Product Introductions, Currency Unions, and the Real Exchange Rate*

Alberto Cavallo
MIT

Brent Neiman
University of Chicago and NBER

Roberto Rigobon
MIT and NBER

December 2012

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 aggregate implications. First, in contrast to the prior literature, we demonstrate that the law of one price holds perfectly within the euro zone for thousands of goods sold by each of the retailers. Second, we find large deviations from the law of one price for these same goods outside of currency unions, even when the nominal exchange rate is pegged. For example, the Danish krone is pegged to the euro but Danish prices differ markedly from those in the euro zone countries. About three-quarters of the cross-sectional variation in law of one price deviations reflects differences in prices at the time goods are first introduced, as opposed to the component emerging from incomplete passthrough or from nominal rigidities during the life of the good. Third, we show that good-level real exchange rates measured at the time goods are first introduced move with the nominal exchange rate. This suggests the importance of real rigidities, in addition to nominal rigidities, for explaining the volatility of the real exchange rate.

JEL-Codes: E3, F3, F4.
Keywords: Law of One Price, Good-level Real Exchange Rate, Purchasing Power Parity, Scraped Internet Prices, Product Life Cycle.

*We thank Mark Aguiar, Fernando Alvarez, Chang-Tai Hsieh, John Huizinga, Pete Klenow, and Emi Nakamura for helpful comments and suggestions. Lionel Potier, Eric Hernandez, and David Sukhin provided outstanding research assistance. This research was supported in part by the Neubauer Family Foundation at the University of Chicago Booth School of Business.

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). The RER is perhaps the most important price in open-economy macroeconomics because its dynamics govern international shock transmission, the co-movement of business cycles, and the optimality of a country's choice of currency regime.[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 focus on three empirical findings. First, we demonstrate that the LOP holds almost precisely within the euro zone for thousands of goods, implying traded RERs approximately equal 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 across countries 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.

Second, we show that there are large magnitude deviations from the LOP 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 differ from those in Sweden, which has a floating exchange rate, and also differ from those in Denmark, which pegs its currency to the euro. We complement this evidence by showing that the LOP with the United States holds far more pervasively for dollarized countries like Ecuador than for otherwise similar countries like Panama, which have their own currency but peg it to the U.S. dollar.

If NER volatility is not the key driver of LOP violations, what is? To answer this question, we introduce a framework to decompose the good-level RER into the RER at the time a good is

[1] Cassel (1918) first used the term "Purchasing Power Parity" (PPP) to describe the condition in which there were 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 persistance and volatility of the RER, what has been termed the "PPP puzzle."

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. We find that the vast 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. For example, they would not expose differences in the RER behavior for pegged countries compared to countries inside the euro zone because after goods are introduced and prices are set, their subsequent dynamics are similar. By contrast, our construction of RERs using price *levels* shows different dynamics for pegged bilateral relationships like Spain-Denmark compared with pairs within a currency area like Spain-Germany.

Third, building on our result that most of the cross-sectional dispersion in the LOP occurs at the time of good introduction, we study how relative introduction prices for countries without a common currency evolve with the NER. Surprisingly, we find that the RER at the time of introduction moves together with the nominal rate. This is evidence against a model in which RER innovations are eliminated at the time of a price change as, after all, a price change inherently occurs when a new product is introduced. It also implies that RER volatility would not be dramatically reduced if product introductions were included in price index construction.

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 conditions under which final good producers worry about price arbitrage. Second, they improve our understanding of traded real exchange rate dynamics, critical for understanding business cycle dynamics in the open-economy. For an important set of products, local currency pricing is the most appropriate assumption in sticky price models, even for periods of time longer than the life of a typical product. This has critical policy implications, including for a country's choice of exchange rate regime and the conduct of monetary policy.[2] 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 – in levels and in changes – may emerge endogenously to the choice of currency regime. Third, our finding that NERs and RERs move together even at the time of product introduction stands in sharp contrast to the pricing behavior in "menu cost" models without variable markups, where all rigidity is due simply to the lack of price adjustment. In this sense, our results are also important for closed-economy

[2]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.

macroeconomic models aiming to understand pricing dynamics and monetary non-neutrality.[3]

Our data include daily prices spanning various subsets of 83 countries and various time periods from December 2008 to July 2012 for all products sold by Apple, IKEA, H&M, and Zara. 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, lagging Samsung and Nokia due to smaller sales of mobile phones. 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, smaller only than Nike and Adidas. Together, their sales exceed $30 billion.

The facts that these companies are among the world's very largest retailers, are headquartered in three different countries, and cover three different industries gives us confidence that inference from our data is appropriately applied to the broader basket of branded and traded goods and is highly relevant for understanding international macroeconomic dynamics. The pricing patterns we identify cannot be oddities associated with a particular firm's, industry's, or country's characteristics. More than 20 percent of U.S. consumption expenditures in goods is on the industries covered in our data.

Studying online prices has the obvious advantage of allowing for the collection of enormous amounts of data at very high frequency. We also 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 confirmed to us that the online and offline prices are identical up to shipping costs and, in limited instances, local taxes or store-specific special promotions. Additionally, as discussed below, we visited the retail stores in the United States to further confirm this to be the case. 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.

[3]Before the euro's introduction in 1999, popular discussion of its potential benefits and academic research often focused on increased competition and the cross-country convergence of prices. For example, Goldberg and Verboven (2005) find evidence of convergence in auto prices after the introduction of the euro, while Parsley and Wei (2008) do not find such evidence in price data on the Big Mac Meal and its ingredients. 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.

An additional advantage of using online data is that it is easier to match products using the information on web pages rather than the bar codes or supermarket product ID codes used in other influential analyses of good-level real exchange rates between the United States and Canada including Broda and Weinstein (2008), Burstein and Jaimovich (2009), and Gopinath, Gourinchas, Hsieh, and Li (2011). For example, 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 greater confidence that composition differences are not important for our key results. By comparison, these other studies typically match less than 5 percent of the total goods.

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 real exchange rate 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, many papers have focused on explanations for LOP deviations or 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 persistance 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), Gopinath, Itskhoki, and Rigobon (2010), and Fitzgerald and Haller (2012). None of this work emphasizes price levels or good-level RERs at the time of product introductions and therefore does not comment on what our analysis suggests is the source for the bulk of LOP violations.

Finally, there are some papers which have looked at disaggregated price data, including in levels.[5] In addition to the papers mentioned above studying prices between the United States and Canada, 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

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

[5]A closely related 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 Borraz, Cavallo, Rigobon, and Zipitria (2012) and Cosar, Grieco, and Tintelnot (2012).

"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 persistance of LOP deviations in the cross-section increase with the importance of the distribution margin. Baxter and Landry (2012) also studies 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 and destruction, but do not report either of the first two findings we focus on in this paper. They report a statistic closely related to our third finding, though do not focus on high frequency movements in the nominal exchange rate as their data is at an annual frequency. Finally, our emphasis that conventionally measured RERs omit the information contained in the price levels when a good is introduced closely relates to work by Nakamura and Steinsson (2012) and Mandel, Gagnon, 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. We restrict our data to the prices for four global retailers where we are able to precisely match prices of identical goods sold in many geographies. We are able to exactly match nearly one hundred thousand unique items across dozens of countries because the firms' web pages organize products using their own company-wide product ID codes. Further, using information from product web addresses (or URLs), we can allocate the products to nestable categories with varying degrees of fineness.[6]

Prices are generally quoted inclusive of taxes and exclusive of within-country shipping costs. The United States is the one large exception, as prices are quoted there exclusive of state-specific sales taxes. We therefore adjust all U.S. prices upward by 6.5 percent to reflect the average combined state and local rates in 2012.[7]

The data include daily prices for the four retailers in some subset of 83 countries during some

[6]For example, the Apple products at http://store.apple.com/us/browse/home/shop_ipad/ipad_accessories/cases can be categorized as cases, in the subcategory of iPad accessories, in the category of iPads.

[7]State sales taxes are charged on internet transactions in the United States when online retailers also have a physical store in the state, as is the case for our retailers in most large U.S. cities. We obtain information on state and local rates for the United States from The Tax Foundation and 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.

subset of the period from December 2008 to July 2012. Table 1 gives a basic description of the country, product, and time coverage in our data. Row (i) indicates that we track prices for about 90,000 products, including 9,000 for Apple, 60,000 for IKEA, 9,000 for H&M, and 11,000 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. We do not have purchase quantities or individual product weights, so all our analyses apportion equal weight to all goods within each store and equal weight to each available store. Each individual IKEA product price, for example, is treated as containing less information than each individual Apple product price. In the Appendix, which can be found on the authors' web pages, we provide details on product, time, and country coverage for each store in our data.

Scraping errors or changes in these companies' web pages occasionally create missing price observations. We interpolate between observed prices with the assumption that prices remain unchanged until a change is observed. This is a reasonable assumption because, as we elaborate below, the prices are highly sticky and do not exhibit high frequency sales behavior seen in other pricing contexts (where a price changes and then returns exactly to its previous value). We exclude the roughly 1 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 web scraping process, our assembly and cleaning of the data, and additional summary statistics, quality checks, and robustness tests are included in the Appendix.

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 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. Fourth, with such a large volume of data that includes multiple product cycles, we can assess and reasonably calibrate the role of product entry and exit. Finally, in measuring prices at a very high frequency, we can more confidently pinpoint

the quantitative role of nominal rigidity in contributing to the RER.

Gopinath, Gourinchas, Hsieh, and Li (2011), Broda and Weinstein (2008), and Burstein and Jaimovich (2009) match identical goods across the United States and Canada to study relative price dynamics. 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, the four stores we focus on sell a majority of their goods in multiple countries, while the matched products in those studies constitute only a very small share of the total observed products.[8]

There are two 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.[9] 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. In fact, this also held true for the one item in those 10 from IKEA which happened to be selling at a discount relative to the previous year's price. 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, listed at the same $99.99 price. In sum, 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.

Next, one might wonder how representative these four retailers are relative to the entire basket of tradable consumption. We may not learn much about the behavior of auto prices or global

[8]Gopinath, Gourinchas, Hsieh, and Li (2011) study 4221 products that are sold by a supermarket chain in both the United States and Canada, which represents only 3.3 percent of the total number of products sold by that store in either the United States or Canada. Broda and Weinstein (2008) use scanner data and find that U.S. city pairs typically have roughly 10,000 matched UPCs, Canadian region pairs have roughly 25,000 matched UPCs, and cross-border pairs typically offer 1,500 matches in their data.

[9]H&M wrote in an email that "H&M website and store prices are the same," other than occasional in-store specials that might not be available on the web site. Zara customer service emailed, "our store and online shop share the same pricing," and IKEA emailed, "IKEA guarantee[s] the same price online for the catalog products." On a phone call, an Apple customer service representative confirmed that prices in Apple's online and retail stores are identical.

commodities from our data, but branded technology, furnishings, and apparel are particularly interesting to study because they are generally produced in the same plant or location, exhibit significant price stickiness, and are often not sold by any other retailers. 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. We use the CPI weights from the U.S. Bureau of Labor Statistics to calculate that, if taken as representative of these three categories of goods, our data cover more than 20 percent of final consumption expenditures on goods. Finally, given our data include four different companies, three different industries, and three different headquarters countries, it is unlikely that our results simply reflect the idiosyncrasies of any particular company, industry, or country. The fact that all four of our retailers exhibit different pricing policies for countries within the euro zone compared with those that simply peg to the euro suggests that a majority, if not large majority, of global firms would do the same.

3 Good-level Real Exchange Rates

We now describe an economic environment that will allow us to more precisely define good-level RERs and to demonstrate why they are informative about aggregate RERs. We show that good-level RERs vary significantly outside of currency unions, even when the exchange rate is pegged. By contrast, the law of one price holds extremely well within the euro zone and holds very frequently when comparing prices in dollarized economies with those in the United States.

3.1 Economic Environment

Consider a world with many countries $i = 1..I$. Each country i has a representative consumer who derives utility from consumption at time t of each of a large number of traded goods z within narrowly defined sectors $n = 1..N$. Let $\Omega_i^n(t)$ denote the set of sector n goods available in country i at time t.

Each good is manufactured in a single plant in some country and sold throughout the world, but shipping the good from the plant to each country requires payment of a good-country specific fixed cost. The union over sectors and countries of these sets will vary over time because of unmodeled product innovations which result in new varieties and cause demand for some products to drop below that required to cover the fixed costs. The set of available varieties might differ across countries within the same time period given heterogeneity in these fixed costs. Preferences are homothetic and symmetric in all countries.

