

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: February 24, 2019

Abstract

Do firms respond to cost shocks by reducing the quality of their products? Using microdata from a large Russian retailer that varies its offerings twice-yearly, we document that ruble devaluations are associated with a reduction in the observed material quality of products imported for resale, but that higher quality goods are also more profitable. We reconcile these facts using a simple multi-product sourcing model that features a demand system with expenditure switching, where more profitable products can be dropped more quickly after a cost shock. The estimated model shows that quality downgrading reduces average pass-through by 6% and has meaningful consequences for welfare. *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 Kalouptsi, 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, Kanef 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? Economists¹ and the business press² have long speculated that companies may reduce the quality of their product offerings instead of raising prices in response to adverse exchange rate movements. This hypothesis complements a long literature on incomplete price pass-through in international finance by providing another margin of adjustment for firms.³

While quality downgrading may offer an explanation for long-run incomplete pass-through, there are two challenges in testing the hypothesis: first, it has been difficult to accurately measure quality; second, any positive evidence of quality downgrading must be reconciled with the quality sorting literature, which shows that higher quality products tend to be more profitable.⁴ Since a cost shock that hits all imports proportionately will typically not change product profit rankings, quality sorting would seem to rule out quality downgrading. Our contribution is to directly test for quality downgrading using new and uniquely granular microdata, as well as to build and estimate a tractable model of product sourcing that can accommodate high quality products being ex ante more profitable—as in the quality sorting literature and our own data—but also dropping out more quickly post shock.

We use novel data from a large Russian online apparel retailer as a laboratory for studying whether quality adjustments are operational during an exchange rate shock. We directly observe the fabric and material composition of hundreds of thousands of individual products offered by the firm, as well as prices, quantities and unit costs. Following [Crozet, Head, and Mayer \(2011\)](#), [Levchenko, Lewis, and Tesar \(2011\)](#), and [Medina \(2018\)](#), who use expert opinions or product descriptions to classify goods as high or low quality, we combine 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 configuration of offered products.

We first confirm that high quality imports tend to be more profitable and more expensive in our data, as in the quality sorting literature. Since the profit ranking of different products does not change in response to a proportional cost shock in a canonical trade model ([Crozet, Head,](#)

¹[Feenstra \(1988\)](#) argues that firms may upgrade their products through changing the design or adding extra features when there is a decline in the quantity sold, in his example as a result of quotas.

²In 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.

³For recent entries on incomplete price pass-through see, for example, [Goldberg and Campa \(2010\)](#), [Gopinath and Itskhoki \(2010a,b\)](#), [Amiti, Itskhoki, and Konings \(2014\)](#), and [Auer, Burstein, and Lein \(2017\)](#).

⁴[Manova and Zhang \(2012\)](#); [Crozet, Head, and Mayer \(2011\)](#)

and Mayer, 2011), quality sorting suggests there should not be quality downgrading.

We then show that high quality imports are dropped more quickly relative to low quality ones within narrow categories after the Russian ruble devaluation increases import costs in 2014. A 1% ruble devaluation causes a roughly 0.35% differential reduction in the fraction of natural fabrics imported versus domestically produced items. The analysis relies on a difference-in-differences strategy with Russian manufactured products as a control group, which rules out common shocks as the explanation for the compositional shift. Quality downgrading is consistent with long-run sticky average prices found in the literature.

Having documented quality downgrading, we next turn to the question of why the firm would react to the exchange rate shock by reallocating towards lower quality products. We rule out “flight from quality” due to falling incomes as the primary mechanism by exploiting a concurrent oil price shock, which affects labor earnings differentially across oil-producing regions of Russia.⁵ Auxiliary regressions indicate no differential pass-through across qualities and document a reallocation of quantities from high to low quality within product categories, suggesting instead that high quality products were dropped because demand for high quality goods was relatively more sensitive to the price increase as in Bems and di Giovanni (2016) or Medina (2018).

To explain the data we build and estimate a model of import sourcing where high quality products can be ex ante more profitable, but can also be dropped more quickly after a cost shock. The key ingredient is a Khandelwal (2010) style logit demand system that supports expenditure switching as in Bems and di Giovanni (2016). This type of demand system has been used successfully to explain how trade responds to differences in incomes, and here we show it can also explain why ex ante more profitable products are disproportionately dropped in response to a proportional cost shock.⁶ Importantly, no income shock is required to generate the product reallocation. On the supply side, including a product affects demand for every other product; this non-separability in sourcing decisions implies a difficult combinatorial discrete choice problem with complementarities as in Antras, Fort, and Tintelnot (2017). By assuming an incomplete information structure within the firm, we retain demand complementarities but dramatically simplify computation.

The recovered model parameters are consistent with our quality classification in the reduced

⁵Such phenomena are well-known in the literature (see Burstein, Eichenbaum, and Rebelo (2005)), and similar mechanisms have been emphasized by Coibion, Gorodnichenko, and Hong (2015), who find that consumers reallocate expenditure across stores in response to economic conditions.

⁶See, for example, Fajgelbaum, Grossman, and Helpman (2011); Levchenko, Lewis, and Tesar (2011); Auer, Chaney, and Sauré (2018).

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.⁷ Low fixed sourcing costs rationalize observed product entry and exit, as well as relatively low sales per product.

The estimated model allows us to decompose the role of quality downgrading in mediating price pass-through. We show that pass through into average prices is roughly 6% lower with quality reallocations compared to a base case with no entry or exit of products. For a high quality product that is replaced with a low quality product, pass-through becomes negative. Complementing the price pass-through results, the model also provides novel insights on how quality mediates the welfare effects of a devaluation. In particular, we show that the bias arising from omitting quality heterogeneity in counterfactuals cannot be signed in general.

This paper contributes to a large literature that explores why pass-through from exchange rate shocks into prices is incomplete. A variety of consistent explanations for incomplete pass-through have been tested using both firm-product [Gopinath and Rigobon \(2008\)](#); [Gopinath and Itskhoki \(2010a,b\)](#) and firm-category (e.g., HS-8 or HS-10) level prices [Knetter \(1989\)](#); [Goldberg, Knetter, et al. \(1997\)](#); [Auer and Chaney \(2009\)](#); [Berman, Martin, and Mayer \(2012\)](#); [Amiti, Itskhoki, and Konings \(2014, 2016\)](#). Our result that quality downgrading can lead to incomplete price pass-through is most applicable to price stickiness within product categories, since our evidence relates to product adding and dropping and not direct replacement. However, the model is consistent with within-firm-product upgrading and downgrading, and thus remains relevant to the within-firm-product long run pass-through findings. Indeed, [Nakamura and Steinsson \(2012\)](#) find that firms often replace products instead of changing prices, giving firms ample opportunity to adjust quality levels.

The present work is also linked to research that focuses on quality sorting of products and quality upgrading. [Manova and Zhang \(2012\)](#) and [Crozet, Head, and Mayer \(2011\)](#) demonstrate cross-sectional quality sorting within firms: high quality products are exported to more destinations and have higher trade values, which in their frameworks is rationalized by the products being more profitable. [Fan, Li, and Yeaple \(2015\)](#), [Bas and Strauss-Kahn \(2015\)](#), [Manova and Yu \(2017\)](#) show that firms may upgrade quality after a trade shock given production function complementarities; their focus is not price pass-through, but rather how trade affects firm level residuals,

⁷Here we recover marginal costs by inverting the demand system, so that these are true marginal costs and not wholesale costs from our data.

either quality or productivity.⁸ [Medina \(2018\)](#) addresses the same focus, but relies on an expenditure switching demand system to induce firms to change their input quality mix in response to an import price shock. While we draw on this literature’s robust finding that higher quality products tend to be more profitable—especially in wealthier countries—we do not speak to the trade literature on how firms produce quality or productivity as our firm purchases its products from wholesalers.

