PRODUCT INTRODUCTIONS, CURRENCY UNIONS, AND THE REAL EXCHANGE
RATE

Alberto Cavallo
Brent Neiman
Roberto Rigobon

ABSTRACT

We use a novel dataset of online prices of identical goods sold by four large global retailers in dozens of countries to study good-level real exchange rates and their aggregated behavior. First, in contrast to the prior literature, we demonstrate that the law of one price holds very well within currency unions for thousands of goods sold by each of the retailers, implying good-level real exchange rates equal to one. Prices of these same goods exhibit large deviations from the law of one price outside of currency unions, even when the nominal exchange rate is pegged. This clarifies that it is the common currency per se, rather than the lack of nominal volatility, that results in the lack of cross-country differences in the prices of these goods. Second, we use a novel decomposition to show that most of the cross-sectional variation in good-level real exchange rates reflects differences in prices at the time products are first introduced, as opposed to the component emerging from heterogeneous passthrough or from nominal rigidities during the life of the good. In fact, international relative prices measured at the time of introduction move together with the nominal exchange rate. This stands in sharp contrast to pricing behavior in models where all price rigidity for any given good is due simply to costly price adjustment for that good.

Alberto Cavallo
MIT - Sloan School of Management
100 Main Street, E62-512
Cambridge, MA 02138
acavallo@mit.edu

Brent Neiman
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
brent.neiman@chicagobooth.edu

Roberto Rigobon
MIT Sloan School of Management
100 Main Street, E62-516
Cambridge, MA 02142
and NBER
rigobon@mit.edu

1 Introduction

For hundreds of years, international economists have taken great interest in cross-country differences in the prices of identical goods (or baskets of goods) when translated into a common currency. The "Law of One Price" (LOP) for traded goods across countries is a fundamental building block of standard models in open-economy macroeconomics. Minor deviations from the LOP are not surprising in a world with barriers to arbitrage such as transport costs. A large literature, however, documents its surprisingly large failure for many traded goods and tries to explain the resulting volatility in the relative price of consumption across countries, or the real exchange rate (RER).[1] This paper uses a novel dataset of online prices for identical traded goods sold in several dozen countries to shed light on the determinants of good-level and aggregate RERs and their dynamics.

We demonstrate that the LOP holds very well within the euro zone for tens of thousands of goods, implying traded RERs approximately equal to one. We show this holds for four different global retailers in three unrelated industries. To the best of our knowledge, this is the first documentation of the LOP holding internationally for a wide variety of differentiated goods, and we show it holds across multiple countries with different and varying tax rates. Physical distance, political and tax territories, language, and culture are all often thought of as forces that segment markets. Our results imply, by contrast, that the choice of currency units is far more important for defining the boundaries between markets for some goods.

Deviations from the LOP are significantly larger for these same products for countries with different currencies, even if their nominal exchange rate (NER) is pegged. For example, prices in the euro zone typically differ from those in Sweden, which has a floating exchange rate, and also from those in Denmark, which pegs its currency to the euro. This clarifies that the common currency per se, and not simply the lack of nominal volatility, is important in reducing cross-country price dispersion. We complement this evidence by showing that the LOP with the United States holds more for dollarized countries like Ecuador and El Salvador than for countries like Hong Kong or Jordan, which have their own currency but peg it to the U.S. dollar.

We introduce a framework to decompose the good-level RER into the RER at the time a good is introduced, a component reflecting price stickiness together with NER volatility, and a residual component due to heterogeneous passthrough which we refer to as reflecting changes in demand.

---

[1]Cassel (1918) first used the term "Purchasing Power Parity" (PPP) to describe the condition in which there are no such cross-country differences in the price of consumption and therefore the RER equals one. See Rogoff (1996) for a history and overview of the high persistence and volatility of the RER, what has been termed the "PPP puzzle."

We find that the majority of LOP deviations occur at the time a good is introduced, rather than emerging subsequently due to price changes or due to price stickiness and NER movements. As a corollary, typical measures of the traded RER that are constructed using only price changes may differ significantly from the underlying object they are designed to capture.

Given the importance of the good-level RER at the time of product introduction, we next study how relative introduction prices evolve with the NER. We find that the RER at the time of introduction largely moves together with the nominal rate. This is evidence against a model in which previous temporary shocks to the RER are fully eliminated at the time of a price change as, after all, a price change inherently occurs when a new product is introduced.

These results are important for a variety of reasons and are relevant for multiple research areas. First, they shed light on the determinants of market segmentation and the pricing problem faced by international retailers. Second, they improve our understanding of traded RER dynamics and carry significant policy implications. For example, the theory of optimal currency areas stresses that a common currency for two countries makes more sense when inflationary shocks in those countries are more synchronized. Our results suggest this synchronization may emerge endogenously to the choice of currency regime.[2] Relatedly, because traded good prices within a common currency area may respond less to country-specific shocks, our results are informative about the nature and efficacy of "internal devaluations." Third, our finding that NERs and RERs move together even at the time of product introduction suggests that local currency pricing may be the most appropriate modeling assumption, even for periods of time longer than the life of a typical product.[3] This result argues in support of models with variable flexible price markups and stands in sharp contrast to the pricing behavior in models where all price rigidity for any given good is due simply to costly price adjustment for that good. In this sense, our results are also important for closed-economy macroeconomic models aiming to understand pricing dynamics and monetary non-neutrality.

Our results are also suggestive of the promise of incorporating consumer psychology, firm organizational frictions, and market norms into macro models of price setting. Firms may equalize prices within currency unions but not outside of them, including in pegged regimes, for fear of

---

[2]Before the euro's introduction in 1999, popular discussion and academic research on its potential impact often focused on increased competition and the cross-country convergence of prices. For example, Goldberg and Verboven (2005) find some evidence of convergence in auto prices after the introduction of the euro, while Parsley and Wei (2008) and Engel and Rogers (2004) do not find such evidence in price data on the Big Mac and other consumer goods. Our results do not on their own indicate whether or not welfare in the euro zone is higher due to the equality of prices. In a model with heterogeneous demand and markups across countries, it is unclear what the removal of barriers to arbitrage implies for overall welfare.

[3]Our empirical results offer further motivation for Berka, Devereux, and Engel (2012), which argues that local currency pricing undermines traditional arguments made in favor of flexible exchange rates.

angering customers that can easily compare prices across borders. In this sense, the Internet itself, which facilitates such price comparisons, may have played an important role in generating these pricing strategies. Alternatively, perhaps establishing pricing departments within large multinational firms involves a very large fixed cost. Firms therefore treat countries that peg their exchange rates differently from those in a currency union because there is a greater likelihood of a future regime change that would then require creating a new pricing procedure. None of these possibilities can on their own explain all of our results, but many of our findings argue that these considerations should feature more prominently in the macroeconomics price-setting literature.

Our data include daily prices for all products sold by Apple, IKEA, H&M, and Zara through their online retail stores in each country. We use the last recorded price in each week to form a weekly dataset that spans various subsets of 85 countries and various time periods from October 2008 to May 2013. At the time of writing, Apple, an American company, is the world's largest company by market capitalization. According to the research firm Euromonitor International (the source for all the market research described in this paragraph), Apple accounted in 2011 for 5.4 percent of the $800 billion global consumer electronics market. This makes it the third largest global firm by sales in that industry. Since at least 2007, more than half of Apple's total retail sales came from online sales. IKEA was founded in Sweden and is the world's largest furniture retailer, accounting for 4.9 percent of the $500 billion global furniture market. H&M, also a Swedish company, and Zara, from Spain, are the world's fourth and third largest clothing retailers respectively. The combined sales of H&M and Zara exceed $30 billion globally.

The pricing patterns we identify cannot be oddities associated with a particular firm's, industry's, or country's characteristics. These companies are among the world's very largest retailers, are headquartered in three different countries, and cover three different industries which account for more than 20 percent of U.S. consumption expenditures in goods. Further, we corroborate using a smaller cross-section of prices that the same qualitative relationship between price dispersion and currency regime also holds for four additional large global retailers in apparel and technology: Adidas, Dell, Mango, and Nike. Together, this gives us confidence that inference from our data is appropriately applied to the broader basket of branded and traded goods and is relevant for understanding international macroeconomic dynamics.

Studying online prices has the obvious advantage of allowing for the collection of enormous amounts of data at very high frequency. Online sales already represent a large and growing share of total global consumption, but we believe our results are no less informative even if a reader cares only about offline sales for these stores. We provide strong evidence that online prices are fully representative of offline prices for all of our goods. The customer service departments for all

four companies state that the online and offline prices are identical up to shipping costs and, in limited instances, local taxes or store-specific special promotions. We also visited several physical retail stores in the United States to confirm this to be the case.

Related Literature

Our work builds on a long literature studying sources of RER movements and relating this movement to the choice of currency regime. Mussa (1986), using aggregate price indices, showed that RER volatility increased markedly with the breakdown of the Bretton Woods system of fixed exchange rates. Engel (1999) demonstrated that movements in the RER did not reflect the relative price within countries of traded and non-traded goods, as in Balassa (1964) and Samuelson (1964). Rather, Engel showed that the bulk of RER volatility comes from movements in the traded-good component, a striking result that holds at horizons ranging from 1 month to 30 years.[4] Motivated in part by this result, much of the literature has focused on explanations for RER movements among traded goods, and we follow in this tradition.

Many papers have focused on the LOP deviations that emerge among traded goods due to movement in the NER in models with price stickiness, as in Devereux, Engel, and Storgaard (2004) and Devereux and Engel (2007). Crucini, Shintani, and Tsuruga (2010) adds sticky information to a sticky price model to match the persistence of good-level LOP deviations. Others have focused on models with exchange rate passthrough and pricing to market even after prices change, including Atkeson and Burstein (2008) and Gopinath, Itskhoki, and Rigobon (2010).

Some papers have looked at disaggregated price data, including in levels. Gopinath, Gourinchas, Hsieh, and Li (2011), Broda and Weinstein (2008), and Burstein and Jaimovich (2009) document large cross-country price differences for identical goods sold in United States and Canada.[5] Crucini, Telmer, and Zachariadis (2005) examined prices across Europe from 1975-1990 for several thousand narrowly defined categories of goods such as "Dried almonds" or a "Record player". They conclude that the distribution of LOP deviations are generally centered around zero and increase in dispersion the less tradable the good is and the more non-traded inputs are used to produce the good. Crucini and Shintani (2008) use similar data to find that the persistence of LOP deviations in the cross-section increases with the importance of the distribution margin.

Finally, Baxter and Landry (2012) also study IKEA products by using 16 years of catalog prices in 6 countries. They detail a rich set of statistics on prices, passthrough, and product creation

---

[4]See also Rogers and Jenkins (1995), which emphasizes the larger role of LOP deviations in the traded sector compared with the relative price of traded and non-traded goods.

[5]A large literature focuses on the contribution international borders make to price dispersion. See, for example, Parsley and Wei (2001) and Engel and Rogers (1996) as well as more recent work including Gorodnichenko and Tesar (2009), Borraz, Cavallo, Rigobon, and Zipitria (2012), and Cosar, Grieco, and Tintelnot (2012).

and destruction. They do not report our finding that the law of one price is far more likely to hold within currency unions, but they do find that LOP deviations for France and Germany move closely together. They additionally note that the scale of LOP deviations is similar for new and continuing goods and use annual data to demonstrate a large covariance between the nominal exchange rate and real exchange rate for goods that are new in both countries that year. These facts are closely related to and consistent with our documentation of the comovement of the RER and NER at the time of product introductions. The implications of this comovement for the measurement of RERs closely relates to Nakamura and Steinsson (2012) and Gagnon, Mandel, and Vigfusson (2012).

2   Scraped Online Prices from Global Retailers

Our dataset is comprised of prices scraped off the Internet by The Billion Prices Project, an academic research initiative at MIT. These pricing data are introduced in Cavallo (2012) and also used in Cavallo and Rigobon (2012), though neither paper compares the prices of identical goods across multiple countries, the focus of this paper. The bulk of our analyses study prices for four global retailers for which we can precisely match prices of identical goods sold in many geographies. We are able to match over one hundred thousand unique items across dozens of countries because the firms' web pages organize products using their own company-wide product ID codes.

The data include daily prices for the four retailers in some subset of 85 countries during some subset of the period from October 2008 to May 2013. Prices are generally quoted inclusive of taxes and exclusive of within-country shipping costs.[6] Table 1 gives a basic description of the country, product, and time coverage in our data. Row (i) indicates that we track prices for over 100,000 products, including more than 11,000 for Apple, 66,000 for IKEA, 14,000 for H&M, and 22,000 for Zara during varying subperiods of the time ranges listed in row (iv). IKEA has significantly more products than the other retailers and Zara covers significantly more countries. Subject to occasional errors in our scraping algorithm, our dataset includes all products sold online by these stores for the relevant countries and time periods.[7]

---

[6]The United States and Canada are the exceptions. We adjust all U.S. and Canadian prices upward by 6.5 and 12 percent, respectively, to reflect the average local tax rates in 2012. We obtain information on state sales tax rates for the United States from The Tax Foundation, on province sales tax rates for Canada from the TMF group, and on VAT rates for other countries from Deloitte. For countries other than the United States and Canada, the same sales (or value added) tax typically applies throughout the entire country. Shipping costs do not appear to be a hidden source of price discrimination. A limited test we conducted for Apple products, for example, showed that the same shipping cost was charged for all destinations within the euro zone.

[7]There are a number of small countries in which Zara has local stores and a country-specific web page, but

Scraping errors or changes in these companies' web pages occasionally create missing price observations. We fill in the gaps between observed prices using the assumption that prices remained unchanged during the missing period. In absence of scraping errors and idiosyncratic in-store specials, our prices include sale prices. The procedure for filling in missing observations may therefore cause us to miss sale prices which then reverted to their same initial values. This concern applies differently for Apple and IKEA compared to the apparel retailers. Apple and IKEA have the most historical data and the longest periods with scraping gaps. In principle, therefore, our treatment of gaps in the data could be of greatest concern for these companies. However, as we elaborate below, their prices are highly sticky and do not typically exhibit high frequency sales behavior seen in other pricing contexts. Though the prices of long-lived apparel items are far less sticky (for example, roughly all apparel items in our data for more than a year have experienced at least one price change), clothing items typically have durations in our data about one-fifth as long as the non-apparel items, reflecting both more recent scraping of their web pages and the importance of seasonality in apparel. As such, though the procedure for filling in gaps in the data may incorrectly omit some temporary price changes, the impact of these errors would likely be short-lived.

We do not have purchase quantities or individual product weights, so all our quantitative analyses apportion equal weight to all goods within each retailer and country pair and equal weight to all country pairs within each retailer. When we aggregate across retailers, we give equal weight to each available store-pair combination. For example, if a country pair has twice as many IKEA goods as Apple products, then each individual IKEA price would be treated as containing less information for that country pair than each individual Apple price. When we aggregate across countries and pairs, however, one store may be given more total weight as its products may be available for more bilateral country pairs. We exclude the roughly 2 percent of goods for which we observe implausibly large price changes or for which the good's relative price across countries is implausibly large. Additional details on the data coverage, web-scraping process, and our assembly and cleaning of the data are included in the Appendix, which can be found on the authors' web pages.

Relative to prior studies that use manufacturing or traded good price indices to understand RER levels or movements, our dataset offers several clear benefits. First, by matching the identical product, we avoid the concern that RER movements misleadingly reflect heterogeneity in the basket of goods or biases that emerge due to the aggregation across goods as highlighted in

_____

does not actually make online sales. Zara representatives confirmed that the online prices also equal the offline prices in these countries.

Imbs, Mumtaz, Ravn, and Rey (2005). Second, by comparing the same product and retailer combination, we can distinguish cross-country pricing differences from cross-chain pricing differences, which Nakamura, Nakamura, and Nakamura (2011) argue explains a large share of total variation in price dynamics. Third, by observing price levels at the date of introduction we are able to reveal what turns out to be the largest component of the RER in our data, a component which is by definition ignored by matched model price indices that are constructed only using observed price changes for continuing goods. Finally, in measuring prices at a very high frequency, we can more confidently pinpoint the role of nominal rigidity in contributing to the RER.

Compared to prior work that matched goods at the barcode level, we emphasize that our data allow for significantly more cross-country comparisons, variation that proves essential to uncover our results on the role of currency unions. Further, a typical large bilateral country pair in our data will have half of the total products available across both countries also available in each country, which gives confidence that composition differences are not driving our key results. By comparison, these previous studies typically match less than 5 percent of the total goods.

There are three primary concerns that may arise from our focus on the online prices of four retailers. First, one might reasonably worry that prices posted online differ from prices paid in physical stores and outlets. Internet transactions are not only a large and growing share of the market, but online prices are highly representative of offline prices. We contacted each of the companies over email or by phone and received confirmation that online and physical stores have identical prices for all four retailers, with only occasional deviations for in-store specials. We also checked this ourselves by sending research assistants to two Apple stores, two H&M stores, two Zara stores, and the only IKEA store near Boston and confirmed for 10 randomly selected items in each store that the online and offline prices (excluding taxes) were exactly identical.[8] There is strong direct and indirect evidence that Internet prices in our data are highly representative of, and typically identical to, prices in physical stores.

Second, one might believe that these companies are unusual in that they typically list a single price per country, a finding in juxtaposition to much of the literature on price dispersion. In fact, if one excludes groceries and related products, the policy of offering a single price per country is the norm for large online retailers. Of the 10 largest U.S. retailers, only Walgreens and Walmart use zip codes to localize prices shown over the Internet.[9] We scraped the webpages of these two

---

[8]In fact, this also held true for the one item in those 10 from IKEA which happened to be selling at a discount. Figure 1(a) shows a screen shot of that product on IKEA's U.S. website, a "HEMNES coffee table, gray-brown." The price is clearly marked as $99.99, and one can see the previous higher price of $119.00 listed above the new price and crossed off with a black line. Figure 1(b) shows a photograph of the price tag of the identical object found in a physical IKEA store on the same day, listed at the same $99.99 price.

[9]According to Deloitte (2013), the remaining 8 "top-10" U.S. retailers include Kroger, Costco, Home Depot,

retailers and found that only 15 percent of the items sold by Walgreens are labeled with "prices vary by store," implying that at least 85 percent of Walgreens' products sold online have the same price throughout the United States. A similar analysis suggests that more than 90 percent of items sold online by Walmart have the same price throughout the country. We further looked at a list of the top 100 retailers by U.S. revenues according to the U.S. National Retail Federation and found that of the 70 retailers that sell online, only 21 require customers to enter a zip code to see local prices, and 13 of these 21 are grocery stores.

