Exchange Rate Shocks and Quality Adjustments*

Daniel Goetz
University of Toronto Mississauga
Rotman School of Management

Alexander Rodnyansky
University of Cambridge

First Version: August 10, 2016
This Version: July 13, 2018

Abstract

How do firms change the quality composition of their traded goods in response to an exchange rate shock? Using unique data from a large Russian retailer that varies its offerings across seasons, we document that ruble devaluations are associated with a reduction in the observed material and fabric quality of goods the retailer imports for resale. Our results indicate that an increase in the retailer's costs, as opposed to a reduction in demand due to shrinking real incomes, is the driving force. We estimate a simple multi-product sourcing model to quantify the welfare impact of quality adjustments and find that preventing firms from downgrading overstates the welfare loss from the 2014 ruble devaluation by 33%, while incorporating cost heterogeneity but ignoring quality has ambiguous effects on welfare changes in general. *JEL Codes*: E30, F14, F31, L11, L15, L16, L81, M11.

*We are grateful to Mark Aguiar, Markus Brunnermeier, Vasco Carvalho, Giancarlo Corsetti, Jan de Loecker, Mikhail Golosov, Bo Honoré, Oleg Itskhoki, Myrto Kalouptsidi, Jakub Kastl, Atif Mian, Virgiliu Midrigan, David Sraer, Claudia Steinwender, Motohiro Yogo, as well as multiple seminar participants at the Barcelona GSE Summer Forum, University of Cambridge, the annual conference of the CEA, Princeton University, Stockholm School of Economics, and University of Toronto for valuable advice. Rodnyansky gratefully acknowledges financial support from the Alfred P. Sloan Foundation, the CME Group Foundation, Fidelity Management & Research, and the MFM initiative. Goetz: Department of Management, Kaneff Centre, Rm 233, 1833 Inner Circle Rd., UTM, Mississauga, ON L5L 1C6, Canada. Email: daniel.goetz@rotman.utoronto.ca. Rodnyansky: Faculty of Economics, Austin Robinson Building, Sidgwick Avenue, Cambridge, CB3 9DD, UK. Email: a.rodnyansky@gmail.com.

1 Introduction

How do firms respond to cost shocks and what are the most relevant margins of adjustment? Trade economists[1] and the business press[2] have long speculated that companies may adjust the quality of their product offerings instead of changing prices in response to exchange rate movements or tariff reforms. This hypothesis complements a long literature on incomplete price pass-through in international trade by providing another margin of adjustment for firms.[3] While recent empirical work has found indirect evidence of the presence of cost shock induced quality adjustment (Ludema and Yu, 2016), finding direct empirical support for this phenomenon has been challenging because of the difficulty in measuring product quality.

This paper addresses the question of whether quality adjustment is an operational margin for firms using new data from an online Russian apparel retailer. We directly observe the fabric and materials content of hundreds of thousands of individual products offered by the firm, as well as prices, quantities and unit costs. Our analysis combines an intuitive restriction on which fabrics are high quality with high frequency changes in firm product stocking to identify the effect of the 2014 Russian currency crisis on the quality composition of offered products.

We draw on insights from the recent literature on non-homothetic demand in trade to understand why a proportional cost shock can lead to quality reallocation. Trade to wealthy countries is biased towards quality:[4] To explain this pattern, a demand system where consumers switch expenditures to higher quality goods as their incomes increase has been proposed (Fajgelbaum, Grossman, and Helpman, 2011). This form of demand system can also imply that when facing a proportional cost increase—as in an exchange rate devaluation—the profits from a high cost, high quality import can shrink disproportionately compared to a low cost, low quality import, leading

[1] Feenstra (1988) argues that quota restrictions led to substantial quality upgrading for U.S. imports of Japanese cars, and that, more generally, firms may upgrade their products through changing the design or adding extra features when there is a decline in the quantity sold as a result of quotas.

[2] In the aftermath of Brexit, the devalued pound was cited as a reason for shrinking candy bars. See, for example, the Financial Times article "Food groups embrace 'shrinkflation' to cope with rising costs" on December 2 of 2016.

[3] For recent entries on incomplete price pass-through see, for example, Goldberg and Campa (2010), Gopinath and Itskhoki (2010a), Gopinath and Itskhoki (2010b), Amiti, Itskhoki, and Konings (2016), Auer, Burstein, and Lein (2017), and Corsetti, Crowley, Han, and Song (2018).

[4] For earlier work on quality bias in trade, see Alchian and Allen (1964) and Hummels and Skiba (2004).

to reallocation. Conversely, the homothetic demand systems otherwise prevalent in the trade literature will predict exit of low-quality varieties in response to a proportional adverse shock (Crozet, Head, and Mayer, 2011).

Our first contribution is to show that the firm responds to the ruble devaluation in late 2014 by reducing the quality of its product offerings in the following season, where quality is proxied by whether a product uses natural fabric inputs as in Medina (2018). A 1% ruble devaluation causes a roughly 0.35% differential reduction in the fraction of natural fabrics in imported products compared to domestically produced items. Variation in domestic and foreign manufacturing origin of the firm's products, as well as regional variation in the economic impact of a concurrent oil price shock, allows us to identify that rising real wholesale costs—and not an income-shock induced "flight from quality"—are the driving force behind the reallocation.

Having documented quality downgrading, we next turn to the question of why the firm would react to a cost shock by reallocating toward lower quality products. Since the reallocation implies high quality goods became less profitable relative to low quality goods after the shock, we examine whether high quality goods experienced less price pass-through and find that they did not, implying that differentially shrinking markups are not responsible. We instead document a reallocation of quantities from high to low quality within product categories, suggesting that high quality goods were differentially sensitive to the price increase. We thus find evidence for non-homothetic demand as in Bems and di Giovanni (2016), and further show that it is responsible for product quality switching.

Our second contribution is to write and estimate a simple structural model of quality choice to understand how quality downgrading mediates the welfare costs of a devaluation. The retailer's problem is to choose the profit maximizing set and qualities of products to stock each season given an assumed fixed per-product cost of sourcing; since demand is not separable across products, this is a combinatorial discrete choice problem with complementarities as in Antras, Fort, and Tintelnot (2017). Instead of using a full-information solution method, we assume that whether and what quality of products are to be offered in a given season are designated to purchasing managers acting independently, who have private information about the fixed costs of sourcing their designated product and who form expectations about the simultaneous choices of other managers. This transforms the sourcing problem into an incomplete information game of entry and discrete (high or low) quality choice, which can be easily estimated (Hotz and Miller, 1993).

On the demand side, we assume that consumers all value quality in the same way but have idiosyncratic tastes across products, implying a logit demand system with both vertical and horizontal differentiation as in Khandelwal (2010) or Fajgelbaum, Grossman, and Helpman (2011). Intuitively, the demand system implies that consumer expenditures will decrease relatively more for the high-cost, high-quality good when price pass-through from the proportional cost shock is almost complete—as it is in our competitive environment—so high quality goods become relatively less profitable to carry, and the relative share of high quality products will decrease.

The recovered model parameters are consistent with our assumptions on quality in the reduced form exercises: all else equal, natural fabric goods sell 12% percent more and have 82% percent higher marginal costs than artificial fabric goods, implying that natural fabrics are both more valued by consumers and more expensive. Low fixed sourcing costs rationalize observed product entry and exit, as well as relatively low sales per product.

Given our parameter estimates, we analyze how welfare losses after the observed devaluation are affected by quality heterogeneity. We first use the estimated model to analyze a setting where the fixed cost of sourcing low-quality goods is prohibitive, so that high-quality products cannot be downgraded—only discontinued. In response to the 2014 devaluation, consumer welfare losses grow from −16.5% in the base case with quality downgrading to −22.0% (a 5.5pp rise), implying that allowing for low-quality substitutes is preferable to exit for consumers.

Second, we show that quality has an ambiguous effect on the welfare cost of the devaluation. Compared to our base case, a model with no quality heterogeneity implies a 0.02pp smaller welfare loss; however, compared to cases with much higher quality shifters than our base case, the model with no quality heterogeneity implies a larger welfare cost. A higher demand shifter increases the surplus loss from the exit of a high quality good, but also decreases the likelihood of exit. Which effect dominates is an empirical question, and thus the bias from omitting quality heterogeneity in counterfactuals cannot be signed in general.

This paper contributes to the literature on the role of quality in governing the response of trade to shocks. One prominent strand, including Levchenko, Lewis, and Tesar (2011), Chen and Juvenal (2015) and Bems and di Giovanni (2016), has found some evidence that the disproportionate drop in the value of trade after the global negative income shock in 2008 was caused by the higher quality of traded goods combined with non-homotheticity of demand.[5] Another strand has shown that firms may choose to upgrade the quality of their exported products, either because exchange rate shocks make exporting to richer countries more attractive (Bastos, Silva, and Verhoogen (2018), see also Fajgelbaum, Grossman, and Helpman (2011)) or because competing with inexpensive imports drives firms to upgrade, as in Medina (2018).[6]

The key modelling requirement that connects our paper to this literature is that consumer expenditures on higher quality goods must be more elastic with respect to price than expenditures on lower quality goods, which is a consequence of non-homothetic demand.[7] As with Bems and di Giovanni (2016), in our framework, an increase in the price of a good will lead to less income being spent on it overall as consumers substitute to the outside option, and given our parameter estimates this effect will be more severe for higher quality goods, implying a greater sensitivity of their profits to price changes.

A key difficulty in the trade literature on quality has been actually identifying which goods are high quality, and then quantifying what that implies for demand. Khandelwal (2010) suggests using a demand residual, while Medina (2018) and Alessandria and Kaboski (2011) make an assumption based on the description of the goods (e.g., pima cotton versus other fabrics, and fresh

[5] The "flight from quality" phenomenon is well-known in the literature (see Burstein, Eichenbaum, and Rebelo (2005)). Similar mechanisms have also been emphasized by Coibion, Gorodnichenko, and Hong (2015), who find that consumers reallocate expenditure across stores in response to economic conditions, and by Argente and Lee (2017), Cravino and Levchenko (2017), and Faber (2014), who study the distributional consequences of various shocks (output, exchange rate, and trade, respectively) across the income distribution.

[6] Other trade shocks that can drive firms to quality upgrade include rising competition from low-wage countries (as in Martin and Mejean (2014)), cheaper intermediate inputs (see Verhoogen (2008), Fieler, Eslava, and Xu (2014) and Bas and Strauss-Kahn (2015)) or access to larger markets (see Bustos (2011), Lileeva and Trefler (2010), and Aw, Roberts, and Xu (2011)).

[7] In their paper on endogenous quality and the terms of trade, Feenstra and Romalis (2014) also require non-homothetic demand. However, not all papers that use quality to explain trade patterns rely on non-homothetic demand; for instance, Crozet, Head, and Mayer (2011) model quality as a demand shifter in a CES framework and use it to explain exporting patterns of French wine producers.

versus frozen fruit). Our paper bridges these approaches by separating out goods into natural and artificial fabrics using their descriptions, but then also quantifying the effect of natural fabrics in a demand regression.

The structural estimation in this paper is closely related to the industrial organization literature that uses static oligopolistic entry models of incomplete information such as Seim (2006) and Ershov (2018) to estimate entry costs and profit parameters. Our adaptation of this framework amounts to a simplified approach for dealing with combinatoric problems with complementarities, which appear in trade contexts (Antras, Fort, and Tintelnot, 2017), and market entry contexts (Jia, 2008).[8] This paper complements other structural IO papers that evaluate exchange-rate shocks in particular industries such as beer (Goldberg and Hellerstein, 2013) and coffee (Nakamura and Zerom, 2010) but which do not allow for quality downgrading or entry and exit.[9] We also connect to Gopinath, Gourinchas, Hsieh, and Li (2011) and Burstein and Jaimovich (2012) insofar as both papers use the decision-making of a single retailer to answer empirical questions in a trade context—in their cases, pricing to market.

Lastly, there is a vast literature on how exchange rate shocks pass through into prices. Ludema and Yu (2016), using the Melitz and Ottaviano (2008) framework, finds indirect evidence that quality changes mediate price pass-through. Chen and Juvenal (2016) find that pass-through decreases with quality after an exchange rate shock, but do not look at quality adjustments. While this paper does not focus on the price dimension of pass-through—indeed, in our data and in the model pass-through for a given product will be nearly complete, heterogeneity in product entry and exit has been recognized as an important mechanism for both transmitting shocks (Alessandria, Kaboski, and Midrigan, 2010b) and measuring pass-through (Nakamura and Steinsson, 2012).

The paper proceeds as follows. Section 2 provides an overview of the data and institutional background used in the paper. Section 3 presents direct evidence on quality downgrading in the Russian online apparel industry. Section 4 describes the structural model and derives the conditions on parameters under which it will predict quality downgrading. Section 5 provides

[8] Solution methods for these complex problems are an active area of research, see Eckert and Arkolakis (2017).

[9] Feenstra, Gagnon, and Knetter (1996) looks at pass-through for cars, and notes that quality adjustments may affect price pass-through numbers.

details on the estimation, recovered parameters, and counterfactuals. Section 6 concludes.

2 Background and Data

Our data comes from a large, online fashion retailer that sells primarily in Russia.[10] The retailer sells clothing, shoes, and accessories. At the retailer-assigned stock-keeping unit (SKU) level, we observe the price, which is constant across Russia but varies month to month, as well as the quantity sold in each province (oblast) in each month.[11] SKUs are comparable to UPCs in that each one describes a specific product—e.g., a particular variety of Adidas running shoe—aggregating only over different colors and sizes of the same product. The data cover January 2012 through September 2015; from September 2014 to March 2015 the ruble devalued by over 50% after holding roughly steady against the U.S. dollar since the early 2000s.

In addition to prices and quantities of SKUs, we observe unit fabric composition and country of manufacture. The fabric composition is what allows us to cleanly measure quality changes, as it provides a visible measure of product quality that is constant pre- and post-shock.

2.1 Store features

The store operates by ordering SKUs at wholesale prices from both large and small brands and then reselling to Russian consumers with a markup. Most SKUs are uniquely associated by the firm with a season, which corresponds to a combination of Fall or Spring and a year. Before a season begins, the firm chooses which brands and SKUs to include. Once the goods start being offered, the firm is free to choose pricing.[12]

We associate the Spring season with the period from March through August, and Fall with September through February of the following year.[13] Figure 1 shows that the majority of revenue

[10] The company is a subsidiary of a publicly traded German enterprise, listed on the Frankfurt Stock Exchange. As of today, the retailer operates in four countries (Belarus, Kazakhstan, Russia, and Ukraine), although the present study focuses exclusively on the largest market, which is Russia.

[11] We also have a disaggregated, consumer level purchase data set that we do not use.

[12] As far as we are aware from interviews with the management team, the firm is not bound by any resale-price maintenance agreements with the manufacturers.

[13] 78% of Spring SKUs and 75% of Fall SKUs are introduced in our designated Spring and Fall months, respectively.

for a season's SKUs happens during the six month window associated with that season. The only slight discrepancy from this pattern occurs in the Fall 2015 season since we only observe 17 full days in September of 2015 after which our data end.[14]

There are two features of the store worth mentioning. First, for most SKUs the firm does all of its stocking up in one initial wave, before the season starts, at a prearranged unit wholesale cost from existing brands. We thus expect any exchange rate pass-through or quality changes to occur with a lag. Second, the product line is almost completely refreshed each season with new SKUs that are associated with the new season, which gives the firm the scope to reallocate fabrics but prevents us from tracking SKUs over long periods.[15]

2.2 Data cleaning and summary statistics

We have price, quantity, material and origin information for 444,629 SKUs spread over 1,583 brands and 26 product groups. The most commonly occurring fabrics are presented in Appendix A. Cotton, polyester, and leather dominate, with at least one of the three present in 50% of SKUs.

