced the observed aggregate changes. In particular, we show that the quality ladder tightens in response to trade shocks as initially low quality firms upgrade their quality relative to other firms, while firms with initially middle and high quality downgrade their output quality. An important qualification is that the origin country of input sourcing is a key determinant in both the uptake of offshoring and resultant decisions regarding quality. Finally, import competition appears to spur entry of higher quality firms and exit of lower quality producers.

*Sharon Traiberman received assistance from a Princeton IES Summer Grant, the International Economics Section at Princeton, the Economics Department at Princeton and Aarhus University. Stephen Redding was enormously generous with his time in advising this work. A special thanks to Henning Bunzel for help with the data. We are also thankful to Diane Alexander, Thomas Chaney, Jan De Loecker, Jonathan Eaton, Oleg Itskhoki, Moshe Katzwer, Amit Khandelwal, Thierry Mayer, Mark Roberts, Esteban Rossi-Hansberg, John Romalis, Jon Vogel and participants in the Princeton Student Trade Workshop, the Princeton Student IO Workshop, IIOC 2014, Paris Trade Seminar, NOITS 2014, Toulouse School of Economics and a BBVA Foundation-Ivie Workshop on Trade, Growth and Product Quality in Valencia for helpful comments and suggestions.

[†]email: vas@asb.dk
[‡]email: straiber@princeton.edu
[§]email: fwa@asb.dk

1 Introduction

Understanding product quality is instrumental to understanding the welfare gains from trade. At the aggregate level, import competition or access to new inputs can increase consumer's choice and lower price but also spur changes in the quality of goods that are offered to consumers. This paper seeks to understand how firms' output quality decisions are affected by changes in trade costs. Our research question is driven by two recent observations in the literature. First, there appears to be a great deal of heterogeneity in the quality of goods across countries within various aggregations of product definitions (Khandelwal, 2010; Hallak and Schott, 2011).[1] Second, there has been an explosion of growth in trades in intermediates, offshoring and supply chain disintegration (Yi, 2003; Feenstra, 2010). Put together, this suggests that in high-income countries, downstream producers may be sourcing from lower quality firms than they had been in the past (including intra-firm trade through FDI). This naturally leads to a question of whether firms' importing of potentially lower quality inputs affects their output quality in an appreciable way.

Some current evidence from middle income countries suggest that access to *high* quality inputs from abroad can help induce quality upgrading (Eslava, Fieler and Xu, 2013). Our paper explores the opposite direction – the sourcing of inputs from low-quality producing countries by a high quality producing country. There is ambiguity in the possible response of quality: access to cheaper, potentially homogeneous inputs and a more competitive environment may lead to upgrading; however, if inputs themselves are differentiated and trade lowers the relative cost of lower quality inputs it may induce quality downgrading.

Estimating the quality of goods presents a host of econometric problems. Product quality is an unobservable and in most datasets used by trade economists there are no observable product characteristics that might act as a proxy.[2] Moreover, there are endogeneity issues since price and quality are normally determined jointly. This has led to a literature that attempts to back out unobservable quality from information on prices and market shares, sometimes with the aid of a structural model. Following Khandelwal (2010), we employ demand models used in the IO literature (e.g. Berry, 1994) to back out quality as a residual of a regression of market shares on price.[3] We exploit a very rich dataset to construct plausibly exogenous instruments that allow us to weaken assumptions that the literature

[1]For example, Hallak and Schott find an average difference in quality between rich and poor countries of .38 log points in 2003, down from 0.67 in 1989.

[2]An exception is Crozet, Mayer and Head (2012) who use direct quality measures of Champagne wines.

[3]Throughout the paper, we will use the term quality to reflect a demand shock or perceived quality from the consumer's point of view, as in Sutton (2012)*: "it 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).*

has made in the past.[4] The structural approach along with our instrumenting strategy allows us to model quality flexibly and separates price effects that may reflect changes to the competition faced by firms and not by changes in physical quality output.

We employ a novel dataset on Danish apparel firms that contain highly disaggregated information on the import and export transactions of firms as well as information on their employees and production. With this data, we empirically document the response of quality, at the firm-product level, to changes in the opportunities for offshoring. We analyze apparel firms before and after China's entry into the WTO as well as the end of the Multi-Fibre Arrangement (MFA) – which led to the dismantling of nearly all quotas and tariffs on apparel in the EU. China's entry in the WTO in December 2001 made it a part oxf the MFA and quotas on apparel and textile imports were slowly phased out, ending completely in 2008. As Denmark is a small country, but a member of the EU and WTO, the specific changes can be viewed as an exogenous change to Danish firms' foreign competition and their offshoring opportunities. As we will document, the industry went through a major change in the aftermath of these events, yielding substantial variation in access to and use of offshoring in the time-series as well as the cross-section. In addition, lowering the MFA induced massive import competition.

Our empirical investigations both confirm and complicate previous work. At the aggregate level, we find that a large shock to trade costs was followed by a concurrent shortening of the quality ladder (i.e., the quality of goods became more similar), and a change in the distribution of quality with more weight lower on the ladder. We also see massive exit of lower end producers and entry of high end producers – suggesting that import competition may force out some low-end goods while spurring specialization in new high-end goods. At the firm level, we find that increased offshoring is associated with a decrease in the quality ladder. We also find that the negative effect of offshoring is particularly strong when sourcing to low-quality countries.

When we begin to allow for heterogeneity in firms' joint offshoring and quality decisions, the story becomes more complicated. 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. We find evidence both at the aggregate and at the firm level that quality ladders tend to tighten and the weight in the tails of the quality distributions tends to shift to the right.

We build a simple model to guide us in our empirical approach. In our framework, firms

[4] A bevy of such datasets has led to the concurrent development of such instruments for different purposes. See e.g. Hummels et al. (forthcoming) and Amiti, Itskhoki and Konings (forthcoming).

endogenously choose their sourcing strategy, their output quality and their price. Firms differ in their ability to produce high quality goods. Moreover, firms at home have an absolute advantage in producing high quality goods relative to firms based abroad. In line with our empirical results, we find that, when facing a trade shock that affects the relative cost of producing at home or abroad, quality differences between firms are declining: firms at the bottom of the quality distribution start offshoring and upgrade the quality of their products, while firms with relatively higher quality that were already involved in offshoring engage in more offshoring and downgrade their quality.

Our paper is related to several recent contributions in the literature. Both Bloom et al. (2012) and Utar (forthcoming) provide strong evidence that increased competition from China led to massive 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) show that quality differences can explain the substantial size heterogeneity observed in many industries, and also that smaller, more focused and higher quality firms were more resistant to the surge of imports from China. Closer to us, Amiti and Khandelwal (2013) extend Khandelwal's 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 about export 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 (2013) 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, none of these papers focus on how firm-level offshoring decisions in advanced economies are related to product quality. Analyzing this relationship is the main contribution of our paper.[6]

[5]See also Martin and Méjean (forthcoming) 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.

[6]In addition to this growing empirical literature there is a theoretical literature on the interaction between offshoring and productivity that suggests offshoring affects output and wages through a myriad of channels that may push in opposite directions (Grossman and Rossi-Hansberg, 2008). While this literature is on productivity and quality, it suggests that the effects of offshoring on certain firm variables may be ambiguous

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 some reduced form evidence about changes that occurred over time. 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 is concentrated predominantly in the medium to high end segment of the fashion industry. Denmark has a well established reputation in producing original design. 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. It also experienced a dramatic growth over the last decade, increasing revenue from DKK 37 billion in 2003 to DKK 56 billion in 2010.[7]

We identify our sample of firms in the apparel industry by looking at all firms that declared having produced at least one type of apparel product in the Survey of Manufacturers (see the next section for a detailed data description). Most firms are specialized in apparel, and we keep all firms with at least 90% of their sales in the apparel industry.[8] This means that most of our firms are what Gereffi (1999) refers to as "Branded Manufacturers" (OMBs) in addition to a few traditional apparel manufacturers. These producers engage in production sharing with their offshore counterparts. I.e., these firms typically engage in outward processing where raw materials or parts are purchased by the firms themselves and then exported for assembly. This type of production is distinct from that of firms dubbed "Branded Marketers." These firms focus solely on design, distribution and marketing while contracting out the entire manufacturing process. These latter firms have become increasingly common and now dominate much of the industry, especially as "fast-fashion" grows in popularity.[9]

---

and requires empirical analysis.

[7]See Ministry of Business and Growth (2013), Denmark at work - Plan for Growth in the Creative Industries and Design.

[8]The distribution of sales is bimodal with a one peak around 90% and another around 1%. For the handful of firms between 1% and 90% of sales in apparel, we spot checked them and used industry codes to identify those firms were engaged in apparel. Leather goods, shoes and bedding do not fall under the 61 and 62 headings, so our spot checks admitted those firms to the sample.

[9]Fast-fashion refers to those firms that compete on responding quickly (or themselves defining) trends. These firms are often concerned with quality as it pertains to the perception of "fashionability" and less

An apparel product is defined as any product in the 2-digit categories 61 and 62 according to the Combined Nomenclature (CN). Table 1 shows the most common products made by our sample of firms. As we can see, the most observed items in our dataset are relatively basic products, although they still incorporate a large Danish design component.

The discussion above hinted at our concept of quality. Quality in apparel is normally broken down into two components – the physical quality of the good (e.g., open-end spinning versus ringspun cotton, thread count, non-bleeding dyes) and the "fashionability" of the item.[10] Branded manufacturers exercise a great deal of control over both of these as they often work on supplying parts to assemblers. Branded marketers, while making contracts with and demands of offshored firms, have less control over the physical quality of the good beyond demands they place on manufacturing firms. Because our firms fall in the former camp, we believe they are best modeled as having a great deal of control over the output quality. What we cannot do in our analysis (and will become clear in later sections) is separate physical quality from fashionability. This can be a particularly big issue when thinking about apparel because "fashionability" can substitute for poor physical input quality. In fact, this is the strategy behind many middle-end wholesalers who engage in "fast fashion" where copying designers and putting clothing out *quickly* and *cheaply* has supplanted more traditional design. The same is true of traditional manufacturers who engage in extensive marketing campaigns. The fashionability discussion comes down to the fact that quality in apparel, unlike for some other goods, is a relative concept as much as an absolute one. I.e., perceived fashionability may depend on the menu of clothing, and changes quickly over time. The trickiness of interpreting quality will come back (with less subtlety) in the empirical section. For now, we highlight that apparel quality is a rich concept that leaves ample room for vertical differentiation – and, importantly, the *source* of quality differences is something on which we remain agnostic. This discussion is particularly relevant for countries like Denmark, which have comparative advantage in design and distribution. With this discussion of the industry in hand, we turn to a brief history of the quota system that governed apparel through the early 2000s.

Starting in the 1970s, most trade in the apparel industry was governed by a series of quotas called the Multi-Fibre Arrangement (MFA), and later the Agreement on Textiles and Clothing (ATC). The MFA was phased out in several stages beginning in 1995 and ending

with physical quality as this is often forgone in order to keep price down and allow for faster response times. For a discussion of "lean-retailing" in general see Evans and Harrigan (2005) or Cachon and Swinney (2011). The canonical example of fast-fashion firms with presence in the US are H&M and Zara.

[10]We 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 Liegey (1993).

in 2005. China entered the WTO in December 2001 and by January 2002 had dismantled many restrictions on its textile and apparel trade and caught up to the transition path of other WTO members. Thus, China's entry into the WTO provided a large, new outsourcing opportunity for Danish firms starting around 2001. While China's entry into the WTO is the largest shock to the Danish textile industry, the phase out of the MFA/ATC in general led to large changes in the industry. In this section we outline a few of the key changes that occurred over the duration of our panel – especially in regards to the changing composition of firms engaged in the apparel trade as well as the import decisions of apparel manufacturers in particular.

From 1997 to 2010, imports of apparel in Denmark grew by 26.5% in real terms. However, the value of *net* imports fell, while in weight terms they continued to rise until very recently. A large part of the drop in value comes from the crisis after 2007. However, even focusing on the years 1997-2007, the value of net imports fell .28 log points in nominal terms while the weight increased .61 log points. Put differently, while the physical number of goods entering Denmark from abroad has grown, their value at port has declined. This pattern is shown in figure 1. Part of this can be explained by the extension of the European Union – as Eastern European countries entered the EU, reporting requirements and prices changed. Moreover, there has been a general (but not extreme) decline in the price index of apparel. However, while these two facts matter, they miss the crucial compositional changes that can explain the stark patterns in imports over time. These patterns can be more fully reconciled if one decomposes imports into those by retailers and branded marketers versus traditional manufacturers and OMB firms (the latter group, recall, is our definition of a manufacturer). For pure importers, the value of net imports increased by 43.2% over the period; however, they decreased by 154.4% for manufacturers. This latter fact stems from the shrinking number of manufacturing firms. Put together, the above points demonstrate a trade environment characterized by changing cost and markup structure as both the cost of goods decline over time and the nature of imports change.

More evidence that these aggregate changes reflect changes in importer composition and import prices comes from the sourcing and exporting patterns of the firms involved. While the former underwent substantial changes, the latter changed very little. The customers of Danish apparel themselves appear to have changed little – in 1997, 85% of Danish apparel exports go to just 7 countries and these same countries constitute 75% of exports by the end of the sample. The countries themselves are all in Scandinavia and Western Europe, consistent with the idea that Denmark specializes in high quality apparel, which it exports to its rich neighbors. This pattern is similar when one breaks exports into those by pure importers and manufacturers.

While the exporting patterns change very little, there are drastic changes in the composition of imports. To discuss the changing composition of imports, it helps to discuss apparel trade at non-manufacturing and manufacturing firms separately. For manufacturing firms, we focus predominantly on imports of raw textiles and apparel (that is to say, assembled goods that the manufacturing firms process or finish) and refer to these as intermediates. These goods constitute over the years anywhere from 40-80% of all imports done by apparel firms and on average constitute 60% of imports with a downward trend.[11] As discussed above, there was a rise in the share of apparel in intermediate inputs for Danish firms.

What is more interesting and relevant for an analysis of the global sourcing chain is the rapid increase in sourcing from Asian countries and in particular China. Figure 2 documents the rise of China in Danish apparel. Figure 3 breaks this down by domestic producers and pure offshorers and importers.[12] 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. When we break things down by domestic producers and not, we see that offshoring to and importing from China climbs to 45% (and imports from China, Hong Kong and India climbs to 60%) for non-domestic firms and climbs from about 4% to 16% for domestic producers. Moreover, in the latter group, we see the collapse of work being done in Central and Eastern Europe (mostly Poland and Lithuania) and its being supplanted by trade with China and Turkey. Overall, there is robust evidence that the end of the MFA and China's entry to the WTO resulted in massive changes in the Danish apparel industry – domestic producers moved their offshoring services from countries nearby to Asia while other firms began both offshoring to and importing directly from Asia.