A key difficulty in the trade literature on quality has been actually identifying which goods are high quality, and quantifying what that implies for demand. In an influential paper, [Khandelwal \(2010\)](#) pioneers using a demand residual, while [Medina \(2018\)](#), [Levchenko, Lewis, and Tesar \(2011\)](#) and [Alessandria and Kaboski \(2011\)](#) make an assumption based on the description of the goods (e.g., pima cotton versus other fabrics, and fresh versus frozen fruit) and [Crozet, Head, and Mayer \(2011\)](#) uses expert opinions. 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 in our structural model. This approach is similar to that of [Auer, Chaney, and Sauré \(2018\)](#), who find quality affects price pass-through, as do [Ludema and Yu \(2016\)](#) and [Chen and Juvenal \(2016\)](#).

Other papers have studied the role of quality in in an international setting. 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. Previous work has also examined the relationship between trade distances and quality ([Alchian and Allen \(1964\)](#), [Hummels and Skiba \(2004\)](#), and [Feenstra and Romalis \(2014\)](#)). 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](#)) or because competing with inexpensive imports drives firms to upgrade, as in [Medina \(2018\)](#).⁹

Finally, 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](#)

⁸For productivity see, e.g., [Bustos \(2011\)](#).

⁹Other 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 \(2011\)](#)).

Zerom, 2010) but which do not allow for quality downgrading or entry and exit.¹⁰ 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.

The paper proceeds as follows. Section 2 provides an overview of the data and institutional background. 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 details on the estimation, recovered parameters, and counterfactuals. Section 6 concludes.

2 Background and Data

Our data come from a large, online apparel retailer that sells across all of Russia.¹¹ The retailer offers clothing, shoes, and accessories. At the retailer-assigned stock-keeping unit (SKU) level, we observe the price, which is constant across Russia but can vary month to month, as well as the quantity sold in each province (oblast) in each month.¹² 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 a product’s inventory, fabric composition, country of manufacture, brand (e.g., Adidas), product group (e.g., shoes), wholesale cost in rubles, and which currency the the firm used to purchase each SKU.¹³ A more precise description of these variables and how they are used in the analysis is provided below.

¹⁰Feenstra, Gagnon, and Knetter (1996) looks at pass-through for cars, and notes that quality adjustments may affect price pass-through numbers.

¹¹The company is owned by 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. The firm is one of two leading online apparel retailers in Russia, wielding significant market power in many of Russia’s regions, and employing more than 4,000 people as of December 2015.

¹²Even though the data is sufficiently granular to facilitate the tracking of purchases for each consumer over time, we aggregate up to the regional level and exploit shocks to local GDP to identify any potential income induced demand reallocation. We find no evidence of an income-shock induced “flight from quality” in section 3.3.

¹³Most imported SKUs are invoiced either in Euros or the U.S. dollar, and the ruble depreciated almost one-for-one against both. The prevalence of dominant currencies in international transactions is consistent with recent evidence from international finance (e.g., see Casas, Díez, Gopinath, and Gourinchas (2017)).

2.1 Store features

The store operates by ordering SKUs at a wholesale cost from both large and small brands and then reselling to Russian consumers with a markup. Most SKUs are uniquely associated by the firm with the Fall/Winter or Spring/Summer season within a year, which are the two main seasons in the fashion industry (Bhardwaj and Fairhurst, 2010). Before a season begins, the firm chooses which brands and SKUs to include, and, once the goods start being offered, the firm is free to choose pricing.¹⁴

We associate the Spring season with the period from March through August, and Fall with September through February of the following year.¹⁵ Figure 1 shows that the majority of revenue 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.¹⁶

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.¹⁷

2.2 Product quality 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 common fabrics are presented in Appendix A. Cotton, polyester, and leather dominate, with at least one of the three present in 50% of SKUs.

We follow Levchenko, Lewis, and Tesar (2011) and Alessandria and Kaboski (2011) and classify products as high or low quality based on their product description, and specifically based on the

¹⁴As 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. We also find that, on average, the retailer charges a markup of two (i.e., doubling wholesale costs) until the goods are put on sale and phased out as the season draws to an end.

¹⁵78% of Spring SKUs and 75% of Fall SKUs are introduced in our designated Spring and Fall months, respectively. 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.

¹⁶Since 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.

¹⁷Related features of the microdata have recently been emphasized in work studying how firms grow through the introduction of new product lines (e.g., Argente, Lee, and Moreira (2018)).




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.

primary material used in the product. To proceed, we first code polyester, plastic polymers, and any fabric with the word “artificial” as low quality. We assume an SKU containing a low quality material 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. Where an artificial fabric appears overwhelmingly as part of a blend and is included to provide a specific property—for instance, elastane, which provides stretchiness—it is coded as high quality. 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 categories is given in Table A.1 in Appendix A.

As in the fabric quality upgrading analysis of Medina (2018), our split reflects that naturally-derived materials such as leather, silk, and cotton have superior attributes compared to fake leather, polymers, and polyester. We verify that our high quality coded products have a larger demand shifter than low quality products using demand regressions in section 5, which maps to the Khandelwal (2010) method of eliciting quality as the demand residual conditional on price in a logit regression.

Table 1 presents summary statistics by product group. The Share column gives the number of SKUs in that group divided by the total number of SKUs offered over the whole sample period,

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.

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

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 end 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

¹⁸The 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 employed about 142,860 workers in 2014.

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

²⁰See Figure A.1 in the Appendix A.

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.

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. Since Table 1 shows that different product groups have very different mean levels of quality, to assess the magnitude of downgrading accurately we will control for reallocation between product groups in Section 3.

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 steep ruble depreciation 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, and financial

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

²²As 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 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.

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

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 identify four empirical facts in our data:

1. High-quality goods are more profitable than low-quality goods.
2. Imported goods experience a greater quality reduction compared to Russian-produced goods, and goods for which quality is more costly to provide experience the greatest quality reduction.
3. Regions in Russia that experience greater income shocks do not differentially reallocate consumption to lower quality goods.
4. High-quality goods do not experience differential pass-through.

Fact 1 implies that our data exhibits the same features as the quality sorting literature where high quality goods are more profitable (Manova and Zhang, 2012). In workhorse models of international trade, this would imply high quality goods would not be dropped after an adverse shock (Crozet, Head, and Mayer, 2011). Facts 2 and 3 establish that the exchange rate shock induces quality downgrading, and rule out an income shock induced “flight from quality” à la Burstein, Eichenbaum, and Rebelo (2005) as the sole explanation for quality downgrading. Fact 4 suggests that differential movements in the relative markups of high and low quality goods cannot explain the disproportionate exit of high quality goods.

3.1 Quality and profitability

Since we observe wholesale costs of a product c_j directly, we can approximate the variable profits of a good j as $\pi_j = q_j(p_j - c_j)$.²⁴ In all following sections, we will refer to high quality prod-

²³This 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.

²⁴Price varies over a product’s life within season; we use sales prices that are actually observed and faced by consumers to compute profits.

Table 3: Mean differences for high quality products

	<i>Dependent variable:</i>		
	$\log(\pi)$ (1)	$\log(q)$ (2)	$\log(p)$ (3)
$Natural_j$	0.046* (0.023)	-0.339*** (0.027)	0.379*** (0.026)
Group FE	✓	✓	✓
Season FE	✓	✓	✓
Observations	305,376	305,376	305,376
R ²	0.365	0.170	0.383

*Note: This table presents coefficient estimates from specification 1. The outcome variables is either the profit, quantity sold, or price of SKU j , in product group g , in season t . Product group and season fixed effects are included. Prices are sales-weighted within SKUs, and standard errors are clustered at the group level. ***, **, * indicate significance at the 0.1%, 1% and 5% levels, respectively.*

ucts interchangeably as “natural,” in line with our classification method. We run the following regression at the SKU-level:

$$\log(y_{jgt}) = \beta \cdot Natural_j + \sum_g \alpha_g \mathbf{D}_g + \sum_t \alpha_t \mathbf{D}_t + \epsilon_{jgt} \quad (1)$$

where y_{jgt} is either the profit, quantity sold, or price of SKU j , in product group g , in season t , \mathbf{D}_g and \mathbf{D}_t are product group and season fixed effects, respectively. Standard errors are clustered at the group level to allow for serial correlation across time. The results are reported in Table 3; high quality goods are found to be about 4.5% more profitable on average. Controlling for brand and product group fixed effects, so that only within brand variation is used, implies a similar estimated magnitude significant at the 0.1% level (see Appendix B).

