

Currency Unions, Product Introductions, and the Real Exchange Rate

Alberto Cavallo
MIT

Brent Neiman
University of Chicago and NBER

Roberto Rigobon
MIT and NBER

Online Appendix

October 2013

This online Appendix includes sections which provide additional details on the construction and content of our data, alternative RER decompositions, and the variance decompositions.

A Data

In this section of the Appendix, we start by describing our pricing data from The Billion Prices Project (BPP) and then describe additional data we use in our analyses.

A.1 Pricing Data

Our dataset is comprised of prices scraped off the Internet by the BPP, an academic research initiative at MIT. In this paper we use weekly aggregations of data collected on a daily basis, from October 2008 to May 2013, from the websites of four global retailers: Apple, IKEA, H&M, and Zara. These firms sell identical goods across dozens of countries using the same company-wide product ID codes everywhere. Using these codes, we are able to match and compare the prices of over one hundred thousand unique products over time and across countries.

The data were collected using a data scraping software that records the price and product information for all goods sold by each of these retailers in multiple locations. The scraping methodology works using scraping “robots” that follow 3 steps: First, the robot locates and downloads all publicly available web pages where product and price information are shown. Second, it analyzes the underlying HTML code to identify product names, product IDs, prices, and descriptions. This is done by using custom characters in the HTML code that identify the start and end of each variable. For example, prices are usually shown with a dollar sign in front of them and two decimal digits at the end. The scraping software is customized with this information, so it can automatically recognize these patterns in the code and locate the relevant data. Third, the data is stored in a panel-format database that contains one record per product per

day. The data scraping method requires a significant initial effort to build the robots, which must be customized for each retailer and country. The advantage, however, is that the robot can be set to run automatically every day and will continue to collect data as long as the format of the page stays roughly the same. For more on the scraping methodology, see AC2012.

To scrape the data, a robot must know the country, the URLs, and the variables that it is supposed to collect. The list of countries was chosen manually to include the United States and United Kingdom from the start, and then gradually more countries were added to the list as the scraping capabilities of the BPP increased. The selection of the additional countries was chosen randomly by students participating in the project during 2008 and 2009. Tables A.1 through A.4 in this Appendix list, for all stores in our data, the countries and earliest dates covered by our data.

Once a country was included in the list for a particular retailer, it remained in our sample from then on except for occasional periods of scraping errors. The URLs, which identify categories of goods, were updated every few months to add new categories of goods sold by each retailer. Because the starting point of the scraping process is a URL, which lists a set of products that belong to a particular category of goods, the set of products scraped each day changes automatically when new products are introduced and removed from the retailer's website. For example, every day the robot built for Apple in the United States looks into this particular URL: http://store.apple.com/us/browse/home/shop_ipad/ipad_accessories/cases. Over time, if Apple offers new "ipad cases" for sale, the robot immediately starts recording their price and product characteristics. Product information is collected until goods disappear from the website (i.e. they do not appear in any of the URLs being monitored). The variables scraped are always the same: each row in the database has a retailer name, country, ID, description, price, and date.

There are gaps in individual price series that occur when the scraping software fails or individual items are temporarily out of stock. We interpolate between observed prices with the assumption that prices remain unchanged until a change is observed. We exclude the roughly 1 percent of price observations for which we observe implausibly large price variations (if a price is more than 4 times or less than 10 percent of the modal price for that good) or for which the good's relative price across countries is implausibly large (when expressed in dollars at the contemporaneous exchange rate it is greater than 4 times the U.S. price or smaller than 10 percent of the U.S. price). We additionally drop the less than 1 percent of remaining goods with a log RER with an absolute value greater than 0.75.