A few studies suggest that in addition to changes in prices and productivity, the end of quotas can induce changes in the quality of exported goods. 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 depending on the 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. Both these results suggest that the MFA shock led China to offer upgrade the quality of its cheaper, already lower-quality goods while focusing less on the high-end segment of the market. In a similar vein, Manova and Zhang (2012) find evidence that Chinese firms that expand and export "use higher quality inputs to produce

[11] This is actually an underestimate, especially later in the sample. Many of our firms have diversified away from consumer apparel into industrial apparel as well as other textile products such as bedding, shoes, bags, etc and also into leather goods. We don't include these in our definition of intermediates.

[12] We do not make much of this distinction between these two subsets of firms, but our sample of non-domestic producers includes those Danish firms that do design at home but offshore production (pure offshorers) as well as those firms that are engaged mostly in retail, whole sale and distribution (traditional importers.)

higher quality goods." Finally, Khandelwal et al. (2013) find that the end of the quota system led to large productivity gains for Chinese firms as more efficient firms entered the industry. The image of China that emerges in this discussion fits in with a general trend of Asian countries that, in the past, have expanded their production activities. Gereffi (1999) details a brief history of the apparel industry in East Asian countries over time and details a general trend of starting out by focusing on low quality, mass production and improving industrial practices over time.

3 Data

We employ several datasets provided by Statistics Denmark that paint a comprehensive picture of the apparel industry in Denmark. The key datasets are the universe of customs transactions (UHDI) as well as production data on all apparel manufacturers that employ at least 10 individuals *or* meet a revenue threshold (VARS). For each firm, we observe all of their product lines at the Combined Nomenclature (CN) 8-digit level.[13] For each product line, we observe the product's revenue value rounded to thousands of DKK and the number of units sold. This allows us to construct unit value which we use as a proxy for price. It is well known that unit values can be a noisy measure of price – even at a highly disaggregated data the product definitions can mask heterogeneity that moves unit prices even when no price movement has occurred. Moreover, in our data set firms are prone to recording errors that are easily spotted. This cleaning is standard practice in this kind of dataset (e.g., Khandelwal (2010)) and is important to identification as the presence of severe outliers is magnified in our fixed effects framework with large clusters. We clean our data in the following way:

1. Removing the top and bottom 5% in prices. This helps remove outliers that most likely represent recording errors or unit-measure errors (e.g, unit values in the pennies or in the hundreds of thousands). At the upper end it helps remove products that are simply not well approximated by our market modeling (e.g., 6000 USD fur coats).

2. Removing those product lines with less than 45,000 DKK (roughly 7500 USD) deflated to the year 2000 price level. This helps avoid rounding errors – because revenue is rounded to the thousands while units are recorded exactly, low revenue firms may end up with the same level of sales reported in our data set but radically different levels of quantity sold.

[13] The Combined Nomenclature is the system for recording trade data used by the EU. The first 6 digits are the same as HS10 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.

3. Removing product-years where the price differs by the median price by more than a factor of 1. This removes only a few observations that would not have been otherwise removed, but in our dataset we find that some product prices will spike in a single year by an order of magnitude from the norm. We assume these are recording errors hence their removal.

4. Removing products lines that exit and reappear for more than 3 years. In particular, we only remove the products *after* reappearance. The reason for this is that if products enter and exit it is difficult to estimate their mean qualities as the product, at a more narrow level, may have changed. For example, we often find persistent large spikes in prices after this kind of entry and exit.

From a third database (FIRE) we observe employment, intermediates use and capital at the firm level. As is usual in multi-product firm datasets, there is no mapping from firm-level inputs to product-level outputs, so this level of disaggregation remains unobserved.

In addition to data on sales and manufacturing inputs, we observe the universe of firms' trade transactions in Denmark by product and destination or origin. The firms in both datasets can be linked together. This allows us to observe the import and export transactions of our apparel firms as well as other firms involved in the apparel industry. The import and export data includes values (without rounding) and quantities so we can construct unit values (a phrase we use interchangeably with price) for these goods as well. Combining these datasets allows us to construct our instruments, as discussed in the next section.

We also merge onto our apparel data, data from the Danish linked employer-employee dataset (LEED). This data covers the universe of worker-firm pairs and contains information on workers' skill (highest educational attainment), age, market experience, occupational tenure, occupation (at the ISCO 3 digit level) and wages. We mostly use it to construct a firm-level measure of average wage and skill intensity.

Our production data panel runs from 1997 to 2010, thus covering China's entry into the WTO, the beginning of the dismantling of the MFA, and the conclusion of this operation in 2005. Our data on trade begins in 1993 as does our data on employment and other firm side variables. Some of our aggregate statistics on trade exploit the full length of the panel, but mostly we focus on the time frame of 1997-2010 so that we can focus on those firms that we know are producing domestically and nothing else.

Finally, we bring in several outside data sources. Data on quotas comes from the EC'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 Conceptual Framework

Before turning to the empirical analysis, we present an illustrative model of a firm that has three key decisions: an output price, an output quality and a sourcing strategy. Our model determines how the firms' "capability" (which we will formalize below) as well as its more standard "productivity" determine these decisions jointly. More importantly, the model yields some predictions of how the distribution of firms' quality (what we will call the quality ladder) changes in response to new offshoring opportunities. To keep the model as parsimonious as possible, we abstract from any exporting behavior of the firm and assume there is a single composite input. Moreover, all propositions regarding falling trade costs are done in a partial equilibrium manner – i.e., we do not assume that domestic prices of inputs themselves change in response to changes in the prices of foreign inputs. Because we only model a single input, this is an extreme model of offshoring where all sourcing is either done domestically or abroad – traditionally one thinks of offshoring as substituting only pieces of the supply chain. We deal with this distinction more carefully in the empirical section, but for the time being we think of the firm's "production" chain as only two separable pieces - design (simply a choice) and production.

4.1 Consumer preferences

Here and in the appendix we provide a conceptual framework robust to many assumptions on utility. In order to match directly to our empirical specifications, we would ideally model a multi-product firm facing a demand system derived from aggregating a nested logit RUM.[14] Unfortunately, the full-blown nested logit model is intractable for two reasons: (1) the firm internalizes its impacts on aggregate shifters because of the finite players assumption; (2) multi-product firms must work with the complete substitution matrix for goods because they internalize a very complex set of cannibalization effects.[15] However, we can overcome these complications if we ignore multi-product firms (or cannibalization effects so that product lines can be separated), and we consider a "limiting" case where firms ignore their own impact on aggregate shifters (so we assume that shares are small enough that the firm ignores the

[14]See Train (1999) or Anderson et al. (1992) for a discussion of RUM models, their estimation and their connection to other utility functions.

[15]This lack of tractability is not surprising given the general problems of solving the firms' problem with logit demand systems (Caplin and Nalebuff, 1991), so that an interior equilibrium is only guaranteed under certain assumptions.

denominator in the nested logit market share formulas).

These assumptions imply a firm that treats demand as if it were generated by the following function[16]:

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

where A is an aggregate demand shifter, v is quality, p is the price and σ modulates consumer sensitivity to price (recall that in logit demand systems elasticities depend on shares in addition to parameters).

4.2 Cost function

Consider a firm j that faces a constant unit cost. Because we have assumed single-product firms, we let j also index goods. Let the set of firms/goods be denoted by $\mathcal{J}$. The quality production function is

$$v_j = \psi^{\alpha}$$

where ψ is the input quality from a single composite input and $\alpha < 1$ implies diminishing marginal returns to input quality.

Input costs are linear in quality[17] and given by:

$$c(\psi) = a_s + \left(b_s/\omega_j^{z_s}\right)\psi$$

where s represents the sourcing strategy of the firm. The firm can either decide to produce at home (H) or abroad (F). Depending on their choice, they face a fixed per-unit cost a_s that the firm pays regardless of quality, and a variable term $b_s/\omega_j^{z_s}$ that the firm pays for higher quality inputs. While b_s is constant for all firms that choose a similar sourcing strategy, the variable cost per unit of input quality diminishes with the firm's ability to produce quality ω_j. This ω_j term is what we refer to as the firm's capability as in Sutton (2012). This parameter is at the heart of our story because, unlike Hicks-Neutral productivity, it does not shift the marginal cost. Instead, it tilts the cost as a function of quality. Moreover, the z_s term sets up the crucial trade-off that firms face in regards to offshoring: whenever $z_s < 1$, firms lose some of their capability to produce high quality efficiently. In the extreme case

[16] Notice that if instead of p we used $\log(p)$ this would become the CES demand system. The use of $\log v$ is a transformation we employ to keep the quality production function consistent across this section and the appendix, but is irrelevant as far as the mapping to estimation is concerned. This connection is explored further by Anderson, de Palma, and Thisse (1987).

[17] It is not difficult to make the costs convex and of the form $a_s + \left(b_s/\omega_j^{z_s}\right)\psi^{\gamma_s}$ with $\gamma_s > 1\ \forall s$. However, no results will depend on this and it adds unnecessary flexibility to the cost function. The additive structure can be rationalized if production is Leontief in two components – a homogeneous piece and a heterogeneous one.

that $z_H = 1$ and $z_F = 0$, the firm is able to utilize it's own productive capacity if it produces in-house but has no ability to curb quality cost if it offshores.[18] The ω_j term is drawn from a distribution Ω and varies across firms. In addition to unit costs, firms pay a fixed cost of operation, f. Finally, we assume that firms have a Hicks-Neutral productivity shifter, λ_j, inherited from a distribution Λ. Thus, the actual cost paid by the firm is $c(\psi)/\lambda_j$.

Putting this together, the firm's profit function is given by,

$$\pi(\psi, p) = A \exp\left\{\alpha \log \psi - \sigma p\right\}(p - c(\psi)/\lambda) - f$$

From the first order conditions one can derive the following equation for quality (independent of the functional form on cost):[19]

$$\frac{\alpha}{\sigma\psi} = c'(\psi)$$

Using our specification of costs, this implies a firm with capability ω sourcing from s chooses quality such that:

$$\begin{aligned}
\psi^* &= \left(\omega^{z_s}\lambda\right)\frac{\alpha}{\sigma b_s} \\
p^* &= \frac{1}{\sigma} + c(\psi^*)/\lambda \\
&= \frac{1}{\sigma/(1+\alpha)} + \frac{a}{\lambda}
\end{aligned}$$

Optimal quality behaves as expected – it is increasing in both capability and productivity and decreasing in the per-unit marginal cost of quality. Moreover, when customers are more price sensitive, quality decreases. This contrasts with the CES case (in which the price parameter does not matter). It arises from the fact that the price elasticity is non-constant – its magnitude is increasing in price. Since a higher quality necessitates a higher cost and thus higher price, it would exacerbate the price sensitivity of consumers. Thus, firms with

[18]To motivate this a bit more, one can imagine that the composite input comes from two inputs – in-house labor and materials. A firm has access to labor-augmenting technology. Thus, if it uses a production plan that is more in-house labor intensive it can exploit said labor-augmenting technology; but if it engages in a more materials-intensive production plan and uses less labor the total productivity effect of the labor-augmenting technology is diminished. Alternatively, this can represent "communication costs" in the vein of Grossman and Rossi-Hansberg (2008).

[19]See Feenstra and Romalis (forthcoming) for a discussion of this condition in the case of a demand system generated by a representative consumer with preferences of the form $U(x_1 v_1^{\alpha_1}, x_2 v_2^{\alpha_2}, ..., x_J v_J^{\alpha_J})$. It can be demonstrated that for a very broad class of utility functions (including CES), one arrives at a particular independent of the utility function. This comes from the fact that when quality and quantity substitute in a multiplicative way the entire problem can be restated in terms of picking quality adjusted quantities ($\tilde{x}_i = x_i v_i^{\alpha_i}$) and prices ($\tilde{p}_i = p_i / v_i^{\alpha_i}$). One reason that the logit model is difficult to work with is precisely because it breaks this multiplicative property, yielding a more complicated set of first order conditions.

already sensitive customers (high σ) want to lower their quality.

Putting this together, the firm's profits are given by,

$$\pi = \kappa_1 \exp\left\{-\alpha \log\left(b_s/(\lambda \omega^{z_s})\right) - \sigma a_s/\lambda\right\} - f$$

where κ_1 is a source and productivity independent constant. The profit function has some aesthetic appeal essentially saying that profits are determined by a weighted sum of the two components of costs – the per-unit cost (a_s) moderated by price sensitivity and the quality-varying cost b_s/ω^{z_s} moderated by tastes for quality.

We will assume that $a_H > a_F$ but $b_F > b_H$. These two restrictions say that the home country has an absolute advantage for high quality input production, but the foreign country has an absolute advantage for low quality input production. To decide whether a firm offshores or produces at home, we only need to compare the exponential term under both options. In particular one offshores if

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

Taking logarithms this reduces to,

$$\alpha\left[(z-1)\log\omega - \log b_F + \log b_H\right] - \frac{\sigma}{\lambda}\left[a_F - a_H\right] > 0$$

Thus one offshores *iff*,

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

Notice that this cutoff is a function of productivity – so that sourcing decisions depend on both the firm's capability and productivity. 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 point made by Holmes and Stevens (2014) that large firms may not be the highest quality firms and vice versa. In our case, conditional on sourcing, larger, more productive firms will have higher quality (assuming $Cov(\omega, \lambda) \geq 0$), but it need not be the case unconditionally. In addition to its dependence on λ, observe that the cutoff is increasing in $a_H - a_F$. Intuitively, if it is more expensive to produce at home than abroad, more firms offshore. Similarly, the cutoff is decreasing in the difference between b_H and b_F.

Turning to a discussion of offshoring and quality, notice that quality is monotonically decreasing in b_s regardless of a_s. On the other hand, any reduction of trade costs induces more offshoring. Thus, a reduction in b_F that nevertheless maintains home's comparative advantage, induces offshorers to upgrade their quality but induces switchers to downgrade. In the case of a reduction in per-unit trade costs, a_F, only switchers show a quality response whatsoever. In time series data we will see firms enter and exit. These patterns of entry and exit can help us understand offshoring and quality, but requires more thinking through the model. Notice that entry and exit decisions depend on aggregate shiters and fixed costs. If trade shocks also impact aggregate shifters (either by creating import competitoin or through creating new export markets) then the overall pattern of entry and exit can be indeterminate. We summarize the above discussion in the following concrete set of propositions that guide our empirical analysis:

Proposition 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 offshoring will not respond.
2. Some firms will begin to offshore and their quality will decrease (relatively more than in the other scenario).
3. Firms of sufficiently high quality will not respond.