Let $p_i^n(z,t)$ denote the log price in local currency of good z in sector n and country i. A first-order Taylor-series approximation around the steady state expenditure weights to the logarithm of the ideal price index in each country is (up to a constant):

$$\hat{p}_i(t) = \sum_{n \in N} \omega_i^n \hat{p}_i^n(t),$$

where

$$\hat{p}_i^n(t) = \sum_{z \in \Omega_i^n(t)} \omega_i^n(z)\, p_i^n(z,t),$$

with ω_i^n denoting the sector's steady state share of aggregate spending and $\omega_i^n(z)$ denoting good z's share of the sector's steady state spending. The log RER $\hat{q}_{ij}(t)$ is defined as the difference between the approximation to the log price index in country i and that in country j after translating all prices into a common currency. 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.

We assume all goods have the same steady-state expenditure shares in all countries in which they are consumed and therefore write:

$$\hat{q}_{ij}(t) = \frac{1}{N} \sum_{n \in N} \hat{q}_{i,}^n(t),$$

with

$$\hat{q}_{ij}^n(t) = \omega_{ij}^n \sum_{z \in \Omega_{ij}^n(t)} q_{ij}^n(z,t) + \left(1 - \omega_{ij}^n\right) \sum_{z \in \Omega_{i-j}^n(t)} \left(p_i^n(z,t) - e_{ij}(t)\right) - \left(1 - \omega_{ij}^n\right) \sum_{z \in \Omega_{j-i}^n(t)} p_j^n(z,t),$$

where ω_{ij}^n is the steady state expenditure share of the goods which are consumed both in countries i and t. We use the notation $\Omega_{ij}^n(t) = \Omega_i^n(t) \cap \Omega_j^n(t)$ and $\Omega_{i-j}^n(t) = \Omega_i^n(t) - \Omega_j^n(t)$. The term $q_{ij}^n(z,t)$ is the log of the good-level RER:

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

and will equal zero when the LOP holds.

3.2 Law of One Price and Floating Exchange Rate Regimes

Figure 2 plots the distribution of the log good-level RERs q_{ij}^n pooling all goods z and weeks t for various country i's with country j fixed as the United States. 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 points 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}^n 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 points 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 U.S. Note that even in China, whose NER with the U.S. dollar has been relatively stable, good-level log RERs diverge significantly from zero. These patterns represent aggregations across all four retailers. 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 Italy and the U.S. for Apple, are missing due to lack of country data for a particular store. 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.

3.3 Law of One Price and Currency Unions

By contrast, we find compelling evidence that the LOP holds nearly perfectly in European countries that share a single currency and holds quite well between countries that use U.S. dollars as their domestic currency.

3.3.1 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, Austria, Germany, Finland, France, Ireland, Italy, Netherlands, and Portugal are members in the euro zone, a single currency area. The prices for tens of thousands of distinct products in those countries are almost always identical and we therefore see huge mass at zero in these histograms (note the differences in scales of the y-axes). It is not the case that consumers in one country can simply order directly from other 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. Additionally, there is no euro zone law mechanically requiring retailers to harmonize prices. This is the most compelling evidence that we are aware of in the academic literature documenting the LOP holding across countries for the same traded differentiated good.

We use Spain as the base country because it, unlike Germany, has prices for all four stores in our data. As we discuss below, however, Zara divides the euro zone into two regions: one with Spain (including the Canary Islands) 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 points 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.

Prices in Denmark, Norway, Sweden, and Switzerland (not shown), by contrast, do not exhibit this same adherence to the LOP. These countries are also parts of Europe with similar geographies and demand structures (and product market regulations), but are outside of the euro zone. Their histograms look closer to that of the United States. Again, note that Denmark, which pegs against the euro but is not in the currency union, has a distribution of good-level RERs with Spain characterized by a broad support.

A large share of these goods that are sold in multiple countries are likely produced in a single plant at the same marginal cost.[10] 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 developed European countries like Spain and Norway. However, companies forgo this markup variation within the euro zone.[11] This implies that the crucial barrier to arbitrage 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 seems highly similar to these differences between Spain and Norway or Sweden. Rather, it implies that companies believe that having to translate a price into different currency units is the most salient friction, even if the different currency units can be translated at a fixed rate as with the pegged value of the Danish krone and the euro.

[10] 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."

[11] We reiterate that these prices are inclusive of sales taxes, which exhibit variation across time and space, further implying that companies are forgoing optimal 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 to 21 percent at the end of our data. These country-specific changes did not produce changes in the degree to which the LOP held for prices inclusive of taxes in 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 almost perfectly for goods sold by Apple, IKEA, and H&M across the euro zone.[12] As mentioned above, Zara's prices are the same within two blocks in the euro zone, but different across these two blocks and across all countries outside of the euro zone. The left two columns of Figure 5(d) underneath the text "vs. Spain" shows that the LOP holds perfectly between Spain and Portugal. The right two columns underneath the text "vs. Germany" shows that the LOP holds perfectly 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 distribution looks similar to that between Spain and the United States. These four companies are not jointly owned, are headquartered in four different countries, and operate in three very distinct global industries. This striking regularity in their manner of international price setting therefore is not a mechanical artifact of joint pricing systems or an integrated organizational structure, an important concern in other settings.

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 almost precisely for the bulk of these products within the euro zone would be too coincidental if these were not in fact identical goods.

3.3.2 Dollarization

Given the high quality and 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

[12] We have also found the LOP holding quite well in the euro zone for other retailers for which we have more limited data. One example is Mango, a global apparel retailer based in Spain with similar characteristics as Zara. Online prices from Mango were used by Simonovska (2011) to study the relationship between prices and per-capita income, though that paper does not report on LOP in the euro zone.

Ecuador and El Salvador, dollarized countries that use actual U.S. dollars as their currency, with the equivalent distributions for Bahrein, Hong Kong, Jordan, Lebanon, Oman, Panama, Qatar, Saudi Arabia, and the United Arab Emirates, countries with their own currencies that are precisely pegged to the U.S. dollar.[13] Of our four stores, we only have data for these countries on Zara products. Each bilateral pair matches roughly 3,500 to 4,000 distinct products. Further, we note that for these smaller countries, even though Zara has local stores, they do not sell online. They do, however, post the prices of their products online, and these prices constitute our data. Zara representatives confirmed the online prices equal the offline prices, even in these countries. Nonetheless, for these reasons, we view this evidence among dollarized and dollar-pegged countries as secondary to our evidence from the euro zone.

Figure 6 shows the distribution of weekly log good-level RERs for these countries relative to the United States for Zara products scraped in the fourth quarter of 2012. All listed countries peg their exchange rates to the U.S. dollar, except for Ecuador and El Salvador, which use actual U.S. dollars as their currency. 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. Among the 9 countries with their own currencies that are pegged to the U.S. dollar, 10 percent of all goods have a log RER with an absolute value less than 0.01. For the two dollarized countries, 40 percent do.

In sum, 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.

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. Differential shipping costs might imply a particular distribution of good-level RERs is shifted for some bilaterals relative to others. Different demand conditions might result in different markups, which might vary differentially over time with cost or exchange rate shocks. Finally, given prices are sticky and NERs are not, patterns in the above histograms might reflect the fact that some goods recently experienced price changes while others did not. For example, there might be large LOP violations between Spain and Norway while there are none between

[13] Both the U.S. dollar and the pegged Panamenian balboa are legal tender in Panama, but Zara's prices are quoted there in balboas.

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.

4.1 Introduction, Demand, and Stickiness

We assume all expenditure weights remain constant over time and focus on decomposing the term $\sum_{z\in\Omega_{ij}^n(t)} q_{ij}^n(z,t)$. Let $i_i^n(z)$ denote the time that good z is introduced in sector n of country i and let $\bar{p}_i^n(z) = p_i^n(z, i_i^n(z))$ 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 $t > i_i^n(z)$ as:

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

where we define $l_i^n(z)$ 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 $\Delta_{i_i^n(z)}^{l_i^n(z)} p_i^n(z)$ 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 it's introduction, then $p_i^n(z,t) = \bar{p}_i^n(z)$.

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

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

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 exchange rate 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 exchange rate ("Nominal Rigidity").

Combining the 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:

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

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 $q_{ij}^n(z,t)$ will change even without movement in local prices ("Nominal Rigidites").

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 $i_i^n(z) = i_j^n(z)$ and have their last price change on the same date $l_i^n(z) = l_j^n(z)$, then (2) reduces to:

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

which has no dependence on country k. It is an undesirable property, however, for the disaggregation of the good-level RER between countries i and j to reflect the exchange rate of a third and potentially unrelated country, so we consider the two special cases when $k = i$ and when $k = j$. If $k = i$, we have:

$$q_{ij}^n(z,t) = \underbrace{\left[\bar{p}_i^n(z) - \bar{p}_j^n(z) + e_{ji}\left(i_j^n(z)\right)\right]}_{\text{Good-Level RER at Introduction}} + \underbrace{\left[\Delta_{i_i^n(z)}^{l_i^n(z)} p_i^n(z) - \Delta_{i_j^n(z)}^{l_j^n(z)}\left(p_j^n(z) - e_{ji}\right)\right]}_{\text{Heterogeneous Demand}} + \underbrace{\left[\Delta_{l_j^n(z)}^t e_{ji}\right]}_{\text{Nominal Rigidities}},$$

while $k = j$ produces:

$$q_{ij}^n(z,t) = \underbrace{\left[\bar{p}_i^n(z) - e_{ij}\left(i_i^n(z)\right) - \bar{p}_j^n(z)\right]}_{\text{Good-Level RER at Introduction}} + \underbrace{\left[\Delta_{i_i^n(z)}^{l_i^n(z)}\left(p_i^n(z) - e_{ij}\right) - \Delta_{i_j^n(z)}^{l_j^n(z)} p_j^n(z)\right]}_{\text{Heterogeneous Demand}} - \underbrace{\left[\Delta_{l_i^n(z)}^{t} e_{ij}\right]}_{\text{Nominal Rigidities}}.$$

We then use as our benchmark disaggregation of the good-level RER the average of these two expressions:

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

where to simplify the exposition we now use the three terms "Introduction," "Demand," and "Stickiness" to represent the three components of the real exchange rate in (3) and write them as:

$$q_{ij}^n(z,t) = q_{ij}^{n,I}(z,t) + q_{ij}^{n,D}(z,t) + q_{ij}^{n,S}(z,t). \tag{4}$$

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. For example, one might instead define the "Introduction" component as the good-level RER when the good is first available in both markets. Unlike equation (4), this would then not use information on the level of the NER at the date when a good is introduced in one country but not the other. One might prefer a more restrictive definition of the "Stickiness" term that includes only nominal movements subsequent to the most recent change in either good price. Or one might prefer weights other than the equal weights used above to combine the expressions when the base country k is chosen to be i versus j. These alternative decompositions and specifications will be studied in the Appendix. Early analysis suggests all results will be robust to these alternatives.

In the top left panels (labeled "a") of Figures 7 and 8 we once again plot histograms of good-level RERs $\left(q_{ij}^n\right)$ for selected bilateral relationships with the United States and Spain respectively, and in the remaining three panels we plot $q_{ij}^{n,I}$ (in panel "b"), $q_{ij}^{n,D}$ (in panel "c"), and $q_{ij}^{n,S}$ (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. As we'll elaborate below, this information is omitted if one studies the RER using matched model price indices. Nominal rigidity or stickiness contributes a moderate amount, particularly in countries like Japan or Mexico, with which the United States had significant exchange rate movements over the period

(and for which we have data spanning a longer period).

Interestingly, the changes in demand channel – the focus of a huge literature – 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 $q_{ij}^{n,D}$ 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, 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 ridigities. Note that Denmark pegs to the euro and therefore has $q_{ij}^{n,S} = 0$ for all goods. But its good-level RER distribution at introduction continues to look completely different from the euro zone countries.

These histograms pool data for all goods over all years and so give a sense for drivers of good-level RERs when combining cross-sectional and time-series variation. We now consider analyses that distinguish between contributors to these two dimensions of variation in the good-level RER distribution.

4.2 Variation in the Cross-Section

To start, we decompose the variance at any date t of the RER as follows:

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

where $\sigma_{ij}^2(t)$ is the variance over sectors n and goods z of $q_{ij}^n(z,t)$. 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:

$$\left(\tilde{\sigma}_{ij}^I\right)^2(t) = \left(\sigma_{ij}^I\right)^2(t) + \sigma_{ij}^{I,D}(t) + \sigma_{ij}^{I,S}(t), \tag{6}$$

where $\left(\sigma_{ij}^I\right)^2(t)$ is the variance over sectors n and goods z of $q_{ij}^{n,I}(z,t)$, $\sigma_{ij}^{I,D}(t)$ equals the covariance over sectors n and goods z of $q_{ij}^{n,I}(z,t)$ and $q_{ij}^{n,D}(z,t)$, and where we define $\left(\tilde{\sigma}_{ij}^D\right)^2$ and $\left(\tilde{\sigma}_{ij}^S\right)^2$

accordingly. This disaggregation equally distributes the contribution of the covariance terms between each of the two relevant terms, an innocuous assumption given covariance terms are quite 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. Each week will contain some mix goods which are newly introduced (where $q_{ij}^n = q_{ij}^{n,I}$), some which just experienced a price change (where $q_{ij}^n = q_{ij}^{n,I} + q_{ij}^{n,D}$), and some which have had a long time pass since the most recent price change or introduction (where $q_{ij}^{n,S}$ 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. Relative to the same calculation that does not subtract the mean, this will minimize the importance of the RER at introduction in total cross-sectional good-level RER dispersion.