A.2 Other Data

In addition to price data, we use information on many country- and country-pair variables. Daily bilateral NERs were obtained from Oanda. We consider as currency unions bilateral pairs among euro zone countries, Andorra, Monaco, and Montenegro as well as bilateral pairs among the United States, Ecuador, El Salvador, and Panama. Dollar pegs include Azerbaijan, Bahrain, China (before June 2010), Honduras (before June 2012), Hong Kong, Jordan, Kuwait, Lebanon, Macao, Kazakhstan, Oman, Qatar, Saudi Arabia, Syria (before September 2011), UAE, Ukraine, and Venezuela. Euro pegs include Bosnia and Herzegovina, Bulgaria, Denmark, Lithuania, and Latvia. These include all countries in our data with exchange rate code “1” from the “course” classification in EICRKR2008. Bilateral distance was taken from The CEPII Gravity Dataset used in KHTMJR2010. Table 3 uses the pcgdp variable for 2007 from the Penn World Tables as our measure of “Abs. Relative Income.” We obtain information on state sales tax rates for the United States from The Tax Foundation, on province sales tax rates for Canada from the TMF group, and on VAT rates for other countries from Deloitte.

B Robustness to Alternative RER Decompositions

In this section of the Appendix, we write three alternative decompositions of the good-level RER into introduction, demand, and stickiness terms. The first changes the definition of the RER at introduction to equal the RER when both goods are first available, while the second changes the definition of the stickiness term to exclude NER movements prior to the most recent price change in either country. The third decomposition combines these previous two adjustments. Our baseline results are robust to these alternative decompositions.

Recall that the decomposition (3) used in the main text is:

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

Remember that good z 's RER $q_{ij}(z, t)$ is only defined in our data once the product is available in both markets, i.e. for $t \geq i_{ij}^*(z) = \max\{i_i(z), i_j(z)\}$.

B.1 Change in RER at Introduction Term

Notice, however, that if $i_i(z) \neq i_j(z)$, NER movements prior to the earliest date at which q_{ij} is defined will matter for the decomposition (3). Given the logic in its derivation and related discussion in the text, we believe this is appropriate. Nonetheless, for robustness, we now adjust the decomposition to eliminate this characteristic.

We write an alternative to (3) where the ‘‘RER at introduction’’ term q_{ij}^I simply equals the RER at the time the good is first available in both countries:

$$\begin{aligned}
 q_{ij}(z, t) = & \underbrace{q_{ij}(z, i_{ij}^*(z))}_{\text{Introduction } q_{ij}^I} - \underbrace{\left[\frac{1}{2} \Delta_{l_i(z,t)}^t e_{ij} + \frac{1}{2} \Delta_{l_j(z,t)}^t e_{ij} \right]}_{\text{Stickiness } q_{ij}^S} \\
 & + \underbrace{\left[\Delta_{i_{ij}^*(z)}^{l_i(z,t)} p_i(z) - \Delta_{i_{ij}^*(z)}^{l_j(z,t)} p_j(z) - \frac{1}{2} \Delta_{i_{ij}^*(z)}^{l_i(z,t)} e_{ij} - \frac{1}{2} \Delta_{i_{ij}^*(z)}^{l_j(z,t)} e_{ij} \right]}_{\text{Demand } q_{ij}^D}. \tag{A.1}
 \end{aligned}$$

With this first alternative decomposition (A.1), the introduction and demand terms change, though the stickiness term remains as before. Note that unless the goods are introduced at the same time (i.e. $i_i = i_j = i_{ij}^*$, in which case this disaggregation is equivalent to the baseline (3)), the ‘‘demand’’ term no longer necessarily equals zero when there are no local currency price changes. Under (A.1), q_{ij}^D will include NER movements during the period in between the two dates of introduction, even if prices are completely sticky.

B.2 Change in Stickiness Term

Next, notice that NER movements in country i may be included in the ‘‘Stickiness’’ term even if they occur before the last price change in country j . Again, we believe this is appropriate, but to demonstrate robustness, we now adjust the decomposition such that the stickiness term only reflects NER movements after the last local price change in either country has occurred.