Proposition 2 (Entry/Exit). *If A stays constant and a_s and/or b_s decrease then low quality firms enter. However, if A moves then the joint distribution of λ and ω determines entry and exit patterns.*

When we turn to the data, it will turn out that nearly all firms engage in at least some offshoring activity. Hence one should think of this model as largely descriptive. To operationalize this proposition in the data we will focus on looking at how the shape of the quality ladder changes and how firms at different positions along the quality ladder respond to new offshoring opportunities. First, quality is a relative statement, so looking over time will require us to see how firms compare to the mean of the distribution. In particular, we modify our predictions about heterogeneity to suggest that lower quality firms ought to increase their offshoring and their quality relative to other firms, while middle and high quality firms ought to offshore even more while moving closer to the mean relative to other firms (downgrading) or having a more muted response. Bringing entry and exit into the situation completes the picture and allows us to discuss higher order moments of the quality distribution. In particular, we also explore the skewness of the distribution.

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. If this ingredient were not there, then the results would be radically different. In our logit set up, the rich interaction between quality and productivity would give way to an empirically false monotonic relationship between productivity, quality and size. In the CES case, things are more damaging – we show in the appendix that in a very general model neutral shifters of revenue do not change firms' optimal choice of quality. On the other hand, some kind of complementarity between sourcing strategy and capability (as in this model) or between capability and design (as in the fixed costs case) can generate a heterogeneous response of firms to new offshoring opportunities. The idea of non-neutral differences in firms' ability to change the quality of a particular variety is what sets vertically differentiated markets apart from those that are horizontally differentiated.

5 Econometric Model

Before moving to the analysis of imported inputs and offshoring we first need to define more clearly our structural estimation procedure for extracting quality from price and sales data. This section outlines our econometric model, including consumer demands, timing assumptions and decision making by firms, as well as the details of mapping our model to data and instrumenting strategy.

5.1 Consumer Demand

We follow the recent work of Khandelwal (2010) and Amiti and Khandelwal (2013) in using the discrete choice framework common in IO and labor to model consumer demand. In particular, assume that consumer i has indirect utility for good (j, t) given by,

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

where δ_{jt} is a common taste for product jt, 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). The GEV distribution allows for more complicated substitution patterns than the extreme value distribution. In particular, it allows for goods to be grouped into non-overlapping "nests." This allows one to model the agent as first picking a nest, then - conditional on their nest - picking a good. Formally, consumer i picks good jt iff

$$V_{ijt} \geq V_{ikt} \forall (kt)$$

Berry (1994) shows that in the limit of a continuum of consumers, the market share for product jt 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 (or groups) and $\mathcal{J}_g$ is the set of products in nest g. In the same paper, he 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. There are $J \times T$ observations here but a total of $(J + 1) \times T + 2$ parameters. Since we can only truly estimate $(\delta_{jt} - \delta_{0t})$, we are free to make one normalization and so we set $\delta_{0t} = 0$. This still leaves the problem unidentified, and so we adopt the practice of splitting the quality parameter into fixed effects and an error term. In particular, we set $\delta_{jt} = \delta^1_j + \delta^2_t + \delta^3_{jt}$ where the first term represents the average quality of good j, the second term represents a secular trend in quality growth and the third term is a product-time deviation. We will treat the last term as a regression error and so we have the new estimating equation,

$$\log s_{jt} - \log s_{0t} = \delta^1_j + \lambda_t - \alpha p_{jt} + \sigma \log s_{j/g} + \delta^3_{jt}$$

where $\lambda_t = \delta^2_t - \delta_{t0}$ is secular growth in quality *relative* to outside good growth – this subtlety will be very important later on. In general, δ^3_{jt} is correlated both with price and the nest share. We will discuss our instrumenting strategy in detail in subsection 5.4. This strategy will depend on our model of firm production, so we turn to that now.

5.2 Firms' Decisions

Since all of our estimation is done off of the demand system, we do not need to make strong assumptions about what happens with production. However, in order to explain our instrumenting strategy, we present a basic outline of the timing of production decisions. We assume that the firm goes through three stages in each period – and only makes decisions in two of them. 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 (1998, 2012). The timing of the shocks is similar to that employed by Ackerberg, Caves and Fraser (2006) – decisions are made after the realization of an initial TFP shock and then there is a second *ex-post* productivity shock.[20] Notice that, as one would expect, price is mechanically correlated with quality through both its impact on the markup and the marginal cost. As usual, we solve this problem using backward induction.

In the final stage of the period (after cost uncertainty has been revealed), firms set their prices and compete. Suppose as in Berry, Levinsohn and Pakes (1995) 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[21], M is the size of the market and s is market share, that depends on

[20]This timing assumption is admittedly not without loss of generality. We believe it fits well with the characteristics of the industry that we examine. In apparel, the design and planning process happens, by definition, before production takes place, while marketing and selling logistics occur after (Frederick and Staritz, 2012). Thus, we think it reasonable to assume that "physical quality" – design and input sourcing – are determined before production cost shocks that price may respond to are realized. We come back to this when we discuss our IV strategy.

[21]BLP models this as a function of the observed characteristics of each specific product w_j and an unobserved component ϖ_j. In the final stage, it is considered as given to the firm. Our assumption is that marginal cost is not known until quality decisions (now endogenous) have been made.

the price vector, as well as the unobserved quality of the good, δ_j.[22] Our first major timing assumption is that firms optimize statically – i.e., they can costlessly move prices and quality each period.

Maximizing profits with respect to price, we get the following FOC:[23]

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

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

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

where we have momentarily used the vector notation to make explicit that the firms' decisions depend on the whole vector of choices. We assume that $mc(\delta_j, \varepsilon)$ depends on quality and some vector of possible cost shifters that the firm has not yet learned. The expectation is over these shifters and of other firms' shifters (since they all jointly determine relative market shares). Firms choose quality to maximize profit, expecting a cost shock in the second stage and price competition in the final stage,

$$E\left(\sum_{j \in \Gamma_f} \left(\frac{\partial p_j}{\partial \delta_k} - \frac{\partial mc(\delta_j, \varepsilon)}{\partial \delta_k}\right) s_j(p, \delta; \vartheta) + \sum_{j \in \Gamma_f} \left(p_j\left(\vec{\delta}, \epsilon\right) - mc(\delta_j, \varepsilon)\right) \frac{\partial s_j}{\partial \delta_j}\right) = 0$$

After this, the shock ε is realized and firms produce. The crucial difference between this model and a model where quality and price are determined simultaneously is the presence of the expectation operator in deciding on quality. Thus, while δ_j will depend on expectations of cost shocks, it will be uncorrelated with particular realizations. This assumption will allow us to exploit the orthogonality between certain 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.

For the eventual estimation of quality growth, we need no further assumptions on produc-

[22]Since this model is for explanatory reasons rather than analytic results, we have allowed δ_j to indicate quality for both the firm and consumer; in general, all we would need is for there to exist a monotonic function $g(\xi_j) = \delta_j$ that maps from the firms' "physical quality" to the consumers' "tastes quality."

[23]While we allow for some general form of competition, as is standard in this literature we assume that the equilibrium is at the point where firms' solve their first order conditions (Caplin and Nalebuff, 1991).

tion. Our assumptions on demand structure are somewhat stronger, but follow the standard in industrial organization. We now turn to a discussion of the data as well as a more detailed look at the precise set of estimating equations and instrumenting strategy that we employ.

5.3 Nest Structure, Trade and Market Size

In the apparel industry, goods are split into knitted and crocheted wear and also woven wear. Our nests ignore this distinction and are based on combining 4 digit Combined Nomenclature codes which are the same as 4 digit HSIC codes. Thus, the nesting structure is based on the *type* of apparel product and ignores construction-method, fabric and weight (when available). The nesting structure respects gender whenever possible. In total there are 16 nests listed in Table 2. In our estimation, we remove the accessories category. This is a matter of over-aggregation *within* an 8 digit code – year to year price and quantity data is very erratic for such a broad category at the firm level. Within each nest, we observe products at the 8 digit level. These are highly disaggregated and normally include the particular type of garment, the material and sometimes 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." 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. As discussed in the section on consumer demand, we break up each variety's quality into a fixed component, an economy-wide time varying component, and a product-time deviation.[24]

Before discussing our instrumenting strategy, some discussion of the outside good is in order. As is typical in the demand literature, the outside good can often be very important for estimation. In our setting, because we use time-fixed effects, the choice of outside good will not matter for 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. This is obviously of great importance to our estimates. For

[24]At this time, our estimation ignores the distinction between domestic sales and foreign sales by firms. If exports and domestic sales are highly correlated, since we use market *share* measures instead of levels-measures, this problem is abated. In an extreme situation, if domestic sales (in quantities) were a constant fraction of total sales (i.e., $q_{ft}^{dom} = q_{ft}\theta_t$) for all firms, then there would be no problem with our estimates. If there is no systematic relationship between share of exports in output and our instruments, then our aggregation of exports and domestic quantities would lead to higher variance of our estimates but no bias. If the share of exports is systematically correlated with our instruments, then there is bias. It is difficult to sign this bias given our instrumenting strategy, however in section 6.1 we assess the validity of our parameters and find that our parameter estimates fall in line with the literature while our quality measures correlate with important firm level variables.

the outside good, we use the total quantity of imports into Denmark. 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. In effect, this fact leads us to focus on looking at how firms respond within time periods and over long differences instead of focusing on year-to-year differences. However, we also discuss several strategies based on movement along a unit-free quality ladder than helps to distinguish relative quality growth of firms.

5.4 Instrumenting Strategy

The standard endogeneity issue in demand estimation is that price will be correlated with the unobservable demand shock. This is also true of the nest share – and in fact, the unobservable is theoretically a direct input into a product's within nest share. Hence, estimating this model relies on locating suitable instruments. The problem of finding plausibly exogenous instruments in the structural framework is that quality and price are chosen concurrently. In fact, many "cost shifters" that an econometrician might identify – e.g., wages – are almost certainly a reflection, at least partially, of the quality of an input. Given the discussion in section 2, we believe that *unanticipated* shocks to costs may actually be plausibly correlated with price, but not with quality. Based on exisiting descriptions of the industry (e.g. Gereffi, 1999; Frederick and Staritz, 2012), we assume that quality decisions are taken when cost shocks affect the firm and that there is a certain delay (frictions) for the firm to adapt quality decisions to these shocks. A similar strategy was followed by Foster, Haltiwanger and Syverson (2008) who used structural estimates of innovations to firm's productivity as instruments. This particular strategy relies on the idea that output and input are homogeneous, and so any differences in productivity truly reflect supply-side shocks. Our environment is one of vertically differentiated goods, and so we attempt to construct cost shocks directly.

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, via exchange rate risk, leads to unanticipated cost shocks to the firm. In particular, we will use forecast errors on exchanges as instruments. Implicitly, we are assuming that a firm's quality is fixed conditional on the choice of a sourcing strategy and that at least some exchange rate risk is passed through in price. The source of variation arises from cross-sectional heterogeneity in import mixes across firms.

To make things more explicit, we model exchange rates as a simple exponential AR(1) process:

$$e_{ct} = e_{ct-1}^{\rho_c} \exp(\mu_c + \sigma_c z_{ct-1})$$

Taking logs this can be expressed as an AR(1):

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

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. The AR(1) was chosen because of the forecasting powers of simple random walks.

After estimation one can construct forecast errors as:

$$\hat{\eta}_{ct} = \epsilon_{ct} - \widehat{E(\epsilon_{ct})}$$

The instrument is given by:

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

where f indexes firms and $s^{imps}_{ft,c}$ is the share of firm f's imports that are from country c.[25] 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.[26] One can construct a similar measure of forecast errors for exports. We construct this measure and include it as a control in our estimation. The reason being that exchange rate shocks to imports act as cost shocks but so do shocks to exports as they act as unanticipated shocks to profitability.[27]

To instrument for the nest share parameters, we use sales weighted averages[28] of the cost shocks across a firm's *competitors* within a nest. This is similar in spirit to the approach used by Berry, Levinsohn and Pakes (1995), who use own product characteristics as instruments for price and average characteristics of firms' competing products as instruments for their

[25] We ignore those countries that are part of the European Exchange Rate Mechanism (ERM 2) as the Danish Kroner is pegged to the Euro (and varies less than 1% around the peg).

[26] We experimented with robust standard errors (as calculated by the user-written Stata command *xtivreg2*) and found that clustering produced larger standard errors; thus we opted for the more stringent strategy.

[27] A secondary reason for including exchange rate shocks to exports is that the effect of a shock to country j's exchange rate on the import side is exactly counteracted by exports (assuming the timing is correct) – thus this acts as an important control. Regardless, we find that coefficient values are not sensitive to inclusion of this variable.

[28] We explored quantity weighted and simple averages of the shocks as well. The coefficients respond some but not in a statistically or qualitatively significant way.

nest share. The instrument[29] is constructed as follows:

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

Aside from the issue discussed above, there are several possible threats to internal validity, and we attempt to address them now. First of all, since sourcing strategies are endogenously determined alongside quality, our first instrument may be invalid. However, since all firms engage in some trade, this problem only occurs if there is a systematic relationship between quality and the exchange rate risk posed by different countries. For example, if low-quality input countries also have higher exchange rate risk than high-quality input countries, the $E(\delta_{it}|\zeta_{ft}^1) \neq 0$. However, even if this were true, this does not mean that exchange rate errors and unobservable quality are not uncorrelated. I.e., $E(\delta_{it}\zeta_{ft}^1) = 0$. This will still be true by our timing assumption and given that forecast errors are mean 0.

The fact that we use *unanticipated* shocks should strengthen the validity of these instruments. Exchange rates can be correlated with quality if exchange rates in levels are informative about sourcing patterns. The use of shocks to other firms as a natural extension of the BLP procedure is also a useful new instrument. Past papers have relied on the use of the number of competitors to identify nest parameters – however, this relies on making assumptions about the response of quality to firm entry and exit. While this assumption is perfectly valid at aggregate levels – such as countries over long time spans – we are focusing on a substantially more granular market. And, as we are interested in examining precisely entrants and exiters in our data, we want to rely on instruments that do not make assumptions about these groups and the markets.[30]

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 – in fact, we assume as much as we are studying quality upgrading. To address both of these concerns, we employ a two-way clustering strategy. Thus, we allow for arbitrary correlation of demand across products within a firm each period, and across time

[29] In the estimation, we interact the nest instrument with a dummy for each nest. I.e., we estimate the responsiveness of the nest instrument on group market share separately for each nest. While the first stage results suggest a strong instrument even without the interaction, the impact of shocks to competitors varies greatly across nests. This is largely due to stark differences in nest size and attempting to fit these effects uniformly causes weak identification in the second stage.