Third, one may wonder about the extent to which our results are representative of the broader basket of tradable goods. Our data may not allow us to better understand the cross-country pricing policies of small firms, to study pricing of commodities or intermediate goods, or to learn about practices in, say, the automobile market. Branded technology, furniture, and apparel goods, however, are particularly interesting to study because they are often produced in one plant or location, are sold in many countries by the same retailer, and exhibit significant price stickiness relative to homogeneous goods.[10] The companies included in our data are among the very largest technology, furniture, and clothing companies in the world and on their own might constitute a non-trivial share of total expenditures on traded goods. Further, we demonstrate with more limited data that the key qualitative patterns we identify also hold in prices from other large global retailers in these sectors including Adidas, Dell, Mango, and Nike. Using CPI weights from the U.S. Bureau of Labor Statistics, we calculate that these three industries cover more than 20 percent of final consumption expenditures on goods.

3   The Law of One Price and Exchange Rate Regimes

We now introduce notation and define good-level RERs. We start by showing that log good-level RERs deviate significantly from zero (i.e. LOP fails) outside of currency unions, even when the NER is pegged. By contrast, we then demonstrate that the law of one price does very well within the euro zone and frequently holds even when comparing prices in dollarized economies with those in the United States. We show this result is robust to restrictions on the value of the goods, to conditioning on other plausible drivers of good-level RERs, to consideration of goods that are not

---

Target, CVS, Best Buy, Lowe's, and Amazon.com. None of these companies requests zip code information before quoting online prices (Kroger is the only one that does not offer any prices online). We also matched online and offline prices for about 200 total items from these stores and found that a bit more than half had identical prices and most deviations were quite small.

[10]IKEA, H&M, and Zara generally sell their own branded goods. Of all the products sold by Apple, however, only a (highly visible) minority are Apple branded items. We demonstrate below that our results also hold for Apple's sales of non-Apple brands.

branded by the retailer, and to the use of prices from several additional retailers for which we have a more limited set of data.

Let $p_i(z, t)$ denote the log price in local currency of good z in country i. We define $e_{ij}(t)$ to be the log of the value of one unit of country j's currency translated into country i's currency. The log good-level RER $q_{ij}(t)$ is defined as the difference between prices in countries i and j after being translated into a common currency:

$$q_{ij}(z, t) = p_i(z, t) - e_{ij}(t) - p_j(z, t).$$

It equals zero when the LOP holds exactly.

Figure 2 plots the distribution of the log good-level RERs q_{ij} pooling all goods z and weeks t for various countries i with country j fixed as the United States. The y-axes indicate the percentage of observations corresponding to each bin of x-axis values. Values concentrated around zero indicate goods which, after being translated into common currencies, have the same price. The histograms include all available weekly relative prices in our dataset other than those exceeding 0.75 log point in magnitude, a set of outliers typically representing about one percent of the total prices. Frequency weights are used so that the total contribution of goods from each store is equalized. The vertical red dotted line indicates the average value (using these same weights) of q_{ij} across all products.

While patterns vary across bilateral relationships, the scale and frequency of LOP deviations are striking. Even when comparing identical branded and tradable products sold by the same firm, one routinely finds goods with prices in other countries that differ from the U.S. price by 0.25 log point or more. The distributions are generally centered near zero, but it is not uncommon to find countries like Japan where prices average nearly 20 percent more than prices in the United States. Note that even in Canada or China, whose NERs with the U.S. dollar have been relatively stable, good-level log RERs diverge significantly from zero. Across all the bilaterals, no individual RER bin accounts for more than 10 percent of the total observations.

Figure 3 shows these same histograms but separately for each of the stores and demonstrates that these patterns are broadly representative. Some bilateral pairings, such as Mexico and the United States for IKEA, are missing due to lack of store data or matched price observations. There are pricing differences across stores, and the dispersion in good-level RERs clearly seems largest for IKEA and smallest in the apparel companies. All, however, exhibit significant deviations from the LOP and share other common regularities such as the higher average prices in Japan.

By contrast, we find compelling evidence that the LOP generally holds in European countries

that share a single currency and also holds quite frequently between countries that use U.S. dollars as their domestic currency.

3.1 The Euro Zone

In Figure 4 we plot the distribution of the log good-level RERs for many European countries (plus the United States) relative to Spain. Together with Spain, countries including Austria, Germany, Finland, France, Ireland, Italy, Netherlands, and Portugal are members in the euro zone, a single currency area. The prices for most of the tens of thousands of distinct products sold in those countries are identical and we therefore see huge mass at zero in these histograms (note the differences in scales of the y-axes).[11] This is the first evidence that we are aware of documenting the LOP holding across countries for a wide variety of identical traded differentiated goods.

Prices in Denmark, Norway, Sweden, and Switzerland (not shown), by contrast, do not exhibit this same adherence to the LOP. These countries are outside of the euro zone, and their histograms look similar to that of the United States. This is despite the fact that they are all parts of continental Europe with similar geographies and demand structures. Denmark and Sweden are members of the European Union and are subject to the same tariffs and product market regulations as is Spain. In addition, Denmark has a strong peg against the euro. This demonstrates that being in the euro zone per se, rather than simply eliminating nominal volatility, is important for good-level RERs.

A large share of these goods are likely produced in a single plant at the same marginal cost.[12] Therefore, the dispersion of good-level RERs in Figures 2 and 4 suggest that companies price to market and have desired markups which differ significantly across countries, even across similar countries like Sweden and Finland. However, companies forgo this markup variation within the euro zone.[13] This implies that the crucial factor defining companies' abilities to charge different

---

[11]It is not the case that consumers in one country can simply order directly from another country's web page. If a shipping address in Madrid is inputted into Apple's German webpage, for example, the customer is either automatically re-routed to Apple's Spanish webpage or is simply not permitted to enter Spain as the country of the delivery address.

[12]For example, Apple's 2011 annual report states (on page 9) that "substantially all of the Company's hardware products are manufactured by outsourcing partners primarily located in Asia."

[13]We reiterate that these prices are inclusive of sales taxes, which exhibit variation across time and space, further implying that companies are forgoing otherwise desirable markup variation within the euro zone. Prices inclusive of tax rates are generally identical in the euro zone even though value added tax rates varied from 19 percent in Germany to 23 percent in Portugal. Similarly, there have been many tax changes in our data, such as Spain's increase from 16 percent in 2008 to 18 percent in 2010 and to 21 percent at the end of our data. These country-specific changes do not appear to have produced changes in the degree to which the LOP held for prices (inclusive of taxes) in the euro zone.

prices may not be shipping frictions, national border effects, or cultural or regulatory boundaries. After all, the differences in physical, cultural, and political distance between Spain and Finland seem highly similar to these differences between Spain and Sweden or Denmark. Rather, it implies that companies believe that having to consider a price with different currency units is the most salient friction, even if the different currency units can be translated at a fixed rate as with the Danish krone and the euro.

We use Spain as the base country because we have more data for it across the four stores we study. Zara, however, divides the euro zone into two regions: one with Spain and Portugal and the other with the remaining euro zone countries other than Greece and Andorra. The LOP generally holds within each of those regions, though prices differ by about 25 percent between the regions (they are lower in Spain/Portugal). This is why there are similar masses of LOP deviations near 0.25 log point in the histograms for most euro zone countries in Figure 4. In this sense, Figure 4 if anything understates the degree to which prices are equalized within the euro zone.

Figure 5 shows that this phenomena is not specific to a particular store and in fact holds for all four of the retailers. The LOP holds very well for tens of thousands of goods sold by Apple, IKEA, and H&M across most of the euro zone. We split the results for Zara into its two euro zone pricing blocks. The left two columns of Figure 5(d) underneath the text "vs. Spain" shows that the LOP holds perfectly for Zara between Spain and Portugal. The right two columns underneath the text "vs. Germany" shows that the LOP holds perfectly for Zara between Germany and euro zone countries other than Spain and Portugal. When data are available comparing prices in Spain and Norway or Sweden or Denmark, however, the LOP never holds to a meaningful extent and the distributions generally look similar to those between Spain and the United States.

We note that, conveniently, this result corroborates our matching algorithm and reduces concern about measurement error. One might have worried that the huge dispersion in good-level RERs between the United States and Spain followed simply from the difficulty in matching identical products. The fact that LOP holds so frequently for the bulk of these products within the euro zone would be too coincidental if these were not in fact identical goods.

3.2 Dollarization

Given the large quantity of data on prices of multiple retailers in Europe, we view the results for the euro zone as the most robust demonstration of the importance of currency unions for LOP. After seeing these results, though, the natural question is whether the euro zone itself is critical for LOP as opposed to common currency areas more generally. We now present results

comparing dollarized countries (i.e. countries that use the U.S. dollar as their currency) with countries that have their own currencies which are pegged to the U.S. dollar. We demonstrate, consistent with the euro zone results, that LOP holds significantly better between dollarized countries than between dollar pegs.

In particular, we compare the distribution of good-level RERs with the United States for Ecuador and El Salvador, countries where prices are quoted and goods transacted in U.S. dollars, with the equivalent distributions for Bahrain, Hong Kong, Jordan, Kazakhstan, Lebanon, Oman, Panama, Qatar, Saudi Arabia, and the United Arab Emirates, countries with their own currencies that are strongly pegged to the U.S. dollar.[14] Figure 6 shows the distribution of weekly log good-level RERs for these countries relative to the United States. The histograms for Ecuador and El Salvador are the only ones that spike distinctly at zero, with substantial mass where the LOP holds almost perfectly. Of the products sold in countries with their own currencies that are pegged to the U.S. dollar, less than 10 percent have a log RER with the United States with an absolute value less than 0.01. Of those sold in Ecuador and El Salvador, more than 40 percent do.[15]

The evidence from dollarized countries corroborates the evidence from the euro zone. Currency unions have striking implications for good-level RERs that do not simply emerge due to the lack of nominal volatility.

3.3 Competition Policy

Can competition policy explain our finding that prices are generally equalized within currency unions? We find this possibility to be highly unlikely for four reasons. First, there is no European Union (EU) law requiring retail prices to be equalized across member countries.[16] According to Bailey and Whish (2012), "The Court of Justice in *United Brands v. Commission* ruled that 'it was permissible for a supplier to charge whatever local conditions of supply and demand dictate, that is to say that there is no obligation to charge a uniform price throughout the EU.'"

---

[14]Most of these countries match at least several thousand items with the United States, with Zara products typically constituting their majority (or entirety). For some of these countries, Zara's web page does not allow for online purchases but rather advertises the prices that customers would pay offline in that country. We label Panama as "Dollarized (Weaker Form)" because both the U.S. dollar and Panamanian balboa coins are legal tender, and Zara's prices are quoted there in balboas.

[15]For this analysis only, we use the U.S. price exclusive rather than inclusive of taxes to highlight that these spikes are then located at precisely zero. As we discuss below, it is intriguing that the LOP holds among euro countries inclusive of taxes and appears to hold exactly for the dollarized countries exclusive of U.S. taxes. In the quantitative work that follows, we instead use values shifted about 0.06 relative to those plotted in these histograms to account for the value after including U.S. taxes. In this sense, one might consider our subsequent quantitative analyses as understating the degree to which LOP holds in the dollarized countries.

[16]We further confirmed this by consulting antitrust lawyers and the appropriate division of the European Commission.

Second, even if firms mistakenly believed there to be such a law, it would apply at the EU level, not at the euro zone level. This would be inconsistent with our finding that all four of our companies generally charge the same price in the euro zone and a different price in Denmark and Sweden, both of which are within the EU. Third, while we show that most goods obey the LOP for most bilaterals in the euro zone, all four of the companies have a large number of exceptions. As we showed in Figure 5(d), Zara charges different prices in Spain and Portugal compared with the rest of the euro zone countries. Fourth, competition policy cannot explain our results for dollarized countries as Ecuador and El Salvador are clearly outside the jurisdiction of U.S. antitrust authorities.

3.4 A Quantitative Analysis

To summarize quantitatively our conclusions on the importance of exchange rate regime for good-level RERs, we start by characterizing the unconditional mean of the absolute value of good-level log RERs by currency regime.[17] We consider different subsets of goods based on the absolute level of prices to demonstrate that our findings are not driven by cheap items. Next, we report the percentage of items for which the LOP holds in each currency regime. Finally, we introduce other observables that likely influence relative prices and run regressions to report the conditional mean magnitude of good-level log RERs by currency regime.

3.4.1 Average Absolute Values of Good-Level RERs

We calculate the average absolute value of the log RER for each good over all weeks t, $\sum_t \frac{1}{\|t\|} |q_{ij}(z,t)|$, and report the unconditional mean of these good-level log RERs in Panel A of Table 2. As above, we use weights that equalize the contribution of each country-pair and store combination. Rows (i) to (iii) report the average values from our full sample. The first column of those rows shows that the typical magnitude of good-level log RERs equals about 8 percent within currency unions, 12 percent for pegged regimes, and 19 percent for floating exchange rate regimes. Consistent with the histograms presented earlier, countries in a currency union have significantly smaller LOP deviations than countries with their own currencies, including those with nominal pegs. Further, the scale of LOP deviations in currency unions are smaller than for pegged regimes, and significantly

---

[17]We consider as currency unions bilateral pairs among euro zone countries, Andorra, Monaco, and Montenegro as well as bilateral pairs among the United States, Ecuador, El Salvador, and Panama. Dollar pegs include Azerbaijan, Bahrain, China (before June 2010), Honduras (before June 2012), Hong Kong, Jordan, Kuwait, Lebanon, Macao, Kazakhstan, Oman, Qatar, Saudi Arabia, Syria (before September 2011), UAE, Ukraine, and Venezuela. Euro pegs include Bosnia and Herzegovinia, Bulgaria, Denmark, Lithuania, and Latvia. These include all countries in our data with exchange rate code "1" from the "course" classification in Ilzetzki, Reinhart, and Rogoff (2008).

smaller than for floating regimes, in all four stores. The gap between average good-level RERs in currency unions compared with floats equals 9 percentage points for IKEA, 10 percentage points for Zara, and 12 percentage points for Apple and H&M.

These important differences in good-level RERs are not driven by cheap items with very low price levels. In rows (iv) to (vi) we re-report these same statistics when calculated only on goods where the average price, after translating into U.S. dollars, exceeds $50, and in rows (vii) to (ix) we repeat the exercise on goods with average prices exceeding $200. The general patterns are highly robust across stores and price levels.

3.4.2 Share of Goods Obeying the Law of One Price

In rows (x) to (xii) of Panel B of Table 2, we report the share of goods with an absolute value of log RER less than 0.01, a level of similarity of prices that we refer to here as obeying the LOP. Across all currency union pairs and stores in our data, more than 60 percent of the goods satisfy the LOP, whereas only 7 percent do so for pegged pairs and less than 5 percent do so for country pairs with a floating bilateral exchange rate. Consistent with the histograms in Figure 5(d), there is some variation across stores in the degree to which currency unions generate identical prices. The effect is weakest in IKEA, where roughly one-third of prices are equalized in currency unions, and is strongest for H&M, where 90 percent of the prices are equalized in currency unions. For all four retailers, however, a substantial percentage of goods obey the LOP in pairs that share a currency, and this percentage is in all cases significantly larger than the equivalent statistic for pegged and floating regimes.

3.4.3 Regression Analyses

Next, we correlate this average absolute value of each good's log RER with indicators of the currency regime, NER volatility, and other potentially important generators of law of one price deviations (the data sources of which are detailed in the Appendix). Table 3 shows results from a regression of the average absolute value of log good-level RERs on (i) an indicator labeled "Outside of Currency Union" which equals zero for pairs in a currency union and one for others, (ii) an indicator equaling one for "Pegged NER" regimes, (iii) a variable capturing the log NER volatility experienced during the life of the good, (iv) the log bilateral distance between each country pair, (v) a variable called "Abs. Relative Income" that equals the absolute value of the log ratio of per-capita PPP GDP, and (vi) a variable called "Abs. Relative Taxes" that equals the absolute value of the difference in value-added tax rates. We also include a dummy variable for each country and for each store. We run these regressions pooling all stores as well

14

as separately for each store (and we weight to equalize the contribution of each store and country pair combination), clustering standard errors by the interaction of store and country pair.

The first column of row (i) of Table 3 reports that goods outside of currency unions, conditional on other observable differences, are expected to have log RERs with absolute values 0.124 higher than equivalent goods within currency unions. The pooled and store-specific regressions all include two columns, one which includes all covariates and country dummies and another which drops all covariates and dummies other than the exchange rate regime variables. The average increase in absolute value of good-level RERs when moving from a currency union to a floating exchange rate, reported in row (i), averages roughly 7 to 13 percent across stores. The effects are precisely estimated, with the small standard errors reported in parentheses.

If pegged exchange rate regimes had the same implications for good-level RERs as currency unions, the coefficients in row (ii) should equal the opposite of those in row (i) since pegged regimes are also considered "Outside of Currency Unions." Indeed, all the coefficients in row (ii) are negative other than one near-zero value, suggesting that LOP deviations are in fact smaller in pegged than floating regimes. The magnitudes of these estimates, however, are typically significantly smaller than those in row (i), confirming that pegged regimes look different from currency unions in terms of their impact on q_{ij}. For example, the first column with the pooled results suggests that currency unions involve LOP deviations that are about 9 percentage points smaller ($0.087 = 0.124 - 0.037$) than with a pegged regime and about 12 percentage points smaller than with a floating regime.

Several other covariates are statistically significant for "All Stores," though none have magnitudes that are economically important relative to the currency regime. Doubling exchange rate volatility on average reduces LOP deviations by about 3 percentage points. Doubling the distance between country pairs on average increases their LOP deviations by about 1.3 percentage point. Relative taxes differentials range from 0 to 0.27, which means that wider tax differentials may plausibly increase LOP deviations by a couple of percentage points. The final three columns run the regressions separately for the set of country pairs with flexible NERs, with pegged NERs, and within currency unions with each other. Outside of currency unions, distance and taxes remain statistically significant determinants of average RER magnitudes. Within currency unions, none of the covariates is statistically significant.

3.5 Goods Not Manufactured by the Distributor

Many of the goods distributed by these companies are also manufactured or branded by them. Relatedly, many of these brands are not carried by other retailers. For example, the large majority

of goods found in IKEA stores are IKEA branded goods, and one cannot buy IKEA products from other furniture retailers. We can get some insight into whether this characteristic influences these companies' pricing strategies by focusing on the case of Apple, where we can observe the distinction between the Apple branded products and the non-Apple products that they sell.[18]

More than half of all products sold by Apple are not Apple branded products. These goods include, for example, Epson printers, Michael Kors travel totes, Canon digital cameras, and various cables and adapters by brands like Apogee, Belkin, and Kanex. When we separately calculate the unconditional average absolute deviations reported in Subsection 3.4.1 for the Apple and non-Apple branded products sold by Apple, we find highly similar results. The average absolute value of log good-level RERs for currency unions equals 0.020 for non-Apple branded goods and 0.032 for Apple branded products. Pegged regimes have log good-level RERs with average absolute values of 0.082 and 0.093 respectively in these two groups, and floats have values of 0.139 and 0.147. The patterns in the share of goods for which LOP holds is also very similar for Apple and non-Apple branded products. Regardless of whether the products are branded by Apple or not, it sells with far reduced price dispersion for currency unions, even relative to pegged regimes.