The bar charts in Figure 9 show the relative importance of prices at introduction, changes in demand, and nominal rigidities for explaining dispersion in good-level RERs in the *cross-section* for several bilateral relationships. Starting for example with the upper-left plot which represents Canada and the United States, we see that the three bars sum to 0.028 log point, equal to the average cross-sectional variance in q_{ij}^n for that bilateral pair. We use the decomposition (5) to measure that 0.023 log point comes from the Introduction term, 0.004 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 exchange rate 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 20 percent of the total cross-sectional variation in these bilateral pairs, leaving roughly 75 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 the variance, 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 exchange rate fluctuations.

5 Frequency of Product Entry/Exit and Price Change

The relative importance of $q_{ij}^{n,I}$, $q_{ij}^{n,D}$, and $q_{ij}^{n,S}$ depends on the rate at which products are introduced and removed from the market as well as the frequency of price changes. For example, if goods are introduced often into the market, this would imply a greater importance in the cross section of the "Introduction" term. If good prices never change due to complete nominal price rigidity, the "Demand" term will by construction always equal zero. Therefore, to better understand the results in Figure 9, as well as to evaluate how the products studied here might relate to the broader basket of traded goods, we now report the frequency of product churning and of price changes in our dataset.

First, we give a measure of the frequency of product entry and exit in our database. We measure product duration within each country. One challenge in measuring the date of a product's introduction and exit is that our data are both left- and right- censored. This is exacerbated by the fact that our scraping algorithm occasionally starts and stops, making it difficult to know whether the initial observation of a good is due to a true introduction or simply reflects failure in earlier periods to record a product that was available. Below, when reporting results at the country or world level, as in rows (i) and (ii), we simply consider a product to have entered and exited if it is somewhere in our data but is not found in the last 30 days available for each store and country combination. When reporting results for matched bilateral pairs, as in rows (iii) to (v), we additionally exclude the small share of goods which first appeared in the two countries more than 15 weeks apart. In future drafts, we plan to report these results when using more stringent filters for the inclusion of goods.

The top panel of Table 2 lists the mean and median product duration in weeks for various categories of goods, after dropping goods with only 1 weekly observation and weighting observations such that the contribution of each store is equalized. Row (i) shows that globally, the average product length equals 27 weeks or about 6 months and the median product life is about half as long. This large difference between the mean and median comes from significant skewness in the distribution, with 5 percent of the goods in fact lasting two years or longer. Row (ii) shows that when we limit the analysis to goods sold in the United States, typical product lives in our data increase, with the mean length equaling 39 weeks or about 9 months and 9 percent of the goods lasting at least two years. We suspect this reflects the higher prevalence of scraping gaps and right-censoring in some of the smallest countries in our data.

The aggregate product duration statistics listed in the "All Stores" column, however, are perhaps less useful than they may appear because they mask significant heterogeneity across

stores. H&M and Zara have such low product durations that they reduce the average reported in the first column.[14] IKEA products sold in the United States had a mean duration of nearly 1 year, with 13 percent of the products lasting for at least 2 years.

Our primary analyses involve products which can be matched in multiple countries, so in rows (iii) to (v) we consider these product life statistics for the set of products in the bilateral pairs plotted in Figures 2 and 4. The basic patterns for the 12 U.S. bilateral pairs, the 4 non euro zone Spanish bilateral pairs, and the 8 euro zone Spanish bilateral pairs are all qualitatively the same. Products from Apple and IKEA have an average life ranging from half a year to 1.25 years, with a large share of goods with much longer lives of two or more years. H&M and Zara product lives are closer to quarters, likely reflecting different seasonal styles in clothing.

In the bottom panel of Table 2, we examine the absolute value of gaps in the time of product introductions for the key bilateral relationships with the United States and Spain. We report the median because the mean value is driven by outliers which are excluded from our main analyses. Apple and IKEA products are frequently offered in the United States before the rest of the world, with typical lags lasting 2 and 8 weeks, respectively, and with a moderate share of products with introduction timing that differs by more than a few months. But more generally, introductions appear to be at least moderately synchronized. For example, the median absolute value of the timing gap for all product introductions in the bilateral pairs involving Spain, whether euro zone or non euro zone countries, equals zero.

We next turn to the frequency of price changes, or stickiness of prices in our data. The top panel of Table 3 lists the percent of products with any price changes during the entire life of the product. For example, the entry in the top left of the panel indicates that of all our products from all stores and countries in the data, 15 percent experience at least one price change at some point. 85 percent of the goods exit the sample with the same price they had at introduction. Given this significant amount of stickiness and given the heterogeneity in product life documented in Table 2, we report this statistic rather than a price change frequency. Among Apple and IKEA, the stores for which we have the longest-lived data, the share of goods with price changes is closer to 25 percent. The significantly smaller overall percentage is driven by the fact that less than 10 percent of all good-country combinations in H&M and Zara exhibit price changes, though this is very likely a reflection of their short tenure in the current version of the data. If we limit the data to goods with at least 1 year in the data, we find that roughly half experienced at least one price change. Though demonstrating a somewhat higher degree of stickiness, this is broadly consistent

[14]This fast turnover was discussed in a recent June 21, 2012 article in Slate titled "Zara Gets Fresh Styles To Stores Insanely Fast."

with results in Gopinath and Rigobon (2008) or Neiman (2010) for differentiated traded good prices.

In the bottom panel of Table 3, we consider only the goods that are offered both in the United States and the key dozen bilateral pairings shown in Figure 2 and the bilaterals with Spain shown in Figure 4. Here we report the share of products which exhibit at least one change in either of the two countries. For example, imagine a white iPhone 3 with certain memory configuration is sold in the United States, Canada, and Japan. This would then count as two entries in "Key Pairs with United States." If a price change occurred for this good in Canada, but not in the United States or Japan, we would characterize 50 percent of these matched goods as having at least one price change – one in the U.S.-Canada pair but none in the U.S.-Japan pair. The high share of matched products in our data which experience at least one price change suggests that the "Demand" channel could have plausibly played a critical role in good-level RER dispersion, though we showed above in Figure 9 that it did not. Further, the fact that stickiness decreases markedly for the apparel countries that exist in the matched data suggests that the exceptional stickiness at the top of Table 3 largely reflects short-lived products in smaller countries that do not match with the United States and therefore do not impact most of our analyses.

In sum, products in our data are very frequently introduced and withdrawn from the market and their prices may be quite sticky. A typical product life for the clothing stores might last only a quarter while the technology or furniture products might last closer to one year, and typical goods might not experience any price changes at all. However, a sizeable share of our data include goods with product lives well in excess of one or two years, and of the matched products for which we study good-level RERs, there are ample examples of price changes. We view the high frequency of product introduction and exit as evidence that more focus should be paid to relative price levels and less to changes. In future versions, we will report our results when restricted to goods with a long product life and some price changes.

6 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. Relatedly, we document that many products have a short product life. As such, it is clear that an understanding of good-level RERs at introduction ($q_{ij}^{n,I}$) is 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 that prices are quite sticky but the distribution of $q_{ij}^{n,I}$ is fixed completely over time with an average value of $\tilde{q}$. In this world, RERs might appreciate or depreciate over time with the nominal exchange rate, 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 that ignore product introduction would miss this hypothetical mechanism for mean reversion of the RER. This is the sense in which taking product introduction prices into account might plausibly have solved (or explained much of) the PPP puzzle.[15]

In fact, we find evidence of the opposite phenomenon. 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 exchange rate 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 good's RER 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.[16] One reason this result is important is that it suggests the high frequency correlation between NERs and RERs cannot simply emerge due simply 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 median of good-level log RERs at introduction $q_{ij}^{n,I}$ for the key bilaterals involving the United States.[17] We separate each of the four stores and mark their medians with each of four markers. The thin black line plots the log bilateral exchange rate e_{ji},

[15] 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.

[16] While this appears to hold among most countries in our data, it strikes us as extremely unlikely this also holds in countries with large trends in their exchange rates. We hope to report on this in the next version.

[17] We drop the very limited number of observations where $|q_{ij}^{n,I}| > 0.75$, which is slightly stronger than the filter $|q_{ij}^{n}| > 0.75$ used to capture outliers in the rest of the paper.

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 where menu costs and nominal volatility account for RER movements, 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, as does the downward movement late in the sample. It is difficult, however, to make conclusions from these rich scatter plots, so we now turn to non-linear fits from these raw data.

First we scale the $q_{ij}^{n,I}$ 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 exchange rate. 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 Germany, France, 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 by 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.

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

$$q_{ij}^{n,I}(z,t) = \gamma_{ij} + \beta e_{ji}(t) + \epsilon_{ij}^{n}(z,t), \tag{7}$$

where good z only appears in the regression in the one period when $t = max\left(i_i^n(z), i_j^n(z)\right)$, where we demean the left-hand side variable for each store and country pair, and where we exclude any good with $|q_{ij}^{n,I}| > 0.75$ or $|q_{ij}^{n}| > 0.75$. 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. 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.

Table 4 reports the coefficients on this regression and, consistent with Figures 10 and 11, shows that the good-level RER closely tracks the NER even at the time of product introductions. For example, looking at row (ii), we see that across all stores, the good-level RER at introduction ($q_{ij}^{n,I}$) moves 0.715 log point for every full log point movement in the NER. In other words, if the bilateral exchange rate with the United States appreciates by 10 percent over the course of the year, one would expect new products to be introduced with relative prices about 7 percent higher than the previous year. The phenomenon holds for Apple products, but less so than for the other stores. IKEA and H&M good-level RERs at the time of product introductions track the NER essentially one-for-one. We cluster the standard errors by retailer-weeks and note that these coefficients are estimated with very high precision.

Regression (7) is very similar to some run in Baxter and Landry (2012) and is reminiscent of exchange rate passthrough regressions, which typically correlate changes in import prices with changes in the bilateral exchange rate between the importer and the exporter. While the relationship of the RER with the bilateral exchange rate is of course related to passthrough (as can be easily seen in the "Changes in Demand" term of (2)), we cannot explicitly comment on exchange rate passthrough here because we do not know the identity of the exporting country for any given good.

For example, imagine a good is produced in Japan and exported to both Spain and the United States, and imagine there is a constant underlying passthrough rate to Spain that equals 0.75 and to the United States that equals 0.25. If prices change only due to exchange rate passthrough, 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 imply identical movements in the dollar-euro exchange rate (a 10 percent euro depreciation), but generate different movements in the good-level RER. This simple example shows that with heterogeneous passthrough rates across markets, the case strongly suggested by the literature, knowledge of the exporting country is required to estimate passthrough.[18]

In sum, we isolate good-level RERs at the time products are introduced and demonstrate that these RERs track closely the NER, in contrast to the predictions of menu-cost models with

[18] Additional complications arise from the lack of information on the source of inputs used in production by the unknown exporter and the large degree of price stickiness in the relatively short panel of data we currently have. In other work, we are exploring panel-based econometric procedures in the hopes of dealing with this issue.

constant markups.

7 Conclusion

Open-economy macroeconomic models require an assumption about international relative prices to comment about such critical topics as optimal currency regimes, the international transmission of shocks, and the benefits of international coordination of monetary policy. Classic models, such as the Balassa-Samuelson model, assume a constant level of the RER among traded goods, though subsequent empirical work has demonstrated the marked violation of this assumption at the good and aggregate level. A voluminous literature has worked to understand the determinants of LOP deviations in the cross-section and time-series, as these deviations imply RER variability.

Using a novel dataset of traded goods with dramatically more products and countries than are covered in many earlier studies, we demonstrate that the LOP holds almost exactly for the vast majority of products sold by four global retailers within the euro zone. Outside the euro zone, even among bilateral pairs with zero nominal volatility, the LOP is flagrantly violated, giving rise to RER volatility. Evidence comparing dollarized to dollar pegged countries is similar. The currency in which prices are quoted is a critical determinant of market segmentation, and for these products is more important than transport costs or tax or taste differences.

Further, in large part since these products have a short life cycle, we show that LOP violations are best understood by measuring relative prices in levels at the time of product introductions, as opposed to focusing on fluctuations due to incomplete passthrough or nominal rigidity. This is particularly important as conventional matched model price indices, the basis for most of the literature on RER movement, are constructed only from observed price changes and exclude the information contained in what we call the "RER at introduction." Finally, we demonstrate that this RER at introduction in fact also moves at high frequency with the NER, which strongly suggests that the root of pricing rigidities is not well captured by models that omit variable markups or pricing complementarities, including those with constant elasticity of substitution demand and menu costs.