[30] Nevertheless, as robustness exercises we have included the log number of competing products in the regression. It does not qualitatively change results but improves the precision of the estimates greatly. This fits with a model in which any short term quality response to entry is either muted or done with sufficient lag for it not to bias the estimates. I.e., concurrent shocks to quality are orthogonal to entry and exit because of timing.

for each period.

6 Results

6.1 Parameter and Quality Estimates Overview

For the sake of comparison, we run an OLS estimate, simple logit model and the full nested logit. The results of the estimation are shown in table 3. First, notice the expected biases in the OLS estimate: both price and nest share are positively correlated with unobserved quality, which drives both coefficients up. In the simple logit model, the estimate of price is pushed up a great deal – this stems from omitting the nest share and imposing overly restrictive substitution patterns. The final nested logit model successfully removes the upward bias and all coefficients are significant – this is even with our fairly conservative clustering strategy which allows for arbitrary auto-correlation in the error term within products over time and also across a firm's products in a given year.[31] Our estimated price coefficient of $-.0077$ falls comfortably in the range of parameters estimated by Khandelwal (2011) for all industries, where the median estimated coefficient was $-.001$ and the IQR for all coefficients across industries was .070.[32]

Before turning to questions about offshoring, we explore the plausibility of our results by looking at several statistics implied by our structural estimates. First, we can back out implied price elasticities that are derived from the nested logit model as follows:

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

where α is the price coefficient and σ the substitution parameter. In the event that $\sigma = 0$, this collapses to the familiar formula for logit demand. Figure 4 contains the density of elasticities implied by our estimates. They are fairly reasonable – the mean elasticity is 1.90 and the median is 1.66. 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 – with women's coats and men's coats (not pictured) containing many of the outliers. This might reflect either model rigidity –

[31] In fact, the most conservative possible clustering strategy would be by firm and allow for arbitrary cross-product-time correlations. We found that in this case our results are *more* precisely estimated. We have chosen to report the results that work the most against us since we still find them plausible and significant.

[32] Khandelwal ran his regression using 2 digit SITC industries to define a goods-market, defined 6-digit products as goods and defined country-product pairs as varieties. Thus, he is estimating a more aggregated system than we are but we take his estimates as a useful benchmark for comparison. See Roberts et al. (2012) for a similar study of the Chinese footwear industry using firm-level export data at the product level.

the same substitution parameter may not be right for all nests. However the assumption keeps things simple without an enormous cost to plausibility. It may also just reflect the idea that men's coats and women's coats have a disproportionately large number of high quality, highly inelastically demanded goods. The range of elasticities is a bit larger in absolute magnitude than those found by Khandelwal, but we believe that makes sense here as our more disaggregated goods might be more substitutable. Importantly, the magnitude of elasticities is highly correlated with quality, which again suggests that the estimated parameters implied by the model display internally valid properties. It is important to note that nothing in our estimation forces these patterns to hold.

In addition to looking at the implied elasticities, we can see how our estimates of quality correlate with unit values, adjusting for various confounders – a common approach in the literature. Table 5 below summarizes the correlation between price, quality, elasticity and size. As expected, quality and price are highly correlated – but imperfectly so. Figure 5 plots this relationship with nest and year means removed. The red line is a lowess fit that is nearly linear and clearly upward sloping. Price and size, measured by employment, are more correlated than quality and size – but all signs are positive. We cannot necessarily establish causality but it speaks to the idea that larger firms can exploit market power in addition to physical quality in order to raise prices. The work of Kugler and Verhoogen (2012) suggests that quality explains the correlation between size and price. One way to see if our results are consistent with this hypothesis is to run their reduced form regression of employment on price controlling for quality. To that end, we run the following regression:

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

where δ_{jft} is our estimate of quality. Here, we purge the regression of product (at the Combined Nomenclature 8 level) level and time fixed effects. For the sake of comparison, we also run a regression with firm-CN8 pair fixed effects – i.e., we also look at the coefficient employing only *within* firm-product variation. The results of this set of regressions are in table 6. We can see here that, when we look at the price-size correlation controlling for observed quality, the coefficient decreases. If we look only at within product-firm variation, we find that the employment effect becomes insignificant while the importance of quality goes up considerably. In either specification, we find quality to be an important piece of the size-price correlation. Our quality estimates are positively correlated with price and size in statistically significant ways and help explain away part of a phenomena that they could not do if they were just noise. We take the collective results above as important proof that the quality estimates derived by the model do capture something non-trivial about firm's products.

It is important to remember that what we measure is not necessarily quality in the sense typically understood. As discussed in section 2, we do believe that firms exercise a great deal of control over consumer tastes for their goods. However, the way this is done is multifaceted – some firms may actually change the physical quality of their good in the sense of improving the material (e.g. thread count, spinning method) or they may alter their distribution to offer a wider variety of goods that change over short intervals. They may also engage in increased marketing or in improving the aesthetic content of their output. We remain agnostic on the source of quality differentiation and, indeed, we believe that most likely firms of a different initial quality are changing their quality by exercising different options available to them. Importantly, for our analysis, this agnosticism does not change the content of our observations. As long as quality reflects real resources of some kind (distribution networks, designers, inputs, etc.), then analyzing how offshoring correlates to our estimated quality measure will reflect important information about firms and their actions.

6.2 Real Correlates of Quality

In this section, we explore how our estimates of quality relate to the firm's observables. In particular, we demonstrate here that higher quality firms tend to pay their workers more, are larger and more skilled. Our purpose here is not to discuss a *theory* of how firms' create quality or *why* the correlations in the data exist. Rather, we wish to demonstrate that our estimates of quality are *correlated* with firm-level observable variables (such as firm size or skill intensity) and do not merely reflect taste shocks.

Since we cannot observe product level use inputs, we aggregate quality at the firm level.[33] Denote by $\tilde{\delta}_{fjt}$ 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,

[33] Aggregation over products presents a problem since many products' quality reflects other factors (which is why we include product fixed effects in previous regressions). To that end, we first remove product-year fixed effects at the 5 digit level. The reason for this level of aggregation is that it is the second most disaggregated level of product codes that is consistent over time. 5-digit product codes represent a good exclusive of size and material. E.g., an 8 digit code will be "Wool men's coat over 1 kg" while the 5 digit good will be men's coats. Unfortunately if *all* of one good are of a higher quality than *all* of another good (say, because of comparative advantage in some good), then this approach will absorb those differences. As an extreme example, if there is only one firm producing a good in a given year, it will automatically have a quality of 0. Nevertheless, since the most popular items are produced by many outifts and aggregating without product fixed effects risks worse aggregation issues, we feel this is the best compromise.

and S_{ft} is the total sales of firm f at time t (i.e. the *sales* weighted mean quality across goods at firm f).[34] The evolution of this firm-level measure is similar to that of the product-level evolution, as discussed below, and a simple regression of product level quality on firm level quality reveals a strong correlation – suggesting this is a good measure of a firm's overall quality, relative to other firms.

With this in mind, we begin our analysis of workers and firms. We begin by running simple regressions to determine the relationship between the log of firm's average wage ($log\overline{w_{ft}}$) and average quality:

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

where β_1 is our object of interest and $\mathbf{X}$ collects control variables. The results of these regressions are in Table 7. As can be seen from the first two regressions, firm quality and average wages are strongly positively correlated – even when we control for productivity using size as a proxy. 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 "quality premium" may be a composition effect. This is, of course, to be expected – higher quality firms seem to be paying more because they are employing better workers. In appendix B we explore the relationship between firm composition, quality and wages in more detail. We demonstrate strong correlations between our measure of quality and wages within occupations and also the composition of occupations; this allows us to conclude that our measure of quality, no matter its imperfections, picks up physical, controllable elements of the firm's production process.

6.3 Quality Evolution and Quality Ladders

To think about how offshoring and the entry of China impacted product quality over time, we need to be able to measure the secular growth in quality. It turns out that, in this methodology, this is an impossible task without making unreasonable assumptions. The details of this will be made clear below. However, as discussed in Section 4, we have predictions about what should happen to the *distribution* of quality over time. By looking at centralized moments of the quality ladder over time, we can still say a great deal about how the growth

[34] We experimented with quantity weighting and achieved similar results. The argument for sales weighting is that quantities may have a product fixed effect wrapped in them (i.e., firms typically sell more under garments than coats) thus overweighting certain products arbitrarily. On the other hand, sales weighting (because of prices) may put arbitrarily high weight on higher quality goods. Neither measure is perfect, but results are robust to each.

of China in the time series and the time series of aggregate quality comove.

To formally explore our predictions from section 4, we first define a good's position on the quality ladder at time t as

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

That is to say, it is the good's quality purged of the time fixed effect and demeaned. Our quality ladder is unitless but cardinal: the magnitude in difference between positions is a measure of the quality difference between products. The changing shape of the quality ladder gives insight into aggregate changes. In order to get a sense of how the distributions of quality change, figure 6 plots the density of the ladder measure at the beginning and the end of our sample. Immediately a few patterns emerge: the right end point shifts in, and there is more mass in the center of the distribution. Harder to see, but a strong fact that we demonstrate below is that the distribution also becomes increasingly left skewed. The issues with interpreting these facts is universal in demand estimation over time and we will discuss this further below. In spite of these facts, we argue that focusing on higher moments of the quality ladder gives us insight into the effects of import competition and offshoring – and which effect dominates in the aggregate. This is drawn into sharper focus when paired with our firm-level analysis in subsection 6.3.

First, we explain the problems with analyzing changes over time and motivate looking at moments besides the first. Recall that time fixed effects serve as secular shifters of quality *relative* to the outside good. Hence, if there is a large shock to the supply of the foreign good (such as quotas being dismantled or tariffs decreasing), then this will be picked up as a negative *demand* shock to domestic goods. This is a weakness of any demand model that regress shares on prices. In particular, this method usually estimates $\delta_{it} - \delta_{0t}$ where δ_{0t} is the outside good's quality. Hence, as long as the share and quality of the outside good is stable over time, then looking at changes in δ_{it} reflects changes in quality. However, if the share of the outside good is changing rapidly, which is likely to be the case in international trade when large liberalizations are often the empirical object of interest, then one cannot easily interpret changes in absolute quality over time.

This does not mean that changes over time are completely meaningless. To understand this, consider the following definition of aggregate *sales* weighted quality:

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

that defines an "industry" level aggregate quality. Ideally, it would pick up secular growth in quality and we could look at quality changes relative to this trend to determine if goods are downgrading or upgrading. However, as we just discussed this is unlikely to be the case. This can be seen in figure 7, which plots the evolution of our aggregate measure of quality. This graph appears to trend upward from 1997 to 2004 or so, then nose-dive around 2004 and 2005 as the MFA came to a close. Then, noisily, it appears to flatten out before beginning to move up again later in the sample – highlighting that the quality and supply of import competition contaminates over-time changes.

To help tease apart the economic forces at play in this quality growth, we use a decomposition of aggregate quality growth that is common in the productivity growth literature. In particular, we break out quality growth into the following components:

$$Qual_t - Qual_{t-1} = \underbrace{\delta_t - \delta_{t-1}}_{\text{Time Trend}} + \underbrace{\frac{N_{entrants}}{N_t}\delta_{it}^{Entrants} - \frac{N_{exits}}{N_{t-1}}\delta_{it-1}^{Exits}}_{\text{Composition}} + \underbrace{\Delta\delta_{it}^{Stayers}}_{\text{Idiosyncratic Growth}} + \underbrace{[Cov(\delta_{jt}, s_{jt}) - Cov(\delta_{jt-1}, s_{jt-1})]}_{\text{Reallocation}}$$

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.

Figures 8-10 show the time fixed effects in *levels*, the composition effects (a growth rate) and again the evolution of the covariance term in levels. One can see that the dismantling of the MFA and fall of the quotas (in as much as these drive the time fixed effects) come out as a massive negative shock in the estimates that is particularly pronounced around 2004 and 2005. The graph of fixed effects demonstrates the problem in looking at estimates over time. On the other hand, the other two figures give us some useful information on aggregate movements. Figure 9 plots the contribution of entry and exit to changes in the quality. One can see that there is a sharp upward trend in this graph: it starts negative and ends the trend positive – reversing sign around the time China enters the WTO. To interpret this figure, think of the composition effect as *where* new entrants enter on the distribution of quality. After the WTO shock, new entrants are of higher quality relative to incumbents and pushed mean quality up. This is in line with the point made in Section 4 that a negative aggregate demand shock to domestic producers should change the entry cutoff in such a way that induced entry of higher quality competitors. It is difficult to determine whether the radical changes in trade barriers that occurred from 2001-2005 were alone responsible for this trend. However, the figure suggests at least some response – that import competition might drive quality upgrading in the aggregate by moving it at the extensive margin.

The final figure plots the evolution of the covariance term between quality and market shares. This term appears to be flat around 2003, then drops considerably and flattens out again in 2005 and onward. A lower covariance between quality and market share may reflect the fact that new entrants are of relatively high quality but of low market share. However, it could also reflect a trend of relative quality *downgrading* among established firms. This latter possibility aligns with our observation that offshoring induces tightening of the quality ladder – in the extreme where all firms only produced one quality there would be zero correlation between market size and output quality as the former would be determined solely by productivity. In some sense then, this decreasing covariance term (the opposite of what one expects for the productivity decomposition) reflects that offshoring allows firms to more easily produce similar quality goods but has a more muted impact on the ability of firms to scale their production. Nevertheless, it is difficult to discern these effects from each other since much of our analysis is about movement along the quality ladder. Both of these stories fit with the predictions of our model – that there is a heterogeneous response to new offshoring opportunities that may induce downgrading by middle quality firms. With the facts of this decomposition in mind, we turn to explaining aggregate changes in the distribution.

We previously stressed that the quality ladder is shortening. A plot of these lengths in figure 11 confirms that this is not specific to our choice of years.[35] We run a structural break test, the results of which are in table 10. The downward trend in ladder length increases after China's entry into the WTO – suggesting that this trade shock plays a role in aggregate trends. The compression 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 have a muted response. It is no doubt true that both forces are important – we have evidence of import competition in the entry and exit patterns; we will show the significance of firms' heterogeneous response in terms of offshoring and qualityin the next subsection. This empirical finding mirrors some of the intuition from the productivity literature that suggests lowest productivity firms exit after a trade shock – in addition to this effect, we show that firms may endogenously compress further the distribution of quality.