We revert to the original decomposition (3), but now define the stickiness term such that it only includes NER changes subsequent to $t_{ij}^*(z, t)$. We do not alter the original definition of the

introduction term, generating:

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

With this second alternative decomposition (A.2), the stickiness and demand terms change, and the introduction term remains as before. Note that, as with (A.1), unless the goods are introduced at the same time (i.e. $i_i = i_j = i_{ij}^*$), the ‘‘Demand’’ term no longer necessarily equals zero when there are no local currency price changes (though price changes will render (A.2) different from (3) even if introductions occur at the same time). Under (A.2), q_{ij}^D will include NER movements during the period in between the two dates of introduction, even if prices are completely sticky.

B.3 Change in Both Introduction and Stickiness Terms

Finally, we consider implementing both of the above changes, so that the RER at introduction equals the RER at $i_{ij}^*(z)$, the later of the introduction dates, and the stickiness term only captures NER movements subsequent to $t_{ij}^*(z, t)$, the most recent price adjustment in either country. We write this decomposition as:

$$q_{ij}(z, t) = \underbrace{q_{ij}(z, i_{ij}^*(z))}_{\text{Introduction } q_{ij}^I} + \underbrace{\Delta_{i_{ij}^*(z)}^{t_{ij}^*(z,t)} q_{ij}(z)}_{\text{Demand } q_{ij}^D} - \underbrace{\Delta_{t_{ij}^*(z,t)}^t e_{ij}}_{\text{Stickiness } q_{ij}^S}. \tag{A.3}$$

With this third alternative decomposition (A.3), all three terms change relative to (3). Unlike the first two alternative decompositions, this third decomposition shares the property of the baseline decomposition that the demand term will only be non-zero if there are local currency price changes.

B.4 Why Did We Choose (3) as Our Baseline?

Though there are some arguments for each of the modifications implemented in decompositions (A.1) and (A.2), and the appearance of (A.3) is dramatically simpler than (3), we prefer the baseline definition for a number of reasons. First, unlike the first two alternatives, the demand term is only non-zero when there is a local currency price change. Second, unlike the third definition, the baseline decomposition more closely associates the timing of price changes with

the NER changes that would plausibly influence them, since companies surely set prices in each country in a way that benefits profits, regardless of what the implication for some arbitrary bilateral RER is.

More generally, our baseline decomposition apportions consistent treatment to pricing behavior in each country independent of what occurs in other countries. For example, if a good is introduced in one country and changes price before its introduction in a second country, (A.1) will ignore the true introduction price. Further, imagine prices always equaled 1 U.S. dollar translated into foreign currencies at the date of actual introduction in each market, but then prices remained fixed. The decomposition in (3) would characterize the RER at introduction as obeying LOP – as we believe would be desirable in this case – while (A.1) would not due to NER movements in between the dates of introduction.

Similarly, the treatment of stickiness in (A.2) can generate vast differences in the importance of stickiness for a given good sold in China and Japan compared to that same good sold in China and Korea. For example, imagine the good never once changes its price in China. The stickiness term can nonetheless equal zero in (A.2) for China-Japan if prices frequently change in Japan. If prices are highly sticky in Korea, the term can be simultaneously tiny in China-Japan but large in China-Korea. Disaggregation (3) smooths such discrepancies out by apportioning roughly equal weight on the stickiness characteristics of the good in each country.

Despite these arguments, for some purposes, these alternative decompositions might be equally or more useful than our baseline choice. We emphasize that all results in this paper hold qualitatively (and generally also hold quantitatively) whether we use our baseline decomposition as in the main text or instead use any of the three alternatives presented above.