Clearly, the pricing behavior documented for these four global retailers need not be representative of all retail sectors. A focus on product introductions is likely unwarranted in the egg or milk product categories, and we doubt that the LOP is any less likely to hold for crude oil in or out of a currency union. For branded manufactured goods that represent a large share of total traded consumption expenditures, however, we provide important evidence on how the behavior of prices at the time of product introduction and the choice of currency regime are critical

determinants of the behavior of the traded-good real exchange rate.

References

Atkeson, A., and A. Burstein (2008): "Pricing-to-market, trade costs, and international relative prices," *The American Economic Review*, 98(5), 1998–2031.

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

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.

Borraz, F., A. Cavallo, R. Rigobon, and L. Zipitria (2012): "Distance and Political Boundaries: Estiimating 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*.

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

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.

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.

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.

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

Engel, C., and J. Rogers (1996): "How wide is the border?," *American Economic Review*, 86, 1112–1125.

Fitzgerald, D., and S. Haller (2012): "Pricing-to-Market: Evidence from Plant-Level Prices," *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.

Gopinath, G., O. Itskhoki, and R. Rigobon (2010): "Currency Choice and Exchange Rate Pass-through," *American Economic Review*, 100, 304–336.

Gopinath, G., and R. Rigobon (2008): "Sticky Borders," *Quarterly Journal of Economics*, 123(2), 531–575.

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.

Mandel, B., E. Gagnon, and R. Vigfusson (2012): "Missing Import Price Changes and Low Exchange Rate Pass-Through," *Working Paper*.

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.

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

Nakamura, E., and J. Steinsson (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.

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.

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

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

		All Stores	Apple	IKEA	H&M	Zara
(i)	# Products, World	89,705	9,078	60,040	9,402	11,185
(ii)	# Products, United States	33,602	4,349	17,597	4,107	7,549
(iii)	# Countries	83	29	20	47	80
(iv)	Time Period	2008:Q4 to 2012:Q3	2009:Q2 to 2012:Q3	2008:Q4 to 2012:Q3	2011:Q3 to 2012:Q3	2011:Q3 to 2012:Q2
(v)	Headquarters		United States	Sweden	Sweden	Spain
(vi)	Industry		Consumer Electronics	Home/Office Furniture	Apparel	Apparel
(vii)	Global Industry Rank		3rd largest	1st largest	4th largest	3rd largest
(viii)	Retail Revs ($Bil.)	≈ 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
Product Duration (weeks)							
(i)	World	Mean	27	33	44	9	6
		Median	13	24	28	8	6
(ii)	United States	Mean	36	39	47	7	3
		Median	21	29	35	8	3
(iii)	Key Pairs with United States	Mean	26	31	55	8	9
		Median	11	20	42	6	8
(iv)	Key Pairs with Spain	Mean	28	28	67	8	10
	(Non Euro Zone)	Median	12	20	47	7	9
(v)	Key Pairs with Spain	Mean	28	27	66	9	11
	(Euro Zone)	Median	13	20	48	8	11
Introduction Gap, Absolute Value (weeks)							
(v)	Key Pairs with United States	Median	1	2	8	0	0
(vi)	Key Pairs with Spain (Non Euro Zone)	Median	0	0	1	2	0
(vii)	Key Pairs with Spain (Euro Zone)	Median	0	0	3	1	0

Table 2: Information about the Product Life Cycle

Notes: First two rows include all products with more than 1 week in data. Rows (iii) through (vii) include matched pairs, which exclude goods that are introduced at dates more than 15 weeks apart in the two countries. H&M and Zara have distinctly shorter product durations than Apple and IKEA. This likely reflects both the greater importance of seasonality in apparel as well as the fact that those stores only exist in our data for about 1 year.

			All Stores	Apple	IKEA	H&M	Zara
Percent of Products with Any Price Changes							
(i)	World	All Products	15	18	30	3	9
		$\geq$ 12 months	48	39	51	-	–
		$\geq$ 24 months	55	42	59	-	-
(ii)	United States	All Products	16	23	21	1	1
		$\geq$ 12 months	32	28	35	-	-
		$\geq$ 24 months	37	25	40	-	-
(iii)	Spain	All Products	25	16	45	1	26
		$\geq$ 12 months	75	64	76	-	-
		$\geq$ 24 months	84	64	88	-	-
Percent of Matched Pairs with Any Price Changes in Either Country							
(iv)	Key Pairs with United States	All Products	51	48	75	40	41
		$\geq$ 12 months	80	54	88	-	-
		$\geq$ 24 months	81	56	93	-	-
(v)	Key Pairs with Spain (Non Euro Zone)	All Products	45	35	70	50	26
		$\geq$ 12 months	89	71	92	-	-
		$\geq$ 24 months	92	73	96	-	-
(vi)	Key Pairs with Spain (Euro Zone)	All Products	16	11	45	1	5
		$\geq$ 12 months	75	60	77	-	-
		$\geq$ 24 months	86	61	94	-	-

Table 3: Information about the Frequency of Price Change

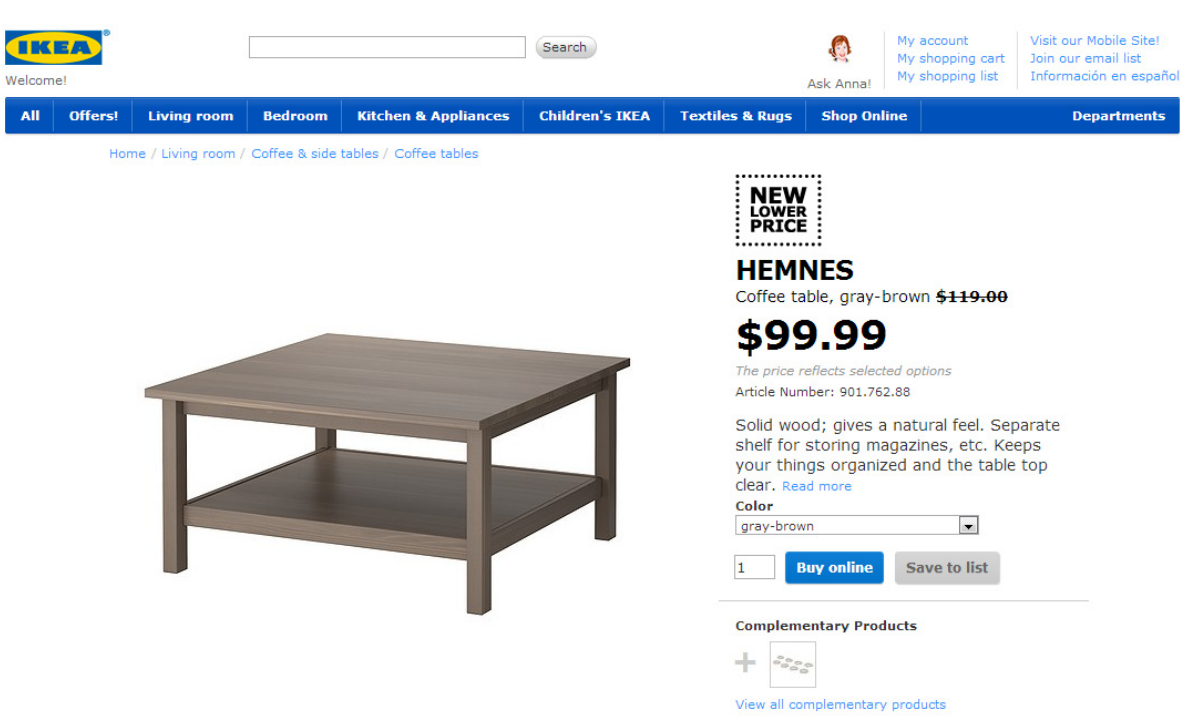
Notes: Results unavailable for $\geq 12, 24$ months for H&M and Zara because those stores only exist in our data for about 1 year. Products with less than 1 week of data are excluded. Each product is separately considered in each country. For example, if one particular Apple product has price changes in Italy but none in Spain, this would be considered as two products, one of which had a price change.

Dependent Variable: Good-Level Log RER at Introduction
Independent Variable: Log Nominal Exchange Rate
Fixed Effects: Country Pair Effects

			All Stores	**Apple**	**IKEA**	**H&M**	**Zara**
(i)	All Bilaterals	Coefficient	0.590	0.485	0.836	0.882	0.772
		Standard Error	(0.008)	(0.012)	(0.029)	(0.006)	(0.011)
		Observations	19,908,201	352,069	872,285	3,318,516	15,365,331
(ii)	All U.S. Bilaterals	Coefficient	0.715	0.617	0.989	1.046	0.747
		Standard Error	(0.025)	(0.030)	(0.048)	(0.027)	(0.052)
		Observations	602,325	25,447	57,576	142,284	377,018
(iii)	All Spain Bilaterals	Coefficient	0.533	0.488	0.735	0.943	0.669
	(Non Euro Zone)	Standard Error	(0.033)	(0.039)	(0.075)	0.023	(0.087)
		Observations	505,766	18,371	66,007	111,368	310,020

Table 4: Comovement of RER at Introduction and Nominal Exchange Rate

Notes: Each good is only included in the regression on a single introduction date. Standard errors are clustered by storeXweek. Unlike other tables and figures, rows (ii) and (iii) include all respective bilaterals, not just the 12 key matches focused on elsewhere. We exclude the very limited number of observations where $|q_{ij}^{n,I}| > 0.75$ or $|q_{ij}^{n}| > 0.75$.



(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 of our stores are identical online and offline, and this is just one example.