In addition to this shortening of the ladder, there is movement from the right end point

[35]We also run a host of robustness checks on measuring ladder length – including looking at standard deviation and the IQR. The compression and results of the MFA shock are robust to these differences.

inward. A plot of the skewness (Figure 12) reveals a sharp trend towards more negative skew. Our model suggests two competing forces that might shape the changing skewness: a muted response at the top end with the heterogeneous response at the bottom and the middle would suggest a thinning right tail with more mass near the mean; on the other hand, highly productive firms that nevertheless produce lower quality goods (productivity versus capability) will remain. If productivity and capability are positively correlated but imperfectly so, there will remain a thin mass of low quality firms at the left tail while many high quality entrants create more mass near the mean and the right end. The skewness result then highlights a tension between the offshoring predictions of our model and more standard predictions regarding entry and exit and suggests that this latter effect may dominate in certain parts of the aggregate story. The focus on higher moments that we have employed here allows us to make this sort of interpretation. However, it is important to note that this is at best suggestive. This serves as a cautionary note for the use of structural demand estimation in trade - the estimates of quality are difficult to compare over time. With this discussion in mind, in the next section we exploit firm-level variation to unpack how offshoring matters for the distribution of quality.

6.4 Firms' Quality Choices and Trade Regime Switches

In this section, we turn our analysis to firm-level variation in offshoring and quality decisions to explore the correlation between the two. We stress that output quality and sourcing decisions are at least partially jointly determined – thus, our regressions are best interpreted as correlations. However, these correlations are still an important step in understanding *how* firm behavior changes in response to offshoring opportunities and whether or not these behavioral changes conform to the predictions of our and others' models. Our model makes a few key predictions about how firms ought to respond to new offshoring opportunities. In particular, we ought to see a negative correlation between offshoring and quality in general. However, and more subtly, we expect that changes in offshoring conditional on being an offshorer already ought to be positive for some firms and negative for others. Finally, for firms that actually *begin* to offshore, we have predictions about how entry patterns into offshoring ought to correlate to initial quality. At first we focus on a firm's total offshoring activity irrespective of source country. Later on, we analyze how the relationship might vary according to the country where imports are coming from. We specifically look at China and exploit changes in quota structure as a large shock to offshoring costs that affected some firms over others.

Our definition of offshoring follows recent examples in the literature (in particular Hummels et al., 2013 and Autor et al., 2013). We first split imports into three rough categories for each firm – apparel (combined nomenclature (CN) headings 61 and 62), direct apparel intermediates (CN headings 52-55), machinery and other imports. Our definition of offshoring is what Hummels et al. call "narrow offshoring" – we focus on apparel imports per employee. That is to say, offshoring is measured by the value of imports of the final good that the firm produces divided by the number of employees at the firm. We feel this is the correct definition of offshoring for at least two reasons. First, it matches up with the intuition that imports of the same good that is supposedly being manufactured is qualitatively different than imports of intermediates. The former represents a decision of the firm to outsource pieces of the supply chain that are *traditionally* done within the boundaries of the firm[36] while the latter represents imports of intermediates that the firm would need to source domestically if not abroad. Second, scaling the imports by the number of employees versus just looking at imports gives a measure of the *intensity* of offshoring *within* the firm. Figures 13 and 14 plot the time series of average offshoring by firms. While there is some volatility in the measure, there is broadly speaking an upward trend in offshoring that is increasing until about 2001, then settles down then begins to climb again around 2004 and 2005 when the MFA ends.[37] A similar pattern emerges when we restrict ourselves to offshoring to China, except that the increase is substantially larger. Moreover, much of the increase after 2007 when the EU's special treatment of Chinese imports came to an end. 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, despite the fact that much manufacturing activity had been sent abroad already, it's clear that firms still moved parts of their production processes abroad.

There are two immediate concerns with this measure of offshoring. First, more productive firms may require less workers to generate the same amount of final output. Thus, our measure of offshoring will be correlated with productivity, which could bias any results we have if quality and productivity are also correlated. Second, we use the value of imports, which is by definition the product of the quantity and the price of the various imported products. Price reflects input quality and so a higher quality firm will mechanically offshore "more" if offshoring is measured in value.

To overcome the first source of bias, we include exports and intermediates as controls for productivity in our regressions. Nearly all models with heterogeneous firms posit a

[36]While we don't observe the boundaries of the firm abroad, discussions with people in the industry as well as the literature on fashion suggests that outsourcing abroad, and not multinational production, is the dominant form of trade in the apparel industry.

[37]One reason for the volatility is that reporting thresholds for imports change for Eastern European countries that entered the EU. This is only mildly problematic since we employ time fixed effects. If one worries some bias is introduced in the cross-section, the robustness of our results with Chinese imports only should alleviate concerns.

positive correlation between exporting and productivity (Melitz, 2003) and many papers find a correlation between the use of imported intermediates and productivity (see e.g. Halpern, Koren and Szeidl, 2005; Kasahara and Rodrigue, 2008).

To address the second source of bias, we did two things. First, we used the wage bill instead of the number of employees in constructing our measure and we find that our results below are ultimately robust to this change in definition. Second, we use a proxy of input quality that we discuss below as a control for the input quality of the firm. This has its limitations since it is a cross-country measure. Note that we posit a decreasing relationship (in the cross-section at least) between offshoring and quality. And so, any positive bias as discussed above actually works *against us*.

As a final caveat, it is important to remember that we can only look at how a firm's quality changes *relative* to other firms. In any differences regression, we are exploring how a product's relative position changed. As discussed in the previous section, we cannot actually identify aggregate changes in quality. This automatically means that, in terms of magnitude, our coefficients will no doubt be smaller than the actual correlations we are interested in. The reason is because our variation is only variation off of trend. However, there is still substantial variation in access to and use of offshoring across products because of quotas and price differences. Moreover, there is substantial heterogeneity in firms' decision to engage in offshoring. This leaves a great deal of variation in the data and thus still leaves plenty of room for quantitative analysis of the correlation between offshoring and quality, in spite of our agnosticism about aggregate changes.

Keeping these issues in mind, we now test the most basic premise of our model – the relationship between offshoring intensity and output quality. We are primarily concerned with regressions of the form:

$$l_{jt} = \alpha_g + \beta \times Offshoring_{ft} + X_{ft}\gamma' + \epsilon_{jt} \tag{1}$$

where l_{jt} is the relative ladder position of good j at time t, α_g is a product (CN 8-digit) fixed effect, $Offshoring$ is our measure of offshoring activity of a firm and X is a set of controls and year fixed effects. Since our measure of offshoring is constant across products within a firm, we follow the same clustering strategy as we did for our original estimation. That is to say, we cluster at the firm-year level. The results of this regression are in table 11. In the first regression, the coefficient on offshoring is positive. This is not surprising in light of our discussion above that our offshoring measure may conflate other aspects of the firm with actual offshoring intensity. In column 2, we report regression results with controls for firm productivity. Once we control for productivity and intermediates, we see that the coefficient changes sign and becomes insignificant. Intermediates on the other hand are positively and significantly related to our quality ladder. Finally, column 3 has the full specification that includes controls for productivity and allows the correlation between quality and productivity to depend on input quality. As for our measure of output quality, input quality is not observable and we create a proxy variable. We use the fact that the literature on trade and quality posits an increasing relationship between the mean quality of a country's exports and the country's mean income (see e.g. Manova and Zhang, 2012). Using this idea, we construct the following proxy for the level of the quality of a firm's offshored inputs:

$$\delta_{ft,input} = \sum_{c} s_{ft,c} \log \left(gdpPerCapita_{c,t} \right)$$

where $s_{ft,c}$ is the share of imports of firm f at time t from country c (where we focus on offshoring imports) and $gdpPerCapita_{c,t}$ is the GDP per capita of country c in year t. A drawback of this method is that does not allow for within-country differentiation. Nevertheless, the results demonstrate that this is already a powerful proxy even if it just relies on variation in firms' sourcing strategies. The measure of quality varies between roughly 8.7 and 9.8 - suggesting that the effect of offshoring varies between roughly between -.044 and -.021. Moreover, when including this variable as control the effect is significant and negative – albeit small in magnitude. The cross-section regression confirms the model's predictions that all else held constant, offshoring tends to be negatively correlated with quality. The important caveat being that this can be tempered (and perhaps reversed) if the sourcing country itself produces sufficiently high quality output. Of course, this latter possibility is not the focus of our work as we are predominantly interested in developed nations offshoring to developing ones. The regressions above also highlight an important but often ignored caveat in regards to assessing the impact of offshoring on many firm level outcomes, as has become increasingly common in the literature. In particular, if firm level variables (e.g., wages) are correlated with output quality then failing to control for quality adequately could produce spurious relationships in the data. Keep in mind that the relationship is ultimately significant even in the presence of our caveats discussed above. That is to say, the small magnitude should not be too disconcerting given the prevalence of offshoring and the difficulty of assessing any absolute measures of quality. Contrary to skepticism, the fact that in an industry that has been globalizing for decades we still find that offshoring matters for output quality (regardless of whether this is a decision variable) underscores the likely importance of vertical differentiation in a complete understanding of globalization. The results also underscore a point that when thinking about trade policy and offshoring, the anticipated effects will not only hinge on country's comparative advantages in terms of productivity along the supply chain, but also their differences in ability to produce different levels of quality.

While we find a negative correlation between quality and offshoring in the cross-section of firms, the more interesting implication of our model and of others in the trade and quality literature is the potential for a heterogeneous response to new offshoring opportunities. To assess how cheaper offshored inputs induce changes in firms' relative position in the quality ladder we run regressions of the form:

$$\Delta l_{jt} = \beta_1 \times \Delta Offshoring_{ft} + \beta_2 \times \Delta Offshoring_{ft} \times l_{t-1} + X_{ft}\gamma' + \epsilon_{jt} \qquad (2)$$

where once again X is a set of controls that includes year fixed effects. We also now include lagged offshoring as a control variable. The idea here is that increasing one's offshoring activity becomes more difficult as the level of offshoring is higher (this would naturally come out of a model of tasks where more complex tasks become more difficult to offshore). To that end, small increases in offshoring might induce large changes in quality (or vice versa) depending on how difficult it is for a firm to offshore the next portion of the supply chain. In addition, looking at changes allows us to explore our model's predictions for heterogeneity in responses. In particular, the β_2 in the above specification tests if movement on the quality ladder depends on one's initial position on it. As a reminder, our model predicts that lower quality firms ought to increase their quality when they offshore more, while middle quality firms ought to downgrade their quality and high end firms ought not to respond much. Thus our model implies that $\beta_1 > 0$ and $\beta_2 < 0$ in the full specification.[38] The results of these regressions are in table 12. In all specifications, the coefficient on growth in offshoring is positive. This suggests that increasing one's offshoring activity tends to increase one's quality. This does not conflict with the negative correlation in the cross-section. Importantly, our model predicts that different kinds of firms will engage in offshoring and take up offshoring opportunities. What is happening is that, in percentage growth, it is mostly low quality firms that increase their offshoring presence and low quality firms that tend to upgrade. The result is the compression of the quality ladder that we noted in the previous section.

In column 4, one can see the results of our regressions that allow for heterogeneity in how growth in offshoring correlates to movement in the quality ladder. We find that, controlling

[38]Technically, it suggests a more complex relationship where there is yet a third quadratic term in lagged quality that is also positive – i.e., the effect of offshoring ought to be increasingly negative but then eventually be positive again. We ran regressions of this sort, but while all the signs were correct we could not separately identify the coefficients in such a large expansion.

for productivity, input quality and initial amount of offshoring, the response of quality to offshoring is eventually negative for middle and higher quality firms while remaining strongly positive for lower quality firms. This need not be the case and in fact, our model demonstrates that what determines the response of firms to offshoring is both the absolute cost of offshoring units and the slope of the cost of offshoring with respect to quality. Depending on how these both move, firms may respond differently. As an important caveat, we again stress that our analysis does not necessarily mean that the absolute level of quality is decreasing. It could indeed be the case that all firms upgrade their quality but that low quality firms upgrade their quality more. That is, in fact, consistent with our model which suggests that aggregate changes in quality are related to the aggregate demand shock and physical productivity changes. Our results here are complementary to those of Bloom et al. (2012). They find that offshoring seems to have a positive but imprecisely measured effect on product innovation as measured by R&D, but find some evidence that offshoring matters for measured productivity. Our results confirm that offshoring seems to induce changes in the products of firms. However, the structural demand estimation allows us to examine products separately from physical measures of innovation such as patents. It also lets us focus not just on physical productivity and process innovation, but on how the good is actually perceived by consumers. Instead, we find a different channel for which offshoring matters – namely that because of costs of output quality in sourcing countries, offshoring can change the output quality ultimately offered by firms.

For the rest of the section, we turn our focus to China – a particularly important player in offshoring discussions. Focusing on China will allow us to explore offshoring to a particular low quality country with its own comparative advantages. Moreover, we can use the fall of the MFA to explore how quality determines selection into offshoring activities. This is particularly important for our analysis because much of apparel had already been offshored by the beginning of our sample. While we already focus on firms that at least source inputs domestically, China's entry to the WTO allows existing offshorers to change their sourcing strategies which means that China's entry to the WTO gives a source of variation even amongst existing players.

As demonstrated above, understanding that input quality varies is an important control in determining the importance of offshoring. While we find it difficult to control for input quality, controlling for the sourcing nation goes a long way in alleviating these issues. Given this fact and also the rising prominence of China in the global economy, we focus our attention there. We augment the regressions above as follows:

$$l_{jt} = \alpha_g + \beta \times ChinaOffshoring_{ft} + X_{ft}\gamma' + \epsilon_{jt} \tag{3}$$

where we have moved our focus to narrow offshoring from China. The first two columns of table 13 repeat our cross-sectional regressions. Again we find negative correlations and, when all the controls are added, a significant negative relationship between offshoring intensity to China and quality. The magnitude is small, reflecting the same caveats we had before. Nevertheless, the sum of our cross-sectional analysis sides with common intuition that a large degree of offshoring is associated with a lower quality product. That being said, we again turn to our more interesting predictions regarding changes and firms' heterogeneous responses. 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 confirm previous findings. Namely, there is a positive relationship between offshoring growth and ladder movement at the low end that becomes muted and eventually changes sign. These results imply that understanding the impact of China could be more nuanced than previously thought. If certain industries compress their quality ladders in response to a China shock while others expand theirs, the implications for other variables such as wages or prices can also vary across industries. Moreover, this means that the predictions of how variables such as wges respond to a China shock depend very much whether the affected industries are better modeled as vertically or horizontally differentiated. This explains some of the cross-industry variation in how variables respond to offshoring shocks from China. It also implies that the effects of trade shocks depend on comparative advantage in differentiated versus homogeneous goods and where different country's lay on the quality ladder.