Since our objective is to code each SKU as either high or low quality, we must decide which fabrics are which quality, and how to treat blends. 32% of products are associated with multiple fabrics, and for most of these we do not have reliable percentage composition information. To proceed, we first code polyester, plastic polymers, and any fabric with the word "artificial" as low quality.[16] We assume an SKU with any low quality component is a low quality product, except SKUs containing polyester, in which case we require that polyester is the only component for it to be low quality.

Our split broadly reflects that naturally-derived materials are high quality and artificial ma-

83% of Spring revenue and 78% of Fall revenue are earned in our designated Spring and Fall months, respectively. Additional graphs of the distribution of Fall and Spring introductions and revenue shares are available in Appendix A.

[14]Since a season's SKUs continue to be introduced beyond the first month of the season, the Fall 2015 revenue share appears low for the final bar of Figure A.2 in Appendix A.

[15]This feature of the microdata has also been recently discovered in work studying how firms grow through the introduction of new product lines (see Argente, Lee, and Moreira (2018)).

[16]Our precise mapping from the 30 most commonly occurring fabrics, present in 97% of SKUs and accounting for all materials in 93% of SKUs, into the high/low quality dummy is given in Table A.1 in the appendix.

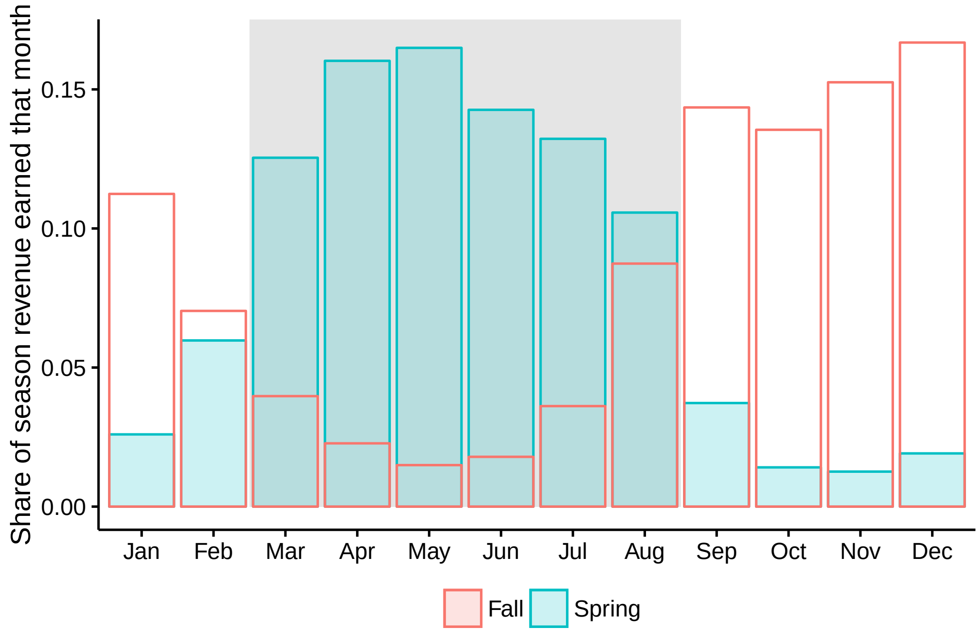

Figure 1: **Monthly revenue shares for SKUs by season**
Note: This figure shows histograms of the distribution of Fall and Spring introductions by revenue. The gray area covers the months we choose to associate with Spring goods of March-August.

terials are low quality, but that blends can also be high quality.[17] We verify in our analysis that our high quality coded products have higher wholesale costs and a positive demand shifter. Our downgrading results also go through if we focus on categories that cannot blend, such as boots, which use either leather or artificial leather. Blend prevalence by fabric type is provided in Table A.1 in Appendix A.

Table 1 presents summary statistics by product group. The Share column gives the total count share of SKUs in that group compared to all SKUs offered over the whole sample period, the Quality column gives the high quality fabric SKU share of each product group, and the Rus. column gives the fraction of Russian manufactured products.[18]

[17] Intuitively blends are often chosen to provide a particular property such as waterproofness, breathability, odor resistance, etc. that are not easily vertically ranked. Elastane, an artificial fabric which can provide stretchiness, is part of a blend in 99.8% of the roughly 51,000 SKUs it is in, and clearly its presence does not indicate low quality.

[18] The Russian apparel industry is made up of numerous manufacturers that tend to be quite labor intensive, with the sector employing around 236,158 workers in medium to large enterprises in 2015 (according to BvD's Amadeus data). For comparison, and according to the U.S. Department of Labor, apparel manufacturers in the United States

Table 1: Cross-sectional summary statistics

Group	Share	Quality	Rus.	Group	Share	Quality	Rus.
Ankle Boots	0.012	0.727	0.091	Outwear	0.060	0.577	0.031
Bags	0.080	0.468	0.060	Sandals	0.019	0.500	0.041
Ballerina Shoes	0.016	0.600	0.039	Scarves	0.022	0.813	0.091
Blazers and Suits	0.011	0.866	0.052	Shirts	0.056	0.769	0.037
Boots	0.039	0.823	0.036	Shoes	0.048	0.787	0.058
Dresses	0.078	0.774	0.117	Shorts	0.018	0.834	0.015
Flip Flops	0.011	0.369	0.068	Skirts	0.020	0.769	0.087
Headwear	0.025	0.894	0.225	Sport Shoes	0.062	0.645	0.014
Heeled Sandals	0.033	0.668	0.057	Sweatshirts	0.032	0.890	0.036
High Boots	0.044	0.775	0.076	Polos	0.114	0.950	0.039
Jeans	0.022	0.988	0.005	Jumpsuits	0.046	0.880	0.051
Knitwear	0.068	0.949	0.039	Underwear	0.016	0.952	0.005
Moccasins	0.018	0.853	0.040	Vests and Tops	0.026	0.793	0.045

Note: This table presents summary statistics by product group. The Share column gives the fraction of SKUs in a group compared to all SKUs offered over the whole sample period, the Quality column lists the high quality fabric SKU share of each product group, and the Rus. column contains the fraction of Russian manufactured products.

Table 2: Time-varying summary statistics

Season	Quality	No. SKUs	Units Sold	Price	Raw Cost	Avg. RUB/USD
2012-03-01	0.816	27,089	339,747	3,874	1,775	31.170
2012-09-01	0.804	33,592	421,807	4,164	1,957	30.840
2013-03-01	0.772	63,584	1,232,188	3,285	1,433	31.947
2013-09-01	0.776	60,638	1,233,759	4,750	1,914	33.225
2014-03-01	0.764	69,945	1,895,759	3,631	1,465	35.324
2014-09-01	0.777	74,885	2,082,531	4,578	1,941	51.704
2015-03-01	0.738	88,122	2,826,627	4,512	1,898	56.898
2015-09-01	0.708	13,100	411,986	4,590	1,983	69.885

Note: This table presents summary statistics at the season level over time. The Season column contains the start date of each respective season, the Quality column lists the fraction of high-quality goods for each season, the number of units sold per season is contained in the fourth column, the average SKU price is in the fifth, the wholesale cost is in the Raw Cost column, and the average U.S. dollar to ruble exchange rate over a season is shown in the last column.

Our panel analysis focuses on the season level SKU stocking choices of the firm, so we aggregate SKUs sales and prices and associate the aggregated values with our assigned time windows. Our baseline results use the first observed price as that SKU's within-season price.[19] Summary statistics at the season level are presented in Table 2. The number of SKUs drops precipitously in the September 2015 season, which reflects the fact that our data ends in September, but SKUs associated with a season continue to be introduced after the first month.[20] Total sales and number of SKUs are on a sharp upward trend, as the firm is expanding during this time period. It is also worth pointing out that the fraction of high-quality products clearly decreases from its previous steady state during the first 2015 season, which is the initial post-devaluation period and is indicative of quality downgrading in the aggregate. While this happens, the unweighted average wholesale cost for this 2015 Spring season rises to 1,898 rubles, far exceeding values of 1,433 and 1,465 rubles for Spring 2013 and Spring 2014, respectively.

2.3 Macroeconomic environment

In 2014, a decline in investor confidence led to a rapid fall in the value of the Russian ruble. Falling confidence in the Russian economy stemmed from two major sources: first, the price of crude oil, a key Russian export, declined by nearly 50% from June 2014 to December 2014; second, the annexation of Crimea in March 2014 precipitated Western asset freezes on Russian energy and banking sectors that were implemented by July 2014.[21] In response, Russia implemented a wide-ranging food import ban against the EU, although no other trade was restricted.

Figure 2 shows how these developments were mirrored in a ruble depreciation of about 60% against the U.S. dollar between July and December 2014. From the vantage point of our firm, which earns revenue in rubles but buys wholesale in foreign currencies, this abrupt movement represents an exogenous cost shock that was fully realized by the time the company was sourcing products for its Spring/Summer 2015 season.[22] Incidentally, the food import ban, oil price shock,

employed about 142,860 workers in 2014.

[19]The results are robust to using a within-season sales-weighted average.

[20]See Figure A.1 in the Appendix A.

[21]See, for example, the New York Times article "Raising Stakes on Russia, U.S. Adds Sanctions" on July 17 of 2014.

[22]As is well-known from the broader exchange rate disconnect puzzle, nominal exchange rates follow a volatile

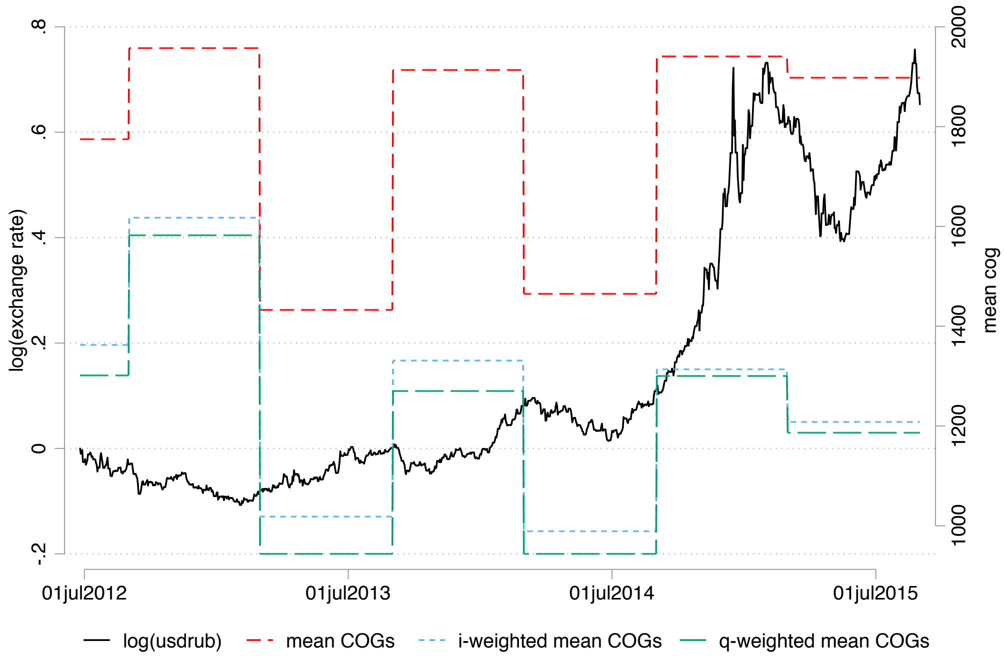

Figure 2: **Cost of goods sold**
Note: This figure shows the normalized log U.S. dollar to ruble exchange rate (black solid line), the mean seasonal (red dashed line), the inventory-weighted mean seasonal (blue short-dashed line), and the purchase quantity-weighted mean seasonal (green long-dashed line) wholesale costs of all SKUs from mid-2012 until 1 Sept 2015.

and financial sanctions on the Russian economy that began in July 2014 may also have represented a substantial income shock to consumers as early as during the Fall 2014 season, which is before any of the quality downgrading is observed.

Besides documenting the exchange rate shock, Figure 2 also provides for an initial look at how the firm responded to the devaluation. A number of patterns are revealed: first, there is a lot of periodicity in the average wholesale cost of goods sold, with Spring/Summer items always being cheaper on average than goods associated with Fall/Winter seasons; second, the steep nominal devaluation at the end of 2014 led to an increase in average wholesale costs during the subsequent Spring 2015 season (mean COGs). Yet costs did not go up nearly as much as one might expect

random walk process that is uncorrelated with macroeconomic fundamentals and is hence largely unpredictable.

under complete pass-through into import prices. Furthermore, inventory-weighted wholesale costs increased even less in percentage terms than unweighted mean costs. This reflects that average stocking quantities per SKU increased in relative terms for cheaper, lower quality goods, which hints at non-homothetic adjustment mechanisms.[23]

3 Reduced Form Evidence

In this section we provide evidence that the firm reacted to the nominal exchange rate shock by reducing the quality of the products it imported for resale. In particular, we verify three empirical facts in the data:

1. Imported goods experienced a greater quality reduction compared to Russian-produced goods, and goods for which quality is more costly to provide experienced the greatest quality reduction. This suggests that the cost shock had an effect on the extensive margin of quality choice beyond any income-induced effect.

2. High-quality goods did not experience differential pass-through, and did experience a differential quantity and expenditure contraction, but only once the firm reallocated offerings.

3. Regions in Russia that experienced greater income shocks did not differentially reallocate consumption to cheaper goods.

3.1 Quality downgrading

We first show that the share of high-quality goods on offer was reduced in response to the exchange rate shock. Our identification strategy is a difference-in-differences (DiD) estimation, where imported SKUs are the treatment group, domestically produced SKUs are the control group,

[23] This pattern is not driven by a large scale removal of high cost goods from the retailer's warehouses (which could be rationalized with consumers moving forward consumption), but rather by a disproportionate amount of stocking-up on low cost goods—the close association between average *quantity*- and *inventory*-weighted wholesale costs confirms this interpretation.

and the fraction of products that are high quality (natural fabric) products is the dependent variable. Intuitively, items manufactured abroad and purchased by the firm in a foreign currency will have a larger increase in ruble costs post-shock than domestically produced items purchased in rubles;[24] if quality adjustment is an important margin for passing through the ruble cost increase, then there will be a significant coefficient for the treatment group dummy post-shock. The DiD framework rules out explanations for quality reallocations that are common across the treatment and control, such as changing tastes or changing commodity/raw fabric costs that are contemporaneous with the devaluation.

In our first specification, we aggregate within seasons to the product group-origin level. For each of the 26 product groups, we will have two observations in each of the eight seasons: the fraction of high quality SKUs for products with a domestic origin, and the high quality fraction for imported products. In order not to impose a timing assumption on when the firm passes through the shock, we run a specification with time-varying treatment effects:

$$natfrac_{grt} = \sum_{t>1} \delta_t \left(nonrus_{gr} \cdot \mathbf{D}_{gt} \right) + \sum_{gr} \alpha_{gr} \mathbf{D}_{gr} + \sum_{gt} \alpha_{gt} \mathbf{D}_{gt} + \epsilon_{grt} \qquad (1)$$

where g indexes a product group (e.g., high boots), r indicates either foreign or domestic manufacturing origin, and t is a season. $natfrac_{grt}$ is the fraction of offered SKUs that use a natural fabric for product group g, source r, in season t, δ_t are the time-varying treatment effects, $nonrus_{gr}$ is a dummy with a value of one for the set of non-Russian products in group g, $\mathbf{D}_{gt}$ are product group-season specific dummies, and $\mathbf{D}_{gr}$ are dummies for each product group-origin combination. Standard errors are clustered at the group$\times$origin level to allow for within-group-origin serial correlation over time. Note that δ_t is omitted for the first season since otherwise the model would be underidentified.

The estimated coefficients δ_t from equation 1 are plotted in Figure 3, along with their associated standard errors. The results indicate the there is no statistically significant differential reduction in quality within product groups for non-Russian goods until the March 2015 season,

[24] We confirm that this is true in pass-through regressions in Section 3.2.