3.6   Evidence from Additional Retailers

While we do not have extensive time-series data that can be matched across countries for retailers beyond Apple, IKEA, H&M, and Zara, we additionally scraped a more limited cross-section of prices for Adidas, Dell, Mango, and Nike.[19] Panels A through C of Table 4 report the statistics offered in Tables 2 and 3 for these additional stores. The share of goods obeying the LOP in currency unions averages closer to 40 than to 60 percent, but the qualitative patterns are essentially identical to those found in the primary data set. Rows (i) to (iii) of the first column show that good-level log RERs in these additional stores have an average magnitude of 0.086 compared with 0.154 for pegs and 0.201 for floats.

---

[18]We separate the products by brand using the first letter of the product ID code. We sampled 100 goods each whose codes began with the letters "H" and "T" and found that all 200 were non-Apple products. Further, all goods we sampled with codes beginning with the letters "M" and "D" were Apple branded goods. The remaining category, codes beginning with "D", appears to include Apple software products.

[19]For Adidas, Dell, and Nike, pairs involving a euro zone country and Denmark are the only peg observations, but more pegged bilaterals exist in the Mango data. Online prices from Mango were used by Simonovska (2011) to study the relationship between average prices and per-capita income.

4   Decomposing the Real Exchange Rate

Above we demonstrated that outside of currency unions, there are marked differences in bilateral good-level RERs, but these LOP deviations and their overall distribution may emerge from multiple sources. For example, there might be large LOP violations between Spain and Norway while there are none between Spain and Portugal because (i) markups are initially set to different levels, (ii) subsequent price changes are of different sizes, or (iii) Spain and Norway have bilateral nominal volatility from the exchange rate while Spain and Portugal do not. We now turn to a disaggregation framework useful for separating out these channels. We start by documenting the short duration of products and the long duration of prices in our data. We next derive a novel decomposition of good-level RERs. Given the short product duration, one important component of the decomposition is the RER at the date of product introduction. Given the long duration or stickiness of prices, another term captures RER variation due to nominal exchange rate movements. Finally, a third residual component captures the impact of differential price adjustment across countries during the life of the good. We demonstrate that the component capturing the RER at the time of product introduction is by far the most salient component of good-level RERs.

4.1   Product Duration and Price Stickiness

Panel A of Table 5 reports the typical life span of products in our data. We list in rows (i) and (ii) the mean and median duration in weeks for all products in any country in our data (each product and country is considered a distinct good) as well as the duration of products in the United States.[20] Durations for the world listed in row (i) are systematically lower than those for the United States listed in row (ii), presumably due to the longer span of our data for the United States. The mean duration of 36 weeks (roughly 3 quarters) is significantly greater than the 15 week median duration, reflecting skewness in the distribution. About 10 percent of the goods for the United States, for example, exist in the data for at least 2 years. These values which pool across stores obscure significant differences among the retailers. IKEA goods on average last 73 weeks (17 months) in the United States compared to average durations of about 1 quarter in the apparel companies. Though the product lives vary, these data suggest that goods are frequently

---

[20]We drop goods which appear for less than 1 month and we weight such that there is an equal contribution from each country and store combination. There is left- and right-censoring plus occasional scraping gaps, which can of course bias downward the duration statistics. We generated alternative estimates that excluded goods introduced on the initial date of scraping, on subsequent dates following scraping gaps, or on days when our algorithm detects a highly unusual volume of new introductions, as well as goods observed on the last date of scraping in each country. Since this restricted set of goods in fact had durations that were about 20 percent shorter (presumably because we excluded the longest-lived types of goods), we simply report the unfiltered results here and note that these duration estimates are likely biased downward.

entering and exiting countries' consumption bundles and suggests that product introductions may be important for understanding RERs.

Next, in Panel B of Table 5, we report the percent of products with any observed price changes. Again, we separately consider in rows (iii) and (iv) all products in our data set as well as only the U.S. data and use the same weighting scheme as in Panel A. Since price changes are of course more likely for longer-lived products, we additionally report this statistic for the set of goods with at least 12 and 24 months of data. Only 15 percent of goods in the data experience price changes. This clearly reflects, however, the very short lives of most apparel products as well as the short tenure of many smaller countries in our data. If instead we focus on goods in our dataset for greater than 12 months, we see that more than half experience at least one price change, including essentially all apparel items with sufficient data. The fact that many goods experience no price changes, however, even when they exist for more than one year in the data, suggests nominal volatility may play an important role in RER volatility.

4.2   Introduction, Demand, and Stickiness

To connect the frequency of product introductions and stickiness of prices to good-level RERs, we now derive a decomposition that will attribute each good-level RER into a component from the time of introduction, a component due to price stickiness, and a residual component owing to differential price adjustment across countries. Let $i_i(z)$ denote the time that good z is introduced in country i and let $\bar{p}_i(z) = p_i(z, i_i(z))$ denote the log of the price at introduction. We assume that prices are characterized by nominal rigidity and so we write the log price of good z in country i at time $t > i_i(z)$ as:

$$p_i(z, t) = \bar{p}_i(z) + \Delta_{i_i(z)}^{l_i(z,t)} p_i(z),$$

where we define $l_i(z, t)$ as the date of the last price change prior to t and where we introduce the multi-period difference operator $\Delta_s^t v = v(t) - v(s)$ for any variable v. The $\Delta_{i_i(z)}^{l_i(z,t)} p_i(z)$ term can be positive or negative and represents the accumulation of one or more price changes. If the good has experienced no price changes since its introduction, then $i_i(z) = l_i(z, t)$ for all t and $p_i(z, t) = \bar{p}_i(z)$.

It will prove convenient to write the price of this good when translated into country k currency units, $p_i(z, t) - e_{ik}(t)$, as:

$$p_i(z, t) - e_{ik}(t) = \underbrace{[\bar{p}_i(z) - e_{ik}(i_i(z))]}_{\text{Price at Introduction}} + \underbrace{\left[\Delta_{i_i(z)}^{l_i(z,t)}(p_i(z) - e_{ik})\right]}_{\text{Cost/Demand Shocks and Passthrough}} - \underbrace{\Delta_{l_i(z,t)}^t e_{ik}}_{\text{Stickiness}}. \quad (1)$$

18

The price of good z expressed in units of currency of some country k at time t can be disaggregated into three terms. The first term on the right hand side of (1) equals the price of good z at the date it was introduced and translated into country k currency units ("Price at Introduction"). The second term captures the extent to which changes in the country i currency price changed along with the NER between countries i and k during a price spell that ended with a price change. We expect price changes in country i to reflect cost or demand shocks as well as the degree to which these shocks are passed through into prices ("Cost/Demand Shocks and Passthrough"). Finally, the country k currency unit price may also fluctuate simply due to the interaction of sticky currency i prices combined with a continuously fluctuating NER ("Stickiness").

Combining expression (1) with the equivalent expression for the same good z in country j, we obtain the following disaggregation of the log good-level RER:

$$
q_{ij}\left(z,t\right) = \underbrace{\left[\bar{p}_i\left(z\right)-e_{ik}\left(i_i\left(z\right)\right)-\bar{p}_j\left(z\right)+e_{jk}\left(i_j\left(z\right)\right)\right]}_{\text{Good-Level RER at Introduction}}
$$
$$
+ \underbrace{\left[\Delta_{i_i(z)}^{l_i(z,t)}\left(p_i\left(z\right)-e_{ik}\right)-\Delta_{i_j(z)}^{l_j(z,t)}\left(p_j\left(z\right)-e_{jk}\right)\right]}_{\text{Changes in Demand}} - \underbrace{\left[\Delta_{l_i(z,t)}^{t}e_{ik}-\Delta_{l_j(z,t)}^{t}e_{jk}\right]}_{\text{Stickiness}}. \quad (2)
$$

One contributor to the log good-level RER at time t is the log good-level RER when the good was first introduced into the two countries ("Good-Level RER at Introduction"). Next, there may be country-specific subsequent demand shocks. Given the assumption that good z is produced in a single plant, production cost shocks on their own cannot influence the RER unless there are also heterogeneous rates of passthrough from the producer country to prices in i and j. For instance, if a 10 percent cost shock is fully passed through to prices in country i but only half of it is passed through to prices in country j, this can generate movement in the good-level RER. Since heterogeneous rates of passthrough without heterogeneity in the underlying production structure reflect heterogeneity in demand conditions, we attribute this second term to demand ("Changes in Demand"). Finally, even when the local currency prices are not moving, the changing exchange rates with k imply $q_{ij}\left(z,t\right)$ will change even without movement in local prices ("Stickiness").

Note that this disaggregation is specific to the choice of country k, though the sum of the terms will be equal for all k. Variation in the disaggregation across countries k is entirely a result of asymmetries in the timing of good introductions and price changes. For example, if both goods are introduced on the same date $i_i\left(z\right)=i_j\left(z\right)$ and have their last price change on the same date

$l_i(z,t) = l_j(z,t)$, then (2) reduces to:

$$q_{ij}(z,t) = \underbrace{[\bar{p}_i(z) - e_{ij}(i_i(z)) - \bar{p}_j(z)]}_{\text{Good-Level RER at Introduction}} + \underbrace{\left[\Delta^{l_i(z,t)}_{i_i(z)}(p_i(z) - p_j(z) - e_{ij})\right]}_{\text{Heterogeneous Demand}} - \underbrace{\Delta^t_{l_i(z,t)}e_{ij}}_{\text{Stickiness}},$$

which has no dependence on country k.

It is an undesirable property for the disaggregation of the good-level RER between countries i and j to reflect the NER of a third and potentially unrelated country, so we consider the two special cases when $k = i$ and when $k = j$. We then use as our disaggregation of the good-level RER an equally weighted average of the two resulting expressions in these two cases:

$$
\begin{aligned}
q_{ij}(z,t) &= \underbrace{\left[\bar{p}_i(z) - \bar{p}_j(z) - \frac{1}{2}e_{ij}(i_i(z)) - \frac{1}{2}e_{ij}(i_j(z))\right]}_{\text{Introduction } q^I_{ij}} - \underbrace{\left[\frac{1}{2}\Delta^t_{l_i(z,t)}e_{ij} + \frac{1}{2}\Delta^t_{l_j(z,t)}e_{ij}\right]}_{\text{Stickiness } q^S_{ij}} \\
&+ \underbrace{\left[\Delta^{l_i(z,t)}_{i_i(z)}p_i(z) - \Delta^{l_j(z,t)}_{i_j(z)}p_j(z) - \frac{1}{2}\Delta^{l_i(z,t)}_{i_i(z)}e_{ij} - \frac{1}{2}\Delta^{l_j(z,t)}_{i_j(z)}e_{ij}\right]}_{\text{Demand } q^D_{ij}}.
\end{aligned}
\tag{3}
$$

We use the three terms "Introduction," "Demand," and "Stickiness" to represent the three components of the RER in (3) and write them as:

$$q_{ij}(z,t) = q^I_{ij}(z,t) + q^D_{ij}(z,t) + q^S_{ij}(z,t). \tag{4}$$

This disaggregation, of course, is not the unique one that allows us to study the relative contribution of the introduction price or nominal rigidities to good-level RERs.[21] We choose the definition in (3) as our baseline because it allocates pricing behavior of good z in country i independently of what occurs in country j. In other words, it uses information on introduction prices and stickiness similarly for a good sold in one country, regardless of the timing of introduction or price changes in the other country of the pair. We consider all new product ID codes appearing for each store in a given country to be a new good. For all subsequent analyses, we restrict our attention to goods introduced to both countries within a 15 week span to ensure an intuitive interpretation of the "Introduction" term.

---

[21]Our results do not change qualitatively if instead we consider any of the following three alternative decompositions. First, we can re-define the Introduction term to be $q_{ij}(z, i^*_{ij}(z))$, where $i^*_{ij}(z) = \max\{i_i(z), i_j(z)\}$ is the later of the two introduction dates, and keep the definition of the Stickiness term unchanged. Second, we can instead re-define the Stickiness term to equal $-\Delta^t_{t^*_{ij}(z,t)}e_{ij}$, where $t^*_{ij}(z,t) = \max\{l_i(z,t), l_j(z,t)\}$ is the most recent observed price change in either country, and keep the definition of the Introduction term unchanged. Third, we can combine both of the previous two adjustments, which implies that all three terms change relative to our baseline definition.

In the top left panels (labeled "a") of Figures 7 and 8 we once again plot histograms of good-level RERs q_{ij} for selected bilateral relationships with the United States and Spain respectively, and in the remaining three panels we plot q_{ij}^I (in panel "b"), q_{ij}^D (in panel "c"), and q_{ij}^S (in panel "d"). Starting with the case of the United States in Figure 7, one immediately notes that the largest share of variation comes from the component at introduction. This information is omitted if one studies the RER using matched model price indices. Nominal rigidity or stickiness contributes a moderate amount.

Surprisingly, given the large attention it has received in the literature, the changes in demand channel contributes only a small amount to international relative prices for these products. This term equals zero by construction when there are no price changes and so the lack of support in the distributions of q_{ij}^D is to some extent equivalent to observing that prices are highly sticky for this class of goods. We will explore this further below, but it will remain a robust conclusion of this paper that differential price change behavior of continuing goods, reflecting differential exchange rate passthrough or other mechanisms, is not the first-order determinant of good-level RERs.

Similar results are seen in Figure 8 for the case of Spain. We saw in Figure 4 that countries outside of the euro zone violated the LOP for goods that were priced identically within the euro zone. In principle, these violations could have reflected LOP violations at introduction, different timing or scales of price changes on existing goods, or could have reflected nominal volatility and price stickiness. In practice, one sees the largest component coming at introduction along with a moderate contribution from nominal rigidities. Note that Denmark pegs to the euro and therefore has $q_{ij}^S = 0$ for all goods. But its good-level RER distribution at introduction q_{ij}^I continues to look completely different from the euro zone countries.

Next, we decompose the variance across good-level RERs at any date t as follows:

$$\sigma_{ij}^2(t) = \underbrace{\left[\left(\tilde{\sigma}_{ij}^I\right)^2(t)\right]}_{\text{Good-Level RER at Introduction}} + \underbrace{\left[\left(\tilde{\sigma}_{ij}^D\right)^2(t)\right]}_{\text{Changes in Demand}} + \underbrace{\left[\left(\tilde{\sigma}_{ij}^S\right)^2(t)\right]}_{\text{Stickiness}}, \tag{5}$$

where $\sigma_{ij}^2(t)$ is the variance over goods z of $q_{ij}(z,t)$. We use tildes in the terms on the right hand side because those terms include not only the variance of each component but also half of the total contribution of the respective covariance terms, an innocuous assumption given that the covariance terms are small. We measure each term on the right hand side of (5) for a selection of countries against the United States and Spain. We perform the decomposition separately for all available weeks with at least 100 matched goods for each bilateral country pair and average across those weeks. As above, we weight in order to equalize the contribution from all store and bilateral

pair combinations in each week. Most weeks will contain some mix of goods which are newly introduced (where $q_{ij} = q_{ij}^I$), some which just experienced a price change (where $q_{ij} = q_{ij}^I + q_{ij}^D$), and some which have had a long time pass since the most recent price change or introduction (where q_{ij}^S may be large). We subtract the mean bilateral RER for each store in order to focus on variation around a mean, rather than changes in the variance due to differences across stores in the mean.[22]

The bar charts in Figure 9 show the relative importance of prices at introduction, changes in demand, and stickiness for explaining dispersion in good-level RERs in the cross-section for several bilateral relationships. The left two columns represent the three terms in the decomposition (5) when measured on the "Full Sample" of data. We return to the right two columns, labeled "Reduced Sample," below.

Starting for example with the upper-left plot which represents Canada and the United States, we see that the three bars sum to 0.027 log point, equal to the average cross-sectional variance in q_{ij} for that bilateral pair. We use the decomposition (5) to measure that 0.021 log point comes from the Introduction term, 0.005 comes from the Demand term, and 0.001 comes from the Stickiness term. Clearly, nominal rigidity explains little of the dispersion of good-level RERs at any given point in time. This likely reflects the fact that movement in the NER is a common shock applying equally to all goods. And looking at the pairs of Spain with Denmark (which is pegged to the euro) and of Spain with France (which is a member of the euro zone), nominal rigidities cannot explain any of the cross-sectional dispersion for countries with a constant bilateral NER.

The "Changes in Demand" term, which reflects unequal price changes (when expressed in a common currency), contributes roughly 25 percent of the total cross-sectional variation in these bilateral pairs, leaving roughly 70 percent due to the relative prices at the time of the goods' introductions, the largest bars which are shaded red. Note that if there is a constant proportional term contributing to LOP deviations, such as a tax or tariff, this cannot explain our result as it would apply equally to all goods and contribute only to the mean, but not to the cross-sectional dispersion, of the good-level RERs.

The relative price at the time of introduction is far more informative about good-level RERs than anything that happens subsequently including price changes and NER fluctuations. We have calculated this decomposition for all bilaterals in our data to confirm that the bilaterals

---

[22]We note that this disaggregation is sensitive to the length of data available. For example, it will by construction attribute all cross-sectional variation to the $\tilde{\sigma}_{ij}^I$ term for the first observed week of any data set. We have looked at the time series patterns of the three terms for several bilateral pairs in our data and find that this time-dependency typically quickly dissipates. We report the simple average across periods of the decomposition, rather than setting rules on which data to exclude, for greatest transparency.

shown in Figure 9 are representative. While the exact breakdown among the σ term varies across countries, most cross-sectional variation is generally attributable to good-level RERs at the time of introduction.

To understand whether our conclusions hold qualitatively for longer-lived products with less price stickiness, we re-calculate these results with a "Reduced Sample" that includes only goods that are in our dataset for each bilateral pair for more than one year and which at some point experience at least one price change in at least one of the countries. Comparing the "Full Sample" results in the first two columns of Figure 9 to their equivalent results from the "Reduced Sample" in the last two columns, it is not surprising that the relative contribution of the "Demand" term increases in all the bar chars. For example, whereas the "Demand" term was about one-fourth as important as the "Intro" term for the Canada-U.S. pair in the "Full Sample," it is nearly half as important in the "Reduced Sample." The "Intro" term remains clearly the most important, however, for all plotted bilaterals. We view the high frequency of product introductions and significant price stickiness as calling for a greater emphasis in empirical work on relative prices in levels rather than in changes.

5   Time Series Variation and Product Introductions

Above, we highlighted the importance of relative prices at the time of product introductions for understanding the cross-section of good-level RERs. An understanding of good-level RERs at introduction (q_{ij}^I) may also be critical for an understanding of the time-series evolution of the RER, which is a key focus of open-economy macroeconomics. A large body of literature including Mussa (1986) and Engel (1999) uses price indices to measure the RER and to note its surprising co-movement at both low and high frequency with the NER. Due to data limitations, these measures are not influenced by the RER at introduction and therefore ignore the component which we above show to contribute the bulk of the variation in good-level RERs. We now show how the RER at introduction varies with the NER in our data.