C Variance Decomposition

In this section of the Appendix, we decompose the variance across good-level RERs at any date t as follows:

$$\sigma_{ij}^2(t) = \underbrace{\left[(\tilde{\sigma}_{ij}^I)^2(t) \right]}_{\text{Good-Level RER at Introduction}} + \underbrace{\left[(\tilde{\sigma}_{ij}^D)^2(t) \right]}_{\text{Changes in Demand}} + \underbrace{\left[(\tilde{\sigma}_{ij}^S)^2(t) \right]}_{\text{Stickiness}}, \quad (\text{A.4})$$

where $\sigma_{ij}^2(t)$ is the variance over goods z of $q_{ij}(z, t)$. We use tildes in the terms on the right hand side because those terms include not only the variance of each component but also half of the total contribution of the respective covariance terms, an innocuous assumption given that the covariance terms are small. We measure each term on the right hand side of (A.4) for

a selection of countries against the United States and Spain. We perform the decomposition separately for all available weeks with at least 100 matched goods for each bilateral country pair and average across those weeks. As above, we weight in order to equalize the contribution from all store and bilateral pair combinations in each week. Most weeks will contain some mix of goods which are newly introduced (where $q_{ij} = q_{ij}^I$), some which just experienced a price change (where $q_{ij} = q_{ij}^I + q_{ij}^D$), and some which have had a long time pass since the most recent price change or introduction (where q_{ij}^S may be large). We subtract the mean bilateral RER for each store in order to focus on variation around a mean, rather than changes in the variance due to differences across stores in the mean.³⁵

The bar charts in Figure A.1 show the relative importance of prices at introduction, changes in demand, and stickiness for explaining dispersion in good-level RERs in the cross-section for several bilateral relationships. The left two columns represent the three terms in the decomposition (A.4) when measured on the “Full Sample” of data. We return to the right two columns, labeled “Reduced Sample,” below.

Starting for example with the upper-left plot which represents Canada and the United States, we see that the three bars sum to 0.027 log point, equal to the average cross-sectional variance in q_{ij} for that bilateral pair. We use the decomposition (A.4) to measure that 0.021 log point comes from the introduction term, 0.005 comes from the demand term, and 0.001 comes from the stickiness term. Clearly, nominal rigidity explains little of the dispersion of good-level RERs at any given point in time. This likely reflects the fact that movement in the NER is a common shock applying equally to all goods. And looking at the pairs of Spain with Denmark (which is pegged to the euro) and of Spain with France (which is a member of the euro zone), nominal rigidities cannot explain any of the cross-sectional dispersion for countries with a constant bilateral NER.

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

³⁵We note that this disaggregation is sensitive to the length of data available. For example, it will by construction attribute all cross-sectional variation to the $\tilde{\sigma}_{ij}^I$ term for the first observed week of any data set. We have looked at the time series patterns of the three terms for several bilateral pairs in our data and find that this time-dependency typically quickly dissipates. We report the simple average across periods of the decomposition, rather than setting rules on which data to exclude, for greatest transparency.

The relative price at the time of introduction is far more informative about good-level RERs than anything that happens subsequently including price changes and NER fluctuations. We have calculated this decomposition for all bilaterals in our data to confirm that the bilaterals shown in Figure A.1 are representative. While the exact breakdown among the σ term varies across countries, most cross-sectional variation is generally attributable to good-level RERs at the time of introduction.

To understand whether our conclusions hold qualitatively for longer-lived products with less price stickiness, we re-calculate these results with a “Reduced Sample” that includes only goods that are in our dataset for each bilateral pair for more than one year and which at some point experience at least one price change in at least one of the countries. Comparing the “Full Sample” results in the first two columns of Figure A.1 to their equivalent results from the “Reduced Sample” in the last two columns, it is not surprising that the relative contribution of the “Demand” term increases in all the bar charts. For example, whereas the “Demand” term was about one-fourth as important as the “Intro” term for the Canada-U.S. pair in the “Full Sample,” it is nearly half as important in the “Reduced Sample.” The “Intro” term remains clearly the most important, however, for all plotted bilaterals. We view the high frequency of product introductions and significant price stickiness as calling for a greater emphasis in empirical work on relative prices in levels rather than in changes.