Note from the quantity regression in Table 3 that high quality goods sell fewer units than low quality goods. Thus, even if there is a per-unit distribution or storage cost in the complete marginal cost, it will not reverse the profit ordering.

3.2 Quality downgrading

We show in this section that the share of high-quality goods on offer was reduced in response to the exchange rate shock. Our identification strategy relies on a difference-in-differences (DiD) ap-

proach, where imported SKUs are the treatment and domestically produced SKUs are the control group. The fraction of products that are high quality (natural fabric) 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;²⁵ if quality adjustment is an important margin for passing through the ruble cost increase, then there will be a negative, significant coefficient for the foreign sourced goods post-shock.

In our first specification, we aggregate within seasons to the product group-origin level.²⁶ 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 fraction of high quality SKUs 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:

$$natfrac_{grt} = \sum_{t>1} \delta_t (nonrus_{gr} \cdot \mathbf{D}_{gt}) + \sum_{gr} \alpha_{gr} \mathbf{D}_{gr} + \sum_{gt} \alpha_{gt} \mathbf{D}_{gt} + \epsilon_{grt} \quad (2)$$

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 , origin r , in season t , δ_t are the time-varying treatment effects, $nonrus_{gr}$ is an indicator with a value of one for the set of non-Russian (imported) 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. The latter sets of indicators are included to account for systematic differences in quality across product groups, as well as for changes in this quality level within groups over time and by origin.

Specification 2 uses only within group-origin variation to identify downgrading. Because the specification includes group-origin and group-season dummies, it is equivalent to running a separate DiD within each product group, using foreign-sourced products as the treatment in each case, and then averaging the treatment effects across product groups. Treatment effects that are the result of seasonal reallocations from high $natfrac$ to low $natfrac$ product groups are therefore ruled out, as are explanations that are common across the treatment and control within a product group, such as changing tastes, changing incomes, or changing commodity/raw fabric costs that are contemporaneous with the devaluation.

²⁵We confirm that this is true in pass-through regressions in Section 3.4.

²⁶While this aggregation helps to transform the data into a tractable format for regression analysis, the results hold for alternative levels of firm stocking choices.




Figure 3: **Quality downgrading**

Note: This figure plots the estimated δ_t coefficients of equation 2 with 95% confidence intervals around them. Fixed effects are at the product group \times country of origin and season level. Standard errors are clustered by group \times origin to allow within-group-origin serial correlation.

The estimated coefficients δ_t from equation 2 are plotted in Figure 3, along with their associated standard errors, clustered at the group \times origin level to allow for within-group-origin serial correlation over time. The results indicate that there is no statistically significant differential reduction in quality within product groups for non-Russian (imported) goods until the March 2015 season, after the peak of the devaluation. That is, there was a significant reduction in the quality of imported products, and it happened on a time frame 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 as part of our identification strategy, and rules out 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 lagged exchange rate movement to play a role:

$$natfrac_{grt} = \delta (nonrus_{gr} \cdot \log(ER_{t-1})) + \sum_{gr} \alpha_{gr} \mathbf{D}_{gr} + \sum_{gt} \alpha_{gt} \mathbf{D}_{gt} + \epsilon_{grt} \quad (3)$$

$\log(ER_{t-1})$ is the average U.S. dollar to ruble 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 express the dependent variable in levels in our baseline specification, but all results go through if we use $\log(natfrac_{grt})$ instead.²⁷

We run equation 3 for three different levels of aggregation: i) one that does not distinguish between product groups at all (no g), so that each season has one observation for the imported high quality fraction and one for the domestic high quality fraction; ii) one where g indicates product groups as in equation 2; iii) and one where g indicates specific brands within a product group.²⁸ 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.

Our base specification in columns (2) and (3) of Table 4 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 high quality offerings. 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), δ is estimated as insignificant, implying that within-brand reallocations are less important for downgrading.²⁹

If the increase in costs from the exchange rate shock—rather than an income shock or a change in the nature of demand—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 3 to vary by product group in our product group level specification:

$$natfrac_{grt} = \sum_g \delta_g (nonrus_{gr} \cdot \log(ER_{t-1})) + \sum_{gr} \alpha_{gr} \mathbf{D}_{gr} + \sum_t \alpha_{gt} \mathbf{D}_{gt} + \epsilon_{grt} \quad (4)$$

For each product group, we recover the quality premium by dividing the average wholesale cost 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

²⁷We also run regressions using the number of high and low quality SKUs instead of the fraction, which we discuss in the robustness section. The results are available in Appendix B.

²⁸For example, Adidas and Puma are two brands within sport shoes, but here a brand will have different fixed effects for all the product groups where it sells items.

²⁹Results for specification 3 using the logged fraction of natural offerings, and results dropping the last season of incomplete data are reported in Appendix B. Both are qualitatively and quantitatively similar to our baseline findings.

Table 4: Differential quality downgrading

	Dependent variable:			
	$natfrac_{grt}$			
	(1)	(2)	(3)	(4)
$nonrus_{gr} \cdot \log(ER_{t-1})$	-0.285** (0.059)	-0.347*** (0.064)	-0.321** (0.115)	0.204 (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 3, aggregating SKUs within non-Russian (imported) or Russian origin in column (1), within product group-origin in columns (2) and (3), and within brand-origin in column (4). The outcome is the fraction of offered SKUs that use a natural fabric for group or brand g , origin r , in season t . $nonrus_{gr}$ is an indicator with a value of one for the set of non-Russian products in group or brand g , and $\log(ER_{t-1})$ is the average exchange rate during season $t - 1$. Standard errors (in brackets) are clustered at product group or brand \times origin level to allow for serial correlation across time. ***, **, * indicate significance at the 0.1%, 1% and 5% levels, respectively.

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 4.³⁰ The strong negative relationship between the costs of providing quality and the amount of quality downgrading supports the hypothesis that costs played a central role in the firm's decision to quality downgrade after the devaluation. Our result that product groups with the highest costs downgrade the most after a proportional increase in input wholesale costs agrees with the evidence in Fan, Li, and Yeaple (2018), who find that firms with the highest costs upgrade the most after a proportional reduction in input prices.

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

³⁰The full regression results from equation 4 are available in Table B.2 in Appendix B.




Figure 4: **Cross-group variation in downgrading**

Note: This figure plots the estimated δ_g coefficients of equation 4. Fixed effects are at the $group \times origin$ and season level. Standard errors are clustered by $group \times origin$ level to allow within-group-origin serial correlation.

lower but still positive in Table 5 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 4 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, rather than downgrading in the treatment group, especially since the control group is relatively small. We perform several checks to address this issue. First, in Appendix B we provide a raw DiD data graph for polymers (Figure B.1), which appear as a rubber and leather substitute in product groups using leather (approx. 40% of total SKUs). 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 the DiD is picking up differential downgrading in the treatment group.

We also run a DiD using only imported goods, treating the logged number of high or low quality SKUs within a group as our dependent variable. The growth in imported natural fabric SKUs is negative and significant compared to imported artificial fabric SKUs, so that Table 4

reflects imports’ natural fabric share is actually shrinking, and not simply growing less quickly than domestic products’ natural fabric share. Full results are reported in Appendix B.