As a benchmark, consider the possibility that prices are quite sticky, leading to volatility in good-level RERs, but that the distribution of q_{ij}^I is stationary with an expected value of $\tilde{q}$. Imagine that all goods frequently enter and exit the market. In this world, the aggregate RER might appreciate or depreciate over time with the NER, but would never wander too far from $\tilde{q}$. This is because every time a product exited, regardless of it's good-level RER at the time of exit, it would be replaced by a new good with an expected good-level RER of $\tilde{q}$. Measures such as matched-model price indices that ignore product introduction would miss this hypothetical

mechanism for mean reversion of the aggregate RER. This is the sense in which taking product introduction prices into account might plausibly have solved (or explained much of) the PPP puzzle.[23]

We do not, however, find evidence supportive of this explanation. As the NER varies, so too does the good-level RER at the time of introduction. For instance, imagine that a typical good is introduced in Spain and the United States at 1 euro and 1 dollar when the NER is hypothetically at parity in January of year 1. If the dollar-euro exchange rate changed to 2 dollars per euro in January of year 2, one might expect a new good to be introduced at 1 euro and 2 dollars at that later date. By contrast, the data indicate that the subsequent introduction would more likely also be priced at 1 euro and 1 dollar, implying the aggregate RER (that combines the old and new goods) moves together with the NER.

Echoing many of the results in the literature on exchange rate passthrough suggesting the prevalence of local currency pricing, our results appear to document the prevalence of local currency introduction pricing too. One reason this result is important is that it suggests the high frequency correlation between NERs and RERs cannot simply emerge due to the lack of price adjustment without some additional pricing complementarities or reference point behavior. After all, a price change must by definition occur at the time of a good's introduction, regardless of what the introduction price level is!

Figure 10 plots the weekly weighted median of good-level log RERs at introduction q_{ij}^I for the key bilaterals involving the United States.[24] We separate each of the four stores and mark their medians with each of four markers. The thin black line plots the log bilateral NER e_{ji}, normalized to zero at the first date. As such, the relative levels of the markers and the black line are not informative, but their time-series movements are. As opposed to sharing none of the time-series properties of the black line, as would be predicted in models with constant markups at the time of introductions or price changes, the markers often appear to move along with the black line.

For example, the upward movement of the red circles representing Apple products early in the sample for Germany and the United States mimics the upward movement of the log NER. This means that as the euro appreciated relative to the dollar in 2009, prices at the time of introduction grew higher in Germany relative to the United States when translated into a common currency. Similarly, there is a downward movement for all of the retailers in the Japan-U.S. sample during

---

[23]This possibility is a cousin of the explanation in Nakamura and Steinsson (2012) that the exclusion of substitution prices from BLS import and export price indices is behind the low levels of exchange rate passthrough in their aggregate indices.

[24]As above, we weight at each date such that the contribution from each store and bilateral pair combination is equalized. We drop the limited number of observations where $|q_{ij}^I| > 0.75$ or $|q_{ij}| > 0.75$.

the large yen depreciation starting in late 2012, represented by the downward movement of the log NER. (The gaps, particularly in the red circles during 2010 and early 2011, are the results of periods without scraping.) It is difficult, however, to make conclusions from these rich scatter plots, so we now turn to non-linear fits of these raw data.

First we scale the q_{ij}^I values for each store by a constant so they have the same mean in early 2012. We do this because we wish to capture the within-store time-series variation in median good-level RERs at introduction as opposed to capturing compositional changes due to stores with different mean LOP deviations entering or exiting our sample. We then use the lowess nonlinear smoother on these data and plot the resulting fitted values as the dashed red line in Figure 11, scaled up or down such that the average value equals that of the log NER. In this sense, there is no information in the levels of either line in the diagram, but the time-series movements are informative. Periods in between observed introductions are interpolated, and therefore long periods lacking introductions appear as straight dashed red lines, such as the interpolations in the middle of the China, France, Germany, and Japan plots.

The comovement at high and low frequency between the red-dashed line and black solid line in Figure 11 is striking. The fitted average values of the RERs at introduction move with the NER. Major secular trends in Canada, China, and Japan are at least partly captured, and higher frequency movements in the NER with euro zone countries, Sweden, and the U.K. are all mirrored by comparable high frequency movements in the log RER at the time of good introductions. Companies appear to price with local currency reference points, even at the time that a new good is introduced and despite large movements in the NER.

To formally quantify this relationship, we run the following regressions:

$$q_{ij}^I(z, i_{ij}^*) = \gamma_{ij} + \beta e_{ji}(i_{ij}^*) + \epsilon_{ij}(z, i_{ij}^*), \tag{6}$$

where good z only appears in the regression in the one period when $t = i_{ij}^*$, where we demean the left-hand side variable for each store and country pair (equivalent to adding store-country-pair dummies), and where we exclude any good with $|q_{ij}^I| > 0.75$ or $|q_{ij}| > 0.75$. An estimated value $\beta = 0$ would imply that goods are introduced at RER levels unrelated to the NER, as would be predicted for instance if the LOP always held at the time prices were adjusted. An estimated value $\beta = 1$ would imply that the RER at the time of good introductions perfectly tracks the NER, as would be predicted if introduction prices were centered around some local currency price target.

Panel A of Table 6 reports the coefficients on this regression when we weight using the inverse

of the number of observations for each store and country pair combination. We consider in rows (i) and (ii) "All Bilaterals" in our data and restrict to bilaterals involving the United States in rows (iii) and (iv). Finally, the second column indicates whether we simply include all goods (that are introduced within 15 weeks in both countries), or a more restricted set of goods that eliminates goods first observed on the first date of scraping for the store-country combination and goods first observed on a date with an unusually large increase in the number of new good observations. This restricted set is labeled "Filtered Intros" and is designed to control for possible misspecification of product introductions due to scraping gaps and errors. We cluster the standard errors by retailer-weeks and note that these coefficients are estimated with high precision.

Consistent with Figures 10 and 11, all four specifications in Panel A show that the good-level RER closely tracks the NER even at the time of product introductions. For example, looking at row (i), we see that across all stores and intros, the good-level RER at introduction (q_{ij}^I) moves about 0.7 log point for every full log point movement in the NER. If we restrict to U.S. bilaterals and the stricter definition of introductions, as in row (iv), the coefficient rises to about 0.9 log point. In other words, if the bilateral NER appreciates by 10 percent over the course of the year, one would expect new products to be introduced with relative prices 7 to 9 percent higher than the previous year (i.e. relative local currency prices move between 1 and 3 percent). The phenomenon holds for Apple products, but less so than it does for the other stores. IKEA, H&M, and Zara good-level RERs for U.S. bilaterals at the time of product introductions, for example, roughly track the NER one-for-one. Regression (6) is very similar to that estimated using data on IKEA products that are "new in both countries" in Table 5 of Baxter and Landry (2012). Their estimate would correspond to a value of 0.64, which is only slightly below our results in rows (i) and (ii) for IKEA and consistent with our qualitative conclusions. We consider these weighted results to be our baseline estimates, but report unweighted results in Panel B. The store-by-store results are similar, but due to the reduced weight on Apple products, the "All Stores" values all increase to levels above 0.8.

In sum, we isolate good-level RERs at the time products are introduced and demonstrate that these RERs track closely the NER. This evidence is similar to findings that exchange rate passthrough is incomplete even conditional on price adjustment and stands in contrast to the predictions of models where all price rigidity for any given good is due simply to costly price adjustment for that good.[25] Furthermore, although the importance of the RER at introduction

---

[25]While the relationship of the RER with the bilateral NER is of course related to passthrough (as can be easily seen in the "Changes in Demand" term of (2)), we cannot explicitly measure exchange rate passthrough in these data because we do not know the identity of the exporting country for any given good. To see why, imagine a good is produced in Japan and exported to both Spain and the United States with constant passthrough rates

in Section 4.2 suggests that some shocks to the true RER between particular countries may not be measured correctly in conventional price indices, the results in Table 6 suggest this problem does not apply as systematically as would be required to resolve the PPP puzzle. The relative prices for these goods move in a way at the time of introduction that closely resembles how they would move if they were existing goods which simply had sticky prices. Our results indicate that, for IKEA, H&M, and Zara, for example, matched-model indices would have roughly no systematic bias due to the omission of introduction prices.

6   Implications for Theories of Price Setting

In this section we discuss how our findings relate to theories of price determination. The basis for our discussion is a simple static model where firms maximize profits associated with the sale of good z in market i by setting the price as a markup $\mu_i(z)$ over the marginal cost of production in the producer's currency $c_i(z)$ plus the distribution d_i and border costs τ_i, all in logs, where the markup is a function simply of the elasticity of demand for the product. Interpreting, for now, countries as markets, and letting e_i^p denote the log exchange rate between the producer country and country i, we can write:[26]

$$p_i(z) = e_i^p(z) + c_i(z) + d_i + \tau_i + \mu_i(z).$$

6.1   Importance of Markup Heterogeneity Across Space and Time

Our results are not supportive of models with perfect competition that rely on different costs of production and distribution to account for differences in $p_i(z)$.[27] Many of the goods in our sample are produced in the same plant at the same marginal cost $c(z)$, regardless of where they are sold. Most would face comparable shipping costs to reach consumers in Chicago or Toronto

---

of 0.75 and 0.25 respectively. If prices change only due to exchange rate shocks, a 10 percent depreciation of the euro relative to the yen with no change in the dollar-yen will produce a 7.5 percent appreciation of the good-level RER between Spain and the United States. Alternatively, a 10 percent appreciation of the dollar relative to the yen with no change in the euro-yen will produce a 2.5 percent appreciation of the same good-level RER. These two scenarios both imply a 10 percent nominal depreciation of the euro but generate different movements in the good-level RER.

[26]Dynamic pricing models further take into account the idea that prices might not change frictionlessly over time, whether due to technological constraints as in time-dependent models or due to adjustment costs as in state-dependent models. Price changes in the dynamic setting reflect expectations about the length of time the price will last and assessments about the future path of the marginal cost and elasticity of demand. Our results shed less light about these dynamic considerations.

[27]Prominent papers in international trade such as Eaton and Kortum (2002) generate price differences from cost differences that emerge due to differences in the sourcing of the same goods and differential international shipping costs. Others including Burstein, Neves, and Rebelo (2003) focus on the importance of the distribution channel in the total cost of reaching consumers.

(i.e. $d_i \approx d_j$). Finally, country-level factors such as tariffs might result in a levels shift of all prices in one country relative to the other (i.e. $\tau_i > \tau_j$ for all z) for that industry, but cannot explain why there are simultaneously many goods with higher prices and others with lower prices in the United States compared with, say, Canada.

Our results are also inconsistent with models where relative price variation over time for any given good emerges primarily due to nominal exchange rate fluctuations Δe_{ij} that occur when that good's price is fixed. Imagine, for example, a standard model with a constant elasticity of demand for each good in each market and no predictable trends in inflation. Production and distribution costs, country-level tariffs, and the distribution across goods of the elasticity of demand are fixed over time. Goods are introduced on the same date into countries i and j with markups $\bar{\mu}_i$ and $\bar{\mu}_j$ and with RERs $\bar{q}_{ij}$, all indexed by good z. Finally, assume prices of existing goods change due to permanent exchange rate shocks Δe_{ij}. In this simple setting, the firm's dynamic pricing problem at time of adjustment is the same as its static problem and the distribution of $\bar{q}_{ij}$ would be unrelated to e_{ij}, in stark contrast to our findings detailed above. This is in line with the finding in Gopinath, Itskhoki, and Rigobon (2010) that passthrough is incomplete even at the time of a price change and is consistent with models generating long-run pricing to market and variable markups.[28] Our results therefore corroborate the importance of markup heterogeneity $\mu_i(z)$ in explaining cross-border price variation.

6.2 Customer and Firm Psychology and the Role of the Internet

The fact that LOP holds well within currency unions but less so among credible and strong pegs points to behavioral theories of price setting in which customer psychology, firm organizational structure, and market norms play important roles. For example, given the above discussion, one might interpret our findings as suggesting the euro zone has eliminated barriers to arbitrage between its members. While barriers may in fact have been reduced with the single currency's introduction and use, we find such an explanation for our results unlikely. We simply do not find it plausible that a German consumer (or arbitrageur) would find it dramatically harder to cross the border and buy a product in Denmark compared to instead doing the same in the Netherlands or Belgium. The only meaningful difference, the transaction cost charged by a currency exchange or a German credit card in order to make a purchase in kroner, is far too small to account quantitatively for the differences found in Tables 2 and 3.

---

[28]These characteristics feature prominently in real exchange rate determination in Krugman (1987) and Dornbusch (1987) and the large subsequent literature which includes Knetter (1993) and Atkeson and Burstein (2008) among many others.

It is possible that firms' pricing decisions are designed to avoid potential consumer anger, which might result from the knowledge that neighboring countries face lower prices. Put more formally, firms may believe that the elasticity of demand for good z in country i is not only a function of the local price but is also a function of the local currency value of foreign prices. How might this account for differences between common currencies and pegs? If countries share a common currency, then it is easy to translate foreign into domestic prices. After all, $exp\{e_{ij}\} = 1$ for pairs in the euro zone while $exp\{e_{ij}\} = 0.1342$ for pairs involving one euro zone country together with Denmark.[29] In fact, the act of joining a currency union or the observation that prices are quoted in the same currency may themselves shift the customer's attention to the fact that there is a relevant international comparison to be made.

Another reason to suspect that consumer psychology may play a role is that in currency unions, we find that *listed* prices are equalized, despite the fact that market norms vary in terms of whether prices are listed inclusive or exclusive of taxes. For example, the histograms in Figure 4 show that prices are often equalized across euro zone countries, despite the fact that European prices are required to include VAT tax rates that differ across countries. Figure 6 shows that prices which include VAT taxes in dollarized countries are often exactly equal to prices in the United States that by norm exclude sales taxes. Firms appear to tolerate markup variation in order to preserve equality in list prices. We view our results as supportive of approaches such as those in Rotemberg (2005), Nakamura and Steinsson (2011), and Bordalo, Gennaioli, and Shleifer (2013a,b). These papers model pricing strategies that take into account elements of consumer psychology such as anger, habit formation, and a disproportionate focus on particular product attributes.

Moving from the customer to the firm side, another interpretation of our results is that the organizational structure of firms is critical to understanding pricing strategies. It is possible that each of these multinational companies has pricing departments that are segmented by currency. If changing such structures is very costly, firms may set up separate departments for pegged regimes simply due to the greater likelihood that they become a floating currency in the future. Standard calibrated price-setting models would have a very difficult time delivering the optimality of such a policy, however, and few papers connect firm organizational structure to pricing dynamics at the macro level.[30]

---

[29]However, this reason alone cannot explain why price differences are often large between the United States and Panama, where the currency is pegged at 1 with the dollar, or between the United States and Canada, where the exchange rate has never deviated from 1 by more than a few percentage points in recent years.

[30]See Zbaracki, Ritson, Levy, Dutta, and Bergen (2004) for a detailed analysis of the price setting process in a large manufacturing firm. Neiman (2010, 2011), Hellerstein and Villas-Boas (2010), and Hong and Li (2013) analyze the implications of vertical organizational structure for price rigidity and the response to exchange rate

Finally, it is interesting to speculate on the role of the Internet in driving firms to equalize prices both within individual countries and in currency unions. We do not have data covering a country that enters or exits the euro zone in sample and therefore all our identification is necessarily from cross-section, not time-series, variation. Further, we have no observations that predate the Internet. Perhaps there is an interaction between Internet penetration and the choice of currency regime. Perhaps joining a currency union carries different implications when technology facilitates cross-country price comparisons. If true, one might expect in the future to observe an increase in within- and across-country price convergence for traded goods, even without reductions in trade frictions.

7 Conclusion

Using a novel dataset of traded goods by four global retailers in dozens of countries, we demonstrate that for tens of thousands of goods, prices are equalized in currency unions and diverge outside of them. The choice of currency regime for these products is the critical determinant of market segmentation, far more important than transport costs or tax or taste differences. Further, we demonstrate that the relative price at the time of product introductions is the most important component of good-level RERs in our data, and this component moves at high frequency with the NER. We estimate that the sectors covered in our data account for more than 20 percent of U.S. consumption on goods, but, clearly, the pricing behavior documented for these four global retailers need not be representative of all retail sectors. Future work should clearly focus on understanding what determines when prices behave like those documented here and when they do not.

These findings have important implications for such critical topics as optimal currency regimes, the dynamics of external adjustment, the international transmission of shocks, and RER measurement. They suggest the repercussions of joining a currency union are far broader than the simple elimination of nominal volatility. Our findings about the RER at introduction imply that the root of pricing rigidities is not well captured by models that omit variable flexible price markups or pricing complementarities, including many of those with monopolistic competition, constant demand elasticities, and menu costs. Most broadly, the patterns we document point to the importance of customer psychology, firm organizational structure, and the internet for price setting, elements which do not yet feature prominantly in most standard macroeconomic models.

---

shocks.

References

ATKESON, A., AND A. BURSTEIN (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98(5), 1998–2031.

BAILEY, D., AND R. WHISH (2012): *Competition Law*. Oxford University Press.

BALASSA, B. (1964): "The Purchasing Power Parity Doctrine: A Reappraisal," *Journal of Political Economy*, 72(6), 584–596.

BAXTER, M., AND A. LANDRY (2012): "IKEA: Product, Pricing, and Pass-Through," Working Paper.

BERKA, M., M. DEVEREUX, AND C. ENGEL (2012): "Real Exchange Rate Adjustment in and out of the Eurozone," *The American Economic Review*, 102(3), 179–185.

BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2013a): "Competition for Attention," Working Paper.

———— (2013b): "Salience and Consumer Choice," Working Paper.

BORRAZ, F., A. CAVALLO, R. RIGOBON, AND L. ZIPITRIA (2012): "Distance and Political Boundaries: Estimating Border Effects under Inequality Constraints," Working Paper.

BRODA, C., AND D. WEINSTEIN (2008): "Understanding International Price Differences Using Barcode Data," Working Paper.

BURSTEIN, A., AND N. JAIMOVICH (2009): "Understanding Movements in Aggregate and Product-Level Real Exchange Rates," *unpublished paper, UCLA and Stanford University*.

BURSTEIN, A., J. NEVES, AND S. REBELO (2003): "Distribution Costs and Real Exchange Rate Dynamics During Exchange-Rate-Based Stabilizations," *Journal of Monetary Economics*, 50(6), 1189–1214.

CASSEL, G. (1918): "Abnormal Deviations in International Exchanges," *The Economic Journal*, 28(112), 413–415.

CAVALLO, A. (2012): "Scraped Data and Sticky Prices," Working Paper.

CAVALLO, A., AND R. RIGOBON (2012): "The Distribution of the Size of Price Changes," Working Paper.

COSAR, K., P. GRIECO, AND F. TINTELNOT (2012): "Borders, Geography, and Oligopoly: Evidence from the Wind Turbine Industry," Working Paper.