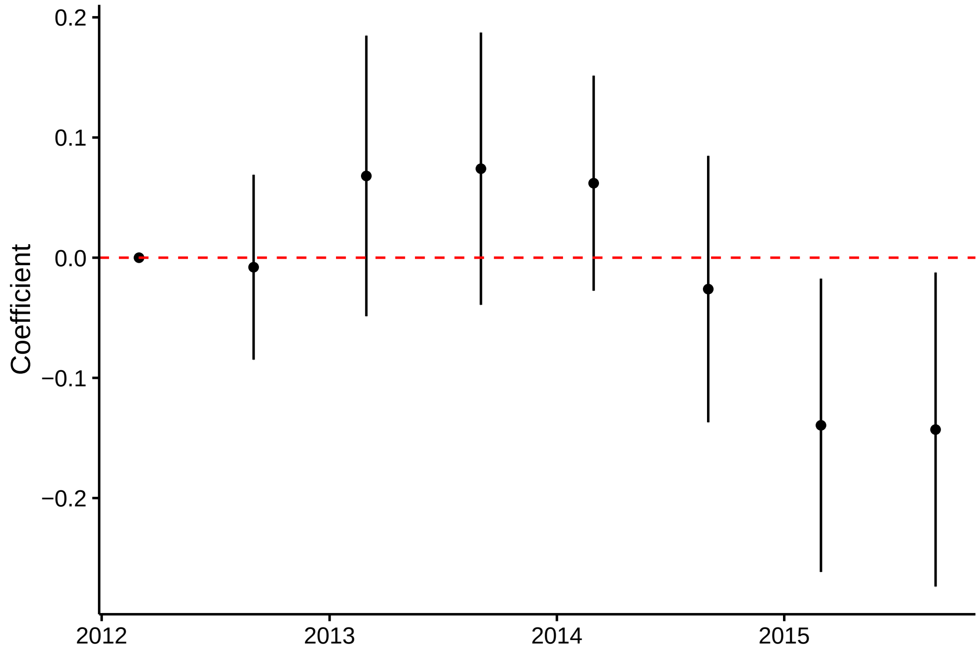

Figure 3: **Quality downgrading**
Note: This figure plots the estimated δ_t coefficients of equation 1 with 95% confidence intervals around them. Fixed effects are at the product group×country of origin and season level. Standard errors are clustered by group×origin to allow within-group-origin serial correlation.

after the peak of the devaluation. That is, this formal timing test of the exchange rate shock suggests that there was a significant effect on the quality of imported products, and that it happened on a timeframe consistent with the firm's one-season-ahead stocking decisions. The lack of a significant treatment effect prior to March 2015 validates the use of domestic products as a control group and provides evidence against a pre-trend as the explanation for the effect.

To quantify the impact of the devaluation on imported products, we next run specifications that allow the magnitude of the exchange rate movement to play a role:

$$natfrac_{grt} = \delta\left(nonrus_{gr} \cdot \log(ER_{t-1})\right) + \sum_{gr} \alpha_{gr}\mathbf{D}_{gr} + \sum_{gt} \alpha_{gt}\mathbf{D}_{gt} + \epsilon_{grt} \tag{2}$$

$\log(ER_{t-1})$ is the average exchange rate during the prior season. The coefficient δ no longer has a t subscript, and can be approximately interpreted as the percent change in $natfrac_{grt}$ that results from a one percent change in the lagged exchange rate.

We run equation 2 for three different levels of aggregation: for no g, so that each season has one observation for the imported high quality fraction and one for the domestic high quality fraction; for g indicating product groups as in equation 1; and for g indicating specific brands within a product group.[25] These specifications are saturated with fixed effects and therefore allow for quality reallocations between product groups, within product groups and between brands, and within brands only for the three regressions, respectively.

The results reported in columns (2) and (3) of Table 3 correspond to the within-product group model, and imply that a one percent devaluation in the prior season leads to a roughly 0.35% reduction in the fraction of offerings that are high quality. In column (1), we recover a negative, significant δ coefficient that is not statistically different from the estimates in (2) and (3), suggesting that reallocation between product groups with different average quality levels is not a key margin for quality downgrading for the firm. In column (4) we estimate a positive, insignificant δ, implying that within-brand reallocations are not an important margin for downgrading.

If the increase in costs from the exchange rate shock is causing quality downgrading, one might expect that for product groups where quality is more expensive to provide, there will be more downgrading. We test this relationship by allowing for the treatment coefficient in equation 2 to vary by product group in our product group level specification:

$$natfrac_{grt} = \sum_{g} \delta_g \left(nonrus_{gr} \cdot \log(ER_{t-1}) \right) + \sum_{gr} \alpha_{gr} \mathbf{D}_{gr} + \sum_{t} \alpha_{gt} \mathbf{D}_{gt} + \epsilon_{grt} \tag{3}$$

For each product group, we recover the quality premium by dividing the average wholesale cost of goods for high versus low quality goods in the seasons prior to March 2015. A value greater than one indicates that high quality goods cost more on average than low quality goods in that product group. For most product groups (20 out of 26), quality is costly.

We plot the estimated coefficients δ_g against the quality premium in Figure 4.[26] The strong negative relationship between the costs of providing quality and the amount of quality down-

[25]For example, Adidias Flip Flops and Adidas Sport Shoes are counted as different brands.

[26]The full regression results from equation 3 are available in Table B.1 in Appendix B.

Table 3: Differential quality downgrading

	Dependent variable:			
	$natfrac_{grt}$			
	(1)	(2)	(3)	(4)
$nonrus_{gr} \times \log(ER_{t-1})$	-0.285^{**}	-0.347^{***}	-0.321^{**}	0.204
	(0.059)	(0.064)	(0.115)	(1.029)
Origin FE	✓			
Season FE	✓	✓		
Group $\times$ Origin FE		✓	✓	
Group $\times$ Season FE			✓	✓
Brand $\times$ Origin FE				✓
Observations	16	395	395	24,820
R^2	0.911	0.692	0.864	0.999

Note: This table presents coefficient estimates from specification 2. The outcome is the fraction of offered SKUs that use a natural fabric for product group g, source r, in season t. $nonrus_{gr}$ is an indicator with a value of one for the set of non-Russian products in group g, and $\log(ER_{t-1})$ is the average exchange rate during season $t-1$. Standard errors (in brackets) are clustered at product group$\times$origin level to allow for serial correlation across time. ***, **, * indicate significance at the 0.1%, 1% and 5% levels, respectively.

grading support the hypothesis that costs played a central role in the firm's decision to quality downgrade after the devaluation.

Quality downgrading robustness

Our identification is based on the assumption that the exchange rate shock does not affect the wholesale cost of Russian-manufactured products as much as foreign-manufactured products. We provide evidence that pass-through from the devaluation into Russian product wholesale costs is lower but still positive in Table 4 in the next section. Since Russian products may use imported intermediates combined with Russian labor this is to be expected, and suggests that our quality downgrading coefficient in Table 3 is a lower bound since the control group experiences a cost shock as well.

One concern is that the treatment effects are driven by quality upgrading in the control group,

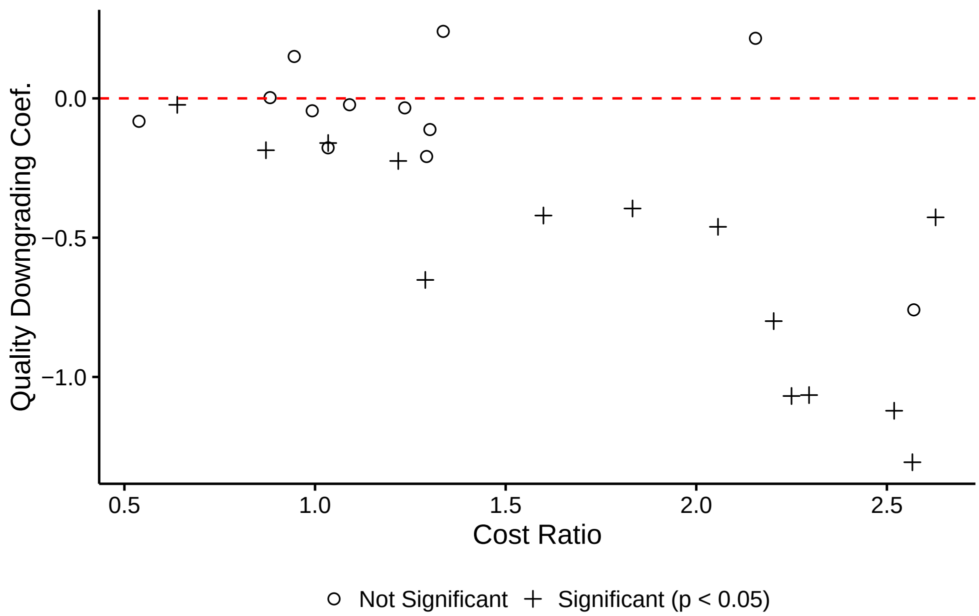

Figure 4: **Cross-group variation in downgrading**
Note: This figure plots the estimated δ_g coefficients of equation 3. Fixed effects are at the product group×country of origin and season level. Standard errors are at the 95% level, clustered by product group×country of origin to allow within-group-origin serial correlation.

rather than downgrading in the treatment group. In Appendix B we provide a raw DiD data graph for polymers (Figure B.1), which mostly appear in a small subset of products using leather such as Boots, High Boots, and Shoes. Looking at this specialized subset helps control for compositional differences in product groups between Russian and non-Russian products. Polymers have a significant presence by end of sample (in 8% of SKUs) and show a clear differential trend, with imports increasing their share while domestic products keep the share roughly constant. This check provides some assurance that our methodology is sound.

A second concern might be that the treatment effects are not driven by quality upgrading, but by the the firm adding one or several large, imported brands around the time the shock hit, that predominantly use artificial fabrics. If the random timing of a large addition of brands were the reason behind the downgrading, however, one would not expect to see the clear relationship between costs of quality and magnitude of downgrading in Figure 4.

3.2 Pass-through and expenditure switching

Having documented quality downgrading in the previous section, in this section we ask why downgrading occurs. If the firm is stocking fewer high-quality goods, then they must have become relatively less profitable; since profit is simply markup multiplied by quantity sold, either high quality markups, quantities, or both must have experienced a relative decline after the shock.

A differential reduction in markups would imply lower pass-through of the shock into high than low quality goods. We run pass-through regressions to determine whether high quality goods experienced a change in relative prices. Since we do not observe most SKUs for longer than one season, our main results are not within SKU; rather, we treat a brand-group-fabric choice as a consistent product over time, while still using SKUs as our unit of observation in the regression. Our specification is:

$$\begin{aligned} \log(y_{jmbgt}) = & \beta_1 \log(ER_{t-1}) + \beta_2 \log(ER_{t-1}) \times Nat_{jmbgt} + \beta_3 \log(ER_{t-1}) \times Rus_{jmbgt} \\ & + \sum_{bgr} \alpha_{bgr} \mathbf{D}_{bgr} + \sum_{mbg} \alpha_{mbg} \mathbf{D}_{mbg} + \epsilon_{jmbgt} \end{aligned} \tag{4}$$

where y_{jmbgt} is either p_{jmbgt}, the first observed price of SKU j of material m for brand b in product group g in season t, or c_{jmbgt}, the constant (within season) wholesale cost of j. ER_{t-1} is the lagged average ruble to U.S. dollar exchange rate, and Nat_{jmbgt} and Rus_{jmbgt} are dummies for whether SKU j has a natural fabric and Russian origin respectively. The specification only uses within material-brand-group variation in prices to identify pass-through.

Results from the regression are presented in Table 4. Pass-through into prices in column (1) is incomplete, as the coefficient on the lagged exchange rate for pass-through into prices is roughly 0.6 and statistically different from 1. However, using the raw data on marginal costs, this imperfect pass-through does not correspond to lowered markups: the pass-through on cost is very similar in column (2).[27] Importantly, the differential change in prices and wholesale costs for high quality goods is not significantly different from zero, implying no differential pass-through for

[27]From discussions with the firm's operations staff, they describe negotiating a "50-50" split of the cost increase (in rubles) with their wholesale suppliers. The coefficient on the lagged exchange rate in column (2) is higher than 0.5, which may reflect that larger brands with more SKUs negotiated higher pass-through into costs.

Table 4: Pass-through coefficients

	Dependent variable:	
	$\log(price)$	$\log(cog)$
	(1)	(2)
$\log(ER_{t-1})$	0.646^{***}	0.626^{***}
	(0.030)	(0.037)
$\log(ER_{t-1}) \times Nat$	0.055	0.010
	(0.029)	(0.035)
$\log(ER_{t-1}) \times Rus$	-0.176^{**}	-0.201^{***}
	(0.058)	(0.049)
Brand $\times$ Origin FE	✓	✓
Brand $\times$ Quality FE	✓	✓
Observations	417,855	393,916
R^2	0.881	0.875

Note: This table presents coefficient estimates from specification 4 at the brand-group-fabric level. The dependent variable is either (1) the first observed price of SKU j or (2) the within season wholesale cost of j. ER_{t-1} is the lagged averaged U.S. dollar to ruble exchange rate, and Nat and Rus are indicators for whether SKU j has a natural fabric and is of Russian origin, respectively. Standard errors (in brackets) are clustered at product group×material level to allow within-group-material serial correlation across time. ***, **, * indicate significance at the 0.1%, 1% and 5% levels, respectively.

these products. While strategic complementarities in price setting can explain some of the price increases among Russian-sourced products following the devaluation, those goods still exhibit significantly lower pass-through than imported items, validating their use in the previous section as a control group that is less exposed to the cost shock.

We address concerns that within material-brand-group selection on low-performing SKUs may be biasing pass-through in Appendix B. We also perform standard within-SKU pass-through regressions on the small set of SKUs we observe for longer than one season in Appendix B.1, and find no evidence of differential pass-through for natural fabric products.

Even with no differential pass-through there may have been differential reduction in demand. With non-homothetic demand, a proportionate price increase can imply a disproportionate reduction in quantity sold of the more expensive, higher quality product. We test whether there was a differential reduction in shares for high quality goods. At the material-group-season level, we run:

$$\log(q_{mgt}) = \sum_t \delta_t \left(Nat_{mg} \cdot \mathbf{D}_t\right) + \sum_{mg} \alpha_{mg}\mathbf{D}_{mg} + \sum_{gt} \alpha_{gt}\mathbf{D}_{gt} + \epsilon_{mgt} \tag{5}$$

where q_{mgt} is the aggregate quantity sold of material m, product group g product purchased in season t. We restrict our sample to imports only. A consumption reallocation away from high quality towards low quality would be reflected in a negative, significant δ_t, starting in March 2015. The results are plotted in Figure 5 and show a relative reduction in the quantity share of high-quality goods right after the steep ruble devaluation. We also estimate the regression using expenditures (price multiplied by quantity sold) as the dependent variable and find very similar results; since we use within product group variation this makes our results comparable to the within group switching in Bems and di Giovanni (2016).

This section highlights that non-homothetic demand plays a key role in the reallocation towards lower quality products, as even with no relative change in prices or markups high quality products disproportionately decrease in quantity purchased. There is also supporting evidence that the quality downgrading was not completely in response to an income shock, since if that were true one might expect some reallocation in Figure 5 towards low quality when the income shock hit in the Fall 2014 season. The fact that significant reallocation only occurred after the firm passed through higher costs into consumer prices suggests that the cost shock played a dominant role in product quality downgrading.