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 of differential consumption reallocation toward 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:

$$\log(Qual_{it}) = \alpha_i + \sum_t \gamma_t \mathbf{D}_t + \sum_t \delta_t (\mathbf{D}_t \cdot Growth_i) + X_{it}' \theta + \sum_t \psi_t (\mathbf{D}_t \cdot X_{it}) + \epsilon_{it} \quad (5)$$

$$\forall i, \forall t \in \{2012m1, \dots, 2015m9\} \setminus \{2014m12\}$$

where $Qual_{it}$ is either the median or mean quality (*natfrac*) 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.³¹ 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 5. Panel (A) of Figure 5 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.

³¹The results are unaffected by inclusion of these additional controls and interaction terms. Appendix B.4 further presents estimates for the median and mean regular prices in region i at time t as alternative outcome variables.

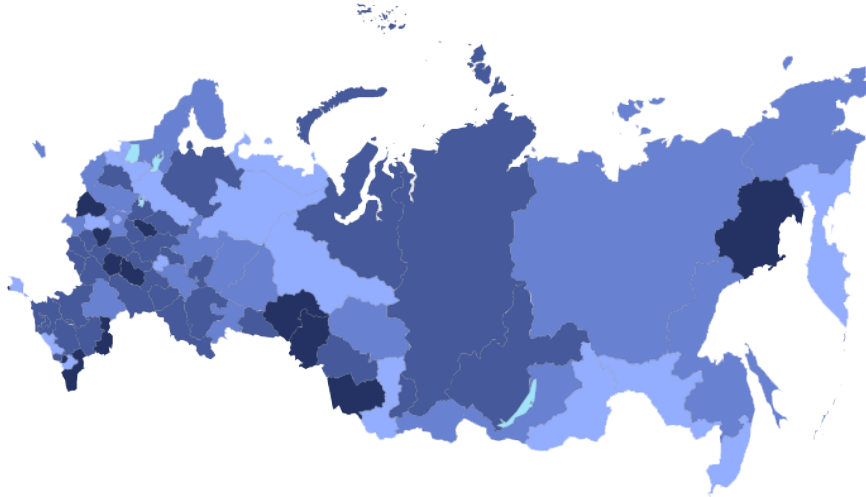
As would be necessary with any DiD estimation approach, this specification also provides evidence on the parallel trends assumption in all outcome variables. That is, in the absence of treatment the unobserved disparities between high- and low-growth regions should be constant over time—the validity 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 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 5 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 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 median and mean quality can only be explained by changes in quantities (purchases).

The findings are summarized in Figure 5, which plots the key estimated parameters of interest, $\hat{\delta}_t$, 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 similarly robust

(A) Regional growth (2015)



(B) Income effect




Figure 5: 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 5 with 95% confidence intervals around them. Results for two distinct outcome variables are displayed over time: the log median regional quality (black), and the log mean regional quality (grey). Time is measured on a monthly basis.

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.

3.4 Price pass-through

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 material-brand-group choice as a consistent product over time through the inclusion of eponymous fixed effects. Meanwhile, we still use SKUs as our unit of observation in the regression. Our specification is:

$$\log(y_{jmbgt}) = \beta_1 \log(ER_{t-1}) + \beta_2 \log(ER_{t-1}) \cdot Nat_{jmbgt} + \beta_3 \log(ER_{t-1}) \cdot Rus_{jmbgt} \quad (6)$$

$$+ \sum_{bgr} \alpha_{bgr} \mathbf{D}_{bgr} + \sum_{mbg} \alpha_{mbg} \mathbf{D}_{mbg} + \epsilon_{jmbgt}$$

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 U.S. dollar to ruble 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 5. 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 wholesale costs, this imperfect pass-through does not correspond to lowered markups: the pass-through on cost is very similar in column (2).³² Importantly, the differential change in prices and wholesale costs for high

³²From 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

Table 5: Pass-through coefficients

	<i>Dependent variable:</i>	
	$\log(\textit{price})$ (1)	$\log(\textit{cog})$ (2)
$\log(ER_{t-1})$	0.646*** (0.030)	0.626*** (0.037)
$\log(ER_{t-1}) \cdot \textit{Nat}$	0.055 (0.029)	0.010 (0.035)
$\log(ER_{t-1}) \cdot \textit{Rus}$	-0.176** (0.058)	-0.201*** (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 6 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 \textit{Nat} and \textit{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 \times material level to allow within-group-material serial correlation across time. ***, **, * indicate significance at the 0.1%, 1% and 5% levels, respectively.*

quality goods is not significantly different from zero, implying no differential pass-through for 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.3. We also perform standard within-SKU pass-through regressions on the small set of SKUs we observe for longer than one season, and find no evidence of differential pass-through for natural fabric products.

Even with no differential pass-through there may have been a differential reduction in demand. With demand that exhibits expenditure switching, a proportionate price increase can imply a disproportionate reduction in quantity sold of the more expensive, higher quality product. Indeed, we find that within product groups, the aggregate quantity sold of high quality products decreases disproportionately more relative to low quality products. Those results are reported in Appendix B.3.

0.5, which may reflect that larger brands with more SKUs negotiated higher pass-through into costs.

4 Structural Model

This section develops and estimates a simple structural model of quality choice. We write a model capable of matching the facts that high quality products are dropped at a faster rate post shock and that demand decreased for these products disproportionately while there was no differential pass-through. The estimated model is then used to assess counterfactuals and the partial welfare implications of quality downgrading.

4.1 Setup

Demand

Each season t there are M_t consumers indexed by i , who choose among products offered during that season and an outside option. They face N_{ht} high quality products and N_{lt} low quality products, each of which is differentiated with a consumer-product specific idiosyncratic demand shock. Consumer 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.³³ 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:

$$s_{jmt}(p_{jmt}, \mathbf{p}_{-jt}, N_{ht}, N_{lt}) = \frac{\exp(q_m + \alpha p_{jmt})}{1 + \sum_{j' \in N_{ht}} \exp(q_h + \alpha p_{j't}) + \sum_{j' \in N_{lt}} \exp(q_l + \alpha p_{j't})} \quad (7)$$

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

We highlight a key feature of the demand system in the following lemma:

Lemma 1. *A proportional increase in both the price of high quality and low quality products, p_h and p_l , will lead to a decrease in s_{jht}/s_{jlt} as long as $p_{jht} > p_{jlt}$.*

If prices increase proportionately, then the more expensive product will experience a greater reduction in market share, as relatively more weight shifts to the outside option. This exactly

³³We follow Medina (2018), Levchenko, Lewis, and Tesar (2011) and Alessandria and Kaboski (2011) and our reduced form in treating quality as a 0-1 dummy corresponding 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 as in Khandelwal (2010).

mirrors the disproportionate substitution away from expensive, high quality goods in [Fajgelbaum, Grossman, and Helpman \(2011\)](#).

Quality choice

We assume that 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.³⁴

Formally, the purchase manager for SKU j 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 choices of other managers, a manager will strategically set prices to maximize profits in a Nash-Bertrand equilibrium:

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

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 and 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 \(2010\)](#). We choose a symmetric equilibrium in the pricing game where any j with quality m has the same optimal price p_{mt}^* .

At time $t - 1$, the manager for j must decide what quality, if any, to source. The profit to j of providing quality m is:

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

where \mathbf{a}_{-jt} denotes the equilibrium entry and quality strategies of all potential entrants, of which

³⁴Models of product sourcing with production or demand interrelationships fall into the class of combinatorial optimization problems ([Antras, Fort, and Tintelnot, 2017](#)). Our model requires demand interrelationships be taken into account, since inward shifting residual demand curves are the only limit on the size of the firm; we thus cannot use the quality sourcing models of [Fan, Li, and Yeaple \(2018\)](#) or [Manova and Yu \(2017\)](#), which rely on single product firms or abstract from interrelationships. Our method implies a tractable sourcing model that is very easy to solve and estimate (< 1 sec to compute an equilibrium, vs. roughly 1 day for [Jia \(2008\)](#)).

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.