CRUCINI, M., AND M. SHINTANI (2008): "Persistence in law of one price deviations: Evidence from micro-data," *Journal of Monetary Economics*, 55(3), 629–644.

CRUCINI, M., M. SHINTANI, AND T. TSURUGA (2010): "Accounting for persistence and volatility of good-level real exchange rates: The role of sticky information," *Journal of International Economics*, 81, 46–60.

CRUCINI, M., C. TELMER, AND M. ZACHARIADIS (2005): "Understanding European real exchange rates," *The American Economic Review*, 95(3), 724–738.

DELOITTE (2013): "Global Powers of Retailing 2013," Report.

DEVEREUX, M., AND C. ENGEL (2007): "Expenditure switching versus real exchange rate stabilization: Competing objectives for exchange rate policy," *Journal of Monetary Economics*, 54, 2346–2374.

DEVEREUX, M., C. ENGEL, AND P. STORGAARD (2004): "Endogenous Exchange Rate Pass-Through when Nominal Prices are Set in Advance," *Journal of International Economics*, 63, 263–291.

DORNBUSCH, R. (1987): "Exchange Rates and Prices," *American Economic Review*, 77(1), 93–106.

EATON, J., AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–1779.

ENGEL, C. (1999): "Accounting for U.S. Real Exchange Rate Changes," *Journal of Political Economy*, 107(3), 507–538.

ENGEL, C., AND J. ROGERS (1996): "How wide is the border?," *American Economic Review*, 86, 1112–1125.

——— (2004): "European product market integration after the euro," *Economic Policy*.

GAGNON, E., B. MANDEL, AND R. VIGFUSSON (2012): "Missing Import Price Changes and Low Exchange Rate Pass-Through," Working Paper.

GOLDBERG, P., AND F. VERBOVEN (2005): "Market Integration and Convergence to the Law of One Price: Evidence from the European Car Market," *Journal of International Economics*, 65, 49–73.

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., O. ITSKHOKI, AND R. RIGOBON (2010): "Currency Choice and Exchange Rate Pass-through," *American Economic Review*, 100, 304–336.

GORODNICHENKO, Y., AND L. TESAR (2009): "Border effect or country effect? Seattle may not be so far from Vancouver after all," *American Economic Journal: Macroeconomics*, (1), 219–241.

HEAD, K., T. MAYER, AND J. RIES (2010): "The erosion of colonial trade linkages after independence," *Journal of International Economics*, 81(1), 1–14.

HELLERSTEIN, R., AND S. B. VILLAS-BOAS (2010): "Outsourcing and Pass-through," *Journal of International Economics*, 81(2), 170–183.

HONG, G. H., AND N. LI (2013): "Market Structure and Cost Pass-Through in Retail," Working Paper.

ILZETZKI, E., C. REINHART, AND K. ROGOFF (2008): "Exchange Rate Arrangements into the 21st Century: Will the Anchor Currency Hold?," Unpublished Data.

IMBS, J., H. MUMTAZ, M. RAVN, AND H. REY (2005): "PPP Strikes Back: Aggregation and the Real Exchange Rate," *Quarterly Journal of Econoimcs*, 120(1), 1–43.

KNETTER, M. (1993): "International Comparisons of Price-to-Market Behavior," *American Economic Review*, 83(3), 473–486.

KRUGMAN, P. (1987): "Pricing to Market When the Exchange Rate Changes," in *Real Financial Linkages Among Open Economics*, ed. by S. Arndt, and J. Richardson, pp. 1355–1389. MIT Press, London.

MUSSA, M. (1986): "Nominal Exchange Rate Regimes and the Behavior of Real Exchange Rates: Evidence and Implications," *Carnegie-Rochester Conference Series on Public Policy*, 25, 117–214.

NAKAMURA, A., E. NAKAMURA, AND L. NAKAMURA (2011): "Price dynamics, retail chains and inflation measurement," *Journal of Econometrics*, 161(1), 47–55.

NAKAMURA, E., AND J. STEINSSON (2011): "Price Setting in Forward-Looking Customer Markets," *Journal of Monetary Economics*.

——— (2012): "Lost in Transit: Product Replacement Bias and Pricing to Market," *American Economic Review*.

NEIMAN, B. (2010): "Stickiness, Synchronization, and Passthrough in Intrafirm Trade Prices," *Journal of Monetary Economics*, 57(3), 295–308.

——— (2011): "A State-Dependent Model of Intermediate Goods Pricing," *Journal of International Economics*, 85(1), 1–13.

PARSLEY, D., AND S.-J. WEI (2001): "Explaining the border effect: the role of exchange rate variability, shipping costs and geography," *Journal of International Economics*, 55, 87–105.

——— (2008): "In Search of a Euro Effect: Big Lessons from a Big Mac Meal?," *Journal of International Money and Finance*, 27(2), 260.

ROGERS, J., AND M. JENKINS (1995): "Haircuts or hysteresis? Sources of movements in real exchange rates," *Journal of International Economics*, 38, 339–360.

ROGOFF, K. (1996): "The Purchasing Power Parity Puzzle," *Journal of Economic Literature*, 34(2), 647–668.

ROTEMBERG, J. (2005): "Customer Anger at Price Increases, Changes in the Frequency of Price Adjustment and Monetary Policy," *Journal of Monetary Economics*, 52(4), 829–852.

SAMUELSON, P. (1964): "Theoretical Notes on Trade Problems," *Review of Economics and Statistics*, 46(2), 145–154.

SIMONOVSKA, I. (2011): "Income Differences and Prices of Tradables," Working Paper.

ZBARACKI, M., M. RITSON, D. LEVY, S. DUTTA, AND M. BERGEN (2004): "Managerial and customer costs of price adjustment: direct evidence from industrial markets," *Review of Economics and Statistics*, 86(2), 514–533.

		All Stores	Apple	IKEA	H&M	Zara
(i)	# Products, World	113,870	11,235	66,455	14,117	22,063
(ii)	# Products, United States	48,133	5,764	18,092	8,195	16,082
(iii)	# Countries	85	39	31	49	79
(iv)	Time Period	2008:Q4 to 2013:Q2	2009:Q2 to 2013:Q2	2008:Q4 to 2013:Q2	2011:Q3 to 2013:Q2	2011:Q3 to 2013:Q2
(v)	Headquarters		United States	Sweden	Sweden	Spain
(vi)	Industry		Consumer Electronics	Home/Office Furniture	Apparel	Apparel
(vii)	Global Industry Rank (2011)		3rd largest	1st largest	4th largest	3rd largest
(viii)	Retail Revs ($Bil. in 2011)	≈ 100	≈ 40	≈ 25	≈ 15	≈ 15

Table 1: Product, Time, and Country Coverage in the Data

Notes: Retail Revenues are calculated using market shares and total industry sales data found in reports by Euromonitor International for 2011. These revenues are smaller than the total revenues listed, for example, in Apple's annual report (which equaled $108 Billion), likely because Euromonitor only considers a subset of each company's sales to be within their specific market definition.

			All Stores	Apple	IKEA	H&M	Zara
Panel A: Average Absolute Values of Log Good-level RERs							
(i)	Full Sample	Currency Unions	0.076	0.023	0.129	0.020	0.102
(ii)	Full Sample	NER Pegs	0.116	0.085	0.146	0.119	0.115
(iii)	Full Sample	Floats	0.187	0.143	0.221	0.145	0.207
(iv)	$(p_i + p_j) > \$100$	Currency Unions	0.065	0.023	0.097	0.005	0.086
(v)	$(p_i + p_j) > \$100$	NER Pegs	0.109	0.081	0.107	0.113	0.111
(vi)	$(p_i + p_j) > \$100$	Floats	0.189	0.144	0.187	0.152	0.205
(vii)	$(p_i + p_j) > \$400$	Currency Unions	0.042	0.022	0.086	0.013	0.097
(viii)	$(p_i + p_j) > \$400$	NER Pegs	0.096	0.078	0.093	0.125	0.118
(ix)	$(p_i + p_j) > \$400$	Floats	0.171	0.151	0.176	0.141	0.270
Panel B: Share of Absolute Value of Log Good-level RERs Less Than 1 Percent							
(x)	Full Sample	Currency Unions	0.613	0.681	0.314	0.911	0.548
(xi)	Full Sample	NER Pegs	0.069	0.140	0.082	0.069	0.064
(xii)	Full Sample	Floats	0.046	0.049	0.039	0.062	0.040

Table 2: Unconditional Moments of Log Good-level RERs by Store, Currency Regime, and Average Price Level

Notes: Panel A of the Table reports unconditional means of the average (across weeks in the data) of the absolute value of each good's RER, separated by the currency regime. We exclude the small number of observations where $|q_{ij}| > 0.75$. Currency regime definitions closely follow Ilzetzki, Reinhart, and Rogoff (2008) and are described in the text. The unconditional mean is reported from our full data set in (i) to (iii), excluding goods with an average price less than \$50 in (iv) to (vi), and excluding goods with an average price less than \$200 in (vii) to (ix). Panel B of the Table reports the share of the goods where LOP violations are less than 0.01 log point. Both panels use the same weights which equalize the contribution from each country-pair and store combination.

| | | All Stores | | Apple | | IKEA | | H&M | | Zara | | All Stores | | |
												Flexible NER	Pegged NER	Currency Unions
(i)	Outside of Cur. Unions	0.124 (0.006)	0.110 (0.005)	0.054 (0.004)	0.120 (0.004)	0.041 (0.006)	0.092 (0.004)	0.091 (0.006)	0.126 (0.004)	0.162 (0.009)	0.105 (0.009)			
(ii)	Pegged NER	-0.037 (0.004)	-0.073 (0.004)	-0.040 (0.011)	-0.058 (0.013)	-0.017 (0.008)	-0.076 (0.010)	0.003 (0.006)	-0.026 (0.007)	-0.053 (0.005)	-0.092 (0.005)			
(iii)	Log NER Volatility	-0.031 (0.006)		-0.017 (0.007)		-0.018 (0.010)		0.001 (0.008)		-0.027 (0.012)		-0.030 (0.006)		
(iv)	Log Bilateral Distance	0.013 (0.002)		0.012 (0.002)		0.015 (0.003)		0.007 (0.003)		0.016 (0.002)		0.019 (0.002)	0.007 (0.003)	0.008 (0.007)
(v)	Abs. Relative Income	0.001 (0.002)		-0.001 (0.006)		0.022 (0.007)		0.003 (0.005)		0.000 (0.002)		0.002 (0.002)	-0.008 (0.005)	-0.034 (0.022)
(vi)	Abs. Relative Taxes	0.078 (0.024)		0.477 (0.024)		0.099 (0.040)		0.049 (0.034)		0.015 (0.030)		0.108 (0.030)	0.288 (0.060)	0.435 (0.267)
	Cty. Dumies:	Y	N	Y	N	Y	N	Y	N	Y	N	Y	Y	Y

Table 3: Regressions of Absolute Value of Log Good-level RERs

Notes: Table reports regressions of the average (across weeks in the data) of the absolute value of each good's RER on a variety of covariates. The column labeled "All Stores" includes store dummies and are clustered by bilateral country pair-store. The columns labeled with store names are clustered by country pairs. Standard errors are reported in parentheses. We exclude the small number of observations where $|q_{ij}| > 0.75$. Currency regime definitions closely follow Ilzetzki, Reinhart, and Rogoff (2008) and are described in the text. All regressions are weighted to equalize the contribution from each country-pair and store combination.

		All Additional Stores	Adidas	Dell	Mango	Nike
Panel A: Average Absolute Values of Log Good-Level RERs						
(i)	Currency Unions	0.086	0.087	0.054	0.112	0.053
(ii)	NER Pegs	0.154	0.172	0.130	0.158	0.103
(iii)	Floats	0.201	0.207	0.139	0.203	0.210
Panel B: Share of Absolute Value of Log Good-level RERs Less Than 1 Percent						
(iv)	Currency Unions	0.377	0.353	0.380	0.332	0.442
(v)	NER Pegs	0.054	0.027	0.041	0.053	0.092
(vi)	Floats	0.049	0.045	0.052	0.041	0.138
Panel C: Regression Results						
(vii)	Outside of Cur. Unions	0.116 (0.005)	0.120 (0.010)	0.086 (0.018)	0.091 (0.011)	0.158 (0.020)
(viii)	Pegged NER	-0.048 (0.008)	-0.035 (0.016)	-0.009 (0.014)	-0.045 (0.009)	-0.107 (0.021)
(ix)	Countries	55	17	18	49	18

Table 4: Results on Absolute Value of Good-Level Log RER for Additional Retailers

Notes: Panel A of the Table reports unconditional means of the average (across weeks in the data) of the absolute value of each good's RER, separated by the currency regime. We exclude the small number of observations where $|q_{ij}| > 0.75$. Currency regime definitions closely follow Ilzetzki, Reinhart, and Rogoff (2008) and are described in the text. The "NER pegs" category for Adidas, Dell, and Nike, however, only includes pairs involving Denmark and euro zone countries. Panel B of the Table reports the share of the goods where LOP violations are less than 0.01 log point. Panel C of the Table reports regressions of the average of the absolute value of each good's RER on currency regime dummies. All panels use the same weights which equalize the contribution from each country-pair and store combination.

			All Stores	Apple	IKEA	H&M	Zara
Panel A: Product Duration (weeks)							
(i)	World	Mean	21	36	42	12	11
		Median	12	26	13	10	12
(ii)	United States	Mean	36	46	73	12	11
		Median	15	31	48	10	11
Panel B: Percent of Products with Any Price Changes							
(iii)	World	All Products	15	27	25	12	7
		$\geq$ 12 months	58	53	61	75	-
		$\geq$ 24 months	63	54	68	-	-
(iv)	United States	All Products	23	32	38	21	3
		$\geq$ 12 months	54	50	55	99	-
		$\geq$ 24 months	60	48	65	-	-

Table 5: Statistics on the Duration of Products and Their Prices

Notes: Panel A reports the mean and median number of weeks in which products appear in our dataset for all countries in the world as well as just in the United States. Due to left and right censoring, these estimates are likely downward biased. Panel B reports the share of goods worldwide and for the United States which experience at least one price change. We report this statistic in the overall data and also in subsamples restricted to goods with at least one and two years of data. Both panels weight to equalize the contribution from each store and country combination.

			All Stores	Apple	IKEA	H&M	Zara
Panel A: Weighted to Equalize Contributions of Store and Country Pair Combinations							
(i)	All Bilaterals	All Intros	0.678	0.414	0.753	0.985	0.798
			(0.007)	(0.010)	(0.028)	(0.004)	(0.011)
(ii)	All Bilaterals	Filtered Intros	0.701	0.325	0.685	0.973	0.778
			(0.006)	(0.010)	(0.025)	(0.003)	(0.012)
(iii)	All U.S. Bilaterals	All Intros	0.693	0.493	0.910	1.021	0.971
			(0.030)	(0.035)	(0.048)	(0.017)	(0.026)
(iv)	All U.S. Bilaterals	Filtered Intros	0.799	0.414	0.936	1.030	1.068
			(0.022)	(0.035)	(0.042)	(0.018)	(0.034)
Panel B: Unweighted							
(i)	All Bilaterals	All Intros	0.827	0.446	0.833	0.991	0.831
			(0.006)	(0.011)	(0.030)	(0.004)	(0.007)
(ii)	All Bilaterals	Filtered Intros	0.829	0.302	0.732	0.978	0.821
			(0.006)	(0.009)	(0.016)	(0.003)	(0.009)
(iii)	All U.S. Bilaterals	All Intros	0.873	0.518	0.980	1.025	0.951
			(0.023)	(0.034)	(0.050)	(0.018)	(0.024)
(iv)	All U.S. Bilaterals	Filtered Intros	0.953	0.395	0.945	1.032	1.053
			(0.016)	(0.032)	(0.038)	(0.019)	(0.025)

Table 6: Regressions of Log RER at Introduction on Log NER

Notes: Table reports coefficient of regressions of the log RER at introduction on the log NER. Each good is only included in the regression on a single introduction date. Standard errors are clustered by store and week combination. We exclude the limited number of observations where $|q_{ij}^{I}| > 0.75$ or $|q_{ij}| > 0.75$. Panel A weights the observations to equalize the contribution of each store and country pair combination, while Panel B does not use any weights.
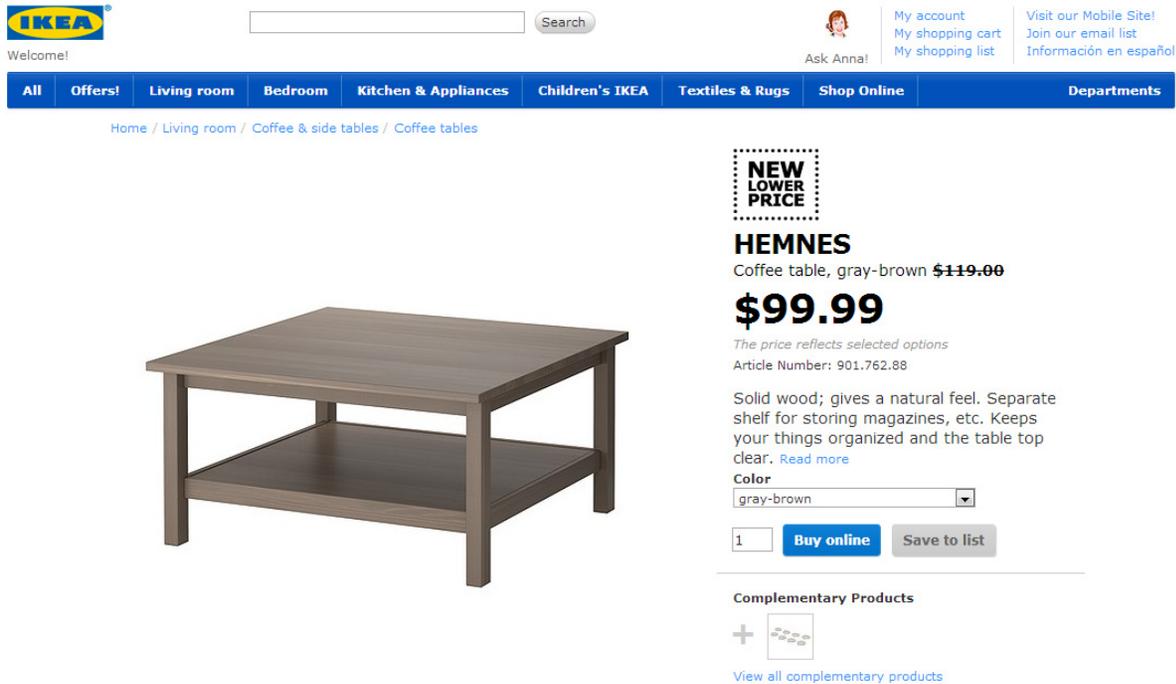
(a) IKEA Online