3.3 Demand channel

One might suspect the observed compositional changes stem from a large demand shift towards cheaper or lower quality goods as a result of an income shock to consumers, rather than a cost shock to apparel manufacturers. In this section we assess the quantitative importance of this mechanism by looking at regions that were more adversely affected during the crisis and comparing their demand patterns to regions that had higher economic growth. We find little evidence

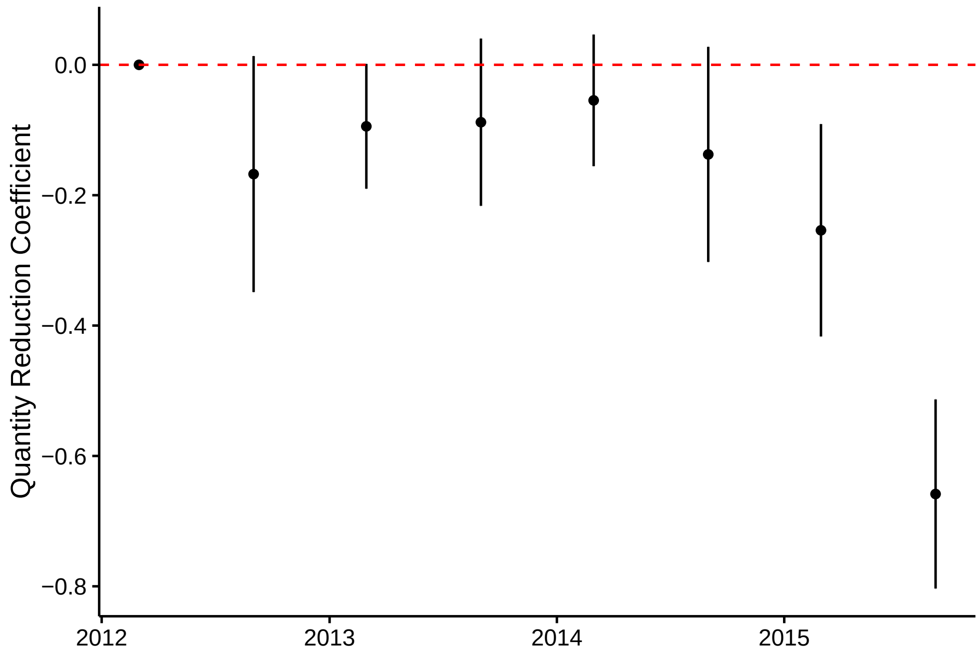

Figure 5: **Differential quantity reduction**
Note: This figure plots the estimated δ_t coefficients of equation 5 with 95% confidence intervals around them. Fixed effects are at the product group×material and product group×season level. Standard errors are clustered at product group×material level to allow within-group-material serial correlation. Results are similar when only using a season, instead of group×season fixed effect.

of differential consumption reallocation towards cheaper goods in Russian regions (oblasts) suffering from extremely low or even negative economic growth in 2015. The basic approach entails a DiD estimation strategy of the following form:

$$\log(Y_{it}) = \alpha_i + \sum_t \gamma_t \mathbf{D}_t + \sum_t \delta_t \left(\mathbf{D}_t \cdot Growth_i\right) + X'_{it}\theta + \sum_t \psi_t \left(\mathbf{D}_t \cdot X_{it}\right) + \epsilon_{it} \qquad (6)$$
$$\forall i,\ \forall t \in \{2012m1, \ldots, 2015m9\} \setminus \{2014m12\}$$

where Y_{it} is either i) the median regular price, or ii) the mean (sales-weighted or unweighted) regular price in region i at time t, α_i are region fixed effects, $Growth_i$ is the nominal regional GDP growth in 2015, $\mathbf{D}_t$ is an indicator for the time period (year-month), with $2014m12$ taken as the omitted category, $(\mathbf{D}_t \cdot Growth_i)$ represents an interaction term between the time indicators and a region's economic performance in 2015, and X_{it} is a matrix of control variables that includes total regional sales (in logs), as well as regional unemployment and income levels. All standard errors are clustered at the region-level to allow for serial correlation across time.

The Russian currency crisis had a vastly differential impact on various regions of the country. This provides for a clean distinction between exposed (low growth) and unexposed (high growth) oblasts that can be utilized when estimating specification 6. Panel (A) of Figure 6 shows a map with geographic regions that grew relatively fast (in dark colors) as well as slowly (in light colors) in 2015. Exclusively devoting attention to oblasts with positive retail sales, the steepest contraction saw regional GDP growth of -10.1% whereas the oblast with the highest growth expanded by 16.1%. The standard deviation of income growth was 3.26 over this period.

As would be necessary with any DiD estimation approach, this specification also provides evidence on the parallel trends assumption in all outcome variables. That is, in the absence of treatment the unobserved disparities between high- and low-growth regions should be constant over time—the validity of the estimation procedure relies on outcome variables that would have continued to develop as they did before the economic shock in all regions. Unless this assumption is valid, the estimated treatment effects would be biased versions of the true impact. As an additional robustness check on the identification strategy, all control variables are interacted with the $\mathbf{D}_t$ indicators to allow for possible heterogeneous responses to negative economic shocks across distinct regions (e.g., poor versus rich oblasts could react differently to the crisis).

The main parameters of interest are the δ_t since they capture the difference between crisis exposed and relatively unscathed regions over time. The estimated fixed-effects model includes leads going back to early 2012 and lags reaching the last available month, September 2015. The specification allows for any causal direction of the findings and assesses if the effects grow or fade over time.

One may also entertain a causal interpretation of the δ_t estimates in equation 6 for other important reasons. Firstly, about 93% of goods sold by the retailer are not produced in Russia, and even when the good is home made it is almost never manufactured in the region under consideration. Hence the specification will not suffer from endogeneity issues typically associated with

(A) Regional growth (2015)

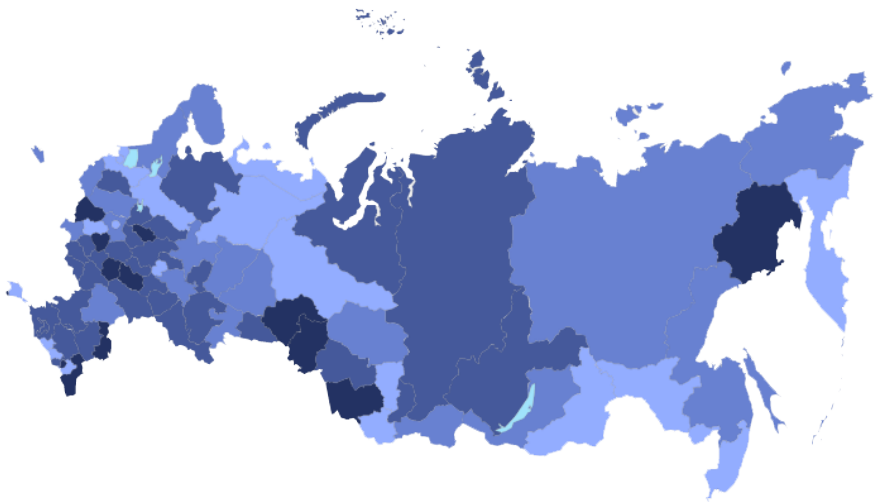

(B) Income effect

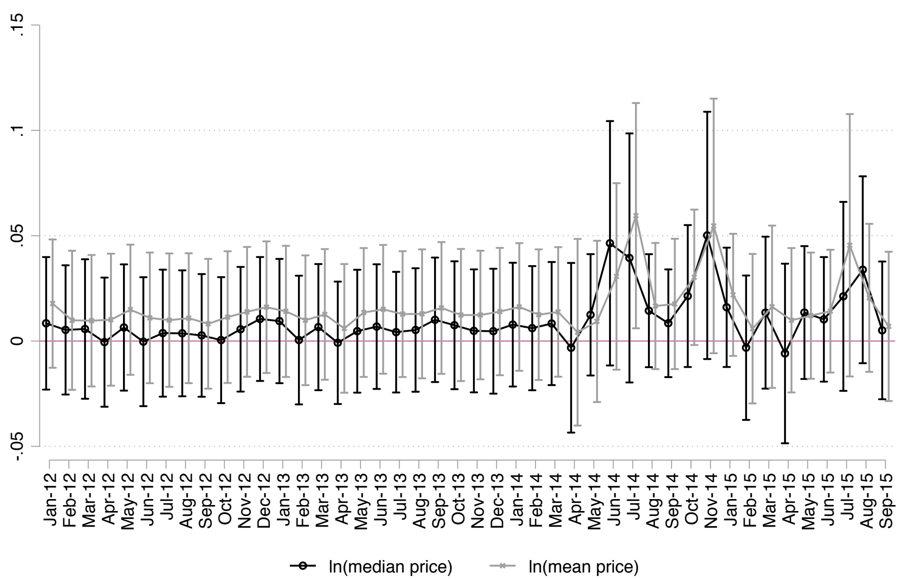

Figure 6: **Demand Channel**
Note: Panel (A) depicts regional GDP growth rates across Russian oblasts in 2015, with darker colors representing higher economic growth; Panel (B) plots the estimated δ_t coefficients of equation 6 with 95% confidence intervals around them. Results for two distinct outcome variables are displayed over time: the log median regional purchase price (black), and the log mean regional purchase price (grey). Time is measured on a monthly basis.

regressions of prices on economic activity. For instance, unobserved productivity innovations for a specific SKU are unlikely to be correlated with local growth rates. In principle, aggregate shocks could lead to simultaneous movements in prices of goods and local economic growth. But since time fixed effects are included, they should eliminate this endogeneity issue too. Finally, the retailer does not price discriminate across geographic regions within Russia and thus any observed divergence in regional (weighted or unweighted) median and mean prices can only be explained by movements in quantities (purchases).

The findings are summarized in Figure 6, which plots the key estimated parameters of interest, $\widehat{\delta}_t$, with 95% confidence intervals around them. As would be consistent with the parallel trends assumption, the estimates in Panel (B) show no robust differences between the positively exposed (high growth) and negatively hit (low growth) regions in the months prior to the onset of Russia's currency crisis. Then, starting around mid-2014, there is increasingly more volatility in the treatment effects for all outcome variables. However, the results are insignificant and hardly moving in the expected positive direction. Together with unreported but also highly robust evidence suggesting no differential effects on total regional sales, this leads us to conclude that income shocks across Russian regions had a marginal role in the observed compositional shifts in the affordable fashion industry and that endogenous amplification channels on the firm-side must be driving most of the quality downgrading.

4 Structural Model

This section develops and estimates a simple structural model of quality choice to understand how quality downgrading mediates the welfare costs of a devaluation. We are interested in whether consumer welfare would have reacted differently to the exchange rate movements if quality downgrading was not possible, if there was no quality dimension to heterogeneous products, or if the value consumers attached to quality increased or decreased. Since we do not observe this variation in the data, we estimate a simple structural model of consumer demand and product sourcing to assess counterfactuals.

4.1 Setup

Each season, purchase managers for each possible individual SKU decide whether to include that SKU in next season's offerings. The manager can decide whether or not she wants the SKU to be a high quality or low quality fabric. The managers take the optimal sourcing strategies of the other purchase managers into account, but otherwise act independently. "SKUs" and "managers" will be used interchangeably.

During each season, each consumer decides what SKU to purchase, if any. SKUs are horizontally differentiated by consumer specific idiosyncratic shocks, and vertically differentiated by quality.

Demand

There are M_t consumers indexed by i, who choose among product offerings during season t and an outside option. In equilibrium, they face N_h high quality products and N_ℓ low quality products, each of which is differentiated with a consumer-product specific idiosyncratic demand shock. i's utility from consuming product j of quality m at time t is given by:

$$U_{ijmt} = q_m + \alpha p_{jmt} + \epsilon_{ijmt},$$

where q_m is the vertical quality shifter and ϵ_{ijmt} is the idiosyncratic portion of utility.[28] We normalize the utility from the outside good to 0 so $U_{i0t} = \epsilon_{i0t}$, and require that ϵ_{ijmt} takes the logit form. With a slight abuse of notation on N_m, the market share of product j of quality m is:

$$s_{jmt}(p_{jmt}, \boldsymbol{p}_{-jt}, N_h, N_\ell) = \frac{\exp(q_m + \alpha p_{jmt})}{1 + \sum_{j' \in N_h} \exp(q_h + \alpha p_{j't}) + \sum_{j' \in N_\ell} \exp(q_\ell + \alpha P_{j't})} \tag{7}$$

We denote the set of available products at time t by $\mathcal{J}_t$.

[28]Although we could recover product-specific qualities as a demand residual as in Khandelwal (2010), we instead follow Medina (2018) and our reduced form in treating quality as a 0-1 dummy correspond to material. In their analysis of the 2008 income shock Levchenko, Lewis, and Tesar (2011) find more evidence of a quality response when using explicit, 0-1 measures of quality instead of demand residuals.

Competition and entry equilibrium

The purchase manager for SKU j first makes an entry and quality decision at time $t-1$, then chooses pricing depending on the competitive environment at time t after entry decisions have been realized. We solve managers' optimal strategies backwards, first taking as given the competitive environment and solving prices, then solving the optimal entry.

Conditional on the set of competitors, managers strategically set prices to maximize profits in a Nash-Berltrand equilibrium:

$$p^*_{jmt} = \arg\max_{p_{jmt}} M_t \cdot s_{jmt}(p_{jmt}, \boldsymbol{p}_{-jt}, N_h, N_\ell) \cdot (p_{jmt} - c_m \cdot ER_{t-1})$$

An SKU j's base marginal cost c_m is in units of foreign currency and is converted to rubles through ER_{t-1}. From the reduced form section, the firm negotiates prices chooses stocks one season in advance, so the effect of the shock will be lagged due to inventory considerations as in Alessandria, Kaboski, and Midrigan (2010a). We choose a symmetric equilibrium in the pricing game where any j with quality m has the same optimal price p^*_{mt}.

At time $t-1$, j must decide whether to enter and if so, what quality to provide. The profit to j of providing quality m is:

$$\pi_{jmt} = \beta \cdot \pi^v_m(\boldsymbol{a}_{-jt}, ER_{t-1}, M_t, \bar{N}) - f_m - \sigma_\epsilon \epsilon_{jmt}$$

$\boldsymbol{a}_{-jt}$ denotes the equilibrium entry and quality strategies of all potential entrants, of which there are $\bar{N}$, which together determine the total number of SKUs of each type that j will compete against at time t. Note that while most subscripts are kept as t to denote that payoff and pricing is realized at time t, entry decisions are made and fixed costs incurred at time $t-1$, so that variable profit is discounted by β. The scale of variable profits are fixed in rubles, so we allow the scale of the variance of ϵ_{jmt} to adjust.

ϵ_{jmt} is an idiosyncratic information shock that is only observed by j. Managers know the distribution G_ϵ and form beliefs about other managers' behavior. In particular, firm k expects that j will choose quality m with probability P_{jmt}, and will choose not to enter with probability

P_{j0t}. Firm j's expected profits from choosing material m are then:

$$\pi^e_{jmt}(\boldsymbol{P}_{-jt}) - \sigma_\epsilon \epsilon_{jmt},$$

where the expectation is taken over all the possible distributions of offered product qualities given strategies $\boldsymbol{P}_{-jt}$. Since ϵ_{-jt} is not observed by j, this is an incomplete information game of entry and quality choice similar to Seim (2006), Augereau, Greenstein, and Rysman (2006) and Ershov (2018).

Assuming that ϵ_{jmt} takes the EV Type 1 distribution, j's probability of choosing quality m is:

$$P_{jmt} = \frac{\exp(\pi^e_{jmt}(\boldsymbol{P}_{-jt})/\sigma_\epsilon)}{1 + \sum_{m'} \exp(\pi^e_{jm't}(\boldsymbol{P}_{-jt})/\sigma_\epsilon)} \tag{8}$$

A Bayesian Nash Equilibrium (BNE) at each time t is a vector of choice probabilities $\boldsymbol{P}_t$ that solves equation 8 so that equilibrium actions are consistent with equilibrium beliefs.