Note that 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, manager k expects that j will choose quality m with probability P_{jmt} , and will choose not to enter with probability P_{j0t} . Manager j 's expected profits from choosing material m are then:

$$\pi_{jmt}^e(\mathbf{P}_{-jt}) - \sigma_\epsilon \epsilon_{jmt},$$

where the expectation is taken over all the possible distributions of offered product qualities given strategies \mathbf{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:

$$P_{jmt} = \frac{\exp(\pi_{jmt}^e(\mathbf{P}_{-jt})/\sigma_\epsilon)}{1 + \sum_{m'} \exp(\pi_{jm't}^e(\mathbf{P}_{-jt})/\sigma_\epsilon)} \quad (8)$$

A Bayesian Nash Equilibrium (BNE) at each time t is a vector of choice probabilities \mathbf{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:

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

Our welfare formula will serve as a useful internal benchmark when we evaluate the consequences of counterfactual devaluations or policies. Since the formula only covers one firm, we do not claim that it represents the full welfare costs of the devaluation.

4.2 Model predictions

We provide intuition for the model’s predictions for how a manager’s choice of products changes in response to a nominal exchange rate devaluation with the following theorem:

Theorem 1. *There exist parameters such that for a given exchange rate, $\pi_h > \pi_\ell$, but (1) a manager’s elasticity of choosing h with respect to the exchange rate is larger than the elasticity for ℓ , and (2) the exchange-rate elasticity of demand is larger for h than for ℓ .*

Proof. See Appendix C. □

The theorem states that it is possible for a high quality good to be more profitable than a low quality one, but still be dropped at a faster rate in response to an exchange rate shock. This result comes from the demand model, and would also be generated by the non-homothetic linear demand curves of Melitz and Ottaviano (2008) or Eckel, Iacovone, Javorcik, and Neary (2015). Both the logit and linear demand can predict disproportionate substitution from a high price, high quality good to the outside option in response to a proportional cost shock affecting all inside goods. The result does not require that markups for the high quality good drop more quickly, but does require that the elasticity of demand for the high quality good is larger, and that consumers actually shift expenditures away from the high quality good as in Bems and di Giovanni (2016). In Appendix C we show that if the expenditure share across h and ℓ goods remains fixed, then even if the elasticity of substitution is higher for h goods than ℓ goods in a CES framework, there will still not be a disproportionate reduction in h offerings.

5 Estimation and Results

This section describes how the model is estimated using the subset of product (or “target”) groups for which quality is costly to provide in the sense of Figure 4.³⁵ We estimate the parameters as a function of the data in three steps: first, demand parameters are estimated; second, the demand system is inverted to 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.

³⁵The six excluded product groups are Jeans, Sweatshirts, Tee-shirts and Polos, Trousers and Jumpsuits, Underwear, and Vests and Tops.

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:

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

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.³⁶ The relationship between market size and total quantity ordered is provided in Figure C.3 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:

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

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

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:

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

³⁶This 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.

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 U.S. dollar to ruble exchange rate in Table 2 lagged one season and normalized by the long run average.³⁷ 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 , and the fixed cost shock variance σ_ϵ . 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:

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

To construct entry probabilities as a function of parameters, we first non-parametrically estimate 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.³⁸ 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](#)).

³⁷To normalize the exchange rate, we divide by the expected value of the AR(1) run on season-level data from 2000-2014.

³⁸We simplify the computation of managers' expected entry profits slightly by ignoring Jensen's inequality; see [Appendix C.3](#) for details.

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 6](#).

The price parameter α implies that average $p/c \approx 3$, which is in the neighborhood of the median of first-period price divided by wholesale cost of 2.4.³⁹ Overestimating margins will lead to an overestimation of fixed costs to rationalize lower participation, but would not 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.⁴⁰ 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 5](#).

The fixed costs are estimated in hundreds of thousands of rubles. At the pre-2014 long run stationary average of 30.75 rubles per U.S. dollar, this implies sourcing costs of \$3,400 and \$5,800 for high and low cost goods, respectively. While at first glance it may seem surprising that high quality goods have a lower fixed cost, one should keep in mind that these are sourcing costs and not development costs as in other work ([Medina, 2018](#); [Ershov, 2018](#)). In their paper on import

³⁹The 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.

⁴⁰The cost shifter is multiplied by the scaling parameter, $\exp(\phi \cdot \beta_1^c) = 82\%$.

Table 6: 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.

sourcing, [Antras, Fort, and Tintelnot \(2017\)](#) estimate that fixed costs tend to be higher for U.S. firms sourcing from low income countries such as China and India—which typically manufacture lower quality inputs—compared to high income countries such as Norway and Germany, which manufacture higher quality inputs.

The model does well in matching the entry and exit data: the correlation between the predicted 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.4](#) in [Appendix C](#).

Importantly, the model is able to match quality downgrading and a disproportionate drop in the profitability of high quality products. [Figure 6](#) shows that with the estimated model parameters, high quality product profits decrease disproportionately quickly compared to low quality profits in response to the devaluation. In turn, there is disproportionate exit of high quality products, and indeed with these parameters low quality products will actually enter over some ranges of the devaluation due to the general equilibrium effect of reduced competition.

5.3 Quality downgrading and price pass-through

To what extent does quality downgrading affect exchange rate pass-through into average prices? The literature typically focuses on pass-through within narrow categories ([Knetter, 1989](#)) or




Figure 6: **Predicted effects of a devaluation**

Note: Profits are expected profits, since the fixed cost of sourcing is stochastic. Profits decrease disproportionately quickly in response to a devaluation for high quality products, leading to their disproportionate exit.

within products (Gopinath and Itskhoki, 2010a). Our product groups are similar to HS-6 or in some cases HS-10 categories, and Theorem 1 shows that the model supports quality downgrading within products, so our pass-through results can easily be applied to those literatures. Our model is estimated and aggregated across all groups to highlight that the mechanism does not rely on any cross-group heterogeneity, and to provide an easy-to-interpret effect magnitude.

We present average predicted prices in Fall 2014 (pre-shock) and Winter 2015 (post-shock) when quality reallocation is allowed and compare them to prices when the number of products is kept constant at Fall 2014 levels. In Fall 2014, the normalized exchange rate rose from 1.15 to 1.67, a 45% increase. Using the structural marginal cost pass-through coefficient $\hat{\beta}_2^c = 0.7$ estimated in Section 5.2, this implies a 32% proportional increase in marginal costs for high and low quality products. From Table 5 in Section 3.4, the 45% exchange rate increase would imply a $45\% \times 0.64 = 28.8\%$ increase in prices within-product.

To solve the equilibrium entry probabilities we use a nested fixed point approach as in Seim (2006).⁴¹ The average price is computed as $(\hat{N}_{ht}\hat{s}_{ht}\hat{p}_{ht} + \hat{N}_{lt}\hat{s}_{lt}\hat{p}_{lt})/(\hat{N}_{ht}\hat{s}_{ht} + \hat{N}_{lt}\hat{s}_{lt})$, where hats indicate predicted values from the model.

⁴¹For 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.

Table 7: Average price pass-through

Model	Average Price (RUB)		$\Delta\%$
	Fall 2014	Spring 2015	
Quality Reallocation	4,575.8	5,021.4	9.7
No Reallocation	4,575.8	5,051.5	10.4

Note: This table presents average pass-through with and without quality adjustments.

The results in Table 7 indicate that quality reallocation plays a role in dampening average price increases, implying a 9.7% increase instead of 10.4% in the model with no reallocation, 0.7 percentage points or roughly 6% lower.⁴²

Any manager may also choose to replace a high quality product with a low quality one, depending on fixed cost draws, as high quality goods become relatively less profitable. While we cannot directly test this scenario in the data as we do not observe when a good is directly replaced—only the aggregate product choices each season—replacement is supported by our model, and others have found evidence of frequent product replacement in the microdata (Nakamura and Steinsson, 2012). The Fall 2014 price of a high low quality good is 4,890 RUB, while the Spring 2015 price of a low quality good is 4,409 RUB. For high quality products that are replaced, our model thus predicts pass-through of -9.8% .