In addition to elaborating on previous results, we can use China's entry to the WTO and the dismantling of the MFA to assess selection into offshoring. Typical models predict that intensity of trade (whether it be exporting, importing intermediates or offshoring) is positively correlated with productivity and thus size. Our model suggests a negative correlation between entry into offshoring to low quality countries and initial quality. In fact, recent literature attempts to explore these two possible sources of heterogeneity (the quality and the productivity margin) and see what this implies for selection into trade (Saravia and Voigtlaender, 2013; Holmes and Stevens, 2013). The sharp drop in quotas from China, paired with our measures of quality and proxies for productivity allow us to identify the degree to which each margin is predictive of entry into offshoring to a low-quality producing country. With this discussion in mind, we run variants of the following probit regression:

$$Y^*_{offshoreToCN} = \beta_1 \times quality_{jt} + \beta_2 \times quality_{jt} \times QuotaUtilization_{jt} + X'_{ft}\gamma + \varepsilon_{jt}$$

where Y^* is the latent propensity to offshore, quality refers to our quality ladder measure and quota utilization is the number of licenses over the quota – i.e., the fill rate. The results of this regression are in table 14. In all specifications, the coefficient on the ladder position is negative. It is significant in the regression that includes all controls. This confirms the common intuition that, once productivity and size are controlled for, quality determines a firm's offshoring patterns. The inclusion of the quota terms sheds light on the response of firms experiencing a large shock. Here the interaction term between ladder and quota is negative. This suggests that higher quality firms are less responsive to drops in quotas than lower quality firms. While many firms have posited such relationships, our papers is one of the few to demonstrate this fact empirically. Our final specification, however, presents a comparison of whether quality or size (here measured by number of workers and amount of intermediates) is more important in determining a firm's decision to offshore. We find that the size effect largely dominates the importance of quality. An apples-to-apples comparison of the magnitudes are difficult because our quality measure has no real units, but even at the maximum levels of our quality measure, the impact on offshoring probabilities is swamped by the coefficients on the size effects. Part of this is because our "size" proxies naturally pick up the fact that firms with large employment shares may be more able to offshore and firms that import many intermediates are naturally more open. Even with these caveats, the disparity in coefficients is quite substantial. In some sense this helps explain the low correlation in the cross-section between quality and offshoring and why it becomes stronger when productivity is controlled for; in particular, it could be that much of the binary decision to engage in offshoring to China might be determined by size considerations while still leaving room for quality of offshored production to play a role. Even given this disparity in magnitudes, a firm's position in the quality ladder is both a statistically and an economically significant predictor of offshoring activity.

7 Conclusion

In this paper, we develop a model of offshoring and quality decisions by firms. We then use detailed information about the products made, imported and exported by Danish apparel firms 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). We use the aggregate responses to MFA to test our model's predictions and as a source of exogeneous variation to the Danish apparel industry's access to foreign input markets. We find that firms' product quality is strongly affected by the change in the competitive environment and offshoring opportunities. We observed a tightening of quality ladder together with a change in the shape of the quality ladder. We also documented how offshoring was associated with a decline in quality, especially when offshoring to countries with lower input quality. However, we also find that there is a good degree of heterogeneity in this response along the quality ladder. Many papers posit such relationships and ours demonstrates robust evidence of such a heterogeneous response. In particular, higher quality firms tend to lower their quality more than already-low-quality firms as they begin to engage in offshoring. Our work therefore suggests that globalization has not only led to more competition, but has also affected the quality of the products designed by Danish firms, as they have seized new opportunities to take advantage of offshoring part of their production process.

Our work suggests a wide range of 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. It would be interesting to determine how much of the cost savings induced by offshoring are passed through into quality-adjusted prices. Also, we have focused here only on those firms that maintain some kind of production operations in Denmark. Future work should bring in wholesale firms and firms that switch out of manufacturing into wholesale; this would allow researchers to analyze how offshoring opportunities not only change quality but the industrial structure and allocation of surplus in vertically differentiated markets. Finally, our estimation procedure relies on income-independent tastes for price and quality. Future work could attempt to estimate richer demand systems that would allow for more precise estimates of quality and allow for realistic welfare calculations.

References

[1] Ackerberg, D., Caves, K. and Frazer, G., 2006. Structural Identification of Production Functions. Mimeo, UCLA.

[2] Amiti, M., Itskhoki, O. and Konings, J., forthcoming. Importers, Exporters, and Exchange Rate Disconnect. American Economic Review.

[3] Amiti, M. and Khandelwal, A., 2013. Import Competition and Quality Upgrading. Review of Economics and Statistics 95, 476-490.

[4] Anderson, S, de Palma, A., and Thisse, J., 1987. The CES is a Discrete Choice Model? Economics Letters 24,139-140

[5] Anderson, S, de Palma, A., and Thisse, J., 1992. Discrete Choice Theory of Product Differentiation, Cambridge, MA: The MIT Press.

[6] Autor, D., Dorn, D. and Hanson, G., 2013. The China Syndrome: Local Labor Market Effects of Import Competition in the United States. American Economic Review 103, 2121–2168.

[7] Bernard, A., Redding, S. and Schott, P., 2007. Comparative Advantage and Heterogeneous Firms. Review of Economic Studies 74, 31–66.

[8] Berry, S., 1994. Estimating Discrete-Choice Models of Product Differentiation. RAND Journal of Economics 25, 242–262.

[9] Berry, S., Levinsohn, J. and Pakes, A., 1995. Automobile Prices in Market Equilibrium. Econometrica 63, 841–890.

[10] Bloom, N., Draca, M. and Van Reenen, J., 2011. Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity. NBER Working Paper # 16717.

[11] Brambilla, I., Khandelwal, A. and Schott, P., 2010. China's Experience Under the Multi-Fiber Arrangement (MFA) and the Agreement on Textiles and Clothing (ATC). In Feenstra, R. and Wei, S.-J. (Eds.), China's Growing Role in World Trade, NBER, 345–387.

[12] Bustos, P., 2011. Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of Mercosur on Argentinian Firms. American Economic Review 101, 304–340.

[13] Cachon, G. and Swinney, R., 2011. The Value of Fast Fashion: Quick Response, Enhanced Design, and Strategic Consumer Behavior. Management Science 57, 778–795.

[14] Caplin, A. and Nalebuff, B., 1991. Aggregation and Imperfect Competition: On the Existence of Equilibrium. Econometrica 59, 25–59.

[15] Crozet, M., Head, K. and Mayer, T., 2012. Quality Sorting and Trade: Firm-level Evidence for French Wine, Review of Economic Studies 79, 609-644.

[16] De Loecker, J., Goldberg, P., Khandelwal, A. and Pavcnik, N., 2012. Prices, markups and Trade Reform. NBER Working Paper # 17925.

[17] Eslava, M., Fieler, A. C. and Xu, D., 2013. Trade, Competition and Quality- Upgrading: A Theory with Evidence from Colombia. Working Paper, Duke University

[18] Evans, C. and Harrigan, J., 2005. Distance, Time, and Specialization: Lean Retailing in General Equilibrium. American Economic Review 95, 292-313.

[19] Feenstra, R., 2010. Offshoring in the Global Economy: Microeconomic Structure and Macroeconomic Implications. MIT Press.

[20] Feenstra, R. and Romalis, J., 2014. International Prices and Endogenous Quality. Quarterly Journal of Economics 129 (2), 477-527.

[21] Frederick, S. and Staritz, C., 2012. Success and Upgrading After the End of the MFA. In Lopez-Acevedo, G. and Robertson, R. (Eds.), Sewing Sucess? Employment, Wages and Poverty Following the End of the Multi-Fibre Arrangement, The World Bank, 87-136.

[22] Gereffi, G., 1999. International Trade and Industrial Upgrading in the Apparel Commodity Chain. Journal of International Economics 48, 37–70.

[23] Grossman, G. and Rossi-Hansberg, E., 2008. Trading Tasks: A Simple Theory of Offshoring. American Economic Review 98, 1978–1997.

[24] Hallak, J.-C. and Schott, P., 2011. Estimating Cross-Country Differences in Product Quality. Quarterly Journal of Economics 126, 417–474.

[25] Halpern, L., Koren, M. and Szeidl, A., 2005. Imports and Productivity. CEPR Discussion Paper # 5139.

[26] Holmes, T. and Stevens, J., forthcoming. An Alternative Theory of the Plant Size Distribution with an Application to Trade. Journal of Political Economy.

[27] Hummels, D. and Alexandre Skiba, A., 2004. Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture. Journal of Political Economy 112, 1384–1402.

[28] Hummels, D., Jørgensen, R., Munch, J. and Xiang, C., forthcoming. The Wage Effects of Offshoring: Evidence from Danish Matched Worker-firm Data. American Economic Review.

[29] Kasahara, H. and Lapham, B., 2013. Productivity and the Decision to Import and Export: Theory and Evidence. Journal of International Economics, 89:297–316.

[30] Khandelwal, A., Schott, P. and Wei, S.-J., 2013. Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters. American Economic Review 103, 2169–2195.

[31] Khandelwal, A., 2010. The Long and Short (of) Quality Ladders. Review of Economic Studies 77, 1450–1476.

[32] Kugler, M. and Verhoogen, E:, 2009. Plants and Imported Inputs: New Facts and an Interpretation. American Economic Review (Papers and Proceedings) 99, 501–507.

[33] Kugler, M. and Verhoogen, E., 2012. Prices, Plant Size, and Product Quality. Review of Economic Studies 79, 307–339.

[34] Liegey, P., 1993. Adjusting Apparel Indexes in the Consumer Price Index for Quality Differences. In Foss, M., Manser, M. and Young, A. (Eds), Price Measurement and Their Uses, NBER, Chapter 6, 209–226.

[35] Manova, K. and Zhang, Z., 2012. Export Prices Across Firms and Destinations. Quarterly Journal of Economics 127, 379–436.

[36] Martin, J. and Méjean, I., forthcoming. Low-wage Country Competition and the Quality Content of Exports. Journal of International Economics.

[37] Melitz, M., 2003. The Impact of Trade on Intra-industry Reallocations and Aggregate Industry Productivity. Econometrica 71, 1695–1725.

[38] Danish Ministry of Business and Growth. Denmark at Work - Plan for Growth in the Creative Industries and Design. Online, February 2013.

[39] Piveteau, P. and Smagghue, G., 2013. A New Method for Quality Estimation Using Trade Data: An Application to French Firms. Mimeo, Columbia University.

[40] Roberts, M., Xu, D., Fan, X. and Zhang, S., 2012. A Structural Model of Demand, Cost, and Export Market Selection for Chinese Footwear Producers. NBER Working Paper # 17725.

[41] Saravia, D. and Voigtlander, N., 2012. Imported Inputs, Quality Complementarity, and Skill Demand. Mimeo, UCLA.

[42] Sutton, J., 1998. Technology and Market Structure: Theory and History. MIT Press.

[43] Sutton, J., 2012. Competing in Capabilities: The Globalization Process. Oxford University Press.

[44] Train, K., 2009. Discrete Choice Methods with Simulation. Cambridge University Press.

[45] Utar, H., forthcoming. When the Floodgates Open: "Northern" Firms' Response to Removal of Trade Quotas on Chinese Goods. American Economic Journal: Applied Economics.

[46] Verhoogen, E., 2008. Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector. Quarterly Journal of Economics 123, 489–530.

[47] Yi, K.-M., 2003. Can vertical specialization explain the growth of world trade? Journal of Political Economy 111, 52–102.

Appendix A: CES version of the model

In this appendix we consider two twists to the model presented in the main text: (1) we build in CES preferences and (2) we build in a slightly different cost structure in order to generalize our main propositions. We explore CES preferences because of their prominence in the trade literature and also because their tractable structure on the demand side allows us to construct a generalized cost side.

To begin, a representative consumer in this model has CES preferences with quality shifters given by:

$$U = \left(\int_{\mathcal{J}} [v(j)x(j)]^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$$

where once again j indexes both a firm and a variety, $\mathcal{J}$ is the set of all varieties, v_j is the quality of variety j, x is the quantity consumed of a given variety and σ is the elasticity of substitution between any two varieties. This admits a well known demand function given by:

$$x(j) = Av(j)^{\sigma-1}p(j)^{-\sigma}$$

where A is an aggregate demand shifter that depends on the ideal price index generated by the CES preferences above and aggregate income. Firms ultimately take A as given.

Now consider a firm j that faces a constant unit cost and the same quality production as in the main text,

$$v_j = \psi^{\alpha}$$

where, as a reminder, ψ is input quality. As is well known, CES demand dictates that price be a constant multiplicative markup over marginal cost. Concentrating out price yields the firm's quality maximization decision:

$$\max_{\psi} \quad A\psi^{\alpha(\sigma-1)}c(\psi,\tau)^{1-\sigma} - f(\psi)$$

where, in an abuse of notation, we have wrapped constants into the aggregate shifter, $f(\psi)$ is a fixed cost and $c(\psi,\tau)$ is a cost function that depends on quality and trade costs τ. For the moment we abstract from a discussion of productivity and capability. We assume that $\tau_F > \tau_H$ and that $c(\psi,\tau)$ is strictly log-supermodular. Notice that there is already a departure from the baseline model since now costs are indexed by a single parameter. WLOG we fix $\tau_H = 0$ and $\tau_F = \tau$. This model of costs does *not* strictly contain the model of section 4 – instead we will introduce firm heterogeneity through fixed costs of design.[39] This model

[39] Nevertheless, all of the analysis quite easily carries through with CES preferences and our original specification of costs. Details of this work are available upon request.

will ultimately have many of the same features and properties of Section 4, but allows for some more precise statements about what we require of the production function to be met.