(b) IKEA in Store

Figure 1: Example of Online and Offline Prices for IKEA

Notes: "IKEA Online" image is a screen shot taken of a product found on IKEA's U.S. website. During that same week, the "IKEA in Store" picture was taken of the price of the identical product at the physical IKEA store located in Stoughton, Massachusetts. With few exceptions, all prices for all four retailers are identical online and offline in all countries, and this is just one example.
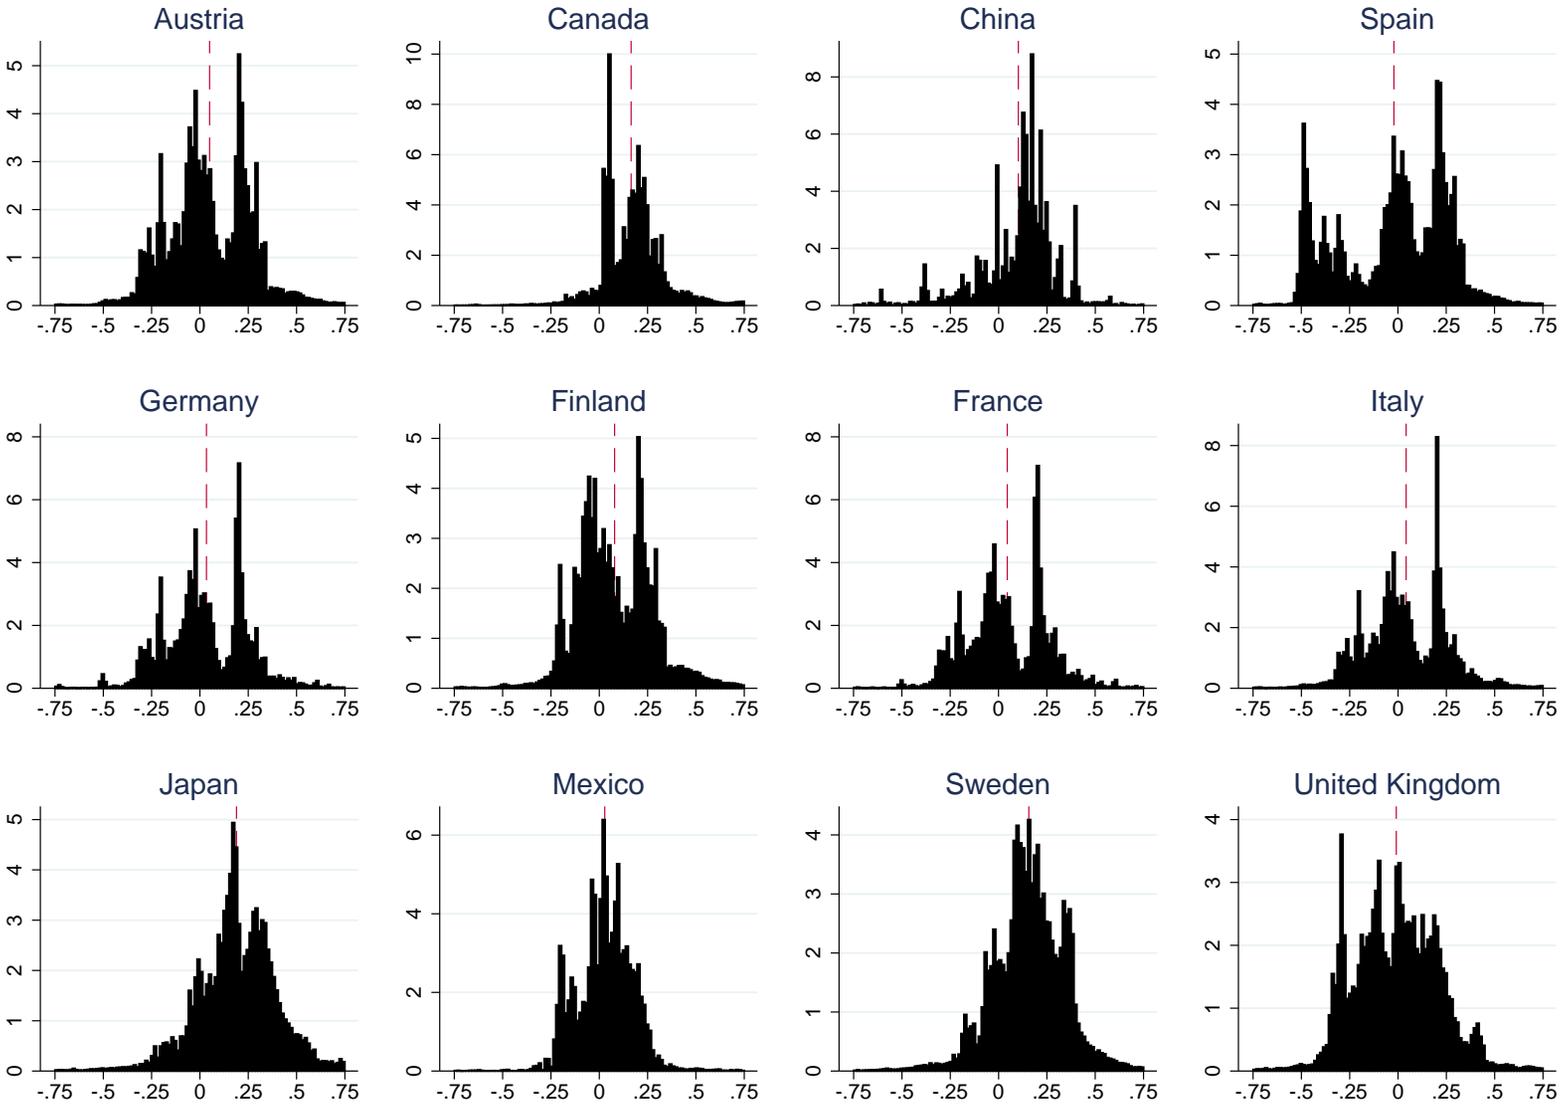
Figure 2: Good-level RERs q_{ij} for Various Countries (i) with the United States (j)

Notes: Figure includes all goods z and all weeks t for which q_{ij} is defined, with United States as country j and the other countries as country i. Histograms include frequency weights such that the contribution from each store is equalized for each country pair. We exclude the small number of observations where $|q_{ij}| > 0.75$. Dashed red vertical lines indicate the weighted average RER. Y-axes plot percents.
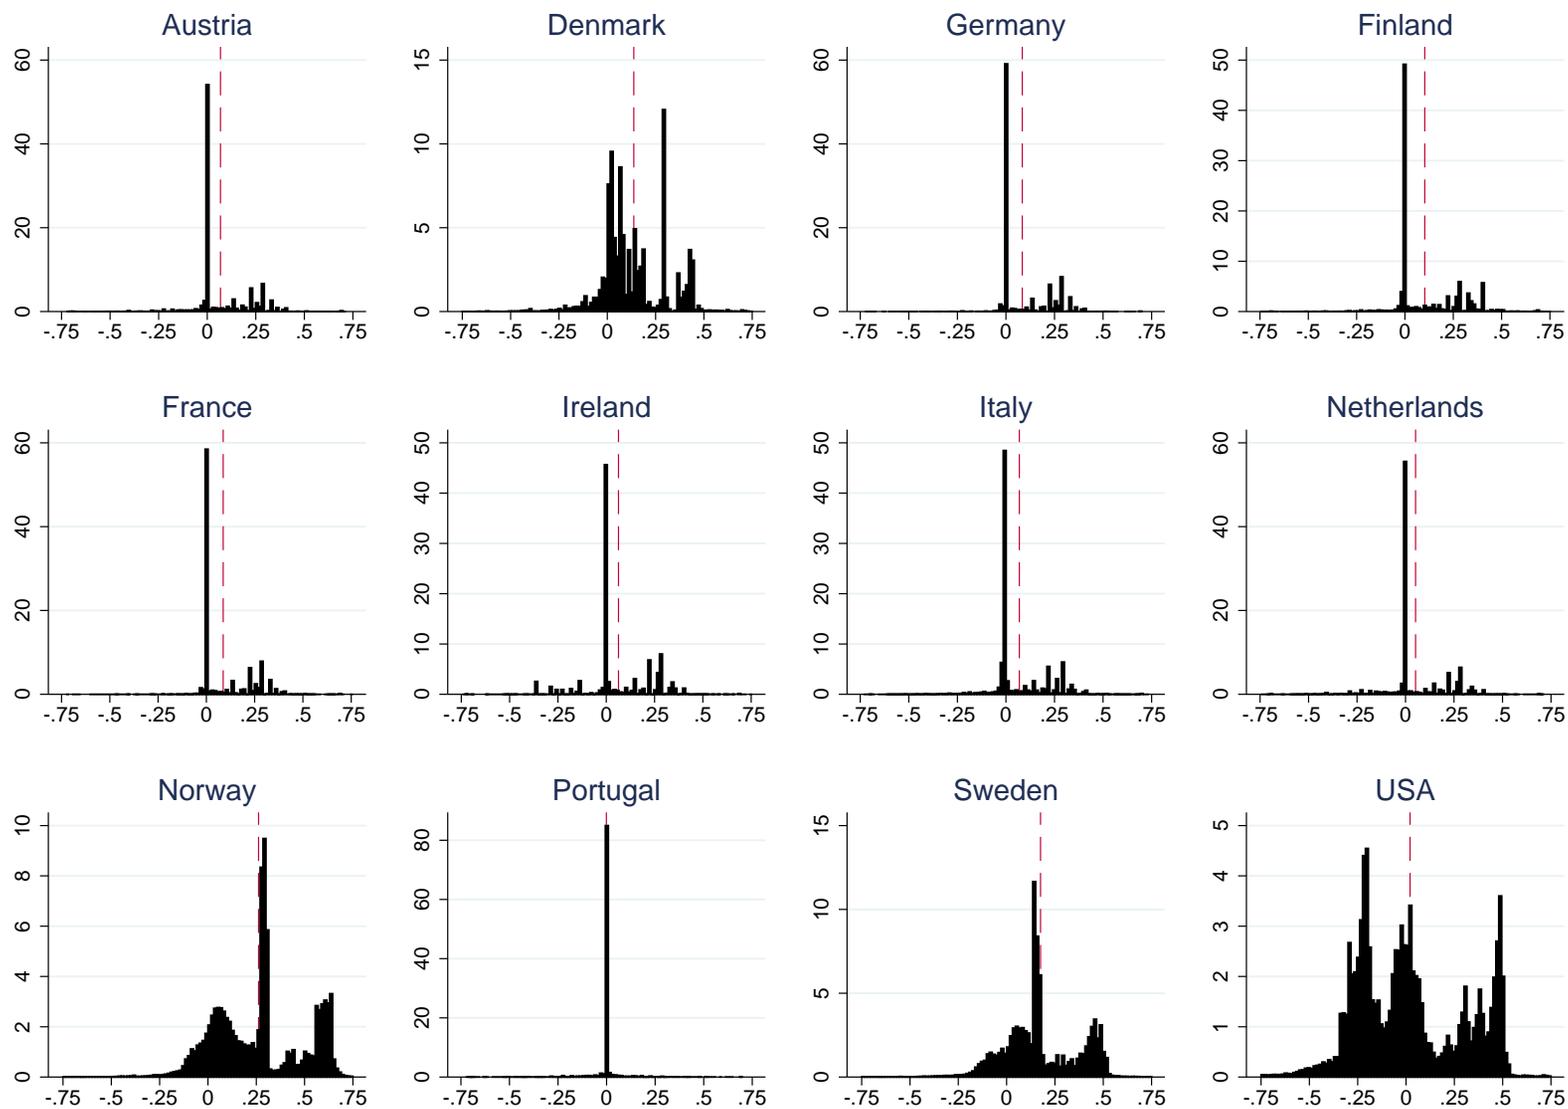
(a) Apple

(b) IKEA

(c) H&M

(d) Zara
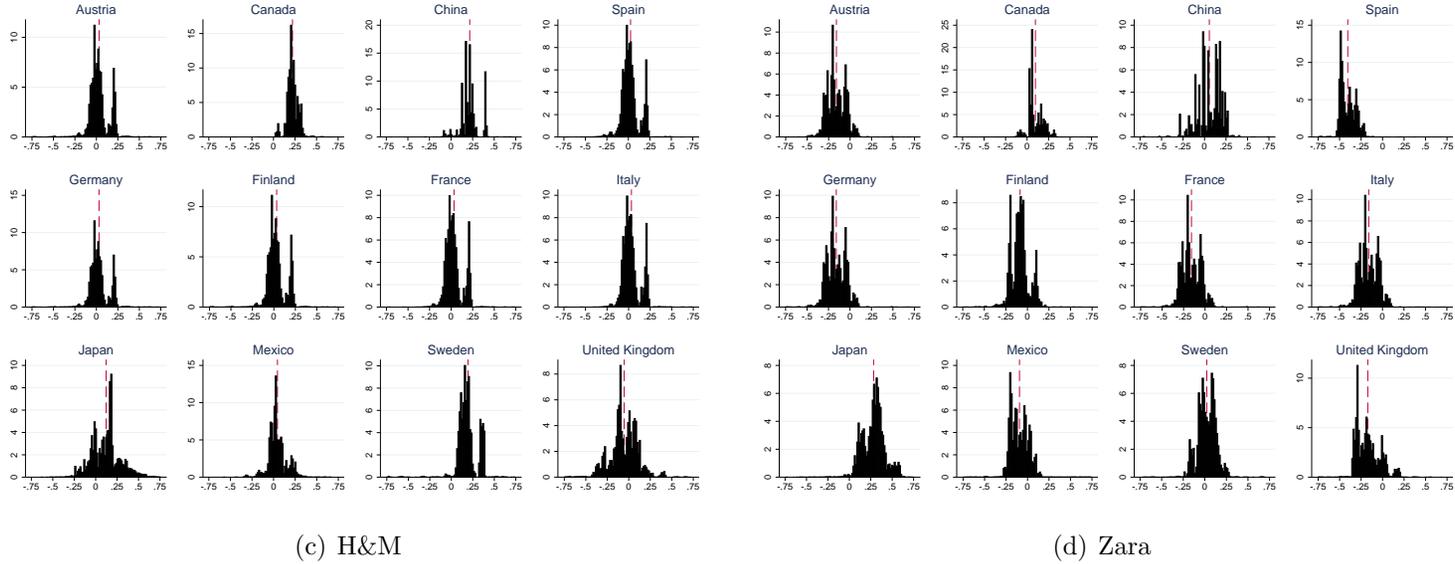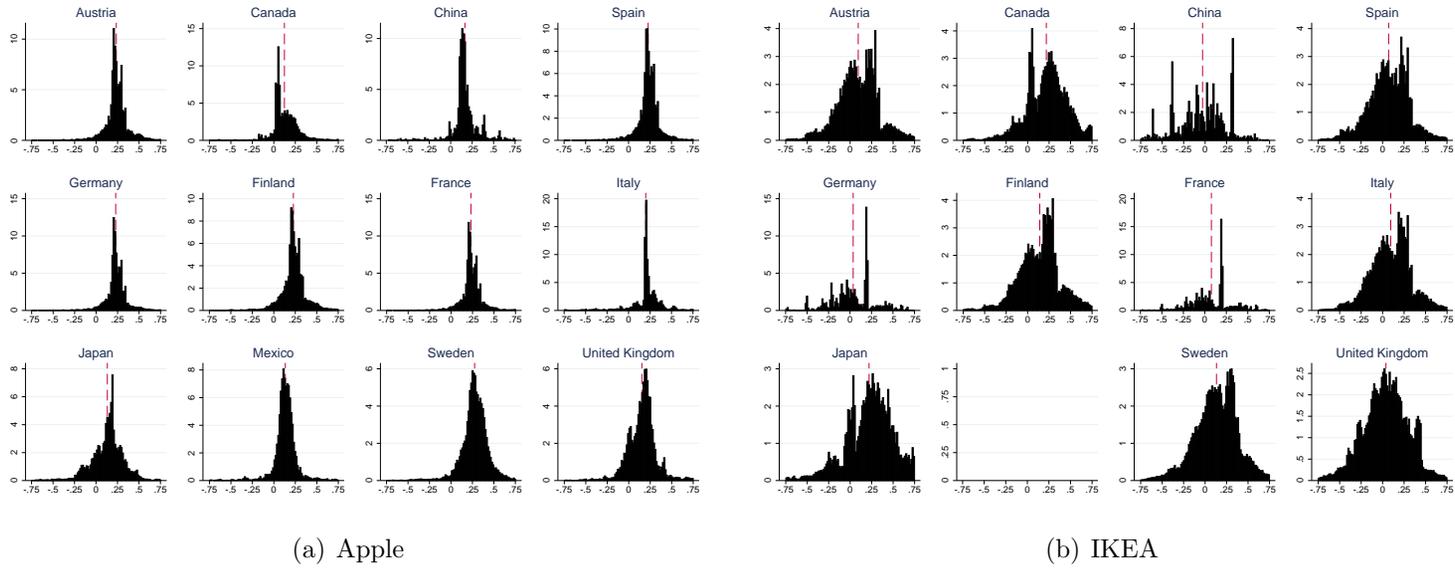
Figure 3: Good-level RERs q_{ij} for Various Countries (i) with the United States (j), by Store