Welfare

Consumer welfare in the model takes the standard logit form. We multiply by market size and divide through by the price coefficient to express welfare in rubles:

$$W_t = M_t \frac{1}{|\alpha|} \log\left(1 + \sum_{j \in \mathcal{J}_t} \exp(\alpha p_{jt} + q_j)\right)$$

4.2 Model predictions

To provide intuition on what to expect in the estimation section, here we derive the model's predictions for how a firm's choice of products changes in response to a nominal exchange rate devaluation.

Theorem 1: Taking the optimal strategies $\boldsymbol{P}_{-j}$ of other firms as given and dropping time subscripts for convenience, the elasticity of firm j's entry probability into quality m with respect to the nominal exchange rate, ER, is:

$$\mathcal{E}_{m,ER} \equiv \frac{\partial P_m}{\partial ER}\frac{ER}{P_m} = -ERc_m \cdot s_m(2 - s_m) \cdot (1 - P_m) < 0,$$

and $\partial(P_h/P_\ell)/\partial ER < 0$ under some parameter conditions, and, in particular, if c_h/c_ℓ is sufficiently large.

Proof: See Appendix C.

The model implies that conditional on offering a product, managers reallocate from high to low quality products.[29] Intuitively, in partial equilibrium the increase in the exchange rate affects firm pricing, and hence potential markups and quantities sold. In our case the markup $p - c$ moves very little for both high and low quality—the large number of SKUs implies the additive logit markup is already close to its lower bound of $1/\alpha$—so the main margin for differential profit reduction is the differential demand response for high and low quality, as in the reduced form.

However, different demand elasticities are not enough to induce differential exit of high quality SKUs: the consumer expenditure shares for high and low quality must be able to adjust. If the share of income a consumer spent on a quality segment remained constant regardless of prices, then any differential exit from that segment would make it less competitive, increasing the attractiveness of entry into that segment—possibly enough to offset the initial differential exit.[30]

5 Estimation and Results

We estimate the model using the subset of product (or "target") groups for which quality is costly to provide in the sense of Figure 4.[31] This section estimates the parameters as a function of data in three steps: first, demand parameters are estimated; second, the demand system is inverted to

[29]IIA in the demand system will imply that any change in the value of the consumer's outside good does not affect the relative shares of h and ℓ goods, so including a time-varying outside option for consumers would not affect the results.

[30]For CES subutility with a Cobb-Douglas aggregator across quality segments, and a higher demand elasticity for high quality, the effects exactly cancel, and there is no differential exit.

[31]The six exluded product groups are Jeans, Sweatshirts, Tee-shirts and Polos, Trousers and Jumpsuits, Underwear, and Vests and Tops.

recover marginal costs; third, the entry and exit model uses demand and cost parameter estimates combined with equilibrium firm strategies to back out fixed costs and the variance of the profit shock.

5.1 Method

5.1.1 Demand model

The model provides an analytic representation of the share of a particular product in equation 7. Taking the log difference between the season sales share of any given product sold that season and the share of the outside option yields:

$$\ln(s_{jmt}) - \ln(s_{0t}) = q_{jm} + \alpha p_{jt} \tag{9}$$

Our data reports quantities, which we transform into shares by making an assumption on the market size. Unique to our online data, in each season we observe the total number of units individuals considered buying but did not—i.e., their shopping carts—which we take as the market size.[32] The relationship between market size and total quantity ordered is provided in Figure C.1 in Appendix C.

In practice, to estimate equation 9 requires the addition of an error term. If the error term is a demand shock observed by the firm, then the OLS coefficient α in equation 9 will be positively biased. We experiment with different estimation strategies and use monthly price and quantity variation to recover $\hat{\alpha}$ independently of quality shifters; details are provided in Appendix C. We then difference out $\hat{\alpha}$ and estimate:

$$\ln(s_{jmt}) - \ln(s_{0t}) - \hat{\alpha} p_{jt} = \beta_0^q + \beta_1^q \mathbf{1}[m(j) = h] + \nu_j^q \tag{10}$$

These coefficients translate into the structural parameters as $q_\ell = \hat{\beta}_0^q$ and $q_h = \hat{\beta}_0^q + \hat{\beta}_1^q$.

[32] This is one way to determine market size in e-commerce industries, and it is especially useful for the largest retailers—as our firm—that are well-known to most of their potential customers. One underlying interpretation is that consumers resort to other stores to obtain the remainder of their initial shopping cart selection.

5.1.2 Costs

We use observed prices and the assumption of Bertrand-Nash competitive price setting to back out baseline marginal costs. In particular, profit maximization implies that:

$$c_{jt} = p_{jt} - \frac{1}{\alpha(1-s_{jt})}$$

We use $\hat{\alpha}$ and observed season-level prices and shares to recover $\hat{c}_{jt}$. To recover the baseline marginal cost we assume $c_{jt} = c_m ER_{t-1}^{\beta_2^c}$, which delivers the estimating equation:

$$\log(c_{jt}) = \beta_0^c + \beta_1^c \mathbf{1}[m(j)=h] + \beta_2^c \log(ER_{t-1}) + \nu_{jt}^c,$$

ER_{t-1} is the mean exchange rate of rubles for U.S. dollars in Table 2 lagged one season and normalized by the long run average.[33] Normalization implies that $c_h = \exp(\hat{\beta}_0^c + \hat{\beta}_1^c)$ and $c_\ell = \exp(\hat{\beta}_0^c)$ are estimated in rubles.

5.1.3 Entry model

The only parameters remaining are the fixed costs of stocking a high cost and low cost good, f_h and f_ℓ. However, to give the model more degrees of freedom to match how products are added and dropped in response to exchange rate fluctuations, we introduce a scaling parameter ϕ that multiplies c_h. That is, $\tilde{c}_h = \phi c_h$. We also introduce a fixed cost $f_{\ell,w}$ for low cost goods during the winter to account for time-of-year fluctuations in the data.

The entry model is thus parametrized by $\theta^s \equiv \{\phi, f_h, f_\ell, f_{\ell,w}\}$. For estimation we maximize the log likelihood function:

$$\mathcal{L}(W, \theta_s) = \sum_t \sum_m \sum_j \log\left(P_{jmt}(\theta^s)\right) \tag{11}$$

To construct entry probabilities as a function of parameters, we first non-parametrically estimate

[33] To normalize the exchange rate, we divide by the expected value of the AR(1) run on season-level data from 2000-2014.

the probabilities as a function of data only as in Medina (2018). We assume $\bar{N}$ is 1.2 times the maximum number of observed SKUs in a season to convert raw entry numbers into probabilities of entry. Using those probabilities as estimates of managers' equilibrium beliefs, we then solve for managers' expected profits and optimal strategies as a function of parameters.[34] This estimation strategy bypasses the difficulties created by multiple equilibria—which is an issue in our entry game with multiple qualities—as long as we assume only one equilibrium is played in the data, which is standard (Hotz and Miller, 1993).

5.1.4 Identification

Identifying the parameters in the demand and cost regressions is straightforward. For the entry model, the fixed costs will be identified by the average probability of entry for each quality of good and the average profitability of each quality. For instance, if the profit of high quality goods is larger on average but the probability of entry is lower, the model will rationalize this feature with a higher fixed cost for high quality goods. Assuming a higher number of potential entrants will lead to a lower probability of entry for each type of product, but will not change the proportions or profit, which will simply lead to higher fixed costs.

Identification of ϕ will depend on whether the baseline c_h and q_h in the data can match the reallocation towards low quality in the March 2015 season of the data. If relatively fewer high quality goods enter after periods of low ruble valuations in the data, then ϕ will increase only if the baseline cost bump $c_h - c_\ell$ is not sufficient to induce the reallocation.

5.2 Results

Results from each stage of the estimation are gathered and presented in Table 5.

The price parameter α implies that average $p/c \approx 3$, which is higher than the median of first-period price divided by wholesale cost of 2.4.[35] Overestimating margins will lead to over-

[34] We simplify the computation of managers' expected entry profits slightly by ignoring Jensen's inequality; see Appendix C.2 for details.

[35] The elasticity may be underestimated due to standard price endogeneity, or because we do not fully capture dynamic demand effects with our months-since-entry dummies.

Table 5: Structural parameter estimates

Type	Parameter	Estimate	SE
Demand	α	-0.32	0.01
	β_0^q	-10.06	
	β_1^q	0.12	0.02
Marginal Cost	β_0^c	6.80	
	β_1^c	0.04	0.01
	β_2^c	0.70	0.01
Entry and Exit	f_h	1.03	0.01
	f_ℓ	1.78	0.01
	$f_{\ell,w}$	0.11	0.01
	σ_ϵ	0.65	0.004
	ϕ	13.94	0.21

Note: This table presents estimation results.

estimation of fixed cost to rationalize lower participation, but does not especially increase the profitability of high versus low quality goods.

The demand shifter for high quality goods $\hat{\beta}_1^q$ is positive, as is the cost shifter $\hat{\beta}_1^c$, giving us that quality is valuable to consumers and expensive for the firm to provide. All else equal, natural fabric goods are expected to sell 12.4% more than artificial fabrics goods, while high cost goods cost 82% more.[36] Pass-through from the lagged exchange rate into marginal costs is 0.70, which is similar to the coefficient recovered from the reduced form regression in Table 4.

The fixed costs are estimated in hundreds of thousands of rubles. At the pre-2014 long run stationary average of 30.75 rubles per USD, this implies sourcing costs of \$3,400 and \$5,800 for high and low cost goods respectively, which rationalizes the fact that low cost goods sell well (given their lower price) but are not as prevalent in the data. The base fixed costs are small and fairly similar, reflecting the fact that in this context the fixed cost is incurred to source the product from a wholesaler—not develop the product from scratch.[37]

The model does well in matching the entry exit data: the correlation between the predicted

[36] The cost shifter is multiplied by the scaling parameter, $\exp(\phi \cdot \beta_1^c) = 82\%$.

[37] In Medina (2018), fixed costs are estimated much larger for both high and low quality fabrics, as the firm's sales are much larger and costs presumably include the cost of development, not just sourcing.

probabilities of entering as a high quality firm and the data is 0.93, and the corresponding correlation for low quality firms is 0.95. The correlation between the ratio of these predicted probabilities and the ratio of the probabilities in the data is 0.75. A full plot of model predictions versus data is provided in Figure C.2 in Appendix C.

5.3 Counterfactuals

The model allows us to answer the question of how welfare would change if firms could not downgrade quality in response to a devaluation, a scenario we do not see in the data. This counterfactual is applicable where there are technological constraints on downgrading, such that only high quality materials are sufficient—for instance, with extreme cold weather gear—or when there are regulations that mandate inputs must be a certain quality for particular products.

We evaluate the change in welfare that would result if the fixed cost of sourcing a low-quality good was prohibitive, so that managers choose between a high-quality good and not entering. Practically, we first assume that the fixed cost of sourcing a low quality good increases by a factor of 10, then simulate the equilibrium probability of entry and the resulting prices pre and post cost shock, and finally compute the ratio of consumer welfare pre (W) and post (W').[38]

Our counterfactual predictions are evaluated using the depreciation that took place at the end of 2014 and which affected the Spring 2015 offerings that were being sourced at that time, holding market size fixed. The normalized exchange rate rose from 1.15 to 1.67, and maps to a 33% increase in marginal costs using the pass-through coefficient of 0.7. To solve the equilibrium entry probabilities we use a nested fixed point approach as in Seim (2006). For the model with no downgrading, high quality goods are the only ones available and a unique entry equilibrium is guaranteed; for the base model counterfactuals, to find optimal entry probabilities we try a range of starting values centered around the empirical probabilities of entry for the Fall/Winter 2014 period and find no evidence of multiple equilibria.

The results are presented in Table 6. In the base model, there is a 4.2 percentage point decrease

[38] This exercise is similar to that in Medina (2018) where she prohibits quality upgrading by increasing the fixed costs of sourcing.

Table 6: Counterfactuals

Model	ΔP_h	ΔP_ℓ	ΔP_{entry}	W'/W
Prices Only				0.869
Base	-0.042	0.006	-0.037	0.835
No downgrading	-0.051		-0.051	0.780

Note: This table presents results from counterfactual simulations. ΔP_h and ΔP_ℓ are the probabilities of entering as a high- and low-quality product, respectively. W'/W gives the welfare change in each case.

in the probability of entering as a high quality product, and a 0.6 percentage point increase in the probability of entering as a low quality product. The baseline probabilities of entry pre-shock are 47.6% and 26.0% respectively, so the loss of high quality products is substantial. The entry of low quality products comes through general equilibrium effects: despite the increase in costs, the substantial reduction in high quality products makes it slightly more profitable to enter as a low quality product on balance. We expect that for larger devaluations, the unconditional probability of low quality entry would also decrease; however, conditional on entry, the probability of entering as a low quality product would still increase.

Welfare changes computed using the base model show that faced with the devaluation experienced in September 2014, consumer surplus decreases in the following season by roughly 16.5%. The model that does not allow quality downgrading would predict a 22.0% decline, a 5.5 percentage point (33%) difference compared to the base case, and a model that prevents firm exit would predict only a 13.1% decline in welfare. Adding an entry/exit margin increases the welfare loss, but offering firms the flexibility to quality downgrade instead of exiting dampens the welfare cost to consumers.

Quality's role in the welfare costs of a devaluation

We are interested in whether eliminating—or increasing—the demand shifter for high quality goods will change the welfare costs of a devaluation. Eliminating the shifter corresponds to a more standard trade model, where costs are the only dimension of product heterogeneity, while

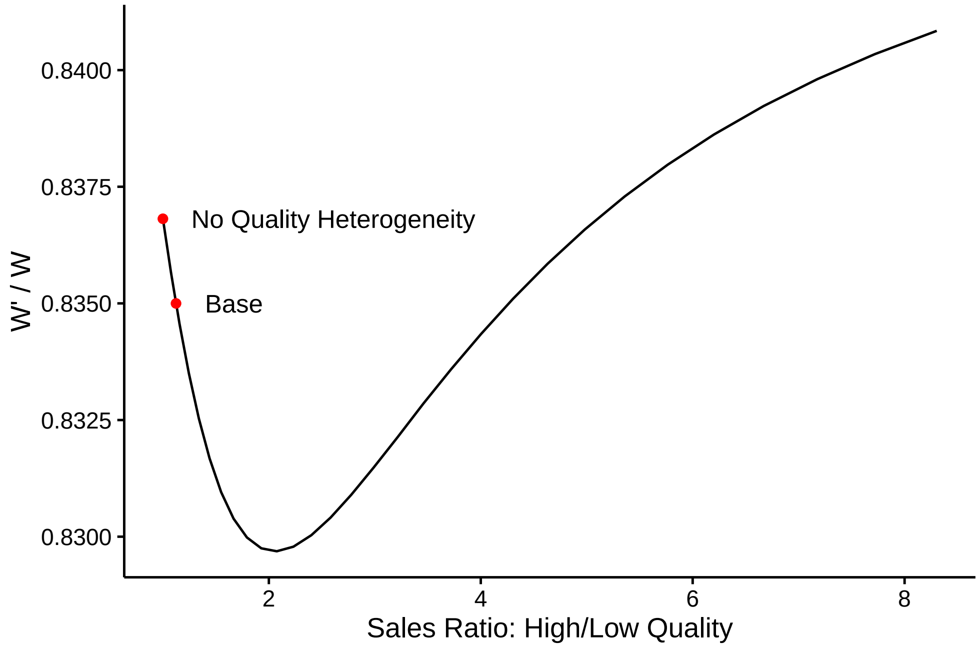

Figure 7: **Changing quality and welfare costs**
Note: This figure plots the welfare cost of the devaluation for different values of the quality demand shifter, holding all other estimated parameters fixed. An x-axis value of one corresponds to no demand increase for high cost goods, i.e., a model with only cost heterogeneity.

increasing the shifter provides insights on industries for which quality is indeed more valued. We simulate equilibrium entry and pricing pre and post cost shock for different values of the quality demand shifter, holding other parameters fixed, and using the same depreciation as for Table 6. We then compute the ratio of consumer welfare pre (W) and post (W') shock for each value of the demand shifter and plot the results in Figure 7.