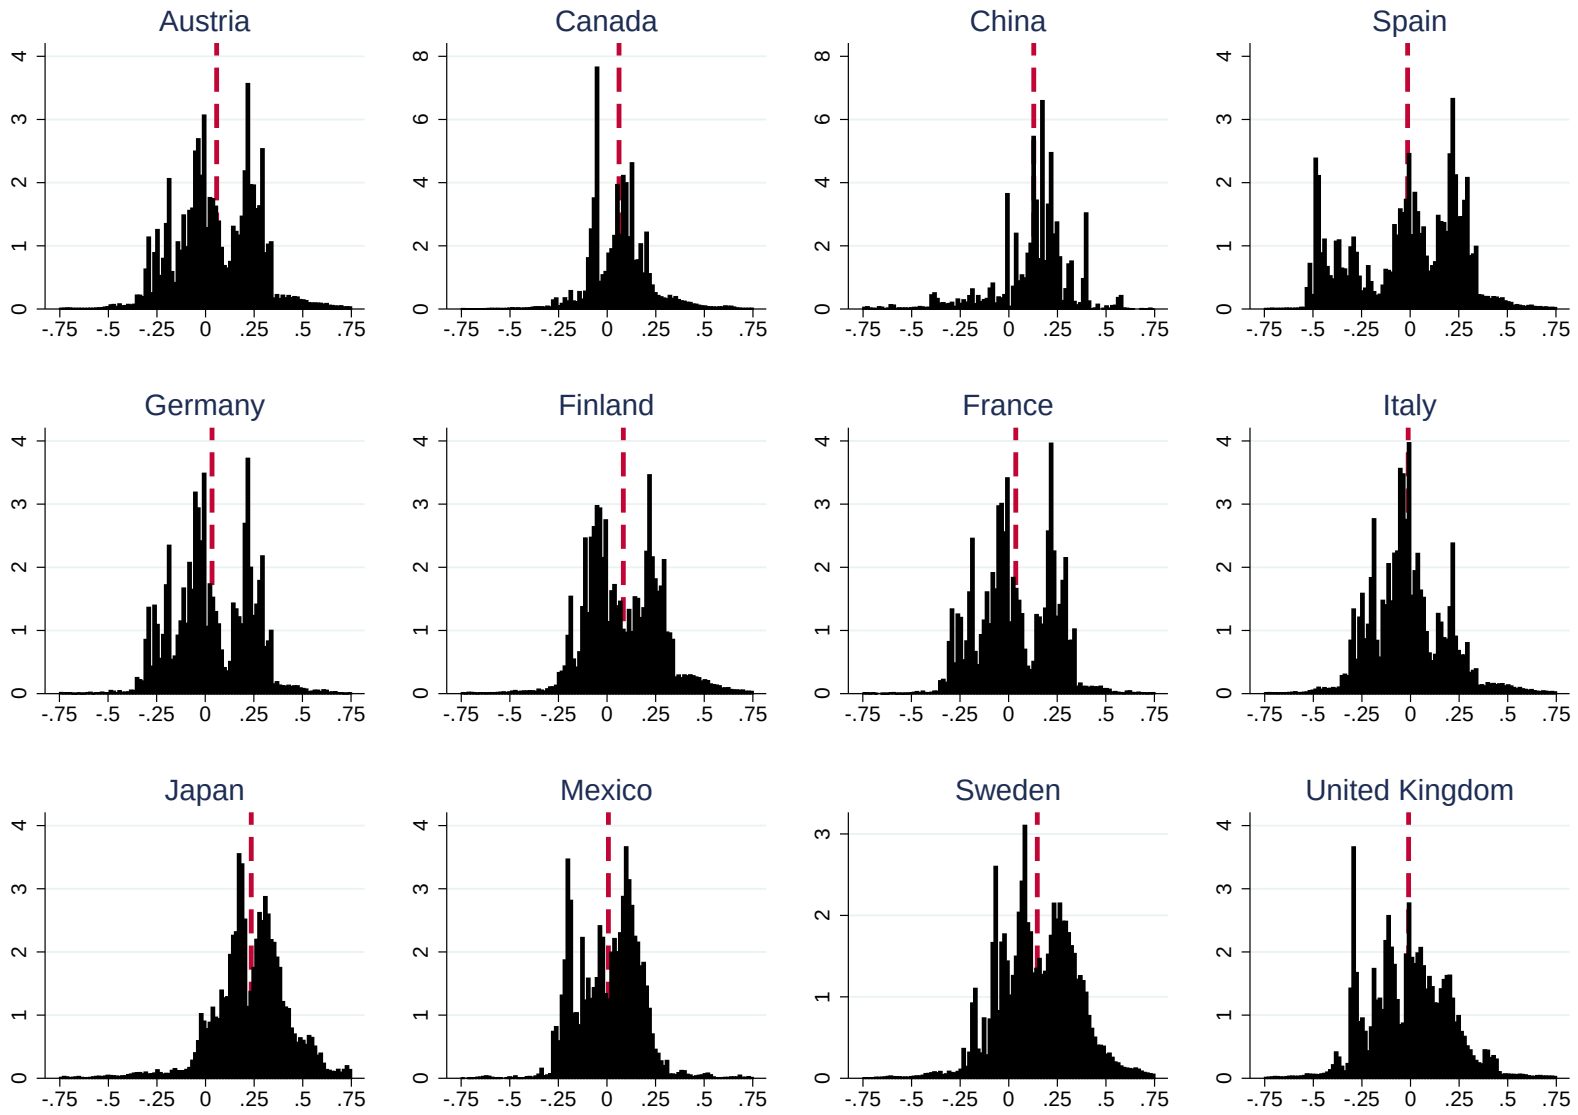
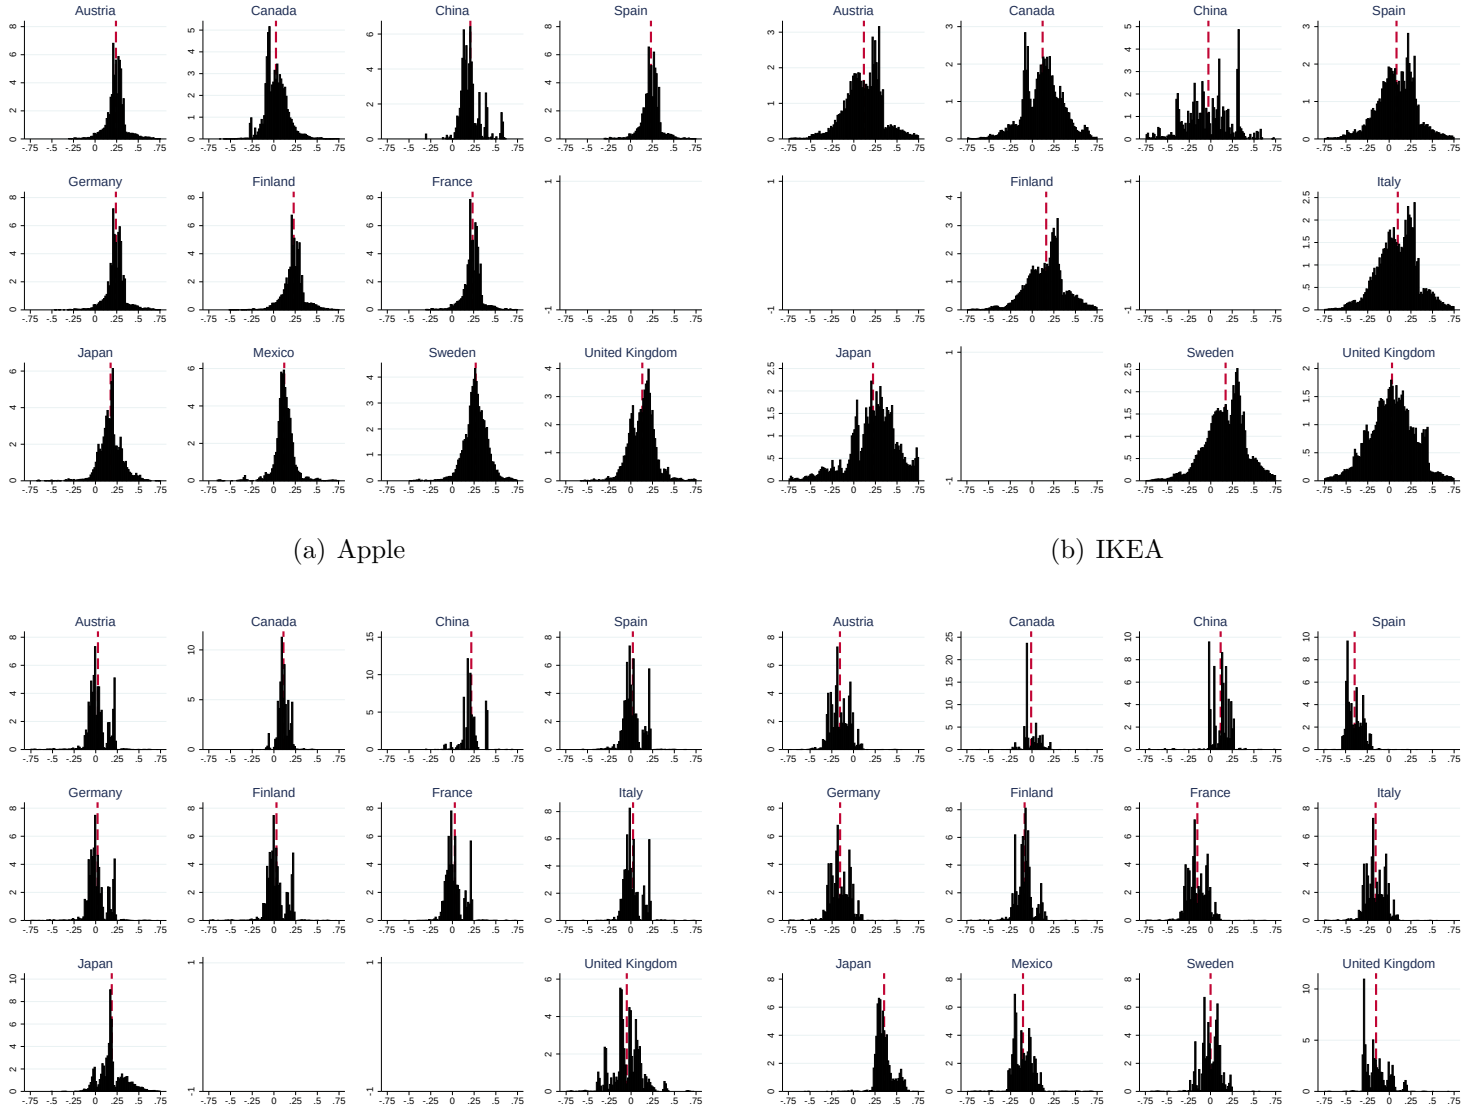


Figure 2: Good-level RERs q_{ij}^n for Various Countries (i) with the United States (j)

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



(a) Apple

(b) IKEA

(c) H&M

(d) Zara

Figure 3: Good-level RERs q_{ij}^n 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}^n is defined, with United States as country j and the other countries as country i. Histograms include frequency weights such that the contribution of goods from each store is equalized. We exclude goods and weeks where $|\left(q_{ij}^n\right)| > 0.75$. These excluded observations represent a very small percentage of total observations. Dashed red vertical lines indicate the weighted average RER. Y-axes plot density.

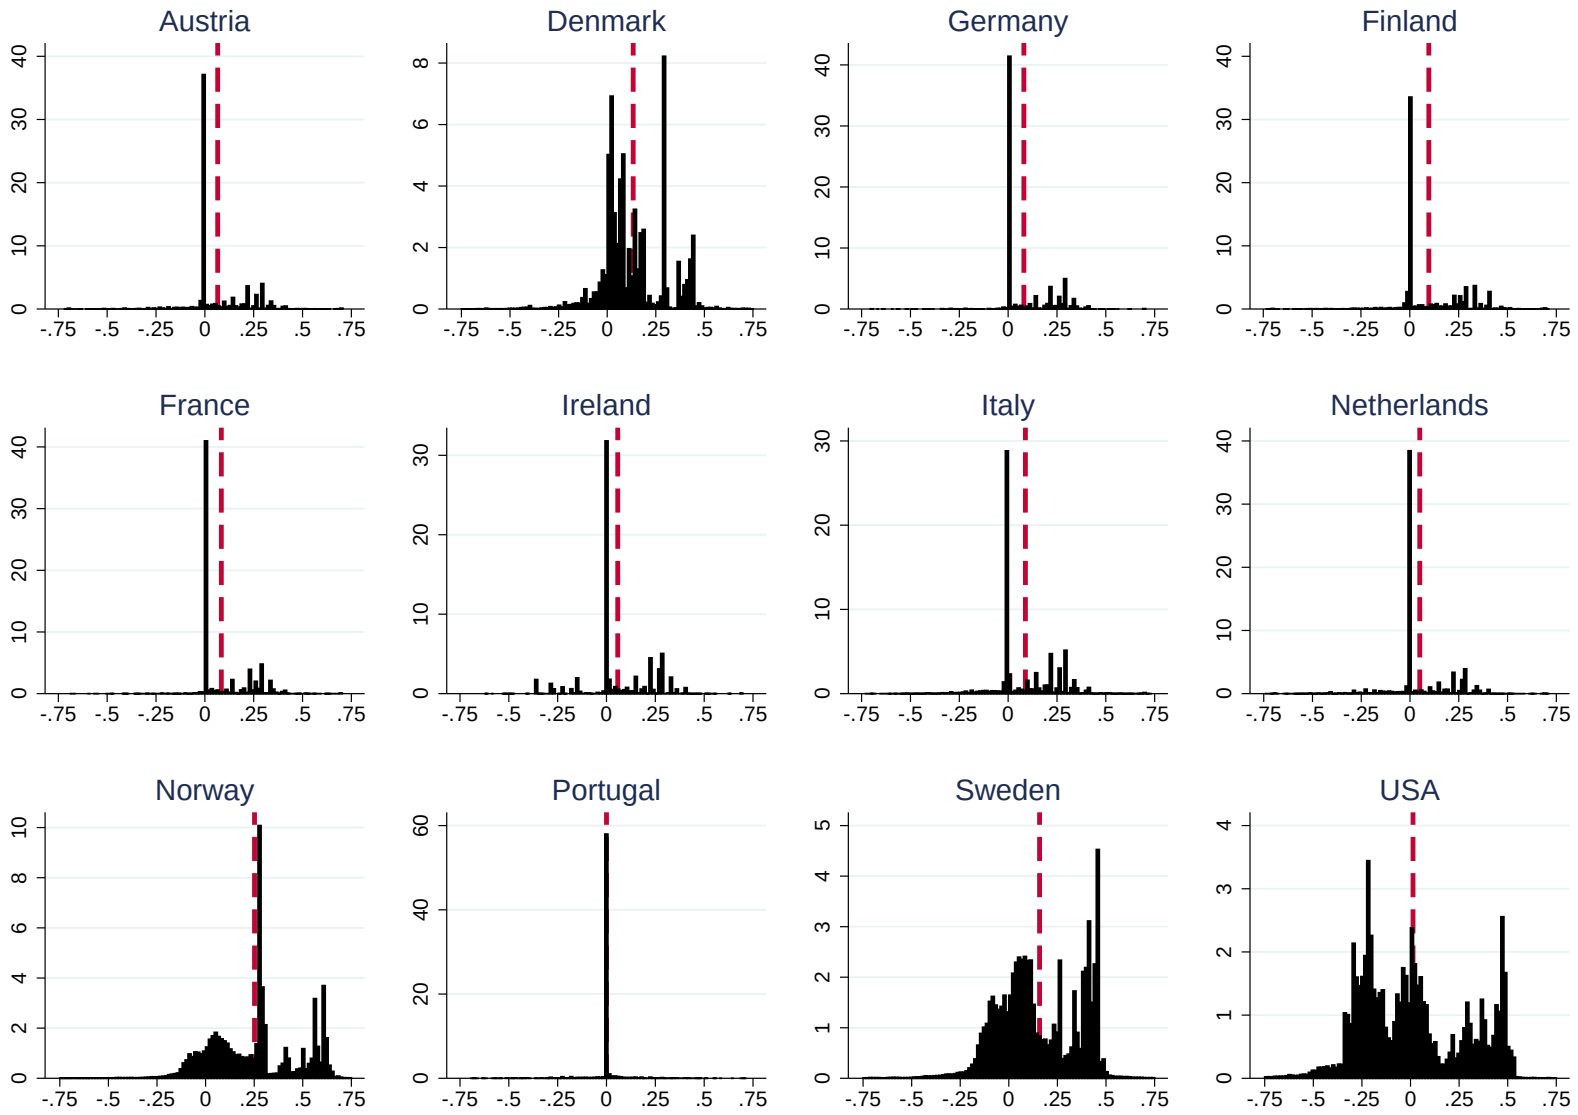
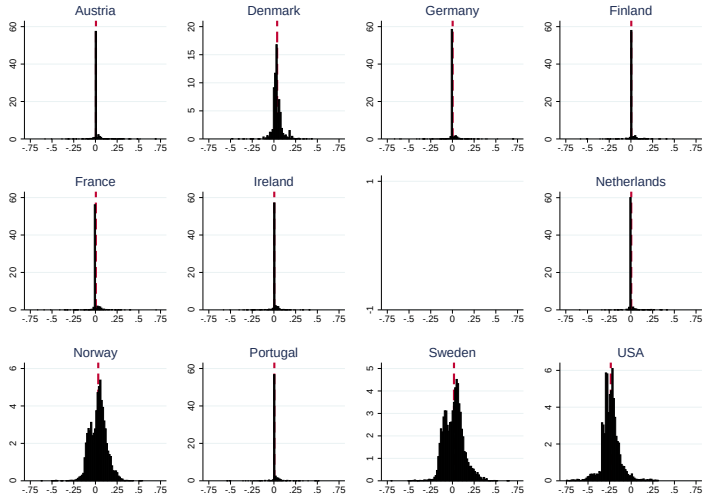
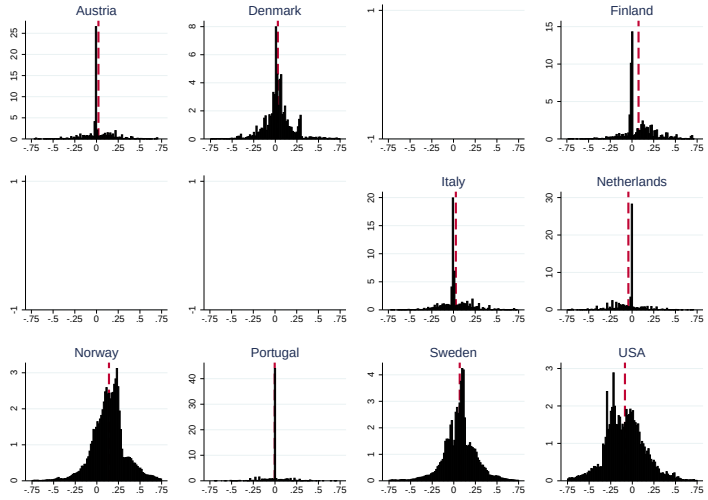