Country	Start Date	Country	Start Date
(1) United States	4/2009	(31) Philippines	4/2009
(2) China	6/2011	(32) Taiwan	4/2009
(3) Japan	4/2009	(33) Czech Republic	10/2012
(4) Germany	4/2009	(34) New Zealand	4/2009
(5) France	4/2009	(35) Hungary	10/2012
(6) Brazil	10/2012	(36) Vietnam	10/2012
(7) United Kingdom	4/2009	(37) Luxembourg	10/2012
(8) Italy	10/2012	(38) Congo	4/2009
(9) Canada	7/2009	(39) Central African Republic	4/2009
(10) Spain	4/2009		
(11) Australia	4/2009		
(12) Mexico	4/2009		
(13) South Korea	4/2009		
(14) Netherlands	4/2009		
(15) Indonesia	4/2009		
(16) Switzerland	10/2012		
(17) Poland	10/2012		
(18) Belgium	10/2012		
(19) Sweden	4/2009		
(20) Norway	4/2009		
(21) Austria	4/2009		
(22) Thailand	4/2009		
(23) Denmark	4/2009		
(24) United Arab Emirates	10/2012		
(25) Malaysia	4/2009		
(26) Finland	4/2009		
(27) Portugal	4/2009		
(28) Hong Kong	4/2009		
(29) Singapore	4/2009		
(30) Ireland	4/2009		

Table A.1: Country and Time Coverage for Apple

Notes: Countries ordered from largest to smallest GDP.

	Country	Start Date		Country	Start Date
(1)	United States	10/2008	(31)	Iceland	11/2009
(2)	China	3/2009			
(3)	Japan	3/2009			
(4)	Germany	9/2012			
(5)	France	9/2012			
(6)	United Kingdom	12/2008			
(7)	Italy	12/2008			
(8)	Canada	2/2009			
(9)	Spain	3/2009			
(10)	Australia	9/2012			
(11)	Netherlands	3/2009			
(12)	Switzerland	2/2011			
(13)	Poland	3/2009			
(14)	Belgium	3/2009			
(15)	Sweden	4/2009			
(16)	Norway	3/2009			
(17)	Austria	3/2009			
(18)	Denmark	3/2009			
(19)	United Arab Emirates	9/2012			
(20)	Malaysia	9/2012			
(21)	Finland	3/2009			
(22)	Portugal	10/2011			
(23)	Hong Kong	9/2012			
(24)	Singapore	9/2012			
(25)	Ireland	9/2012			
(26)	Taiwan	9/2012			
(27)	Czech Republic	2/2009			
(28)	Romania	9/2012			
(29)	Hungary	9/2012			
(30)	Slovakia	3/2009			

Table A.2: Country and Time Coverage for IKEA

Notes: Countries ordered from largest to smallest GDP.

	Country	Start Date		Country	Start Date
(1)	United States	12/2011	(31)	Israel	2/2012
(2)	China	9/2011	(32)	Singapore	9/2011
(3)	Japan	9/2011	(33)	Ireland	2/2012
(4)	Germany	2/2012	(34)	Czech Republic	2/2012
(5)	France	9/2011	(35)	Romania	2/2012
(6)	United Kingdom	2/2012	(36)	Hungary	2/2012
(7)	Italy	2/2012	(37)	Qatar	2/2012
(8)	Canada	2/2012	(38)	Kuwait	2/2012
(9)	Russia	2/2012	(39)	Montenegro	2/2012
(10)	Spain	9/2011	(40)	Slovakia	2/2012
(11)	Mexico	9/2012	(41)	Croatia	2/2012
(12)	South Korea	2/2012	(42)	Oman	2/2012
(13)	Netherlands	2/2012	(43)	Luxembourg	2/2012
(14)	Turkey	2/2012	(44)	Bulgaria	2/2012
(15)	Switzerland	10/2011	(45)	Slovenia	2/2012
(16)	Poland	2/2012	(46)	Lebanon	2/2012
(17)	Belgium	2/2012	(47)	Jordan	2/2012
(18)	Sweden	9/2012	(48)	Latvia	7/2012
(19)	Saudi Arabia	2/2012	(49)	Bahrain	2/2012
(20)	Norway	2/2012			
(21)	Austria	10/2011			
(22)	Thailand	7/2012			
(23)	Denmark	2/2012			
(24)	Greece	2/2012			
(25)	United Arab Emirates	9/2011			
(26)	Malaysia	7/2012			
(27)	Finland	9/2011			
(28)	Portugal	2/2012			
(29)	Hong Kong	10/2011			
(30)	Egypt	2/2012			