The first order condition with respect to quality is now given by,

$$\frac{d\pi}{d\psi} : (\sigma - 1)\, r(\psi, \tau) \left[\frac{\alpha}{\psi} - \frac{\partial c / \partial \psi}{c(\psi, \tau)} \right] = f'(\psi)$$

where $r(\psi, \tau)$ refers to the non-fixed part of profit. If fixed costs are not a function of quality then this reduces to a simple FOC, reminiscent of that in section 4. If $f''(\psi) > 0$ then a *sufficient* condition for profit to be strictly concave is $\frac{d}{d\psi}\, (\sigma - 1)\, r(\psi, \tau) \left[\frac{\alpha}{\psi} - \frac{\partial c/\partial \psi}{c(\psi,\tau)} \right] \leq 0$. This requires some restrictions on the elasticity of costs with respect to quality and the elasticity of the derivative with respect to quality. For example, for constant elasticity costs, $A\psi^{\gamma}$, the restriction is that $(\sigma - 1)(\alpha - \gamma) \in (0, 1]$. Assuming these restrictions hold then the profit function is strictly concave. We assume that an Inada condition holds such that $\lim_{\psi \to 0} r'(\psi) = \infty$ and $\lim_{\psi \to \infty} r'(\psi, \tau) = 0$. In the case of constant elasticity costs it is sufficient to assume that $(\sigma - 1)(\alpha - \gamma) < 1$. This guarantees a unique maximum at some ψ^* exists.

Now we can use the implicit function theorem to deduce how quality responds to trade cost changes. First of all,

$$\frac{\partial \psi}{\partial \tau} = -\frac{\partial^2 \pi / \left(\partial \psi \partial \tau\right)}{\partial^2 \pi / \partial \psi^2}$$

From the second order condition, the denominator is strictly negative. Hence, the sign of this derivative will be equal to the sign of the numerator. The derivative of which is given by,

$$\frac{\partial^2 \pi}{\partial \psi \partial \tau} = (1 - \sigma) \left(f'(\psi) \frac{\partial c}{\partial \tau} \frac{1}{c} + r(\psi, \tau) \frac{\partial \log c(\psi, \tau)}{\partial \psi \partial \tau} \right)$$

Since $f'(\psi) > 0$ and $\partial c / \partial \tau > 0$, log supermodularity of marginal cost is a sufficient condition to ensure that the sign of the derivative is negative. Thus, lowering trade costs induces quality upgrading. Moreover, this automatically implies that if a firm produces at home the quality will be higher *ceteris paribus* than if it offshores.

To induce heterogeneity in sourcing options we allow firms to differ in their ability to do design:

$$\pi(\psi; \tau, \omega) = A\psi^{\alpha(\sigma-1)} c(\psi, \tau)^{1-\sigma} - f(\psi, \omega)$$

where we assume that $-f$ is supermodular and $\partial f/\partial\omega < 0$ while $\partial f/\partial\psi > 0$. Now we have,

$$\frac{\partial\psi}{\partial\omega} = -\frac{\partial^2\pi/\left(\partial\psi\partial\omega\right)}{\partial^2\pi/\partial\psi^2}$$

Once again, all that matters for the sign is the numerator given by,

$$\frac{\partial^2\pi}{\partial\psi\partial\omega} = -\frac{\partial f(\psi,\omega)}{\partial\psi\partial\omega}$$

which, because of supermodularity of $-f$, implies that the optimal choice of ψ is increasing in ω. Finally, suppose there are two countries that firm can source from. In this case $\frac{\partial(\pi_H-\pi_F)}{\partial\omega} > 0$ because $-f$ is supermodular, $\psi_H > \psi_F$, and the envelope theorem. We assume that there is a unique $\omega^* \in [\underline{\omega},\bar{\omega}]$ such that

$$\pi^*(\tau_H,\omega) = \pi^*(\tau_F,\omega)$$

$$A\psi_H^{\alpha(\sigma-1)}c(\psi_H,0)^{1-\sigma} - f(\psi_H,\omega) = A\psi_F^{\alpha(\sigma-1)}c(\psi_F,\tau)^{1-\sigma} - f(\psi_F,\omega)$$

Then this ω^* will be a cutoff value such that firms with $\omega < \omega^*$ offshore and those above produce domestically. Now assume that τ falls. It's easy to see that the cutoff will rise – so that more firms will begin to offshore. Hence, the drop in trade costs will induce an increase in quality by low quality producers who are already offshoring, a decrease in quality by those middle quality firms that begin to offshore and a muted response by high quality firms who continue not to offshore. Thus, the assumptions on costs above would yield a proposition exactly like that of section 4. As a final aside, one can show that $\partial\psi/\partial A \propto f'(\psi)/A$. Hence, A is irrelevant in the absence of fixed costs but a decrease in A should lower quality of all domestic producing firms if there are fixed costs. This also relates to our Section 4 model. In particular, it demonstrates that with CES preferences in the absence of fixed costs a neutral shifter of marginal costs would not impact the optimal choice of quality. In this model with fixed costs, if we were to add a neutral productivity shifter then our analysis *conditional* on productivity would follow as before and we would find, similar to section 4, that quality and sourcing decisions are now carved out in productivity and capability space, with their joint distribution determining the joint distribution of sourcing, quality and size.

Appendix B: Further evidence on the link between workforce and product quality

As an additional test, we split workers into cells and then see how the sizes of these cells vary with a firm's quality position relative to other firms. The cells are based on workers' skill level (proxied by completion of at least 2 years of higher education) and their position in the firm. For firm position, we divide workers into blue collar, mid-skill service and white collar jobs.[40] Table 8 contains the results of this exercise. The group differences are clear – higher quality firms tend to employ workers in more professional roles, and in every professional group, a larger fraction is skilled. This demonstrates that at least some of the quality premium we estimate is due to compositional differences.

The link between average wage and quality does not affect workers the same way. In table B1, we run the same specification as in Table 7 separately for each skill group. Results suggest that, conditional on skill, only high skilled workers enjoy a quality premium. One reason for this assymetry may be that skilled workers are more directly involved in design, sourcing and distribution – key elements of quality.

To explore this further, we link our estimates of quality to data on a panel of the universe of workers at apparel firms. To assess whether or not higher quality firms pay higher wages we run panel regressions of the form:

$$\log w_{ift} = \beta_0 + \beta_2 \log \tilde{\delta}_{ft} + \beta_2 \log Emp_{ft} + \beta_3 \mathbf{X}_{ift} + \varepsilon_{ift}$$

where i indexes workers and s refers to a worker's skill. Because we include size (as a proxy for productivity) and quality, we do not include firm fixed effects, however we cluser errors at the firm-year level to allow for arbitrary correlation between workers at a firm.

no discussion of table 8

Table 9 (B2) contains the results of this analysis. The results reflect the aggregate results. In particular, there appears to be no difference in the wages accruing to low skilled workers at apparel firms. On the other hand, there is a positive correlation between quality and and wages for skilled workers. Notice however that this correlation becomes closer to 0 after worker observables are controlled for (but returns with job fixed effects). This backs up the idea that the quality premium is largely picking up differences in the composition of workers

[40] We define white collar workers as workers in professions with a DISCO code in major group 1 (Legislators, Senior Officials and Managers), 2 (Professionals) and 3 (Technicians and Associate Professionals). We define blue collar workers as workers in professions with a DISCO code in major group 7 (Craft and Related Trades Workers), 8 (Plant and Machine Operators and Assemblers) and 9 (Elementary Occupations). Finally we have a third group of service workers in groups 4 (Clerks), 5 (Service Workers and Shop and Market Sales Workers) and 6 (Skilled Agricultural and Fishery Workers).

across firms. In particular, higher quality firms hire more skilled workers in higher positions, and also hire workers with better worker attributes: in a series of not-shown regressions, we find that market experience, and not age or occupational tenure, is the most positively correlated with quality. From this section, we can conclude that while our measure of quality is noisy and imperfect, it clearly picks up physical differences across firms. Moreover, these physical differences are markers of higher quality inputs (higher wages, more experience, more skilled workers, etc.). On the other hand, there is still unexplained variation in our quality measures – suggesting that taste shocks, uncontrolled by firms, are a large component of our estimates. This does not diminish our results going forward, but highlights common drawbacks to the structural estimation literature – measured shocks are often a conflation of several factors and care needs to be taken in their interpretation.

Appendix C: 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 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	

Table 3: Demand Estimation for Domestic Apparel

	OLS	IV: Logit	IV: Nested Logit
Dep Var:	$\log(s_{jft}/s_{0t})$	$\log(s_{jft}/s_{0t})$	$\log(s_{jft}/s_{0t})$
p_{fjt}	$-.00013$	$-.02129^{*}$	$-.00768^{*}$
	(-1.16)	(-1.82)	(-1.89)
$\log s_{jgft}$	$.901^{***}$		$.321^{***}$
	(95.32)		(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	8,378	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%. All estimation involving two-way clustering done using Stata's *xtivreg2*.

Table 4: Detail on 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

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

Table 6: Estimating the Size-Price Correlation

	(KV)	(1)	(2)
Dep Var: $\log P_{jft}$			
$\log Emp_{ft}$	.1130***	.0989***	.0123
	(2.92)	(2.72)	(.44)
δ_{fjt}		.0509***	.1189***
		(12.92)	(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%.

Table 8: The Composition of Firms Producing Different Quality Outputs

Lowest Quality Quantile	Blue Collar	Mid-Level	White Collar
Low Skill	7.4	5.4	34.3
High Skill	24.4	11.5	17.0
Frac. High	.76	.68	.33

Highest Quality Quantile	Blue Collar	Mid-Level	White Collar
Low Skill	4.4	6.6	34.3
High Skill	21.9	13.9	18.9
Frac. High	.83	.68	.35

Table 7: Correlation Between Wages and Firm Quality

	(1)	(2)	(3)	Low Only	High Only
Dep Var: $\log w_{jft}$					
δ_{fjt}	.0245**	.0224*	.0155	.0149	.028**
	(2.20)	(1.73)	(1.25)	(.92)	(2.26)
$\log Emp_{ft}$		.0055	.0053	.0076	.0079
		(0.32)	(0.33)	(.35)	(.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%.

Table 9: Correlation Between Wages an Firm Quality

	Low Skill			High Skill		
	(1)	(2)	(3)	(4)	(5)	(6)
Dep Var: $\log w_{jft}$						
δ_{fjt}	.0007	-.0083	.0022	.0142**	.0065	.146**
	(.10)	(-1.27)	(.34)	(2.37)	(1.06)	(2.52)
$\log Emp_{ft}$	-.0170	.0050	-.0028	$-.0175$***	-.0054	-.0012
	(-2.45)	(.81)	(-.42)	(-3.11)	(-.98)	(-.23)
Controls		$\surd$	$\surd$		$\surd$	$\surd$
FE	No	No	Job	No	No	Job
Cluster:	Firm-Year	Firm-Year	Firm-Year	Firm-Year	Firm-Year	Firm-Year
	(804)	(804)	(804)	(803)	(803)	(803)
R^2	.0160	.2499	.3715	.0194	.0772	.3372
N	14871	14871	14871	16550	16550	16550

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

Table 10: Offshoring and Quality Ladder Position in the Cross-Section

	Length Measures			Skewness
	$l_{max} - l_{min}$	$l_{p99} - l_{p1}$	σ_l	–
t	$-.130$	$-.054$	-.009	-.027
	(.002)	(.001)	(.0001)	(.0001)
t post-MFA	$-.255$	-.114	-.016	-.027
	(.002)	(.001)	(.0002)	(.0001)

Table 11: Offshoring and Quality Ladder Position in the Cross-Section

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

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

Table 12: Offshoring and Ladder Movement - Overall and Heterogeneous Effects

	Dependent Variable: Δl_{jt}			
	(1)	(2)	(3)	(4)
$\Delta \log(Offshoring)$	.0596***	.0620*	.0623*	.0885***
	(2.03)	(1.92)	(1.92)	(2.69)
$\Delta \log(Offshoring) \times l_{jt-1}$				$-.0827$***
				(-3.22)
$\delta_{ft,input}$			$-.0202$	$-.0103$
			(-.52)	(-.27)
$\log(Intermediates)$		.0069	.0065	.0074
		(.83)	(.78)	(.99)
$\log(Exports)$		.0107	.0099	.0125
		(.77)	(.72)	(.95)
$\log(Offshoring)_{t-1}$		$-.0275$	$-.0269$	$-.0228$
		(-1.29)	(-1.26)	(-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%.

Table 13: Offshoring and Ladder Movement - Overall and Heterogeneous Effects

	Dependent Variable: l_{jt}		Dependent Variable: Δl_{jt}		
	(1)	(2)	(3)	(4)	(5)
$\log(Offshoring_{CN})$	$-.0230$	$-.0290^{*}$			
	(-1.49)	(-1.90)			
$\Delta \log(Offshoring_{CN})$			$-.0044$	$-.113$	.0055
			(-.30)	(-.70)	(.32)
$\Delta \log(Offshoring_{CN}) \times l_{jt-1}$					$-.0521^{***}$
					(-3.57)
$\log(Intermediates)$		.0246		.0053	.0106
		(1.16)		(.30)	(.71)
$\log(Exports)$		$.1373^{***}$		.212	.0192
		(3.03)		(1.27)	(1.16)
$\log(Offshoring_{CN})_{t-1}$				$-.0141$	$-.0141$
				(-1.58)	(-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%.

Table 14: Probit Regression: Probability of Offshoring Projected on Quality

	(1)	(2)	(3)	(4)
δ_{jt}	$-.0313^{*}$	$-.0305$	$-.0233$	$-.0757^{***}$
	(-1.66)	(-1.62)	(-1.22)	(-3.74)
$quotaFill$		$-.2189^{***}$	$-.2016^{***}$	$-.1765^{**}$
		(-3.46)	(-3.20)	(-2.36)
$\delta_{jt} \times quotaFill$			$-.1161^{**}$	$-.0976^{*}$
			(-2.56)	(-1.84)
$\log(Employees)$				$.5736^{***}$
				(5.01)
$\log(Intermediates)$				$.1115^{***}$
				(2.61)
Fixed Effecs:	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%.