5.4 Welfare counterfactuals

The model allows us to answer the question of how welfare would change if the firm 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').⁴³

⁴²The marginal cost is roughly one-third of price and markups in the logit model are additive, so the price pass-through is roughly one-third that predicted in the reduced form.

⁴³This exercise is similar to that in Medina (2018) where the author prohibits quality upgrading by increasing the

Table 8: 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 counterfactual simulation results. ΔP_h and ΔP_ℓ are the probabilities of entering as a high- and low-quality product, respectively. W'/W is the welfare change in each case.

Our counterfactual predictions are evaluated using the same Fall 2014 depreciation as used in the previous section. The results are presented in [Table 8](#).

In the base model, there is a 4.2 percentage point decrease 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 reduction in competition due to fewer 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, the relative probability of high to low quality entry would still decrease.

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

fixed costs of sourcing.




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 8. 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 will underpredict the true welfare costs of the nominal devaluation (as reported in the first row of Table 8 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.⁴⁴

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 nonmonotonic. The U-shape is the result of two countervailing forces: as the benefit of quality increases, it be-

⁴⁴Using 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 plot the welfare loss as a function of the quality shifter in Appendix C and show it can be up to 1.5pp larger.

comes 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 zero, 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 as an innately suitable laboratory to dissect how firms respond to cost shocks. 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 toward 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 expenditure switching in demand systems. Using a simple structural model of multiproduct sourcing, the paper shows how allowing goods to be heterogeneous in both quality and cost, and letting firms quality downgrade, offers more nuanced predictions of the welfare effect of a devaluation.

Our study looks at the effects of the exchange rate shock on quality holding consumer preferences fixed. Yet 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 further future reallocation. We leave those questions regarding the long-run demand 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 (2010): “The Great Trade Collapse of 2008–09: An Inventory Adjustment?,” *IMF Economic Review*, 58(2), 254–294.
- AMITI, M., O. ITSKHOKI, AND J. KONINGS (2014): “Importers, Exporters, and Exchange Rate Disconnect,” *American Economic Review*, 104(7), 1942–78.
- (2016): “International Shocks and Domestic Prices: How Large Are Strategic Complementarities?,” *Working Paper*.
- ANTRAS, P., T. C. FORT, AND F. TINTELOT (2017): “The Margins of Global Sourcing: Theory and Evidence From US firms,” *American Economic Review*, 107(9), 2514–64.
- 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*.
- AUER, R., AND T. CHANEY (2009): “Exchange Rate Pass-through in a Competitive Model of Pricing-to-Market,” *Journal of Money, Credit and Banking*, 41, 151–175.
- AUER, R., T. CHANEY, AND P. SAURÉ (2018): “Quality Pricing-to-Market,” *Journal of International Economics*, 110, 87–102.
- 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.
- BERMAN, N., P. MARTIN, AND T. MAYER (2012): "How do Different Exporters React to Exchange Rate Changes?," *Quarterly Journal of Economics*, 127(1), 437–492.
- BHARDWAJ, V., AND A. FAIRHURST (2010): "Fast Fashion: Response to Changes in the Fashion Industry," *International Review of Retail, Distribution and Consumer Research*, 20(1), 165–173.
- 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.
- CASAS, C., F. DÍEZ, G. GOPINATH, AND P.-O. GOURINCHAS (2017): "Dominant Currency Paradigm: A New Model for Small Open Economies," *Working Paper*.
- 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 Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications," *American Economic Review*, 105(3), 993–1029.
- 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.
- ECKEL, C., L. IACOVONE, B. JAVORCIK, AND J. P. NEARY (2015): "Multi-product Firms at Home and Away: Cost- versus Quality-based Competence," *Journal of International Economics*, 95(2), 216–232.
- ERSHOV, D. (2018): "The Effects of Consumer Search Costs on Entry and Quality in the Mobile App Market," *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.
- FAN, H., Y. A. LI, AND S. R. YEAPLE (2015): "Trade Liberalization, Quality, and Export Prices," *Review of Economics and Statistics*, 97(5), 1033–1051.
- (2018): "On the Relationship Between Quality and Productivity: Evidence from China's Accession to the WTO," *Journal of International Economics*, 110, 28–49.
- 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.
- FIELDER, 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.
- GOLDBERG, P. K., M. M. KNETTER, ET AL. (1997): "Goods Prices and Exchange Rates: What Have We Learned?," *Journal of Economic Literature*, 35(3), 1243–1272.
- 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.
- GOPINATH, G., AND R. RIGOBON (2008): "Sticky Borders," *Quarterly Journal of Economics*, 123(2), 531–575.

- 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.
- KNETTER, M. M. (1989): "Price discrimination by US and German exporters," *American Economic Review*, 79(1), 198–210.
- 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.
- MANOVA, K., AND Z. YU (2017): "Multi-product Firms and Product Quality," *Journal of International Economics*, 109, 116–137.
- MANOVA, K., AND Z. ZHANG (2012): "Export Prices Across Firms and Destinations," *Quarterly Journal of Economics*, 127(1), 379–436.
- 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 Exporting: 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

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.




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.

B Reduced Form Evidence

B.1 Profit and quality

We run the following regression on the entire set of pre-shock products (Fall 2014 and earlier) and report the results in Table B.1:

$$\log(y_{jbg t}) = \beta \cdot \text{Natural}_j + \sum_{bg} \alpha_{bg} \mathbf{D}_{bg} + \sum_t \alpha_t \mathbf{D}_t + \epsilon_{jbg t} \quad (\text{B.1})$$

where $y_{jbg t}$ is either the profit, quantity sold, or price of SKU j , in product group g , in season t , \mathbf{D}_{bg} and \mathbf{D}_t are brand \times product group and season fixed effects, respectively. Standard errors are clustered at the brand \times product group level to allow for serial correlation across time. The results are similar to before: high quality goods are significantly about 4.4% more profitable, and sell at a 5.2% higher price on average.

Table B.1: Mean differences for high quality products

	<i>Dependent variable:</i>		
	$\log(\pi)$ (1)	$\log(q)$ (2)	$\log(p)$ (3)
Natural_j	0.044*** (0.012)	-0.007 (0.010)	0.052*** (0.010)
Brand \times Group FE	✓	✓	✓
Season FE	✓	✓	✓
Observations	305,376	305,376	305,376
R ²	0.624	0.592	0.881

*Note: This table presents coefficient estimates from specification B.1. The outcome variables is either the profit, quantity sold, or price of SKU j , in product group g , in season t . Product group and season fixed effects are included. Prices are sales-weighted within SKUs, and standard errors are clustered at the brand \times group level. ***, **, * indicate significance at the 0.1%, 1% and 5% levels, respectively.*

B.2 Quality downgrading

Number of SKU quality downgrading regressions

In this section we run regressions to assess quality downgrading using the logged raw number of SKUs as a dependent variable, instead of the high quality share of SKUs. Regressions in this

section are at the season and product group level.

The first regression drops Russian produced goods, and does a DiD analysis using the natural material category as a treatment group. Let m index quality, with $m = 1$ indicating high quality and $m = 0$ low quality. Using only imported products and taking the log number of high or low quality SKUs as the dependent variable, we run the following specification:

$$\log(N_{mgt}) = \delta (\text{natural}_{mg} \cdot \log(ER_{t-1})) + \sum_{mg} \alpha_{mg} \mathbf{D}_{mg} + \sum_{gt} \alpha_{gt} \mathbf{D}_{gt} + \epsilon_{mgt}, \quad (\text{B.2})$$

Table B.2: 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.

where $natural_{mg}$ is a dummy equal to 1 for high quality products in group g . If relatively more low quality goods than high quality goods were introduced after the cost shock, then δ should be estimated negative, and the quality downgrading finding in the main paper should not be the result of quality upgrading in the control group.