Figure 4: Good-level RERs q_{ij}^n for Various Countries (i) with Spain (j).

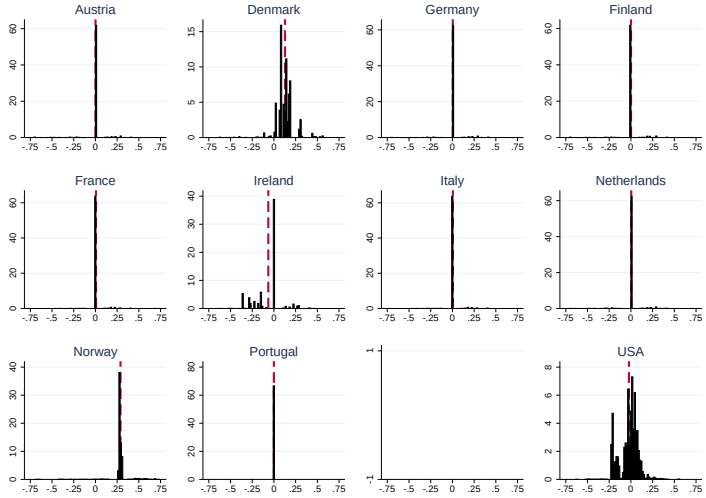
Notes: Figure includes all goods z and all weeks t for which q_{ij}^n is defined, with Spain as country j and the other countries as country i. Histograms include frequency weights such that the contribution of goods from each store is equalized. We exclude goods and weeks where $|\left(q_{ij}^n\right)|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



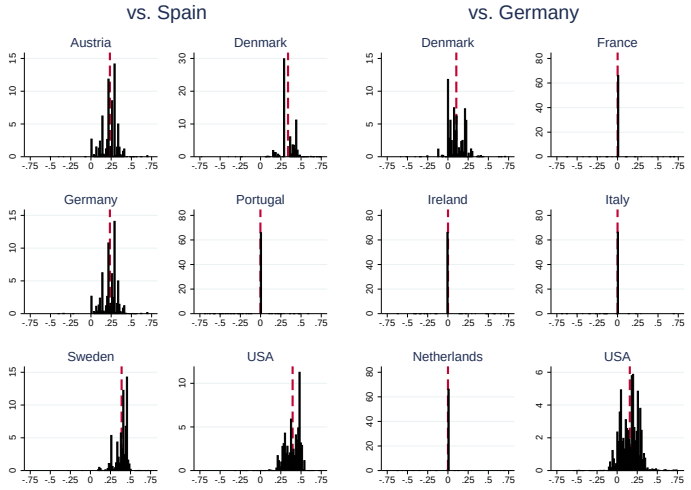
(a) Apple



(b) IKEA



(c) H&M



(d) Zara

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

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

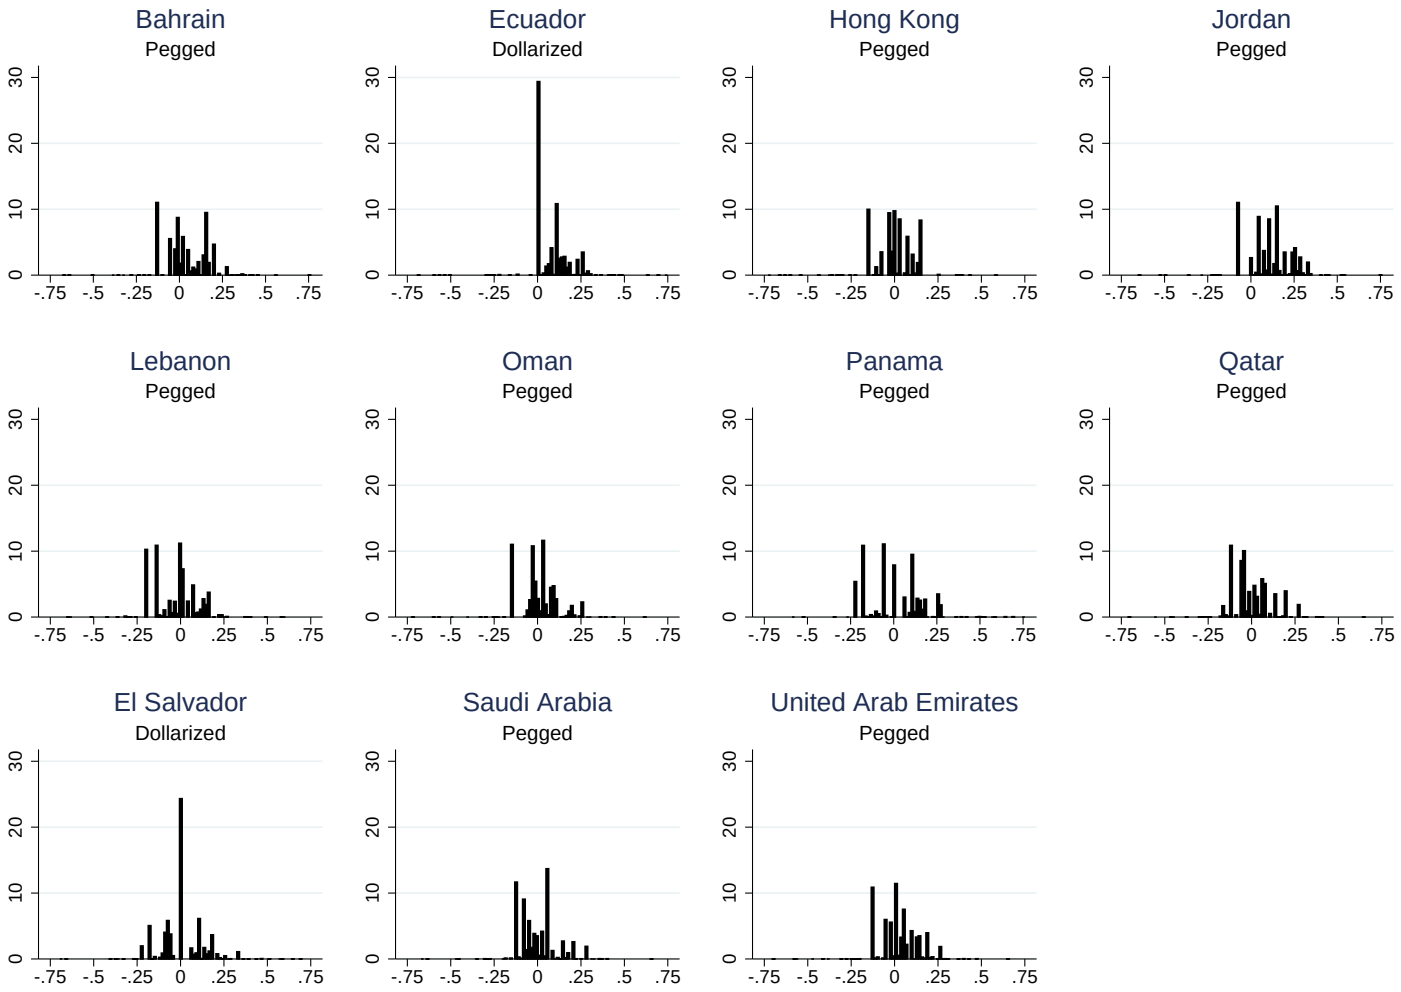
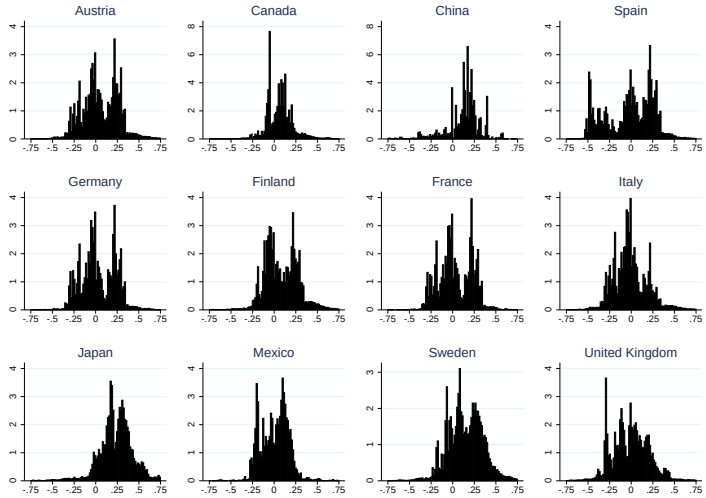
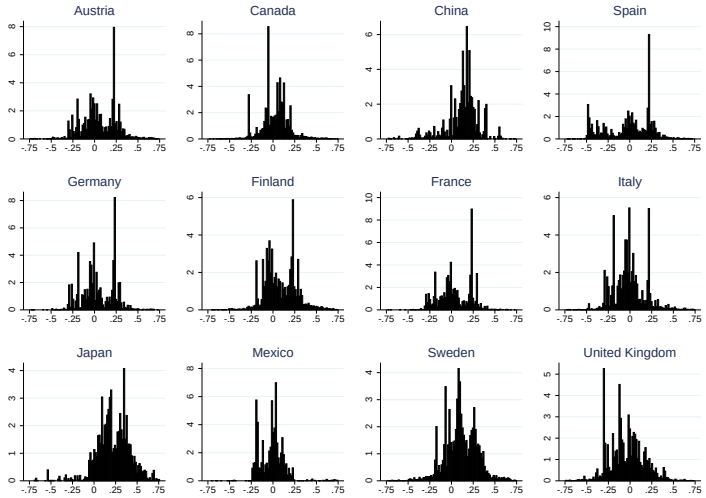


Figure 6: Good-level RERs q_{ij}^n for Various Countries (i) with the United States (j), Zara only

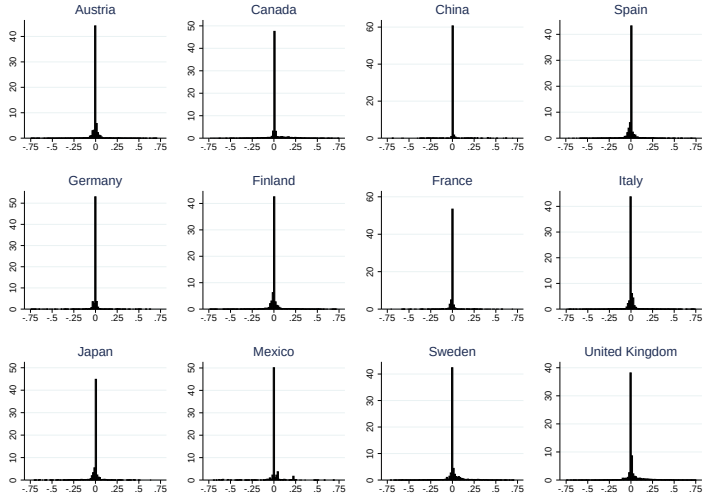
Notes: Figure includes all goods z and all weeks t for which q_{ij}^n is defined, with United States as country j and the other countries as country i. We include 9 countries with an exchange rate peg with the dollar as well as two countries that are dollarized (i.e. actually use the U.S. dollar as their currency), Ecuador and El Salvador. Only goods/prices from Zara are included because we lack data for the dollarized countries for the other stores. We goods and weeks where $|\left(q_{ij}^n\right)| > 0.75$. These observations represent a very small percentage of total observations. We note that for some of these countries, unlike the larger countries presented elsewhere, Zara's web page advertises product prices but does not allow for actual online purchases. In such cases, according to the company, online prices equal those in physical retail stores in the country. Y-axes plot density.