Appendix D: Figures

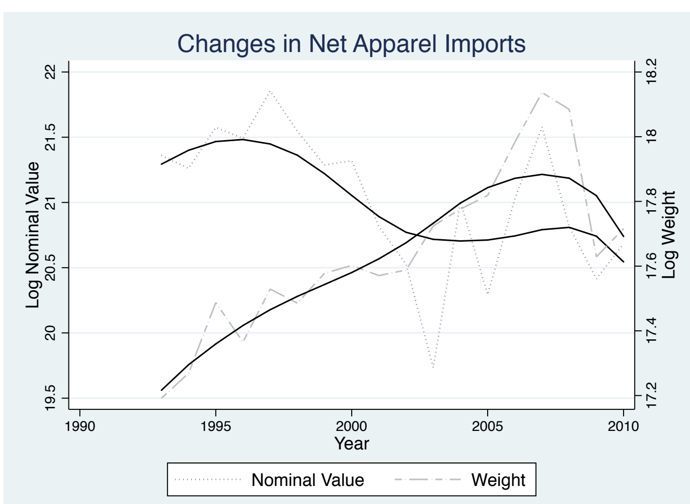

Figure 1: Time Series of Danish Apparel Import

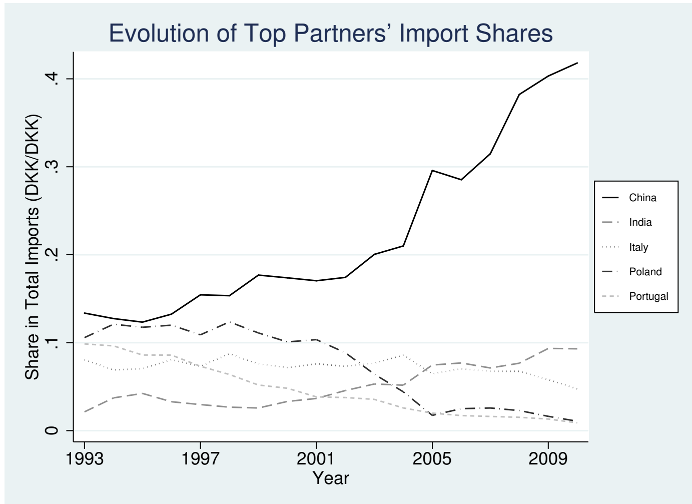

Figure 2: Changes in the Danish Apparel Industry

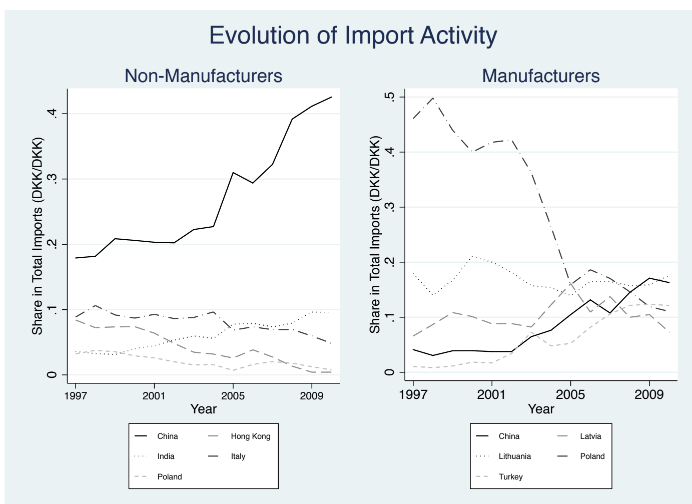

Figure 3: Growth of Chinese Share in Apparel Imports

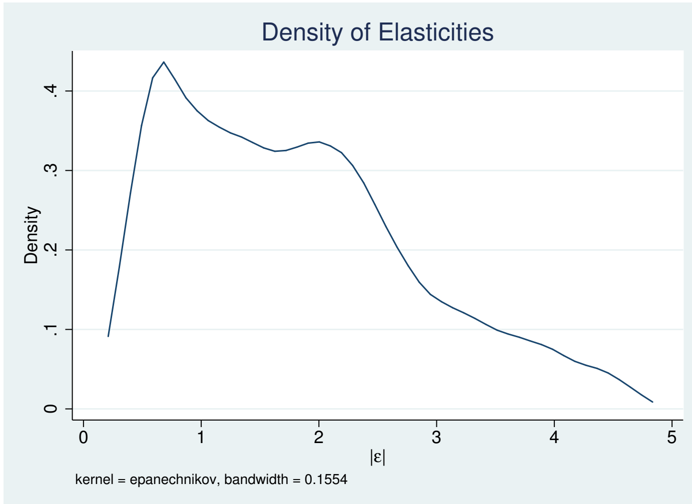

Figure 4: Density of Elasticities

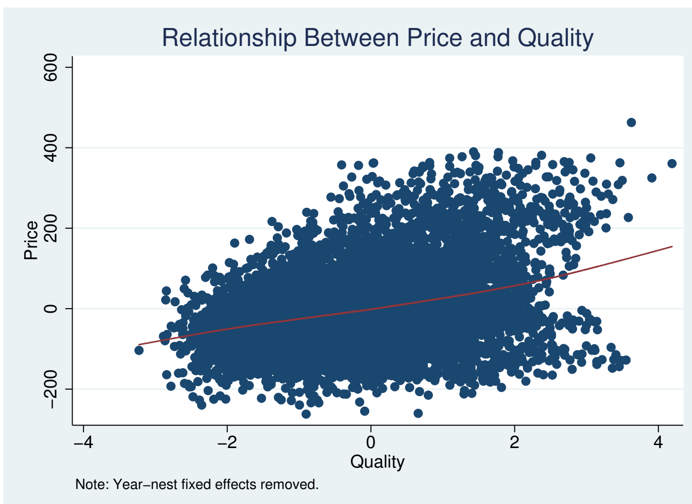

Figure 5: Price versus Quality

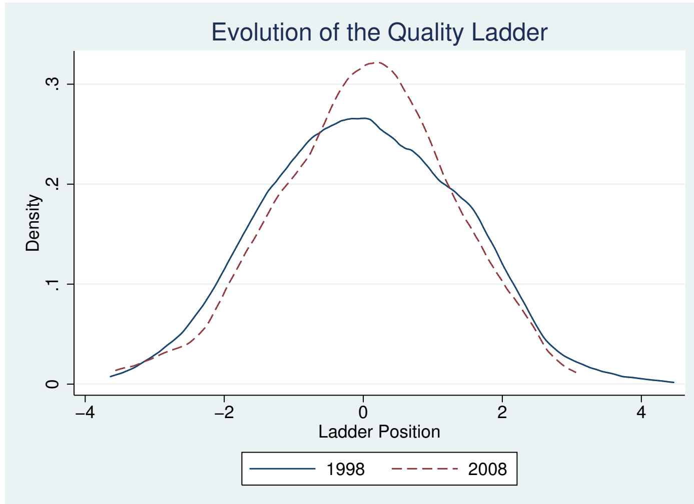

Figure 6: Evolution of Quality Ladders

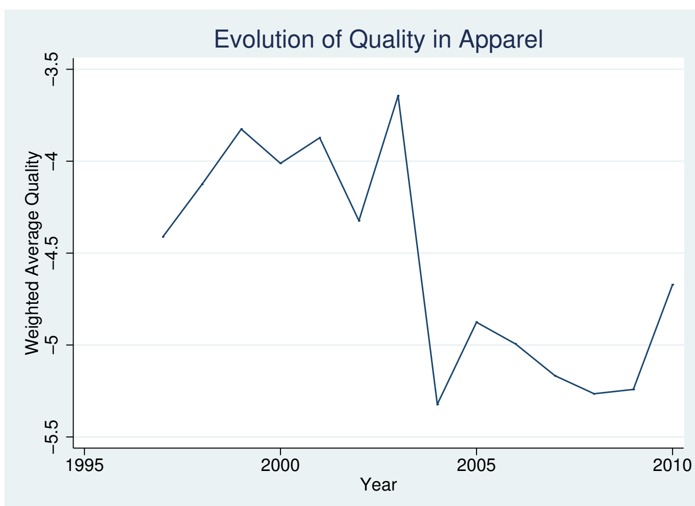

Figure 7: Evolution of Quality

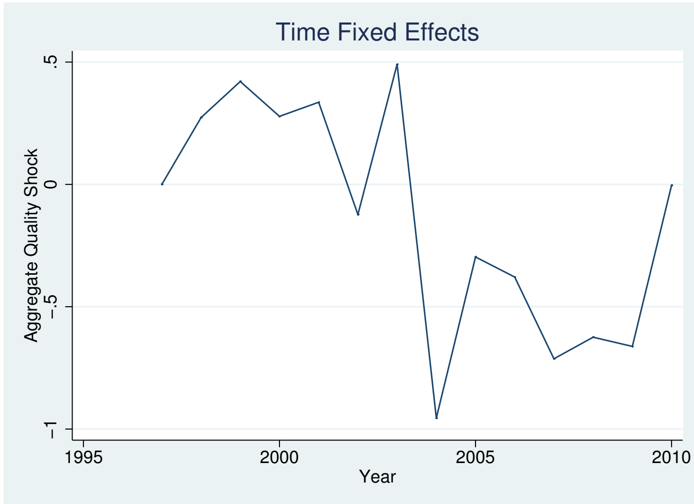

Figure 8: Series of Time Fixed Effects

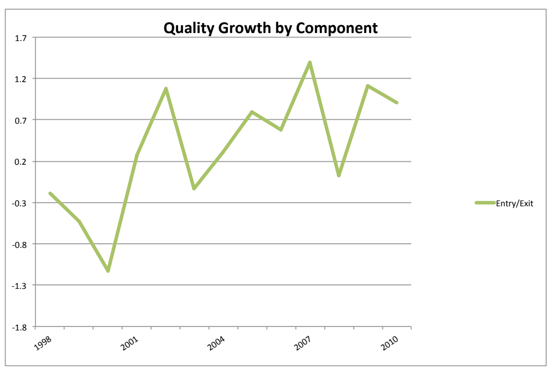

Figure 9: Entry and Exit Component of Quality Growth

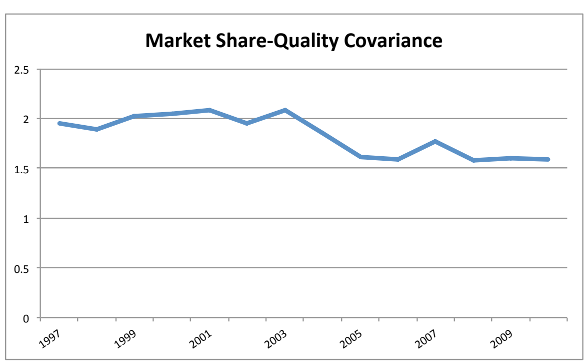

Figure 10: Market Share-Quality Covariance Evolution

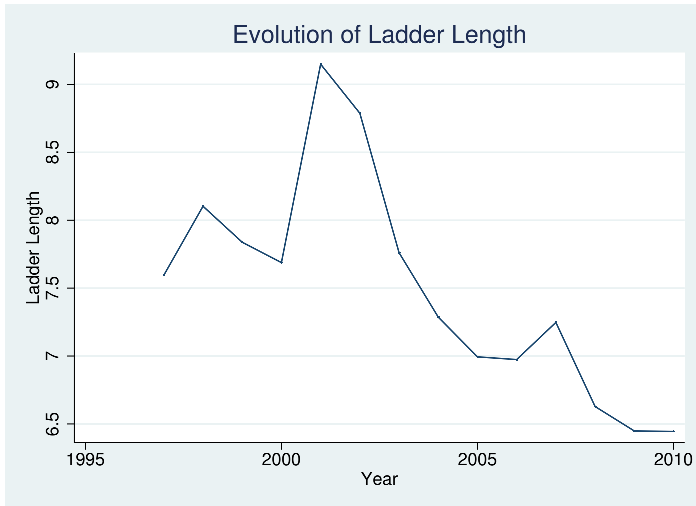

Figure 11: Evolution of Ladder Length

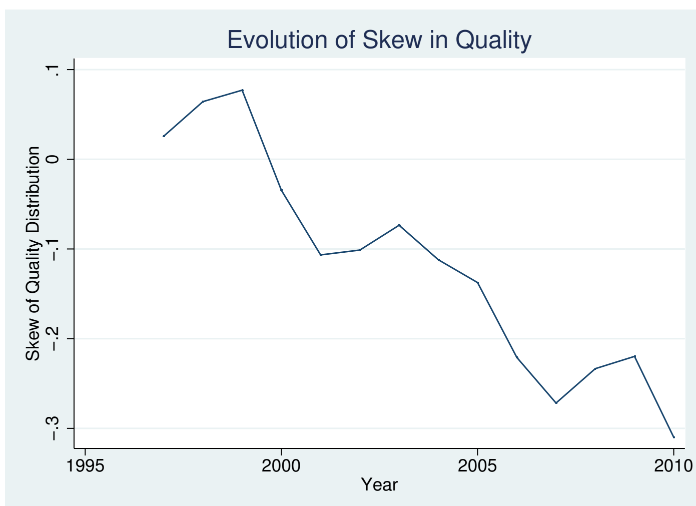

Figure 12: Evolution of Skew in Quality Ladder Distribution

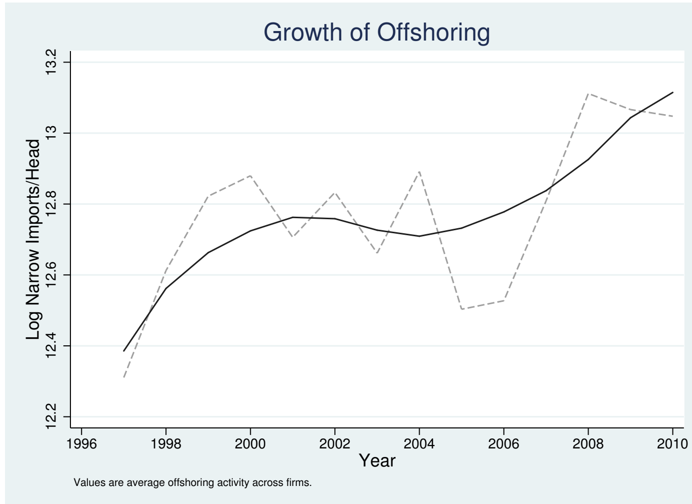

Figure 13: Evolution of Offshoring Activity

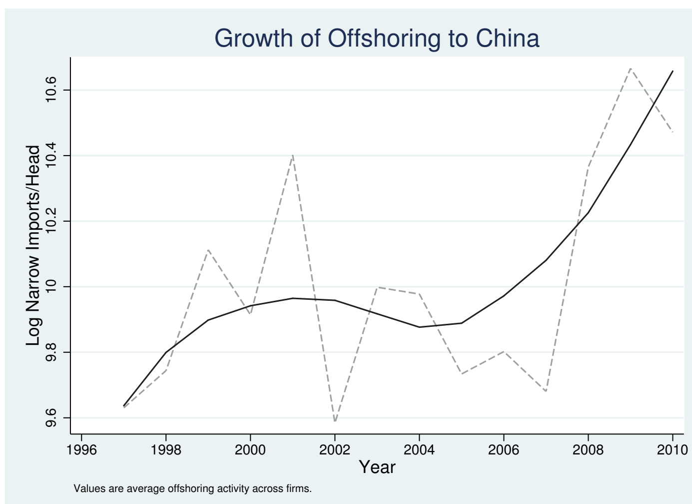

Figure 14: Evolution of Offshoring Activity in China