The second regression keeps Russian produced goods and is therefore a triple-difference specification, with natural, imported goods as the treatment group:

$$\log(N_{mgrt}) = \delta (natural_{mg} \cdot nonrus_{gr} \cdot \log(ER_{t-1})) + \sum_{mgr} \alpha_{mgr} \mathbf{D}_{mgr} \quad (\text{B.3})$$

$$+ \sum_{mgt} \alpha_{mgt} \mathbf{D}_{mgt} + \sum_{grt} \alpha_{grt} \mathbf{D}_{grt} + \epsilon_{mgrt}$$

This is a triple difference regression: compared to the differential movement of Russian high quality goods relative to Russian low quality goods, δ will be negative if the decrease in imported high quality goods relative to low quality goods is *lower* than for Russian goods. The estimation results are reported in Table B.5. We find that both prediction are borne out in the data and thus lend additional support to the cost shock generated quality downgrading mechanism.

Table B.3: Differential quality downgrading robustness: logged dependent variable

	Dependent variable:			
	$\log(\text{natfrac}_{grt})$			
	(1)	(2)	(3)	(4)
$\text{nonrus}_{gr} \cdot \log(ER_{t-1})$	-0.360** (0.072)	-0.662*** (0.133)	-0.595** (0.212)	0.095 (1.909)
Origin FE	✓			
Season FE	✓	✓		
Group \times Origin FE		✓	✓	
Group \times Season FE				✓
Brand \times Origin FE				✓
Observations	16	393	393	22,945
R ²	0.915	0.647	0.853	0.999

*Note: This table presents coefficient estimates from specification 3, but with a logged dependent variable, aggregating SKUs within non-Russian (imported) or Russian origin in column (1), within product group-origin in columns (2) and (3), and within brand-origin in column (4). The outcome is the fraction of offered SKUs that use a natural fabric for group or brand g , origin r , in season t . nonrus_{gr} is an indicator with a value of one for the set of non-Russian products in group or brand g , and $\log(ER_{t-1})$ is the average exchange rate during season $t - 1$. Standard errors (in brackets) are clustered at product group or brand \times origin level to allow for serial correlation across time. ***, **, * indicate significance at the 0.1%, 1% and 5% levels, respectively.*

B.3 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

Table B.4: Differential quality downgrading robustness: dropped final season

	<i>Dependent variable:</i>			
	<i>natfrac_{grt}</i>			
	(1)	(2)	(3)	(4)
<i>nonrus_{gr}</i> · log(<i>ER_{t-1}</i>)	-0.237* (0.088)	-0.355*** (0.074)	-0.348* (0.156)	0.244 (1.254)
Origin FE	✓			
Season FE	✓	✓		
Group × Origin FE		✓	✓	
Group × Season FE				✓
Brand × Origin FE				✓
Observations	14	347	347	23,423
R ²	0.858	0.695	0.864	0.999

*Note: This table presents coefficient estimates from specification 3, but dropping the last season (2015-09), aggregating SKUs within non-Russian (imported) or Russian origin in column (1), within product group-origin in columns (2) and (3), and within brand-origin in column (4). The outcome is the fraction of offered SKUs that use a natural fabric for group or brand g , origin r , in season t . $nonrus_{gr}$ is an indicator with a value of one for the set of non-Russian products in group or brand g , and $\log(ER_{t-1})$ is the average exchange rate during season $t - 1$. Standard errors (in brackets) are clustered at product group or brand × origin level to allow for serial correlation across time. ***, **, * indicate significance at the 0.1%, 1% and 5% levels, respectively.*

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 material-brand-season level to check for differential reductions in price and cost dispersion of a brand's high quality SKUs:

$$\begin{aligned}
 CV_{mbgt}^x = & \beta_1 \log(ER_{t-1}) + \beta_2 \log(ER_{t-1}) \cdot Nat_{mbgt} + \log(ER_{t-1}) \cdot Rus_{mbgt} \\
 & + \sum_{bgr} \alpha_{bgr} \mathbf{D}_{bgr} + \sum_{mbg} \alpha_{mbg} \mathbf{D}_{mbg} + \epsilon_{mbgt},
 \end{aligned} \tag{B.4}$$

Table B.5: Number of SKUs quality downgrading results

	<i>Dependent variable:</i>	
	$\log(N_{mgt})$	
	(1)	(2)
$\log(ER_{t-1}) \cdot natural_{mg}$	-1.218*** (0.271)	
$\log(ER_{t-1}) \cdot natural_{mg} \cdot nonrus_{gr}$		-1.109** (0.412)
Group \times Quality FE	✓	
Group \times Season FE	✓	
Group \times Origin \times Quality FE		✓
Group \times Origin \times Season FE		✓
Group \times Quality \times Season FE		✓
Observations	416	732
R ²	0.983	0.90

*Note: The outcome is the log number of high or low quality SKUs. $natural_{mg}$ is an indicator equal to 1 for high quality products in group g , $nonrus_{gr}$ 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 group \times origin level to allow for serial correlation across time. ***, **, * indicate significance at the 0.1%, 1% and 5% levels, respectively.*

where $\beta_2 \neq 0$ would indicate a differential effect of the exchange rate on the coefficient of variation 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.6 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.

Table B.6: No change in within-brand-fabric price dispersion

	<i>Dependent variable:</i>	
	CV(p) (1)	CV(cog) (2)
$\log(ER_{t-1})$	-0.006 (0.012)	-0.006 (0.013)
$\log(ER_{t-1}) \cdot Nat$	-0.016 (0.014)	-0.012 (0.015)
$\log(ER_{t-1}) \cdot Rus$	-0.010** (0.004)	-0.012*** (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 B.4 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.*

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:

$$\Delta \bar{p}_{i,t} = \beta_1 \Delta_{\tau_1} e_t + \beta_2 \Delta_{\tau_2} e_{t-\tau_1} + \eta_i + \epsilon_{i,t} \quad (\text{B.5})$$

where i indexes the SKU, t stands for the date, the outcome variable, $\Delta \bar{p}_{i,t}$, 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 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

Table B.7: Within-SKU pass-through

	<i>Dependent variable: $\Delta \log(p_{i,t})$</i>			
	(1)	(2)	(3)	(4)
$\Delta_{\tau_1} \text{ usdrub}_{i,t}$	0.993*** [0.279]	0.921** [0.409]		
$\Delta_{\tau_2} \text{ usdrub}_{i,t-\tau_1}$	0.649*** [0.203]	0.553 [0.410]		
$\Delta_{\tau_1} \text{ usdrub}_{i,t} \cdot \text{Nat}$		0.894 [0.975]		
$\Delta_{\tau_2} \text{ usdrub}_{i,t-\tau_1} \cdot \text{Nat}$		-0.410 [0.923]		
$\Delta_{\tau_1} \text{ eurrub}_{i,t}$			0.500* [0.270]	0.383 [0.383]
$\Delta_{\tau_2} \text{ eurrub}_{i,t-\tau_1}$			0.461** [0.217]	0.190 [0.437]
$\Delta_{\tau_1} \text{ eurrub}_{i,t} \cdot \text{Nat}$				0.948 [0.766]
$\Delta_{\tau_2} \text{ eurrub}_{i,t-\tau_1} \cdot \text{Nat}$				-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 B.5. 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.*

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.7 reports the results from estimations of regression B.5. 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.

Differential quantity reduction

We test whether there was a differential reduction in shares for high quality goods. Using similar units of observation as for the above pass-through regressions, at the material-group-season level, we run:

$$\log(q_{mgt}) = \sum_t \delta_t (Nat_{mg} \cdot \mathbf{D}_t) + \sum_{mg} \alpha_{mg} \mathbf{D}_{mg} + \sum_{gt} \alpha_{gt} \mathbf{D}_{gt} + \epsilon_{mgt} \quad (\text{B.6})$$

where q_{mgt} is the aggregate quantity sold of material m , product group g , 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 B.2 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 differential demand responses play 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 B.2 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.




Figure B.2: **Differential quantity reduction**

Note: This figure plots the estimated δ_t coefficients of equation B.6 with 95% confidence intervals around them. Fixed effects are at the product group \times material and product group \times season level. Standard errors are clustered at product group \times material level to allow within-group-material serial correlation. Results are similar when only using a season, instead of group \times season fixed effect.