Notes: Figure includes all goods z and all weeks t for which q_{ij} is defined, with United States as country j and the other countries as country i. All goods in each histogram are equally weighted. We exclude the small number of observations where $|q_{ij}| > 0.75$. Dashed red vertical lines indicate the weighted average RER. Y-axes plot percents.
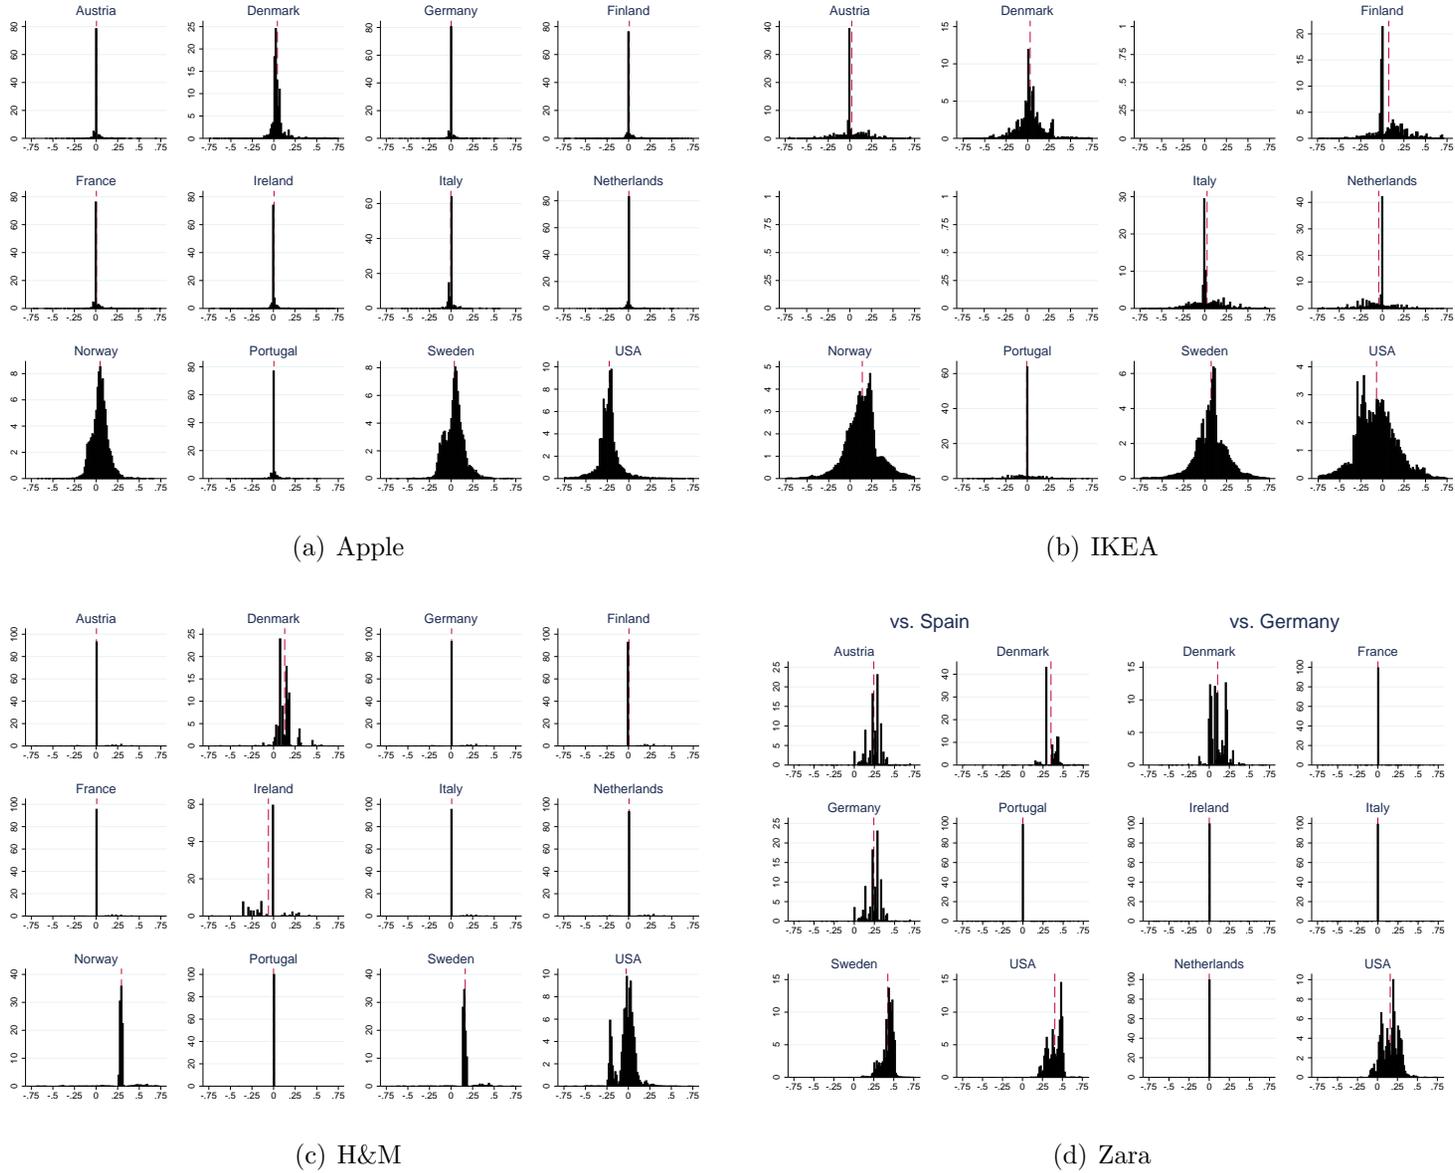
Figure 4: Good-level RERs q_{ij} for Various Countries (i) with Spain (j)

Notes: Figure includes all goods z and all weeks t for which q_{ij} is defined, with Spain as country j and the other countries as country i. Histograms include frequency weights such that the contribution from each store is equalized for each country pair. We exclude the small number of observations where $|q_{ij}| > 0.75$. Dashed red vertical lines indicate the weighted average RER. Y-axes plot percents.
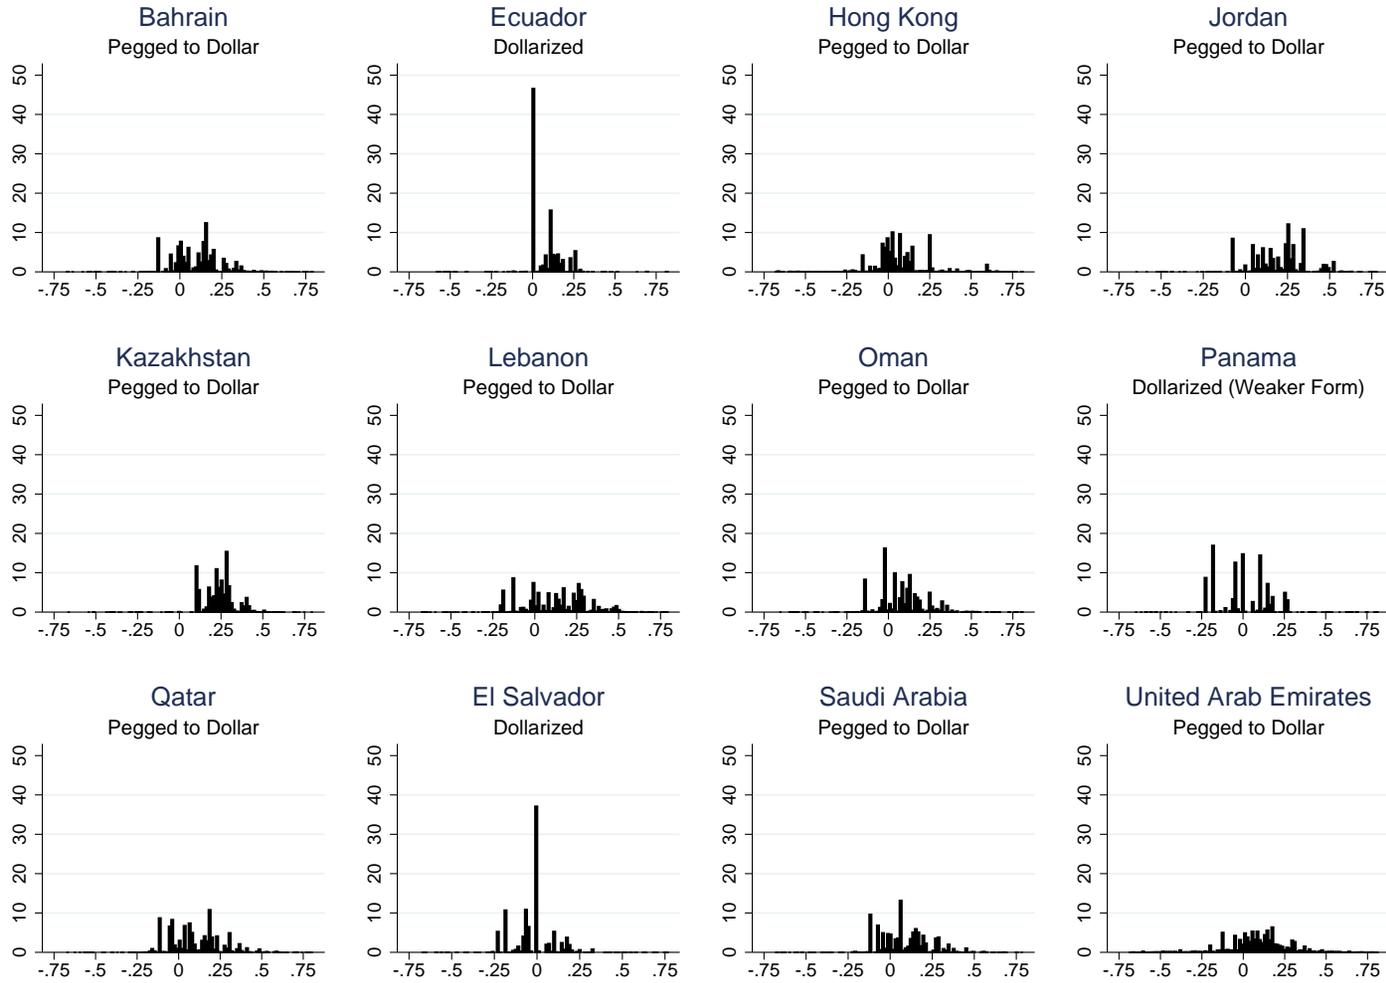
(a) Apple

(b) IKEA

(c) H&M

(d) Zara

Figure 5: Good-level RERs q_{ij} for Various Countries (i) with Spain (j), by Store

Notes: Figure includes all goods z and all weeks t for which q_{ij} is defined, with Spain as country j and the other countries as country i. All goods in each histogram are equally weighted. We exclude the small number of observations where $|q_{ij}| > 0.75$. Dashed red vertical lines indicate the weighted average RER. Y-axes plot percents.

Figure 6: Good-level RERs q_{ij} for Various Countries (i) with the United States (j)

Notes: Figure includes all goods z and all weeks t for which q_{ij} is defined, with United States as country j and the other countries as country i. We include 9 countries with exchange rate pegs with the dollar as well as three countries that use the U.S. dollar as their currency. Ecuador and El Salvador only have the dollar as legal tender, and prices in our data are quoted in dollars for these countries. We label Panama as "Dollarized (Weaker Form)" since dollars and balboa coins are both accepted as legal tender, and prices in our data are quoted in balboas for Panama. We exclude the small number of observations where $|q_{ij}| > 0.75$. Y-axes plot percents.
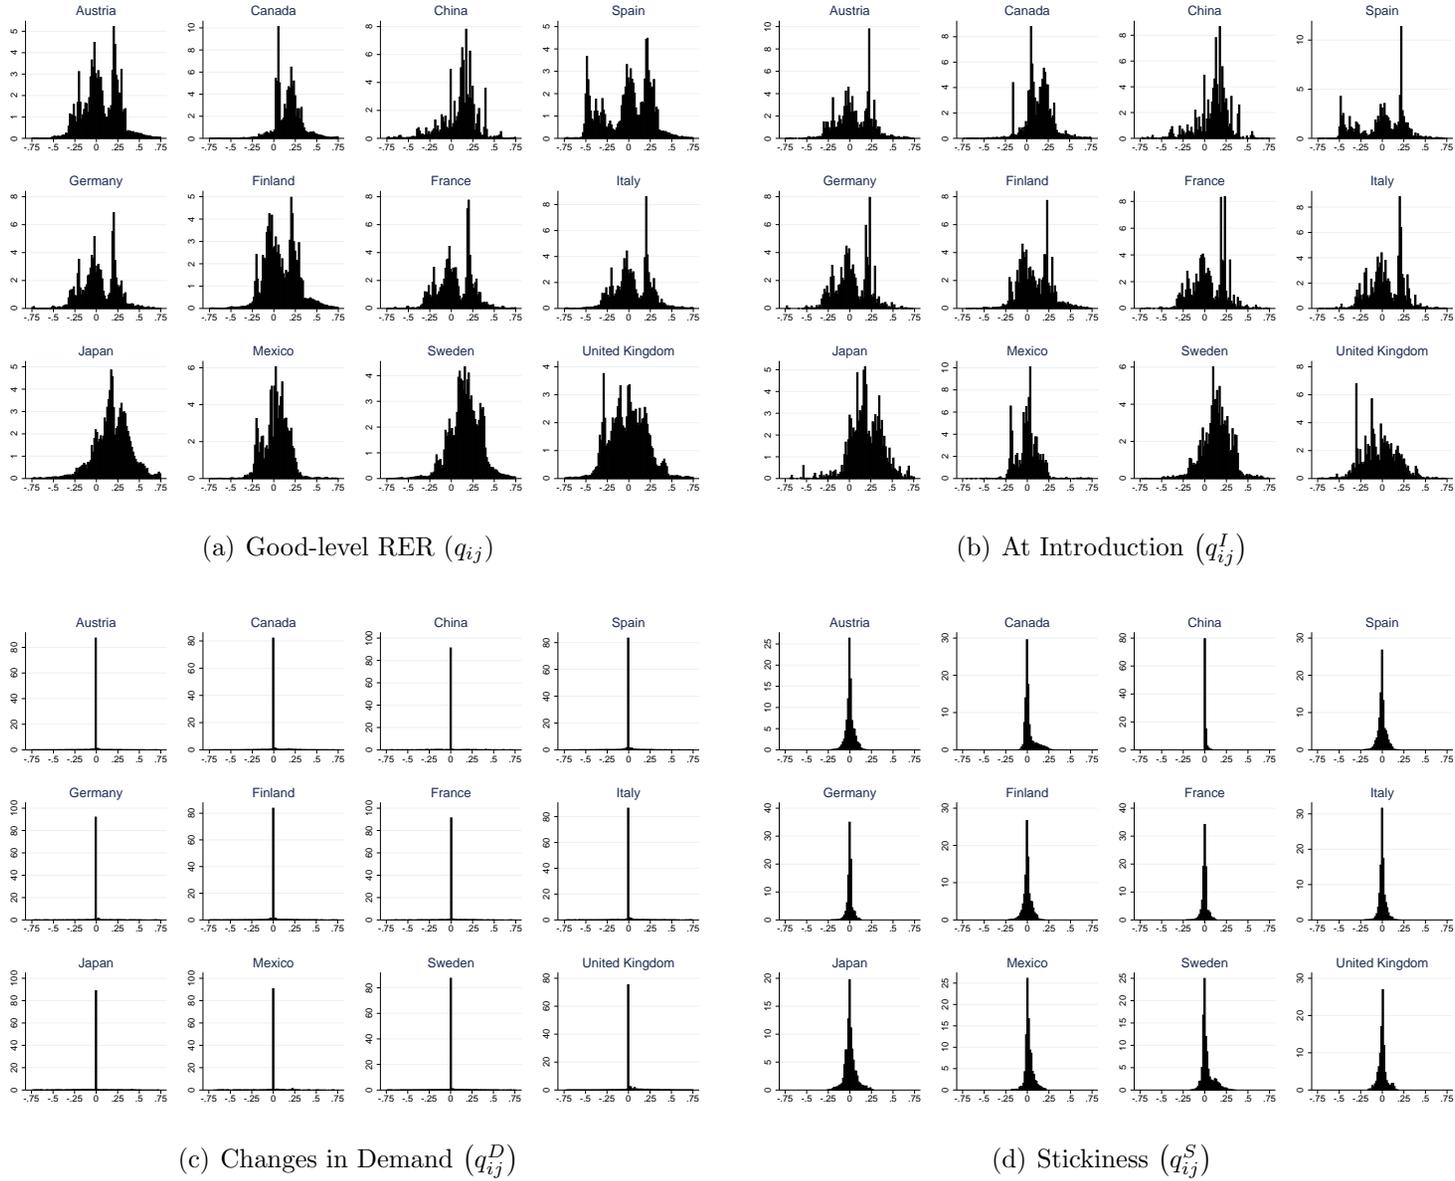
(a) Good-level RER (q_{ij})

(b) At Introduction (q_{ij}^{I})

(c) Changes in Demand (q_{ij}^{D})

(d) Stickiness (q_{ij}^{S})

Figure 7: Good-level RER Decomposition $q_{ij} = q_{ij}^{I} + q_{ij}^{D} + q_{ij}^{S}$ for Various Countries (i) with the United States (j)

Notes: Figure includes all goods z and all weeks t for which q_{ij} is defined, with the United States as country j and the other countries as country i. Histograms include frequency weights such that the contribution from each store is equalized within each country pair. We exclude the small number of observations where $|q_{ij}| > 0.75$. Y-axes plot percents.
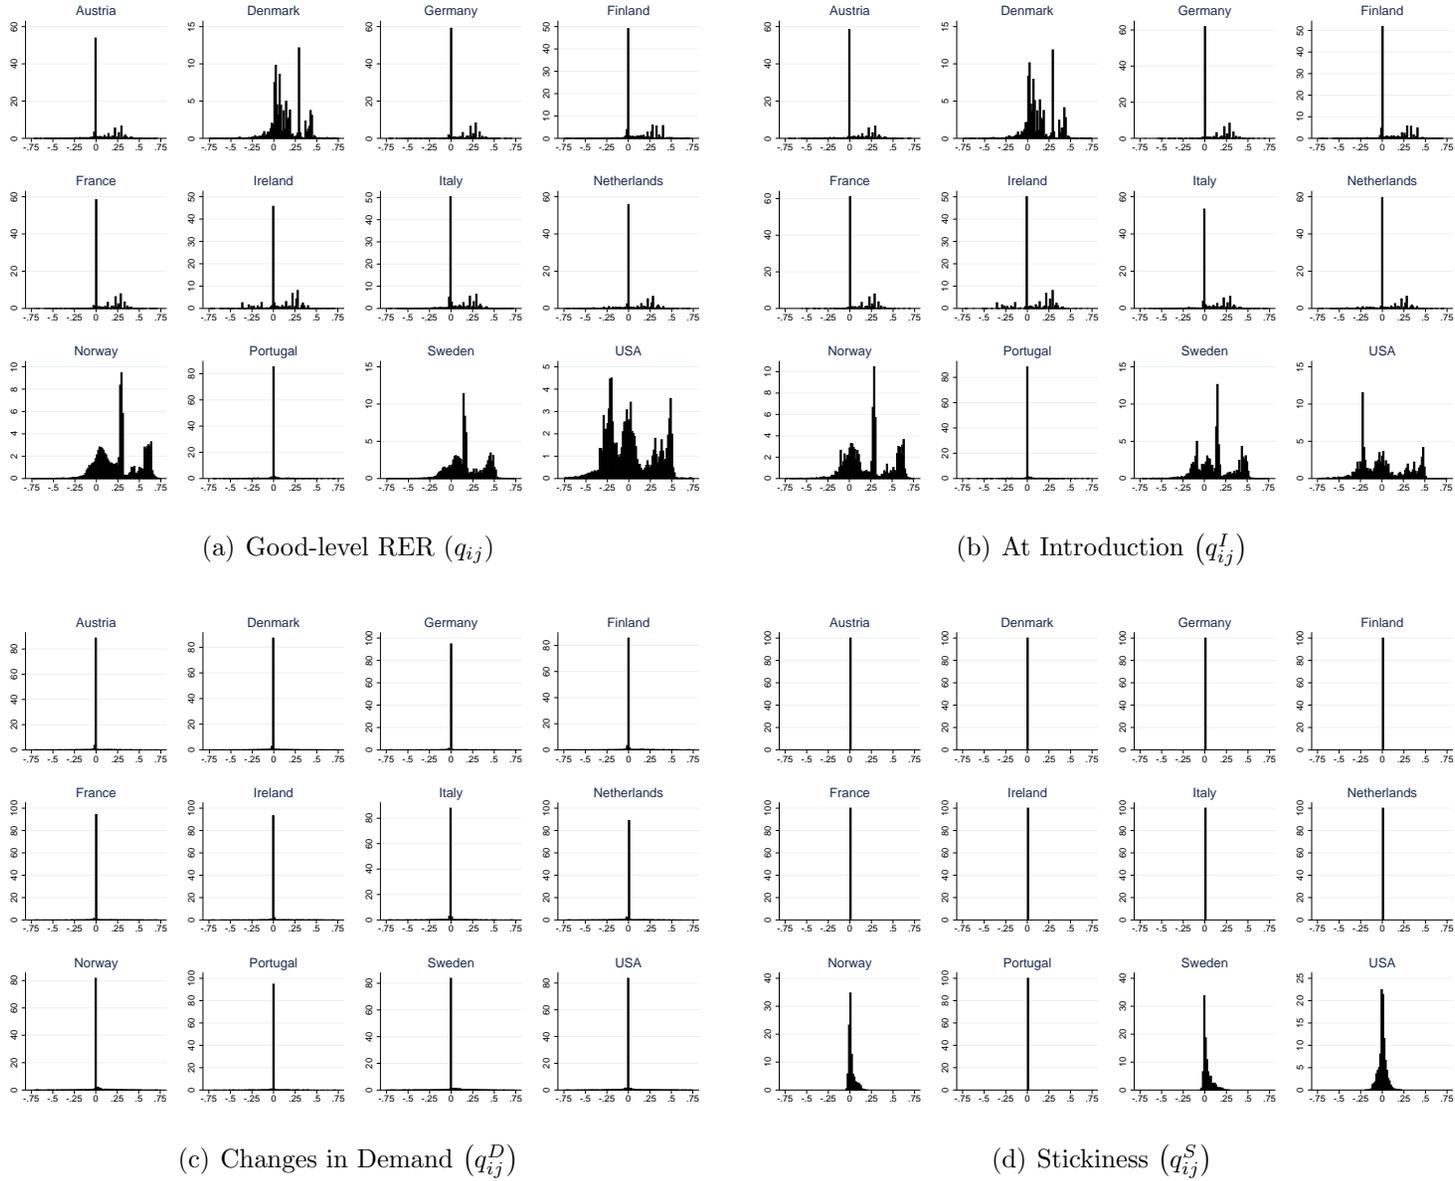
(a) Good-level RER (q_{ij})