A model with no quality heterogeneity (highlighted in Figure 7 will underpredict the true welfare costs of the devaluation (as reported in Table 6 and highlighted in Figure 7). For our estimated parameters the error is slight; the baseline model with its relatively small demand shifter only predicts a 0.2 percentage point greater reduction in welfare compared to the model with no quality heterogeneity. For a demand shifter equal in magnitude to the cost shifter the welfare reduction would be 0.7pp greater.[39]

[39] Using a demand shifter of $\hat{\beta}_1^c \times \hat{\phi}$. For a high quality/low quality cost ratio of 2.7, the maximum in Table 4, we

Interestingly, the effect of increasing the demand shifter from 0 (where the sales ratio of high/low quality is 1, all else equal) on the welfare cost of the devaluation is not monotonic. The U-shape is the result of two countervailing forces: as the benefit of quality increases, it becomes less likely that a product will be dropped in response to the devaluation because quality will have a buffering effect on profits; however, for those goods that are dropped, the welfare cost to consumers of losing those products is increased. For our parameters, as the shifter increases from 0 quality products continue to be dropped at a fast rate in response to the devaluation, and the increased quality of the goods being dropped makes consumers worse off overall. Eventually, the buffering effect of quality takes over and the decrease in the drop rate counterbalances the increased welfare loss from dropping.

In general, the counterfactual suggests that a model with only cost heterogeneity may either overstate or understate the welfare loss from a devaluation, depending on the relative strength of the two effects of quality at the estimated parameters. Signing the bias from omitting quality during devaluations or tariff shocks may therefore not be possible ex ante.

6 Conclusion

We use a novel and unique online retail dataset that spans Russia's enormous currency depreciation in late 2014 to dissect how firms respond to cost shocks and to study the most relevant margins of adjustment. We document that changes to product quality figure prominently in the micro-transmission following exchange rate shocks. The data shows that there is a reallocation towards relatively low quality goods in response to the ruble devaluation and that an increase in firm costs, not a reduction in income, is the primary driver of this quality downgrading. Our paper complements a long literature on incomplete exchange rate pass through by showing direct evidence of another margin of adjustment for firms, and introduces an endogenous firm reallocation margin to the literature on non-homothetic demand and expenditure switching. Using a simple structural model of multiproduct sourcing, the paper shows how allowing goods to be

plot the welfare loss as a function of the quality shifter in Appendix C and show it can be up to 1.5pp larger.

heterogeneous in both quality and cost, and letting firms quality downgrade, offers more nuanced predictions of the welfare effect of a devaluation.

While our study looks at the short run effects of the exchange rate shock on quality, the long run effects may also be substantial. For instance, reductions in quality may deplete firms' relationship capital with consumers, leading to larger long-run demand elasticities and less reallocation; conversely, consumers' tastes may adapt to the suddenly more-prevalent low quality goods, implying yet more future reallocation. We leave those questions regarding the long-run consequences of adjusting quality in response to cost shocks for future research.

References

Alchian, A. A., and W. R. Allen (1964): *University Economics*. Wadsworth.

Alessandria, G., and J. Kaboski (2011): "Pricing-to-Market and the Failure of Absolute PPP," *American Economic Journal: Macroeconomics*, 3, 91–217.

Alessandria, G., J. P. Kaboski, and V. Midrigan (2010a): "The Great Trade Collapse of 2008–09: An Inventory Adjustment?," *IMF Economic Review*, 58(2), 254–294.

——— (2010b): "Inventories, Lumpy Trade, and Large Devaluations," *American Economic Review*, 100(5), 2304–39.

Amiti, M., O. Itskhoki, and J. Konings (2016): "International Shocks and Domestic Prices: How Large Are Strategic Complementarities?," *Working Paper*.

Antras, P., T. C. Fort, and F. Tintelnot (2017): "The Margins of Global Sourcing: Theory and Evidence From US firms," *American Economic Review*, 107(9), 2514–64.

Argente, D., and M. Lee (2017): "Cost of Living Inequality during the Great Recession," *Working Paper*.

Argente, D., M. Lee, and S. Moreira (2018): "How Do Firms Grow? The Life Cycle of Products Matters," *Working Paper*.

Auer, R., A. Burstein, and S. M. Lein (2017): "Price and Consumption Responses to Large Exchange Rate Shocks: Evidence from the 2015 Appreciation in Switzerland," *Working Paper*.

Augereau, A., S. Greenstein, and M. Rysman (2006): "Coordination Versus Differentiation in a Standards War: 56k Modems," *RAND Journal of Economics*, 37(4), 887–909.

Aw, B. Y., M. J. Roberts, and D. Y. Xu (2011): "R&D Investment, Exporting, and Productivity Dynamics," *American Economic Review*, 101(4), 1312–44.

Bas, M., and V. Strauss-Kahn (2015): "Input-Trade Liberalization, Export Prices and Quality Upgrading," *Journal of International Economics*, 95(2), 250–262.

Bastos, P., J. Silva, and E. Verhoogen (2018): "Export Destinations and Input Prices," *American Economic Review*, 108(2), 353–92.

Bems, R., and J. di Giovanni (2016): "Income-Induced Expenditure Switching," *American Economic Review*, 106(12), 3898–3931.

Burstein, A., M. Eichenbaum, and S. Rebelo (2005): "Large Devaluations and the Real Exchange Rate," *Journal of Political Economy*, 113(4), 742–784.

Burstein, A., and N. Jaimovich (2012): "Understanding Movements in Aggregate and Product-Level Real Exchange Rates," *Working Paper.*

Bustos, P. (2011): "Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms," *American Economic Review*, 101(1), 304–40.

Chen, N., and L. Juvenal (2015): "Quality and the Great Trade Collapse," *Working Paper.*

——— (2016): "Quality, Trade, and Exchange Rate Pass-Through," *Journal of International Economics*, 100, 61–80.

Coibion, O., Y. Gorodnichenko, and G. H. Hong (2015): "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications," *American Economic Review*, 105(3), 993–1029.

Corsetti, G., M. Crowley, L. Han, and H. Song (2018): "Markets and Markups: A New Empirical Framework and Evidence on Exporters from China," *Working Paper.*

Cravino, J., and A. A. Levchenko (2017): "The Distributional Consequences of Large Devaluations," *American Economic Review*, 107(11), 3477–3509.

Crozet, M., K. Head, and T. Mayer (2011): "Quality Sorting and Trade: Firm-Level Evidence for French Wine," *Review of Economic Studies*, 79(2), 609–644.

Eckert, F., and C. Arkolakis (2017): "Combinatorial Discrete Choice," *Working Paper.*

Ershov, D. (2018): "The Effects of Consumer Search Costs on Entry and Quality in the Mobile App Market," *Working Paper.*

Faber, B. (2014): "Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices," *Working Paper.*

Fajgelbaum, P., G. M. Grossman, and E. Helpman (2011): "Income Distribution, Product Quality and International Trade," *Journal of Political Economy*, 119(4), 721–765.

Feenstra, R. C. (1988): "Quality Change Under Trade Restraints in Japanese Autos," *Quarterly Journal of Economics*, 103(1), 131–146.

Feenstra, R. C., J. E. Gagnon, and M. M. Knetter (1996): "Market Share and Exchange Rate Pass-Through in World Automobile Trade," *Journal of International Economics*, 40(1), 187–207.

Feenstra, R. C., and J. Romalis (2014): "International Prices and Endogenous Quality," *Quarterly Journal of Economics*, 129(2), 477–527.

Fieler, A. C., M. Eslava, and D. Xu (2014): "Trade, Skills, and Quality Upgrading: A Theory with Evidence From Colombia," Discussion paper, National Bureau of Economic Research.

Goldberg, L. S., and J. M. Campa (2010): "The Sensitivity of the CPI to Exchange Rates: Distribution Margins, Imported Inputs, and Trade Exposure," *Review of Economics and Statistics*, 92(2), 392–407.

Goldberg, P. K., and R. Hellerstein (2013): "A Structural Approach to Identifying the Sources of Local Currency Price Stability," *Review of Economic Studies*, 80(1), 175–210.

Gopinath, G., P.-O. Gourinchas, C.-T. Hsieh, and N. Li (2011): "International Prices, Costs, and Markup Differences," *American Economic Review*, 101(6), 2450–2486.

Gopinath, G., and O. Itskhoki (2010a): "Frequency of Price Adjustment and Pass-Through," *Quarterly Journal of Economics*, 125(2), 675–727.

——— (2010b): "In Search of Real Rigidities," *NBER Macroeconomics Annual*, 25(1), 261–310.

Hotz, V. J., and R. A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60(3), 497–529.

Hummels, D., and A. Skiba (2004): "Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture," *Journal of Political Economy*, 112(6), 1384–1402.

Jia, P. (2008): "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry," *Econometrica*, 76(6), 1263–1316.

Khandelwal, A. (2010): "The Long and Short (of) Quality Ladders," *Review of Economic Studies*, 77(4), 1450–1476.

Levchenko, A. A., L. T. Lewis, and L. L. Tesar (2011): "The "Collapse in Quality" Hypothesis," *American Economic Review: Papers & Proceedings*, 101(3), 293–297.

Lileeva, A., and D. Trefler (2010): "Improved Access to Foreign Markets Raises Plant-Level Productivity for Some Plants," *Quarterly Journal of Economics*, 125(3), 1051–1099.

Ludema, R. D., and Z. Yu (2016): "Tariff Pass-through, Firm Heterogeneity and Product Quality," *Journal of International Economics*, 103, 234–249.

Martin, J., and I. Mejean (2014): "Low-Wage Country Competition and the Quality Content of High-Wage Country Exports," *Journal of International Economics*, 93(1), 140–152.

Medina, P. (2018): "Import Competition, Quality Upgrading and Expoting: Evidence from the Peruvian Apparel Industry," *American Economic Review (forthcoming).*

Melitz, M. J., and G. I. Ottaviano (2008): "Market Size, Trade, and Productivity," *Review of Economic Studies*, 75(1), 295–316.

Nakamura, E., and J. Steinsson (2012): "Lost in Transit: Product Replacement Bias and Pricing to Market," *American Economic Review*, 102(7), 3277–3316.

Nakamura, E., and D. Zerom (2010): "Accounting for Incomplete Pass-Through," *Review of Economic Studies*, 77(3), 1192–1230.

Seim, K. (2006): "An Empirical Model of Firm Entry With Endogenous Product-Type Choices," *RAND Journal of Economics*, 37(3), 619–640.

Verhoogen, E. A. (2008): "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector," *Quarterly Journal of Economics*, 123(2), 489–530.

Online Appendix: Not For Publication

A Data

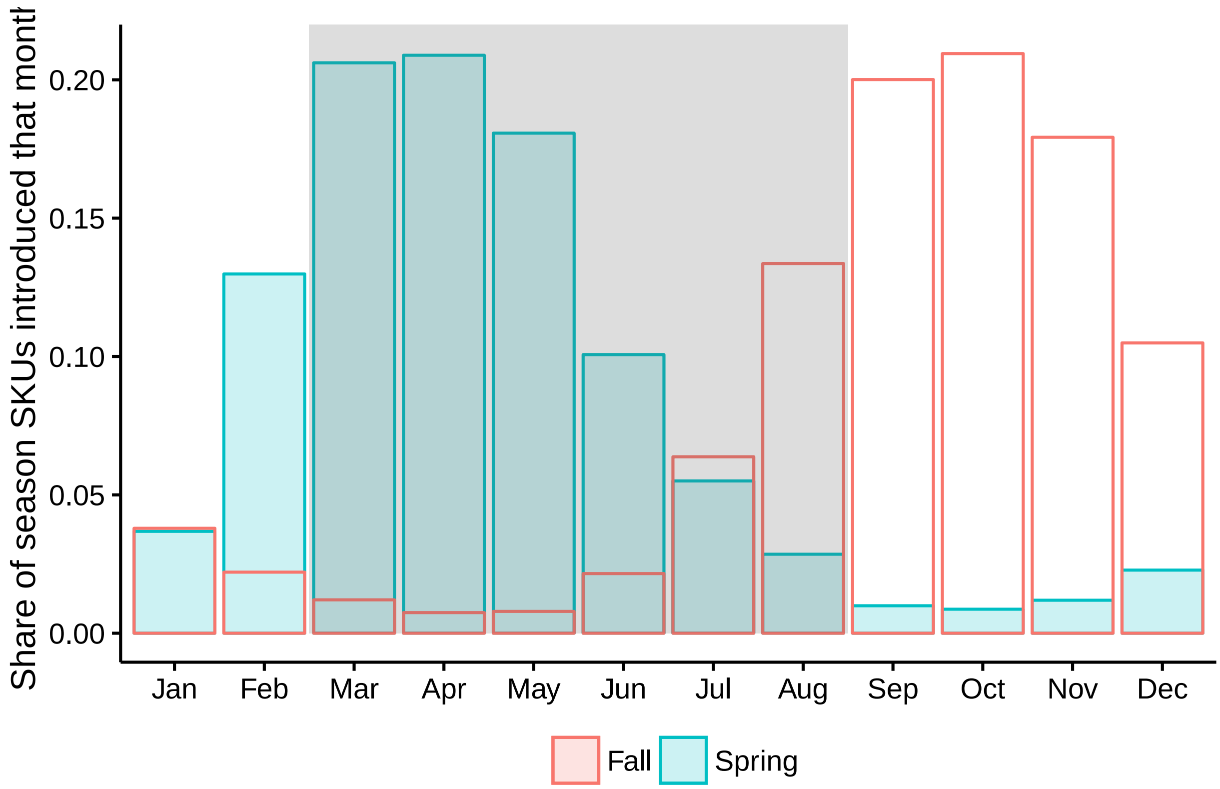

Figure A.1: **Month of first appearance for new SKUs by season**
Note: This figure shows histograms of the distribution of Fall and Spring introductions by month. The gray area covers the months we choose to associate with Spring goods of March-August.

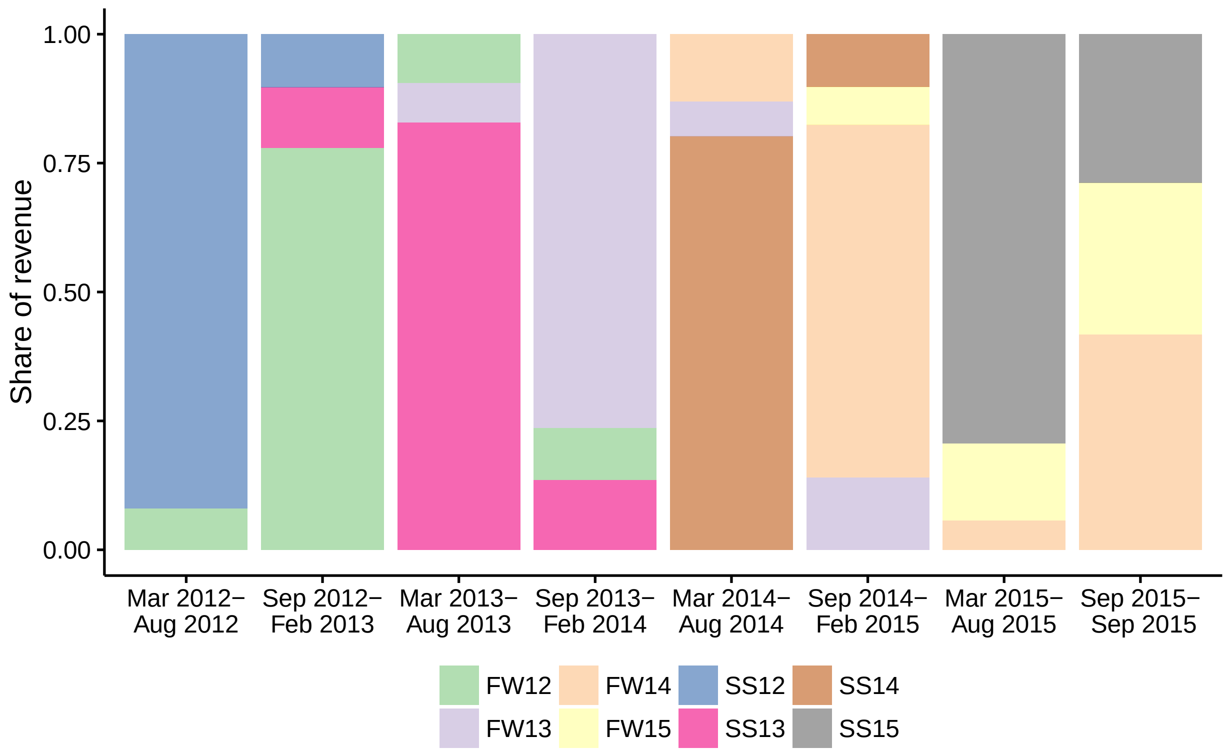

Figure A.2: **Overlapping generations of goods**
Note: This figure plots the revenue shares (between 0 and 1) for each generation of goods over subsequent Fall and Spring seasons.