B.4 Demand channel robustness

Prices as outcome variables

Regression model 5 is estimated using the median and mean regular prices in region i at time t as the outcome variables instead. The results are displayed in Figure B.3. Again, the parallel trends assumption seems to hold. The estimated $\hat{\delta}_t$ are somewhat more volatile but insignificant, and not moving in the expected positive direction. This suggests that product quality downgrading is driven by an endogenous amplification channel on the part of the firm rather than by an income-induced “flight from quality” phenomenon originating from consumers.




Figure B.3: **Income effect**

Note: This figure plots the estimated δ_t coefficients of equation 5 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.

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_{mt}^e , 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 $\partial P_m / \partial \gamma$:

$$\frac{\partial P_m}{\partial \gamma} = -c_m \cdot s_m(2 - s_m) \cdot P_m(1 - P_m) \quad (\text{C.1})$$

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 consider when the ratio of elasticities will be less than one:

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

Using $s_h > s_\ell \Leftrightarrow s_h(2 - s_h) > s_\ell(2 - s_\ell)$ for $s_h, s_\ell \in (0, 1)$ and the logit structure of demand, we have

$$\frac{\mathcal{E}_{h\gamma}}{\mathcal{E}_{\ell\gamma}} < 1 \Leftrightarrow \frac{c_h \exp(q_h + \alpha p_h^*)}{c_\ell \exp(q_\ell + \alpha p_\ell^*)} \frac{1 - P_h^*}{1 - P_\ell^*} > 1 \quad (\text{C.2})$$

The primitives are c_m, q_m and f_m ; the marginal costs, qualities, and fixed costs of providing each of the qualities $m \in \{\ell, h\}$, respectively. Although there is no closed form solution for p_m^* , $\partial p_m^* / \partial c_m > 0$ and $\partial P_m^* / \partial c_m < 0$.

It is clear that since there are many degrees of freedom (i.e., many primitives that can be varied), equation C.2 can be easily satisfied: for instance, with $q_h = q_\ell, c_h = c_\ell$, and $f_h < f_\ell$.

We are interested in whether it is possible for the high price, high quality good ($c_h > c_\ell$ and $q_h > q_\ell$) to be ex ante more profitable ($\pi_h > \pi_\ell$), but nonetheless dropped at a faster rate after the shock, which will occur if equation C.2 is satisfied. To do so requires finding only one set of parameters at which this is true, which is straightforward. For $f_h = f_\ell = 1, c_h = 3 > c_\ell = 2$, and $q_h = -8 > q_\ell = -9$, we plot profit and entry probabilities for high and low quality products as a function of the exchange rate parameter γ (as γ increases, the ruble devalues and marginal costs increase). We also assume a total market size M of 5,000,000 units, total potential entrants as 50,000, $\alpha = -0.32$, and $\sigma_\epsilon = 1$ which are in line with our later structural estimation.

C.1 CES demand with free entry

In this section we show that a CES demand model with fixed expenditure shares can match that (1) there is no differential pass-through for high-quality goods, and (2) quantities of high quality




Figure C.1: Simulated profit for high and low products
Note: This figure plots simulated profits for high and low quality products.




Figure C.2: Simulated entry probabilities for high and low products
Note: This figure plots simulated entry probabilities for high and low quality products.

goods are more responsive to a price change, as in our empirical results. However, we show that this model cannot replicate the differential reduction of high quality goods after the exchange

rate shock. A reallocation of expenditure shares as in [Bems and di Giovanni \(2016\)](#) is necessary to achieve quality downgrading.

Suppose we have a Cobb-Douglas utility function with CES aggregators over varieties of high and low quality products:

$$U = \left(\int_0^{N_h} m_h(\omega)^{\sigma_h} q_h(\omega)^{\frac{\sigma_h-1}{\sigma_h}} \partial\omega \right)^{\alpha \cdot \frac{\sigma_h}{\sigma_h-1}} \left(\int_0^{N_\ell} m_\ell(\omega)^{\sigma_\ell} q_\ell(\omega)^{\frac{\sigma_\ell-1}{\sigma_\ell}} \partial\omega \right)^{(1-\alpha) \cdot \frac{\sigma_\ell}{\sigma_\ell-1}}$$

where N_h is the number of h goods, m_h is the quality shifter for h goods, σ_h is the elasticity of substitution between horizontally differentiated varieties of h , and α is the fixed income share going to type h goods (similarly for ℓ variables). Suppose further that in the CES aggregators, $\sigma_h > \sigma_\ell$ as in [Medina \(2018\)](#), so that demand for the high quality product is more elastic with respect to price.⁴⁵

If a representative consumer maximizes with respect to a budget constraint, we recover the following demand functions:

$$q_h = \frac{\alpha Y m_h^{\sigma_h} p_h^{-\sigma_h}}{P_h^{1-\sigma_h}}, \quad q_\ell = \frac{(1-\alpha) Y m_\ell^{\sigma_\ell} p_\ell^{-\sigma_\ell}}{P_\ell^{1-\sigma_\ell}}$$

where $P_h^{1-\sigma_h} = N_h p_h^{1-\sigma_h} m_h^{\sigma_h}$ with symmetric products (similarly for ℓ). Notice that the elasticity of demand with respect to price is $-\sigma_h$ for a high quality good and $-\sigma_\ell$ for a low quality good, so given $\sigma_h > \sigma_\ell$ demand is more responsive for high quality goods, as in our empirical results. Moreover, the elasticity of price with respect to the cost shock is equal for both high and low quality goods, since optimal prices are a multiplicative markup of costs:

$$p_h = \frac{\sigma_h}{(\sigma_h - 1)} \gamma c_h, \quad p_\ell = \frac{\sigma_\ell}{(\sigma_\ell - 1)} \gamma c_\ell$$

Solving for $\pi_h = q_h(p_h - \gamma c_h) - f_h$ and setting it equal to zero, where γ is the cost shifter, yields the equilibrium number of firms with free entry (similarly for ℓ):

$$N_h = \frac{\alpha Y}{\sigma_h f_h}, \quad N_\ell = \frac{(1-\alpha) Y}{\sigma_\ell f_\ell},$$

which does not depend on γ . Not only is there no differential reduction in N_h in response to a

⁴⁵[Medina \(2018\)](#) uses CES aggregators, but they are linearly additive instead of entering in a Cobb-Douglas format. This allows expenditure shares to vary with income, which is key for her result.

cost increase $\gamma > 1$, the number of firms of both types is completely flat in γ .

C.2 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.

C.3 Expected profit approximation

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

$$\mathbb{E}[\pi_m^v(\mathbf{a}_{-jt}, \cdot)] = \sum_{N_\ell, N_h | N_{\ell t} + N_{ht} \leq \tilde{N}_t} \frac{\tilde{N}_t!}{N_{\ell t}! N_{ht}! (\tilde{N}_t - N_{\ell t} - N_{ht})!} P_{\ell t}^{N_{\ell t}} P_{ht}^{N_{ht}} (1 - P_{\ell t} - P_{ht})^{\tilde{N}_t - N_{\ell t} - N_{ht}} \cdot \pi_m^v(N_{ht}, N_{\ell t}, \cdot)$$




Figure C.3: **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.

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

$$\mathbb{E}[\pi_m^v(\mathbf{a}_{-jt}, \cdot)] \approx \pi_m^v(\tilde{N}_t \hat{P}_{ht}, \tilde{N}_t \hat{P}_{lt}, \cdot),$$

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.4 Model fit and counterfactuals




Figure C.4: **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.5: **Welfare loss with alternative parameters**

Note: This figure plots the welfare cost of the devaluation for various quality demand shifters, assuming the cost of high quality is 2.7 times that of low quality products, whose cost is held fixed at the estimated level. A value of one on the x-axis means no demand increase for high cost goods.