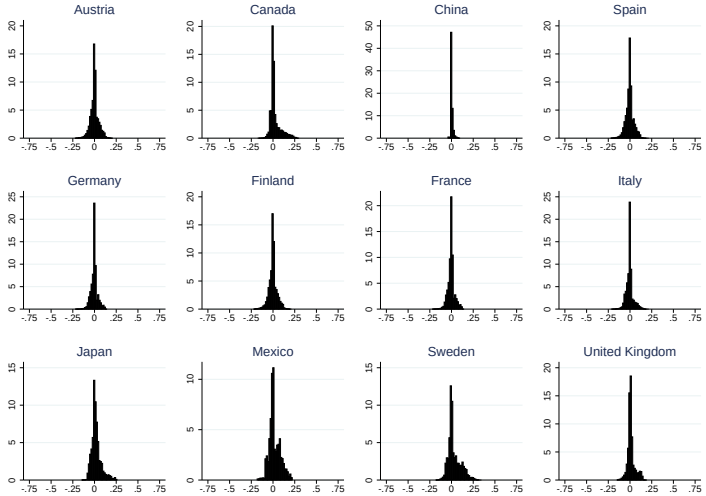
(a) Good-level RER $\left(q_{ij}^{n}\right)$



(b) At Introduction $\left(q_{ij}^{n,I}\right)$



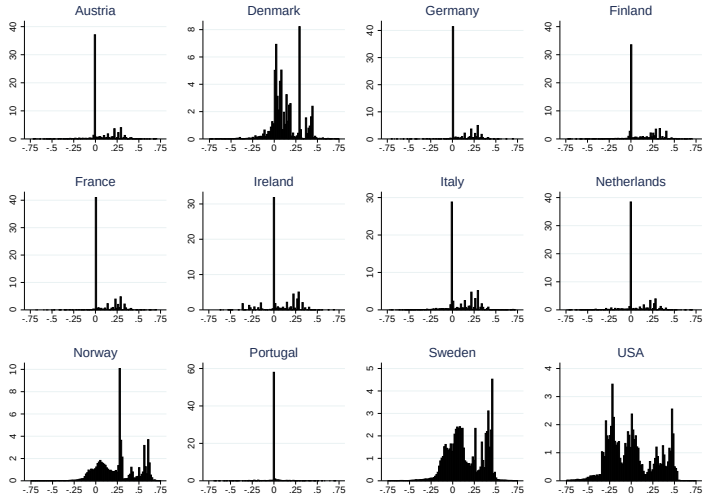
(c) Changes in Demand $\left(q_{ij}^{n,D}\right)$



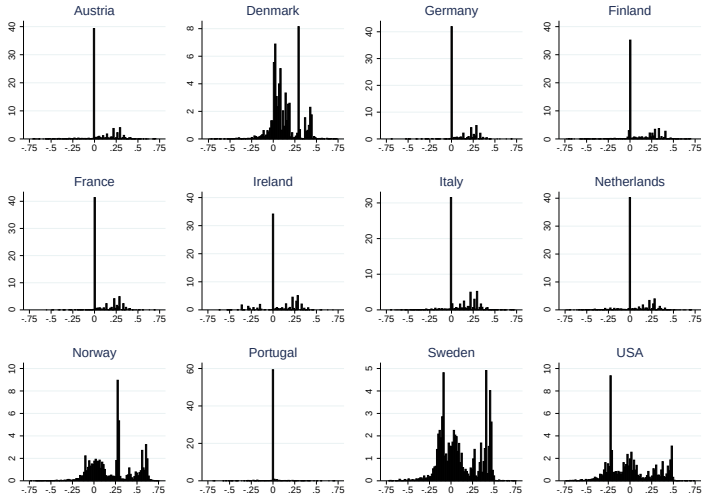
(d) Nominal Rigidites $\left(q_{ij}^{n,S}\right)$

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

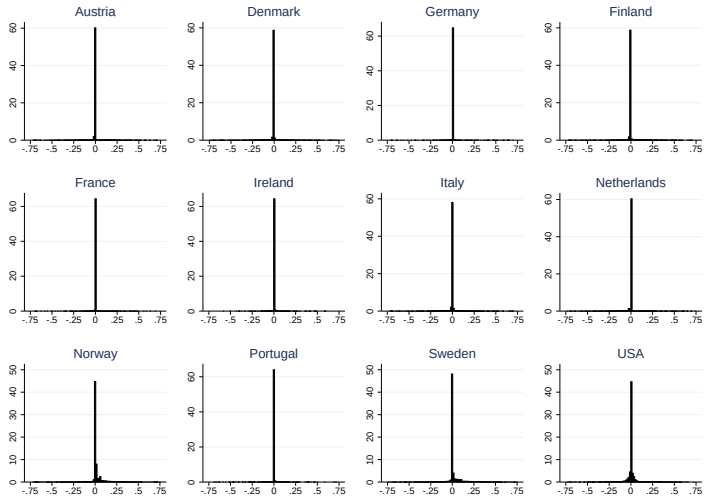
Notes: Figure includes all goods z and all weeks t for which q_{ij}^{n} is defined, with the United States as country j and the other countries as country i. Histograms include frequency weights such that the contribution of goods from each store is equalized. We exclude goods and weeks where $|\left(q_{ij}^{n}\right)| > 0.75$. These excluded observations represent a very small percentage of total observations. Y-axes plot density.



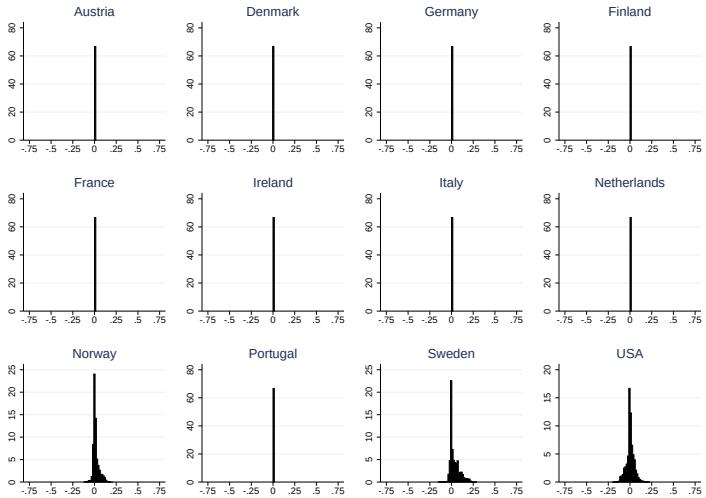
(a) Good-level RER $\left(q_{ij}^{n}\right)$



(b) At Introduction $\left(q_{ij}^{n,I}\right)$



(c) Changes in Demand $\left(q_{ij}^{n,D}\right)$



(d) Nominal Rigidites $\left(q_{ij}^{n,S}\right)$

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

Notes: Figure includes all goods z and all weeks t for which q_{ij}^{n} is defined, with Spain as country j and the other countries as country i. Histograms include frequency weights such that the contribution of goods from each store is equalized. We exclude goods and weeks where $|\left(q_{ij}^{n}\right)| > 0.75$. These excluded observations represent a very small percentage of total observations. Y-axes plot density.

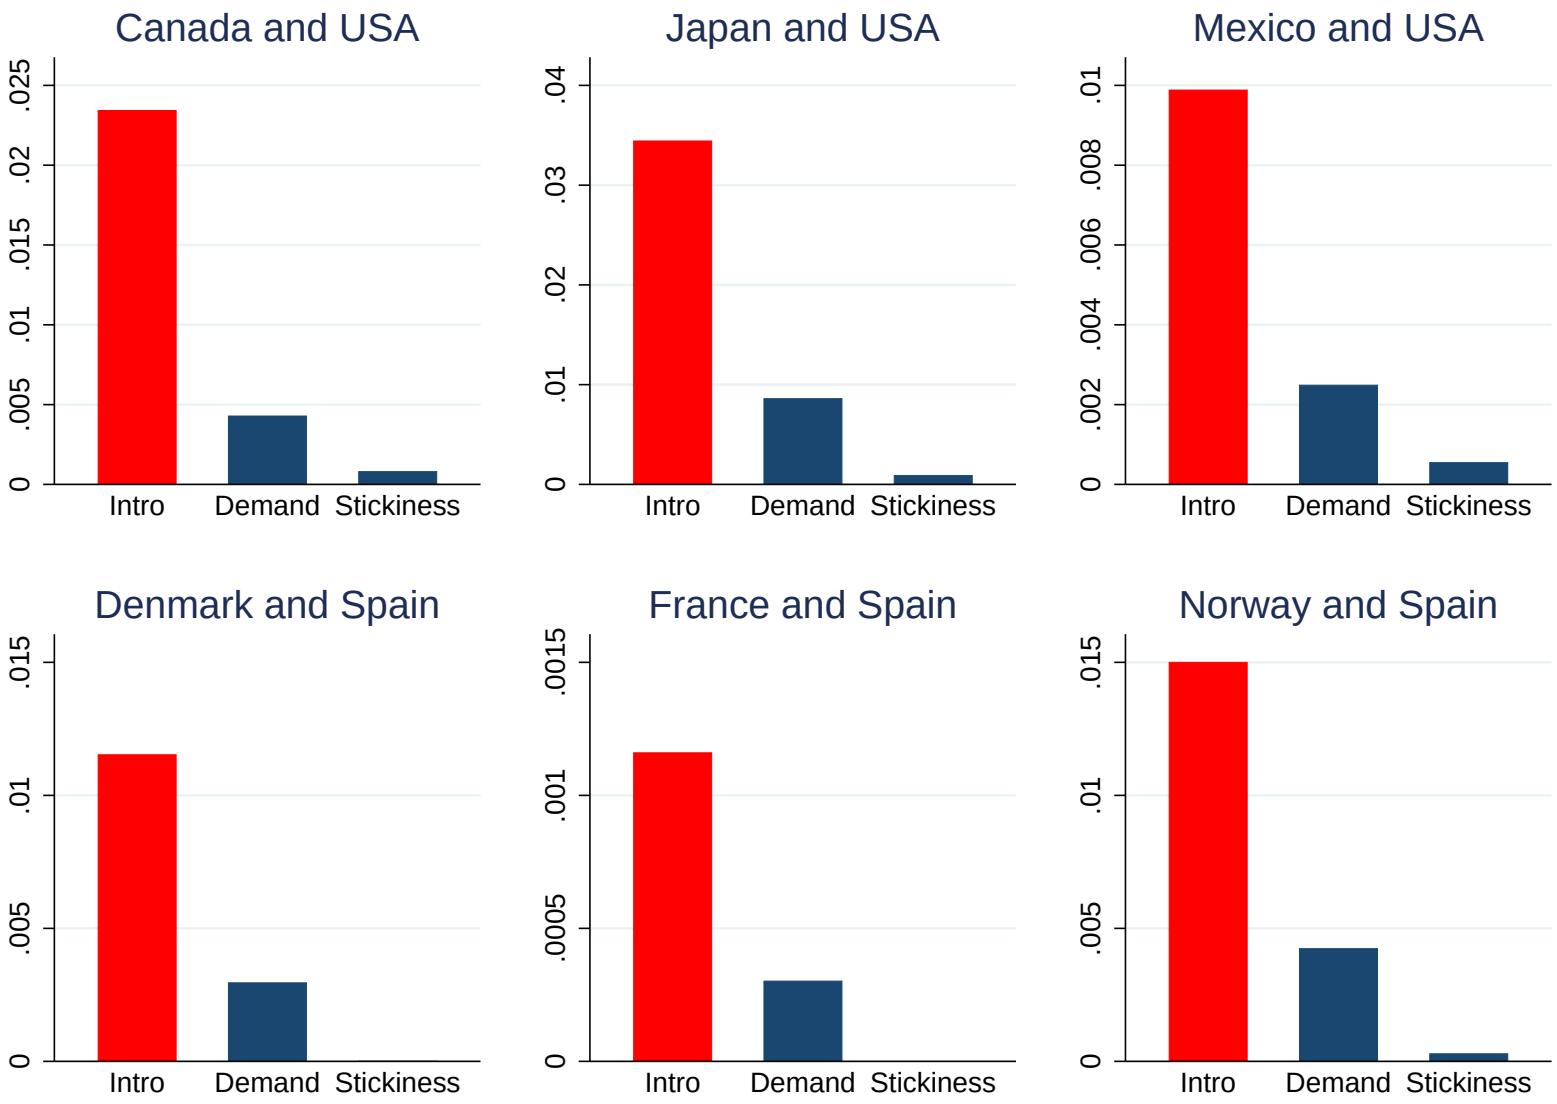


Figure 9: Decomposing the Cross-Section of Good-Level RERs q_{ij}^n for Selected Bilateral Pairs.