Table A.1: Material quality mapping

Material	High Quality	Num. SKUs	Blend Fraction
Cotton	1	140,665	0.508
Polyester	0	104,400	0.653
Leather	1	71,173	0.057
Elastane	1	51,757	0.999
Viscose	1	42,806	0.774
Nylon	1	31,613	0.814
Artificial Leather	0	28,637	0.062
Polymer	0	27,614	0.323
Textile	1	17,618	0.334
Acrylic	0	17,480	0.657
Wool	1	17,411	0.842
Suede	1	10,344	0.028
Spandex	1	8,089	1
Nubuck	1	4,776	0.004
Velour	1	4,046	0.0002
Silk	1	4,024	0.450
Artificial	0	3,256	0.233
Lycra	1	2,751	0.998
Linen	1	2,745	0.765
Rubber	1	2,729	0.715
Angora	1	2,111	0.998
Modal	1	1,924	0.866
Artificial Suede	0	1,900	0.001
Cashmere	1	1,678	0.931
Split	1	1,511	0.001
Artificial Nubuck	0	933	0.002
District	1	852	0.826
Mohair	1	767	0.982
Acetate	0	676	0.934

Note: This table presents the mapping from the 30 most commonly occurring fabrics, 97% of SKUs and accounting for all materials in 93% of SKUs, into the high/low quality dummy.

B Reduced Form Evidence

Table B.1: Heterogeneous downgrading coefficients

Group	Cost Ratio	Coef.	SE	p-val
Ankle Boots	2.571	−1.404	0.152	0
Bags	2.155	0.409	0.204	0.045
Ballerina Shoes	2.296	−1.065	0.430	0.013
Blazers And Suits	1.235	0.153	0.076	0.044
Boots	2.057	−0.383	0.171	0.025
Dresses	1.218	−0.258	0.063	0.00004
Flip Flops	1.833	−0.395	0.084	0.00000
Headwear	1.090	0.139	0.276	0.614
Heeled Sandals	2.250	−1.068	0.209	0.00000
High Boots	2.567	−1.114	0.309	0.0003
Jeans	0.639	−0.056	0.024	0.018
Knitwear	1.034	−0.120	0.057	0.036
Moccasins	2.628	−0.427	0.073	0
Outwear	1.293	−0.625	0.224	0.005
Sandals	2.203	−0.800	0.317	0.012
Scarves	1.599	−0.659	1.090	0.546
Shirts	1.301	−0.145	0.117	0.212
Shoes	2.519	−1.038	0.264	0.0001
Shorts	1.336	0.241	0.225	0.285
Skirts	1.034	−0.116	0.194	0.551
Sport Shoes	1.289	−0.609	0.413	0.140
Sweatshirts	0.993	−0.019	0.068	0.778
Tee−Shirts And Polos	0.945	0.537	0.066	0
Trousers And Jumpsuits	0.871	−0.130	0.054	0.017
Underwear	0.538	−0.051	0.050	0.302
Vests And Tops	0.882	−0.150	0.072	0.036

Note: This table presents estimated quality downgrading coefficients across various product categories along with their levels of statistical significance.

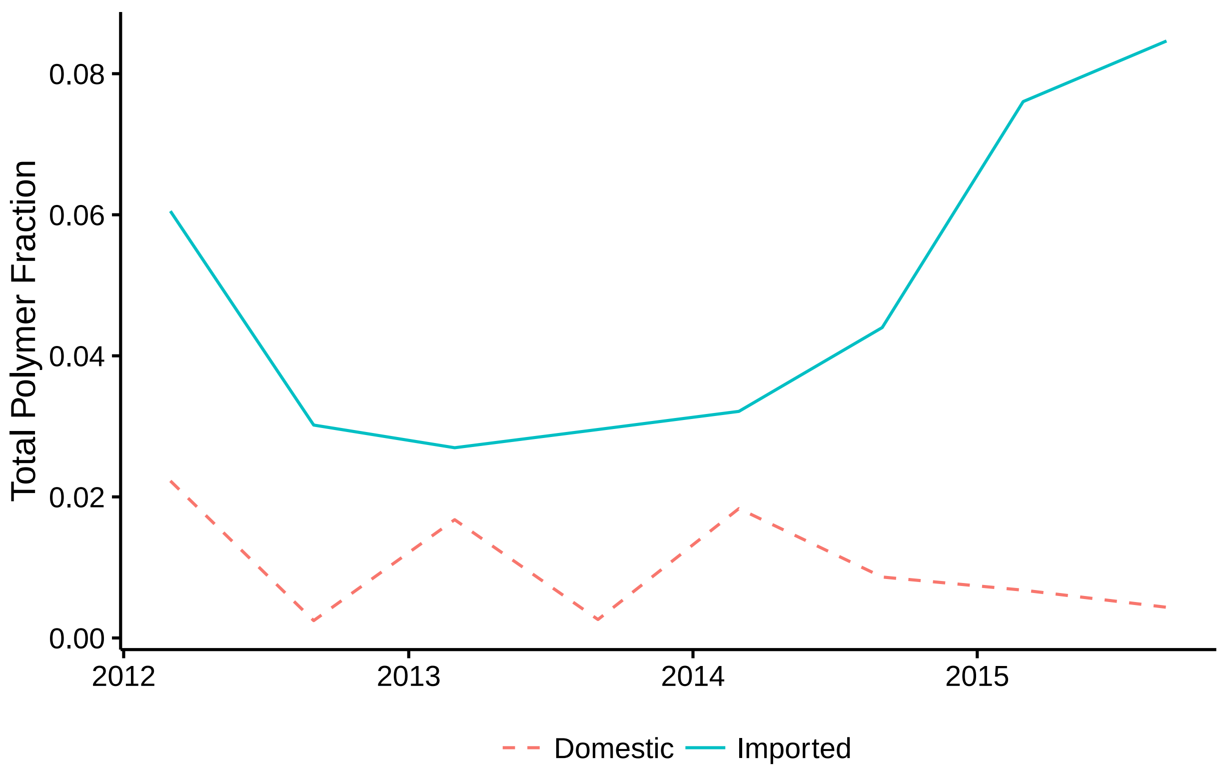

Figure B.1: **Polymer presence by manufacturing origin**
Note: This figure shows the fraction of SKUs where "polymer" is listed as a component over time by domestic (red dashed line) and imported (blue solid line) goods.

B.1 Price pass-through and quantity switching

Differential pass-through dispersion

A concern with the main price pass-through regressions is that since we are not measuring price changes within SKUs, but within material-brand-groups, there may be differential selection of products after the exchange rate shock in a way that biases the results. For instance, if there are different types of high quality products for a particular brand, and if some of them reduce markups more in response to the devaluation, it stands to reason that those high quality goods would drop out by more as they become less profitable. Our regression would thus find more pass-through for high quality goods than there should be.

We evaluate the role within-brand-material SKU heterogeneity plays by checking the second moments of the price and wholesale cost distributions for high and low quality goods. Suppose demand is such that a brand's least expensive high quality goods have more scope for incomplete pass-through compared to its other high quality goods; if the markup contraction makes these goods unprofitable to stock after the cost shock, then the coefficient of variation for a brand's high quality goods' prices (σ_p/μ_p) should decrease, as lower priced SKUs from the bottom of the brand's price distribution of high quality SKUs drop out. The coefficient of variation for a brand's high quality goods' prices would also decrease if it is a brand's most expensive high quality goods that have more scope for incomplete pass-through. If the coefficient of variation for a brand's high quality goods prices does not decrease after the cost shock, then even if there is heterogeneity in pass-through within-brand-material it will not bias the average pass-through regressions through selection.

We run the following specification at the fabric-brand-season level to check for differential reductions in price and cost dispersion of a brand's high quality SKUs:

$$
\begin{aligned}
\mathrm{CV}^x_{mbgt} =& \beta_1 \log(ER_{t-1}) + \beta_2 \log(ER_{t-1}) \times Nat_{mbgt} + \log(ER_{t-1}) \times Rus_{mbgt} \\
&+ \sum_{bgr} \alpha_{bgr} \mathbf{D}_{bgr} + \sum_{mbg} \alpha_{mbg} \mathbf{D}_{mbg} + \epsilon_{mbgt},
\end{aligned} \tag{12}
$$

where $\beta_2 \neq 0$ would indicate a differential effect of the exchange rate on the coefficient of varia-

Table B.2: No change in within-brand-fabric price dispersion

	Dependent variable:	
	CV(p)	CV(cog)
	(1)	(2)
$\log(ER_{t-1})$	-0.006	-0.006
	(0.012)	(0.013)
$\log(ER_{t-1}) \times Nat$	-0.016	-0.012
	(0.014)	(0.015)
$\log(ER_{t-1}) \times Rus$	-0.010^{**}	-0.012^{***}
	(0.004)	(0.003)
Brand $\times$ Origin FE	✓	✓
Brand $\times$ Quality FE	✓	✓
Observations	21,533	21,429
R^2	0.815	0.772

Note: This table presents coefficient estimates from specification 12 at the fabric-brand-season level. The dependent variable is either (1) the within brand-fabric coefficient of variation of prices or (2) the same but for wholesale costs. ER_{t-1} is the lagged averaged U.S. dollar to ruble exchange rate, and Nat and Rus are indicators for whether SKU j has a natural fabric and is of Russian origin, respectively. Standard errors (in brackets) are clustered at the Brand$\times$Origin and Brand$\times$Quality-level to allow for serial correlation across time. ***, **, * indicate significance at the 0.1%, 1% and 5% levels, respectively.

tion of either the prices or wholesale costs for fabric type m for brand b in season s, and $\beta_1 \neq 0$ indicates a baseline effect of the exchange rate on dispersion, and can be estimated when the fixed effects do not control for season. Results in Table B.2 show no significance for β_2, implying that the dispersion in prices and costs did not change differentially for high quality goods. Moreover, β_1 itself is not significantly different from zero, suggesting no effect of the cost shock on the baseline within-brand pricing dispersion. These results suggests that differential dropping of low margin, high quality goods in response to the cost shock is not biasing our pass-through results.

Micro-dynamics of price adjustments

Conditioning on price adjustments, the next section shows that within-SKU pass-through is very high for imported goods. Even though the number of products that live across seasons is small relative to the overall volume, one can use those observations to ask if natural items experienced any differential exchange rate pass-through.

At the SKU-level, we estimate pass-through into prices of exchange rate shocks realized during the most recent period of price non-adjustment and of those that were realized prior to the previous price adjustment. As discussed in the literature (Gopinath and Itskhoki (2010a)), in the absence of real rigidities, all adjustment should take place at the first instance of price change and hence the coefficient on the exchange rate change prior to the previous price adjustment should be zero. More precisely, the following regression is estimated:

$$\Delta \overline{p}_{i,t} = \beta_1 \Delta_{\tau_1} e_t + \beta_2 \Delta_{\tau_2} e_{t-\tau_1} + \eta_i + \epsilon_{i,t} \tag{13}$$

where i indexes the SKU, t stands for the date, the outcome variable, $\Delta \overline{p}_{i,t}$, is the change in the log ruble price of a good, *conditional on price adjustment*, and $\Delta_{\tau_1} e_t \equiv e_t - e_{t-\tau_1}$ is the the cumulative change in the log of the nominal exchange rate over the duration when the previous price was in effect (denoted as τ_1). Analogously, τ_2 denotes the duration of the previous price of the firm so that $\Delta_{\tau_2} e_{t-\tau_1} \equiv e_{t-\tau_1} - e_{t-\tau_1-\tau_2}$ is the cumulative exchange rate change over the previous period of non-adjustment, i.e., the period prior to the previous price change. Solely within-SKU variation is exploited via the inclusion of good-specific fixed effects, η_i, and standard errors are clustered at the SKU-level to allow for serial correlation across time.

Table B.3 reports the results from estimations of regression 13. The number of SKUs is much smaller than in previous regressions due to the fact that there are very few goods that live across seasons. Still, the findings in columns (1) and (3) show that pass-through high after the cost shock. Compared to the Euro, the estimated coefficients are larger and more significant for the U.S. dollar to ruble exchange rate. This is because most trade is invoiced in U.S. dollars rather than in Euros. Columns (2) and (4) present very similar results, but allowing for exchange rate pass-through to differ across natural versus non-natural SKUs, which means that the model is augmented with interaction terms between the exchange rate change and the natural dummy. None of the multiplicative terms are statistically distinguishable from zero, suggesting yet again that pass-through does not vary across high quality and low quality goods.

Table B.3: Within-SKU pass-through

	Dependent variable: $\Delta \log(p_{i,t})$			
	(1)	(2)	(3)	(4)
Δ_{τ_1} usdrub$_{i,t}$	0.993***	0.921**		
	[0.279]	[0.409]		
Δ_{τ_2} usdrub$_{i,t-\tau_1}$	0.649***	0.553		
	[0.203]	[0.410]		
Δ_{τ_1} usdrub$_{i,t} \times Nat$		0.894		
		[0.975]		
Δ_{τ_2} usdrub$_{i,t-\tau_1} \times Nat$		-0.410		
		[0.923]		
Δ_{τ_1} eurrub$_{i,t}$			0.500*	0.383
			[0.270]	[0.383]
Δ_{τ_2} eurrub$_{i,t-\tau_1}$			0.461**	0.190
			[0.217]	[0.437]
Δ_{τ_1} eurrub$_{i,t} \times Nat$				0.948
				[0.766]
Δ_{τ_2} eurrub$_{i,t-\tau_1} \times Nat$				-0.272
				[0.935]
SKU FE	✓	✓	✓	✓
Observations	1,391	1,055	1,391	1,055
No. SKUs	1,126	839	1,126	839
R^2	0.028	0.035	0.009	0.023

Note: This table presents pass-through coefficient estimates at the first and second rounds of price adjustment, respectively, estimated from regression 13. The outcome variable is the change in the log ruble price of a good, conditional on price adjustment. All specifications include SKU fixed effects and standard errors [in brackets] are clustered at the SKU-level to allow for serial correlation across time. The estimation results are based on daily observations between Jan 1, 2014 and April 1, 2015. ***, **, * indicate significance at the 1%, 5% and 10% levels, respectively.