(b) At Introduction (q_{ij}^{I})

(c) Changes in Demand (q_{ij}^{D})

(d) Stickiness (q_{ij}^{S})

Figure 8: Good-level RER Decomposition $q_{ij} = q_{ij}^{I} + q_{ij}^{D} + q_{ij}^{S}$ for Various Countries (i) with Spain (j)

Notes: Figure includes all goods z and all weeks t for which q_{ij} is defined, with Spain as country j and the other countries as country i. Histograms include frequency weights such that the contribution from each store is equalized within each country pair. We exclude the small number of observations where $|q_{ij}| > 0.75$. Y-axes plot percents.
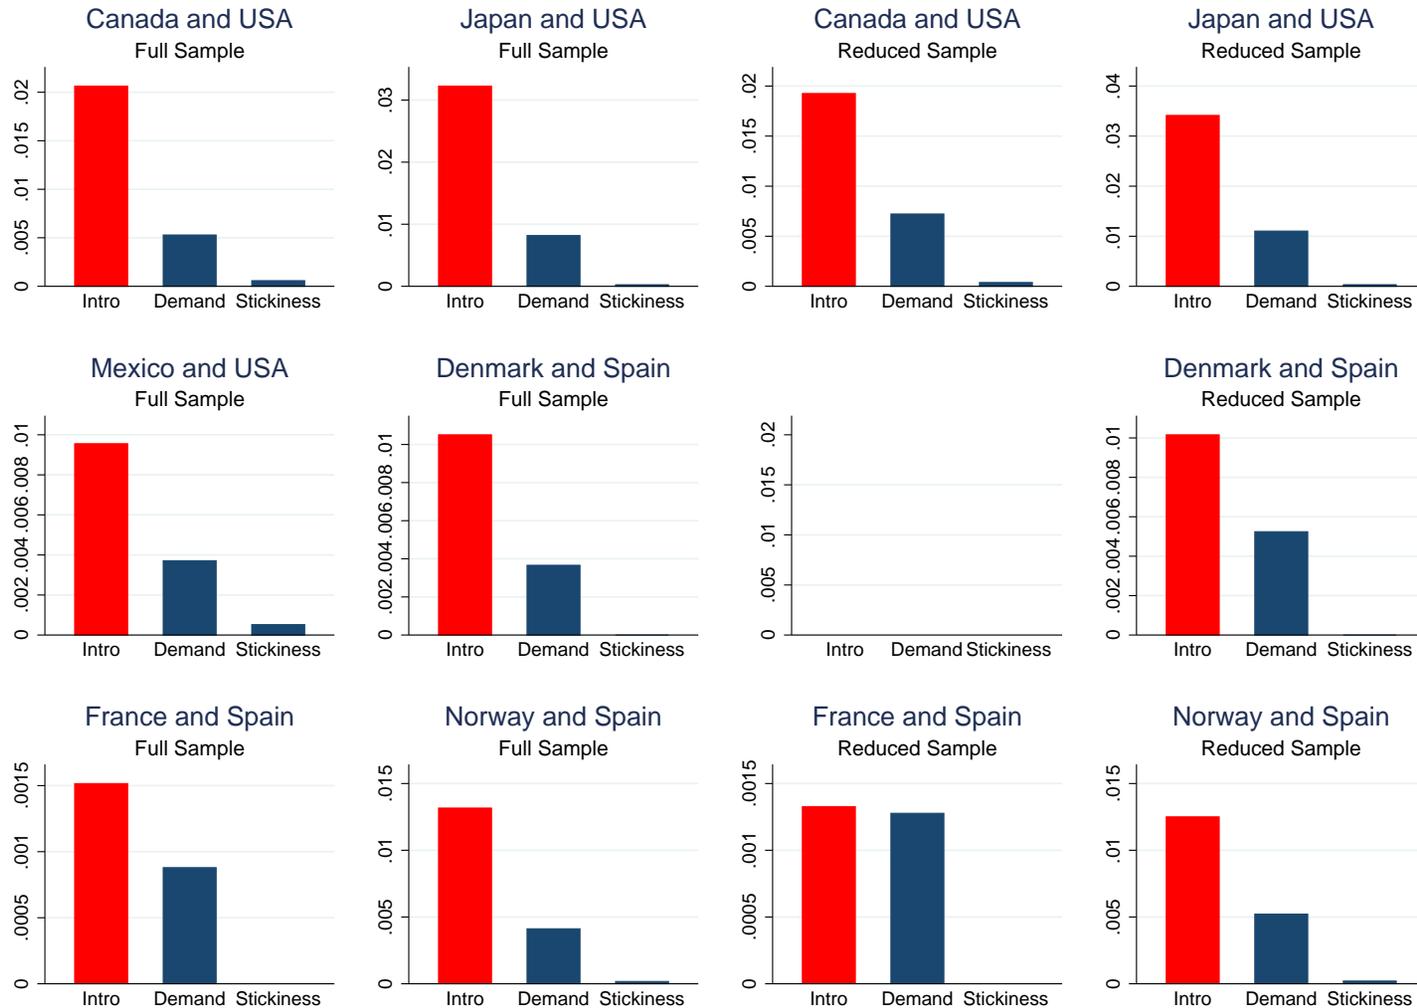
Figure 9: Decomposing the Cross-Section of Good-Level RERs q_{ij} for Selected Bilateral Pairs

Notes: Figure plots the three terms from the cross-sectional decomposition (5). The decomposition in the left two columns ("Full Sample") is calculated for each country pair at each date that contains at least 100 goods and then the results are averaged across all available dates. The decomposition in the right two columns ("Reduced Sample") is equivalently calculated but on the subsample of goods that ultimately have at least one price change in one of the countries and that remain in the sample for at least one year. Weights are used to equalize the contribution of all stores within each country pair. We exclude the small number of observations where $|q_{ij}| > 0.75$.
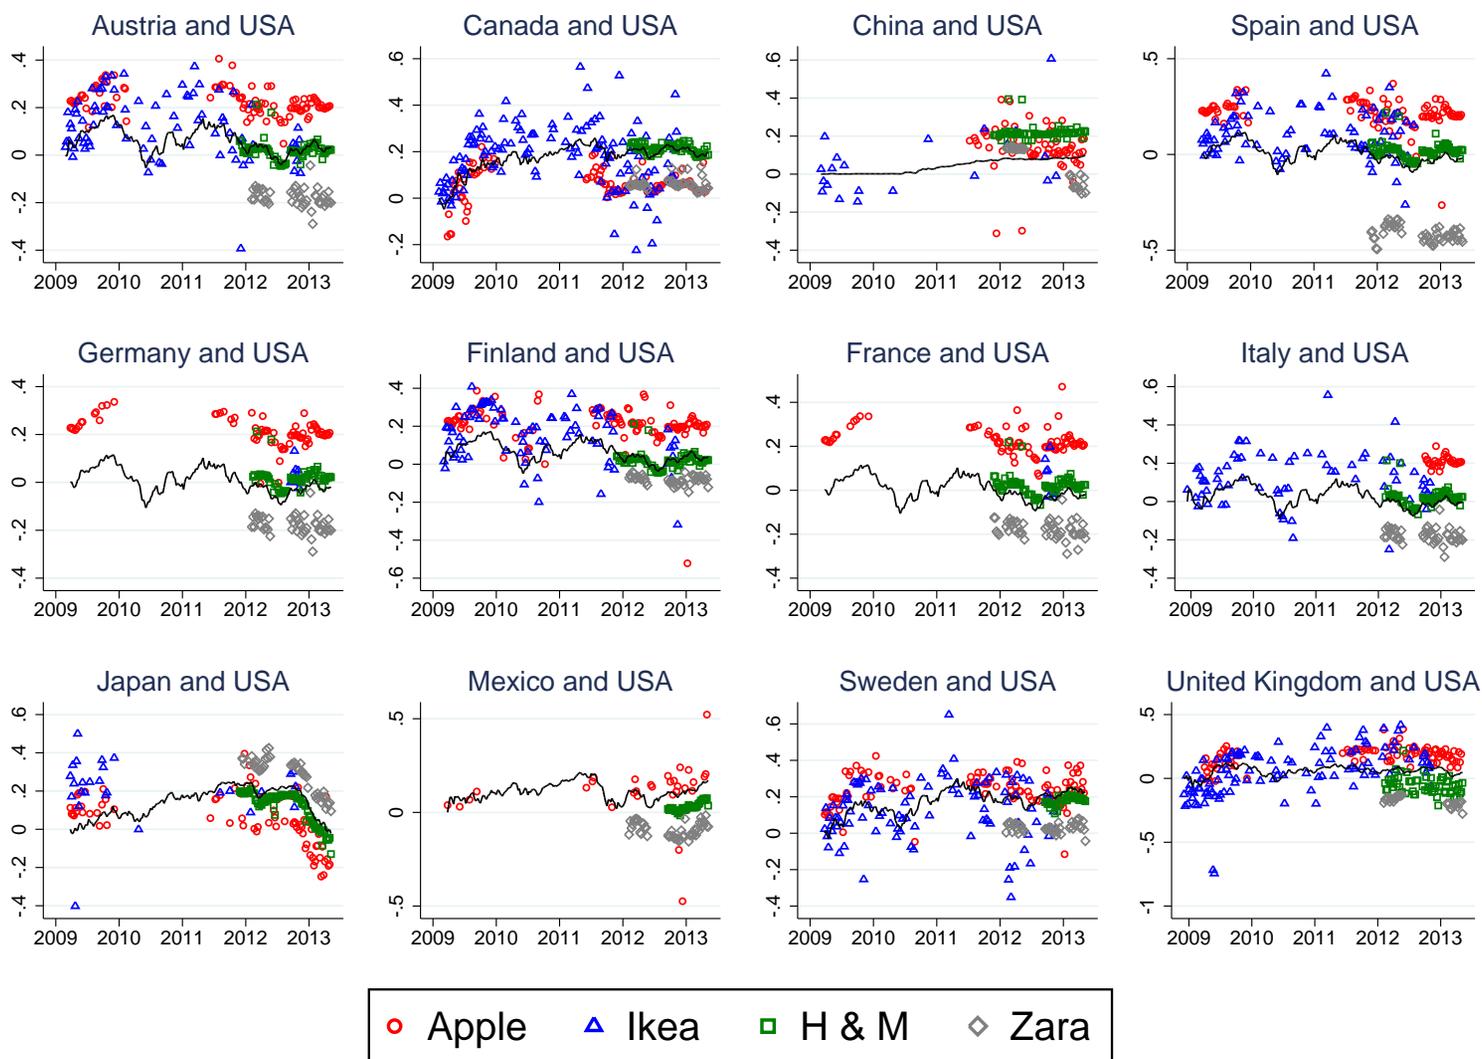
Figure 10: Evolution of Good-Level RERs at Introduction (q_{ij}^I) and the NER, Raw Data

Notes: Figure plots median log good-level RER at the time of introduction for each week and store combination for each bilateral relationship shown, weighted such that the contribution from each store is equalized for each bilateral pair. The black line plots the log NER, normalized to equal zero in the beginning of the sample. The figure is therefore informative about the time-series comovement between the RER and the NER, but not about the level. Any given good contributes (at most) to only one point in the figure. We drop the very limited number of observations where $|q_{ij}^I| > 0.75$ or $|q_{ij}| > 0.75$.
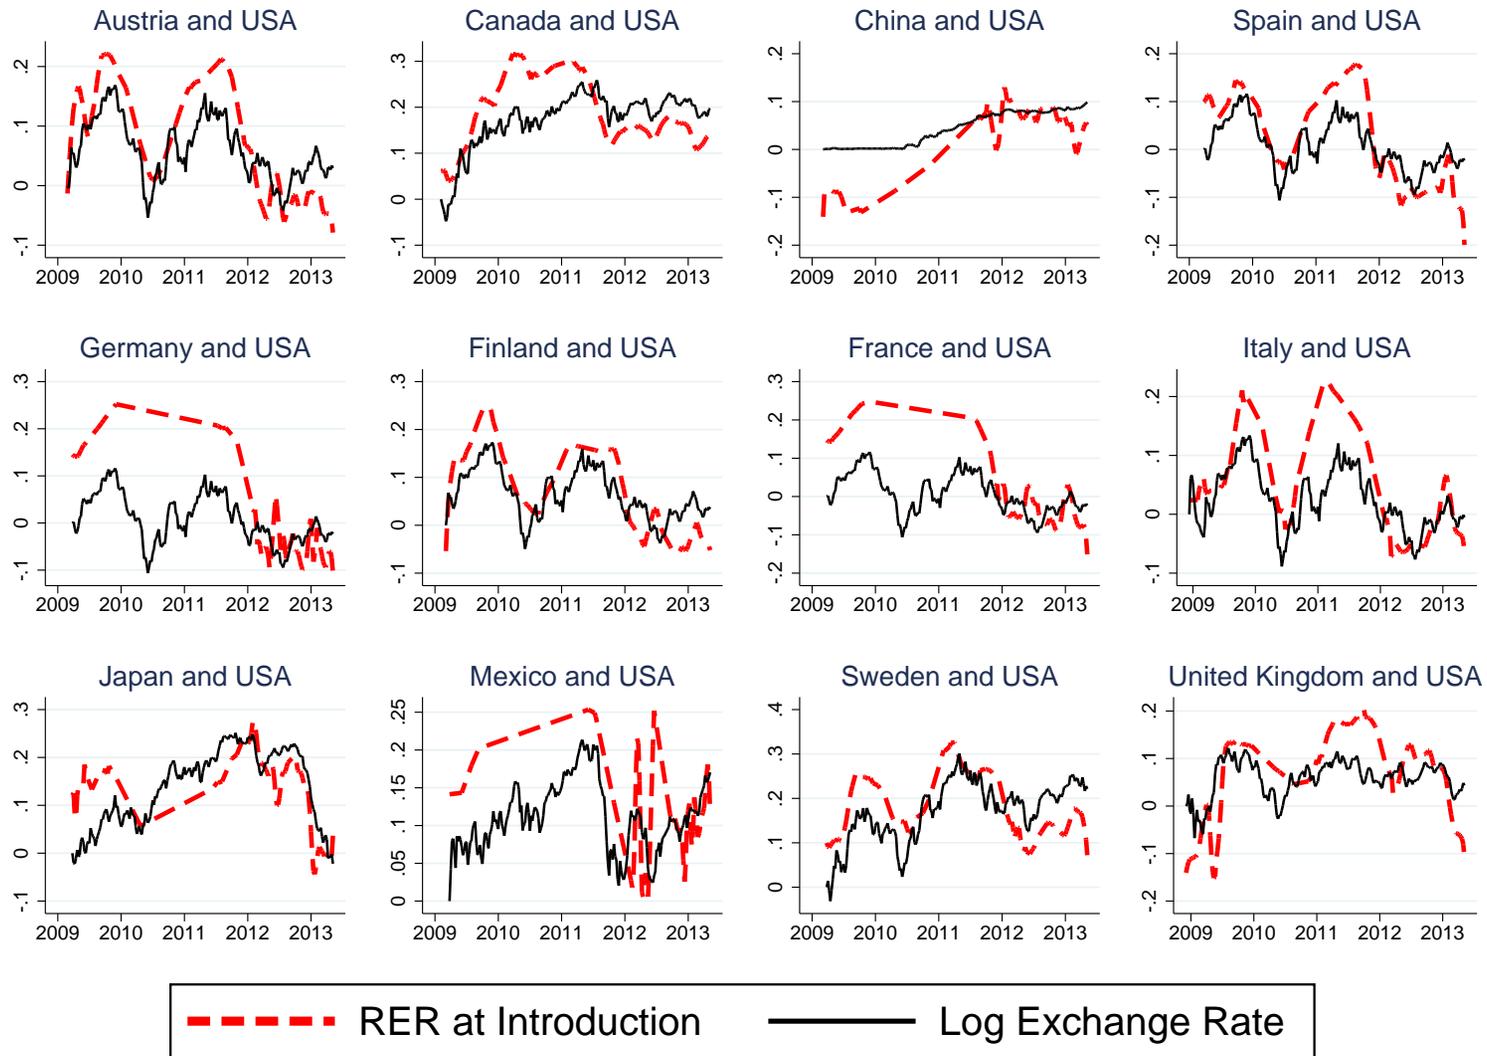
Figure 11: Evolution of Good-Level RERs at Introduction (q_{ij}^{I}) and the NER, Lowess

Notes: Figure plots with a dashed red line the non-linear fit (using Stata's "lowess" command with a bandwidth of 0.1) of the weighted median log good-level RER at the time of introduction for each week and store combination for each bilateral relationship, as shown in Figure 10. The black line plots the log NER, normalized to equal zero in the beginning of the sample.