Notes: Figure plots the three terms from the cross-sectional decomposition (5). The decomposition 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. Weights are used to equalize the contribution of all stores. We exclude goods and weeks where $|\left(q_{ij}^n\right)| > 0.75$. These excluded observations represent a very small percentage of total observations.

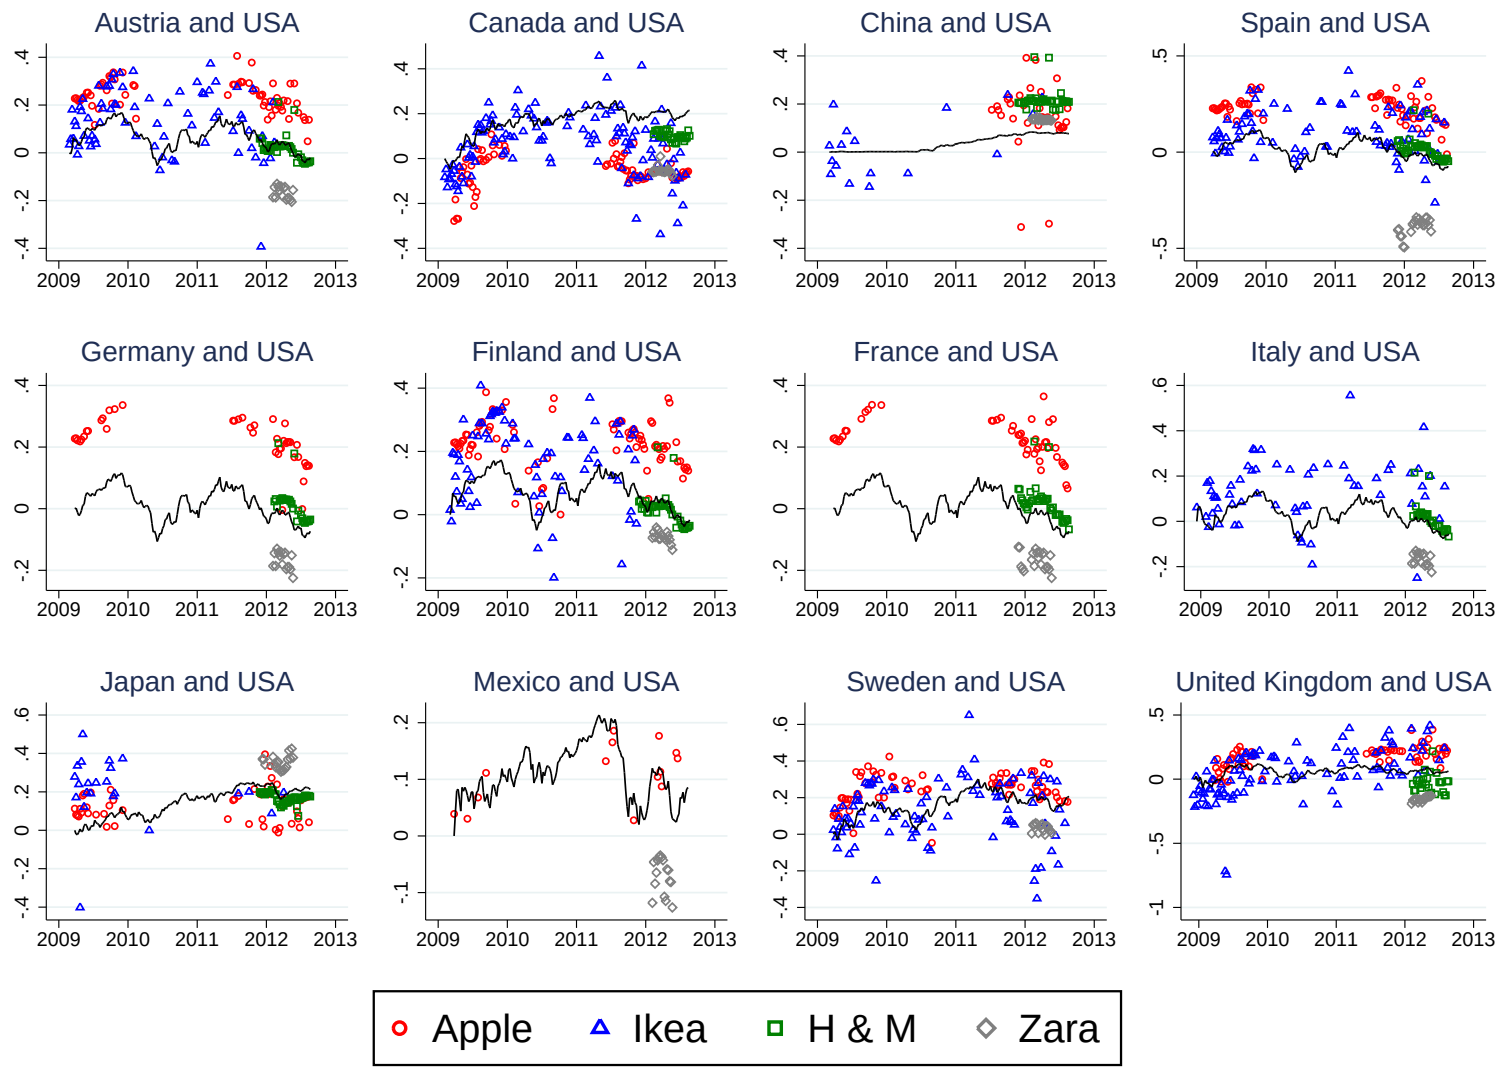


Figure 10: Evolution of Good-Level RERs at Introduction ($q_{ij}^{n,I}$) and the Nominal Exchange Rate, Raw Data

Notes: Figure plots median log good-level RER at the time of introduction for each week and store combination for each bilateral relationship shown. 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 $|q_{ij}^{n,I}| > 0.75$, which is slightly stronger than the filter $|q_{ij}^{n}| > 0.75$ used to capture outliers in the rest of the paper.

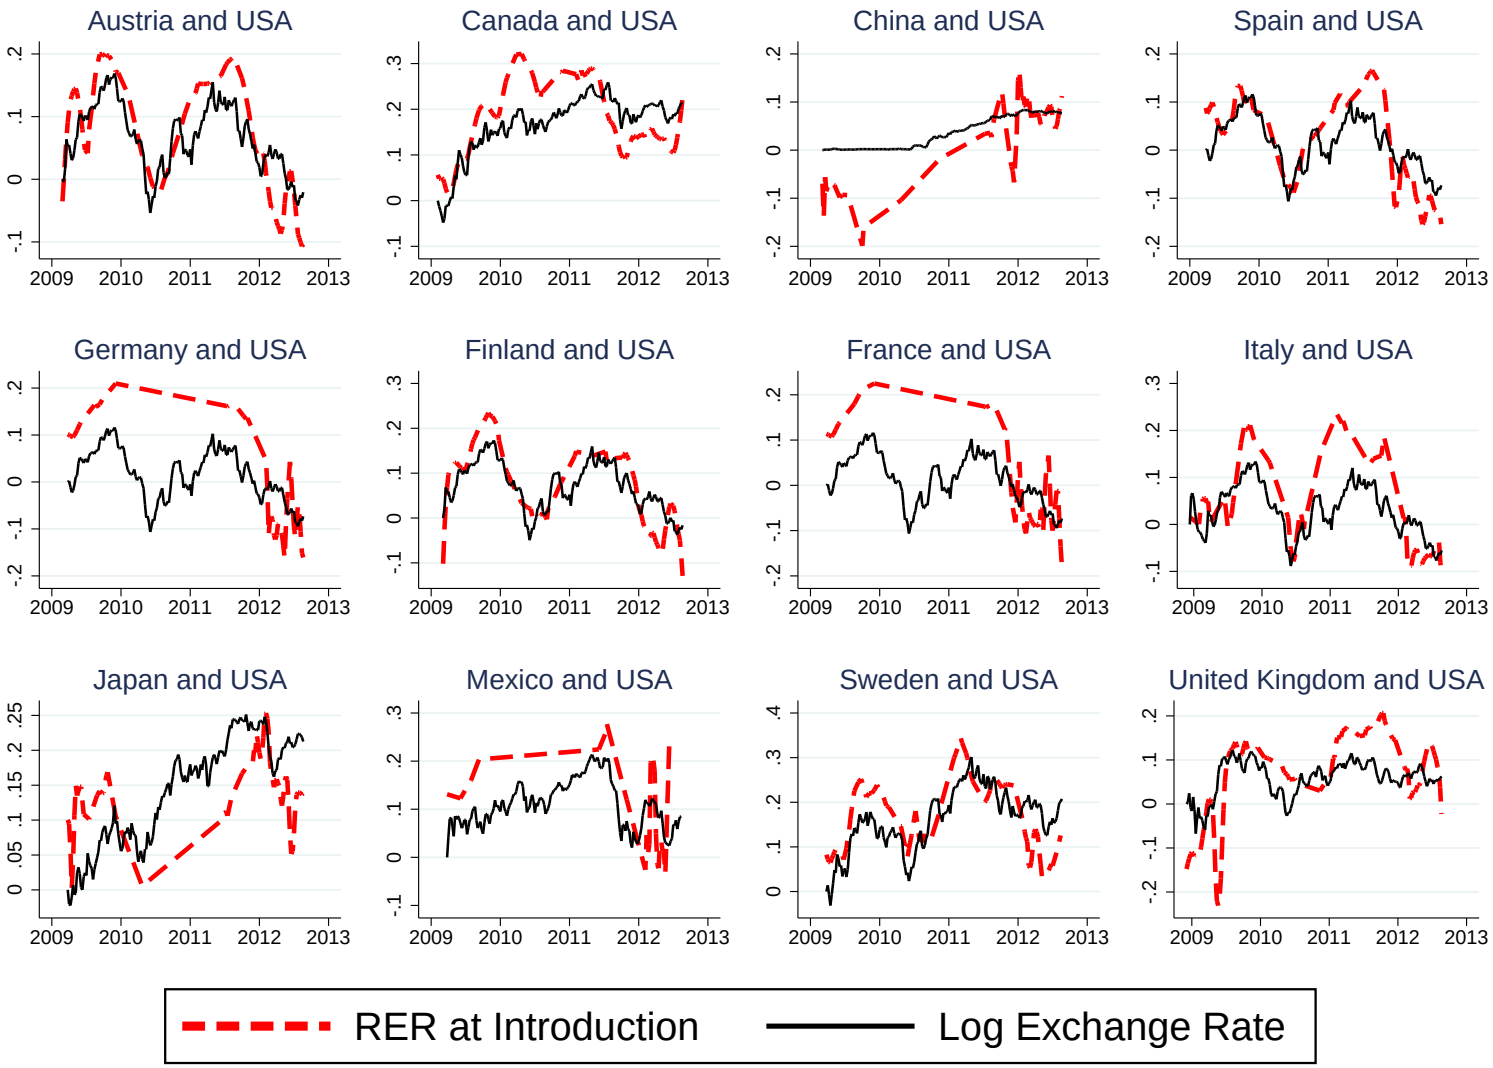


Figure 11: Evolution of Good-Level RERs at Introduction ($q_{ij}^{n,I}$) and the Nominal Exchange Rate, Lowess

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 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. The comovement of the red dashed line and the black line suggest that even at the time of product introductions, when "menu costs" should be irrelevant, companies price with a local currency stickiness. Any given good contributes (at most) to only one point in the figure.