C Structural Model

We drop j subscripts with the understanding that the strategies and prices of opponent firms $-j$ are being held constant. Denote the exchange rate by γ and variable profit π^e_{mt}, and recall that $P_m = \exp(\pi_m - f_m)/(1 + \exp(\pi_\ell - f_\ell) + \exp(\pi_h) - f_h))$

$$\begin{aligned}
\frac{\partial P_m}{\partial \gamma} &= \frac{\partial \pi_m}{\partial \gamma} P_m(1 - P_m) \\
\frac{\partial \pi_m}{\partial \gamma} &= \frac{\partial s_m}{\partial \gamma}(p_m - \gamma c_m) - s_m c_m \\
\frac{\partial s_m}{\partial \gamma} &= \alpha \frac{\partial p_m}{\partial \gamma} s_m(1 - s_m)
\end{aligned}$$

The optimally set price p_m solves $\partial \pi_m/\partial p_m = 0$, which implies $p^*_m = \gamma c_m - 1/\alpha(1 - s_m)$. Taking implicit derivatives with respect to γ gives $\frac{\partial p^*_m}{\partial \gamma} = c_m(1 - s_m)$. Recursively substituting the expressions into each preceding line yields the expression for $\partial P_m/\partial \gamma$:

$$\frac{\partial P_m}{\partial \gamma} = -c_m \cdot s_m(2 - s_m) \cdot P_m(1 - P_m) \tag{14}$$

from which the elasticity $\mathcal{E}_{m\gamma} \equiv \frac{\partial P_m}{\partial \gamma}\frac{\gamma}{P_m}$ follows simply. It is straightforward to show that if $\mathcal{E}_{h\gamma} < \mathcal{E}_{\ell\gamma}$, then $\partial(P_h/P_\ell)/\partial\gamma < 0$. We are interested in how changing (c_h, q_h) relative to (c_ℓ, q_ℓ) affects the likelihood and severity of quality downgrading. We consider when the ratio of elasticities will be less than one:

$$\frac{\mathcal{E}_{h\gamma}}{\mathcal{E}_{\ell\gamma}} < 1 \Leftrightarrow \frac{c_h}{c_\ell}\frac{s_h(2 - s_h)}{s_\ell(2 - s_\ell)}\frac{1 - P_h}{1 - P_\ell} > 1 \tag{15}$$

Since $s_m \in (0, 1)$, as long as $s_h > s_\ell$ it will be true that $s_h(2 - s_h) > s_\ell(2 - s_\ell)$. Using the logit structure of demand,

$$\frac{s_h}{s_\ell} = \frac{\exp(q_h + \alpha p^*_h)}{\exp(q_\ell + \alpha p^*_\ell)}. \tag{16}$$

Which is increasing in q_h. Although there is no closed form solution for p^*_m, $\partial p^*_m/\partial c_m > 0$.

Finally, since profitability of the high type good is increasing in q_h and decreasing in c_h, P_h exhibits the same behavior. We therefore have that increasing q_h both makes it easier to satisfy Equation 15 (by increasing s_h) and more difficult (by increasing P_h). Increasing c_h has similar dual effects.

Through the degree of freedom afforded by the fixed costs f_m that also affect P_m, we can easily find conditions under which the ratio will decrease in γ. For instance, suppose $c_h > c_\ell$ and $q_h > q_\ell$, such that $s_h > s_\ell$, and $f_h > f_\ell$ so that $P_h = P_\ell$; then clearly Equation 15 will be satisfied.

C.1 Demand Estimation

In the model, prices are static within a season. However, as discussed in the data section, we observe price and consumption variation within a season across months, and indeed this is the primary source of our price variation for a product since products only last one season. We therefore run a demand regression at the monthly level with product (j) and month (τ) fixed effects to recover the price coefficient α:

$$\ln(s_{j\tau}) = \alpha p_{j\tau} + \kappa_j + \kappa_\tau + h(j,\tau) + \xi_{j\tau}$$

Note that we do not need to include the outside share as it is time-varying only, and therefore incorporated into κ_τ.

Since demand is dynamic, prices are lowered over time but demand does not necessarily increase—purchasing a product late in the season for which it is intended (e.g., buying winter boots in March) decreases utility from the purchase. The function $h(j,\tau)$ outputs how many months it has been since a product j's first introduction; each number of months since introduction is allowed to have a different intercept. We do not instrument for price for two reasons. First, unobserved product-specific characteristics and dynamic demand are the main sources of unobserved heterogeneity, and both are controlled for. Second, there is no good candidate for an instrument: the exchange rate only affects the initial stock up of product and not month-to-month prices, while the wholesale cost is not time-varying.

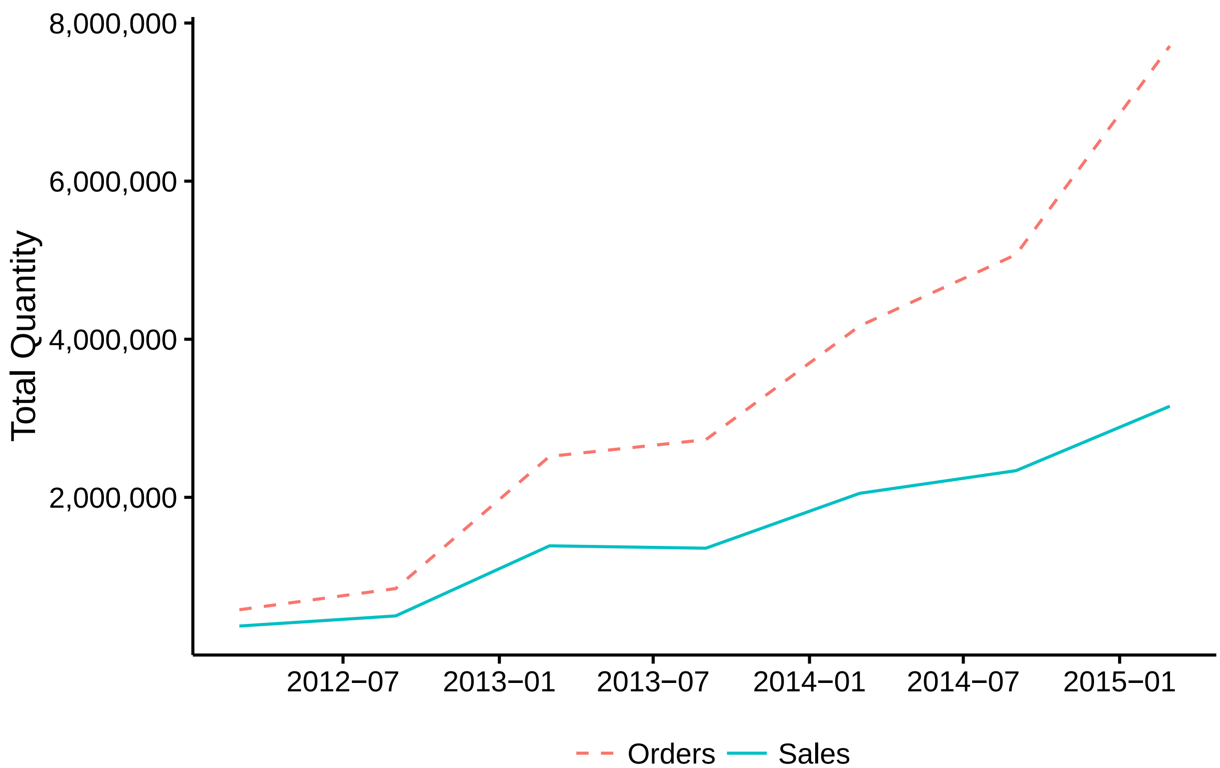

Figure C.1: **Orders and sales**
Note: This figure shows the total quantity ordered by consumers (red dashed line) as well as the total quantity actually sold to consumers (blue solid line) over time.

C.2 Expected Profit Approximation

Formally, $\pi^{e}_{jmt}(\hat{\boldsymbol{P}}_{-jt}, \theta_s) = \mathbb{E}\left[\pi^{v}_{m}(\boldsymbol{a}_{-jt}, \cdot)\right] - f_m$, where the expectation is over the multinomial distribution:

$$\mathbb{E}\left[\pi^{v}_{m}(\boldsymbol{a}_{-jt}, \cdot)\right] = \sum_{N_\ell, N_h | N_{\ell t} + N_{ht} \leq \tilde{N}_t} \frac{\tilde{N}_t!}{N_{\ell t}! N_{ht}! (\tilde{N}_t - N_{ht} - N_{\ell t})!} \times$$
$$P_{\ell t}^{N_{\ell t}} P_{ht}^{N_{ht}} (1 - P_{ellt} - P_{ht})^{\tilde{N}_t - N_{ht} - N_{\ell t}} \cdot \pi^{v}_{m}(N_{ht}, N_{\ell t}, \cdot)$$

Since N_{ht} and $N_{\ell t}$ are typically quite large, we approximate the expectation of the profit with the profit of the expectations as in Ershov (2018). This implies

$$\mathbb{E}[\pi^{v}_{m}(\boldsymbol{a}_{-jt}, \cdot)] \approx \pi^{v}_{m}(\tilde{N}_t \hat{P}_{ht}, \tilde{N}_t \hat{P}_{\ell t}, \cdot),$$

which is straightforward to calculate. Simulations using the multivariate normal approximation to the multinomial and integration using sparse quadrature suggest the error from violating Jensen's inequality is not substantial.

C.3 Model Fit and Counterfactuals

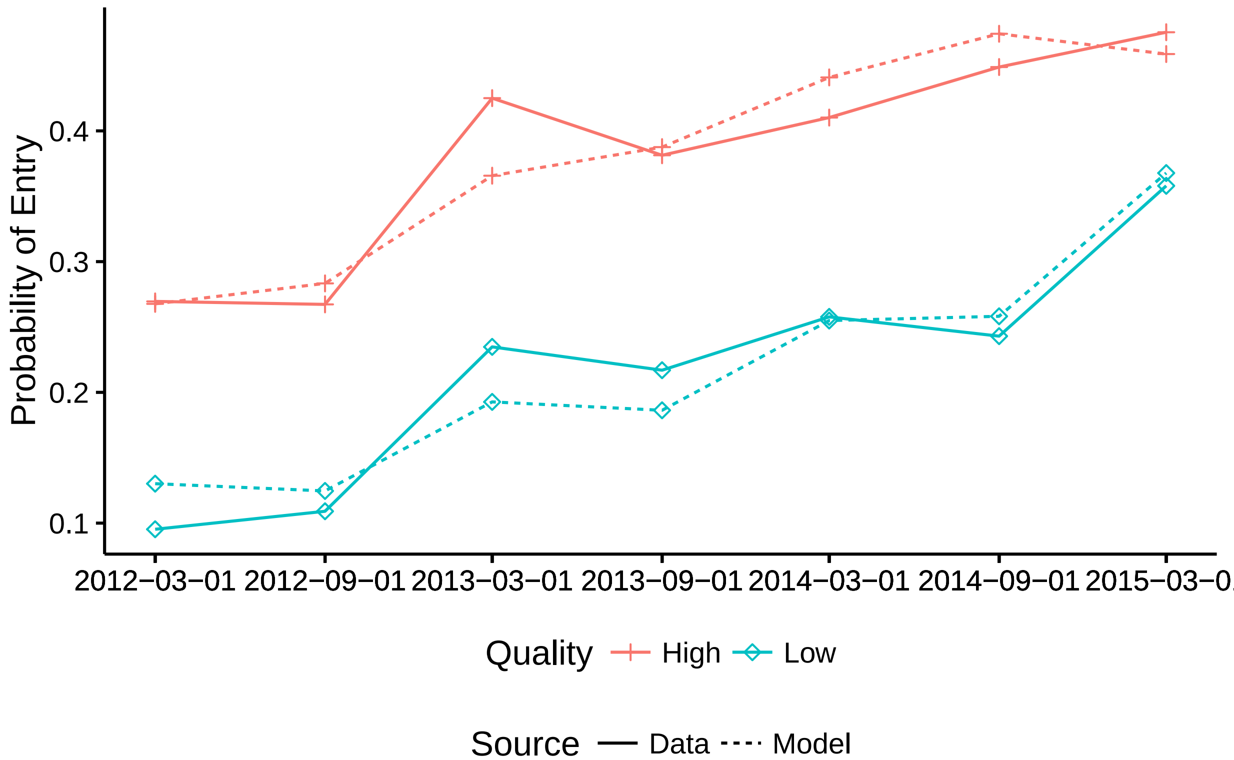

Figure C.2: **Structural model predicted probabilities of entry**
Note: This figure shows the model predicted probability of entry for high (dashed red line, crosses) and low (dashed blue line, diamonds) quality goods over time, and their relationship to the corresponding probabilities of entry in the data (solid lines).

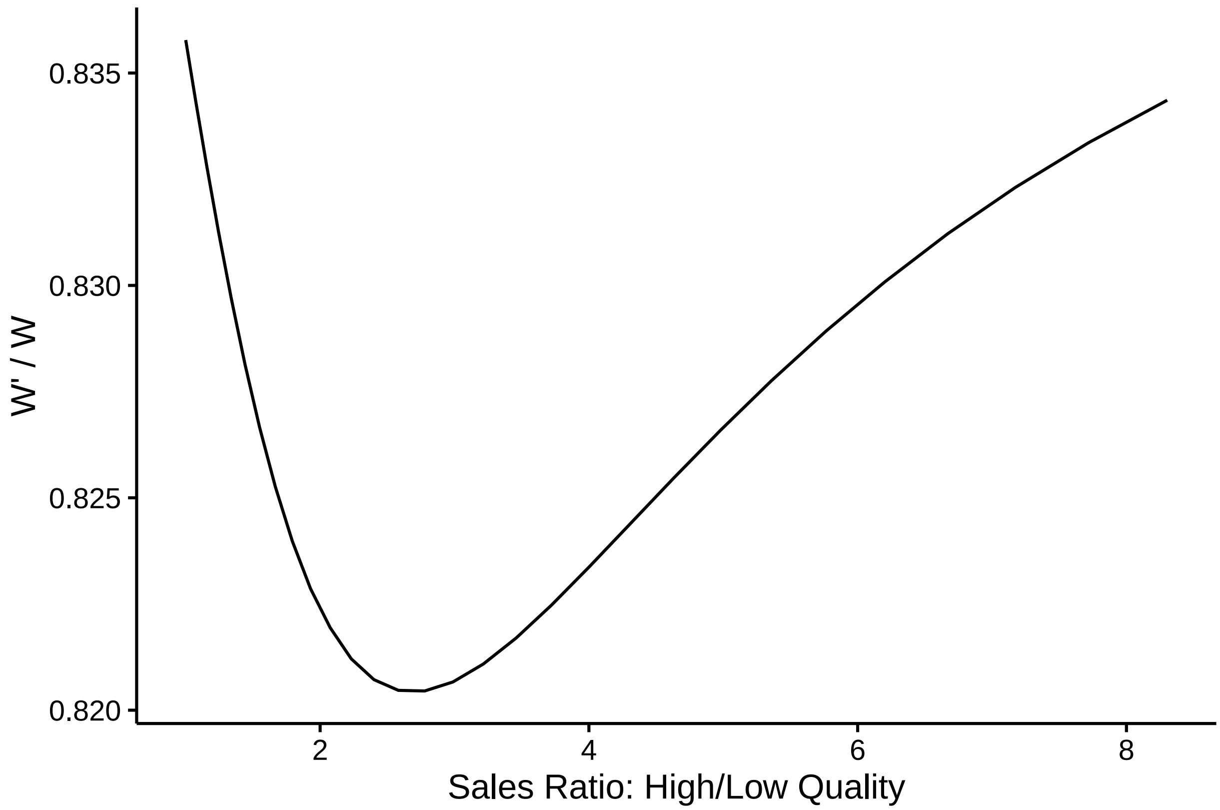

Figure C.3: **Welfare loss with alternative parameters**
Note: This figure plots the welfare cost of the devaluation for different values of the quality demand shifter, assuming the cost of the high quality product is 2.7 times the low quality product, whose cost is held fixed at the estimated level. An x-axis value of one corresponds to no demand increase for high cost goods, i.e., a model with only cost heterogeneity.