Table A.3: Country and Time Coverage for H&M

Notes: Countries ordered from largest to smallest GDP.

	Country	Start Date		Country	Start Date		Country	Start Date
(1)	United States	12/2011	(31)	Finland	2/2012	(61)	Serbia	2/2012
(2)	China	2/2012	(32)	Portugal	9/2011	(62)	Lithuania	2/2012
(3)	Japan	12/2011	(33)	Hong Kong	2/2012	(63)	Croatia	10/2011
(4)	Germany	2/2012	(34)	Egypt	2/2012	(64)	Macao	2/2012
(5)	France	9/2011	(35)	Israel	10/2011	(65)	Panama	2/2012
(6)	United Kingdom	2/2012	(36)	Singapore	10/2011	(66)	Jordan	2/2012
(7)	Italy	2/2012	(37)	Ireland	2/2012	(67)	Latvia	2/2012
(8)	India	10/2011	(38)	Philippines	2/2012	(68)	Cyprus	2/2012
(9)	Canada	2/2012	(39)	Taiwan	2/2012	(69)	Bahrain	2/2012
(10)	Russia	2/2012	(40)	Czech Republic	2/2012	(70)	El Salvador	2/2012
(11)	Spain	9/2011	(41)	Romania	2/2012	(71)	Estonia	2/2012
(12)	Mexico	2/2012	(42)	Kazakhstan	2/2012	(72)	Bosnia and Herzegovinia	4/2012
(13)	South Korea	2/2012	(43)	Ukraine	2/2012	(73)	Honduras	2/2012
(14)	Netherlands	2/2012	(44)	Hungary	2/2012	(74)	Iceland	2/2012
(15)	Turkey	2/2012	(45)	Qatar	2/2012	(75)	Georgia	4/2012
(16)	Indonesia	2/2012	(46)	Kuwait	2/2012	(76)	Malta	2/2012
(17)	Switzerland	2/2012	(47)	Morocco	2/2012	(77)	Monaco	2/2012
(18)	Poland	2/2012	(48)	Slovakia	2/2012	(78)	Montenegro	2/2012
(19)	Belgium	2/2012	(49)	Croatia	2/2012	(79)	Andorra	2/2012
(20)	Sweden	2/2012	(50)	Syria	2/2012			
(21)	Saudi Arabia	2/2012	(51)	Ecuador	9/2012			
(22)	Norway	2/2012	(52)	Oman	2/2012			
(23)	Venezuela	2/2012	(53)	Luxembourg	2/2012			
(24)	Austria	2/2012	(54)	Azerbaijan	2/2012			
(25)	Thailand	10/2011	(55)	Dominican Republic	2/2012			
(26)	Denmark	2/2012	(56)	Bulgaria	2/2012			
(27)	Greece	10/2011	(57)	Slovenia	2/2012			
(28)	United Arab Emirates	10/2011	(58)	Tunisia	2/2012			
(29)	Colombia	2/2012	(59)	Guatemala	2/2012			
(30)	Malaysia	2/2012	(60)	Lebanon	2/2012			

Table A.4: Country and Time Coverage for Zara

Notes: Countries ordered from largest to smallest GDP.

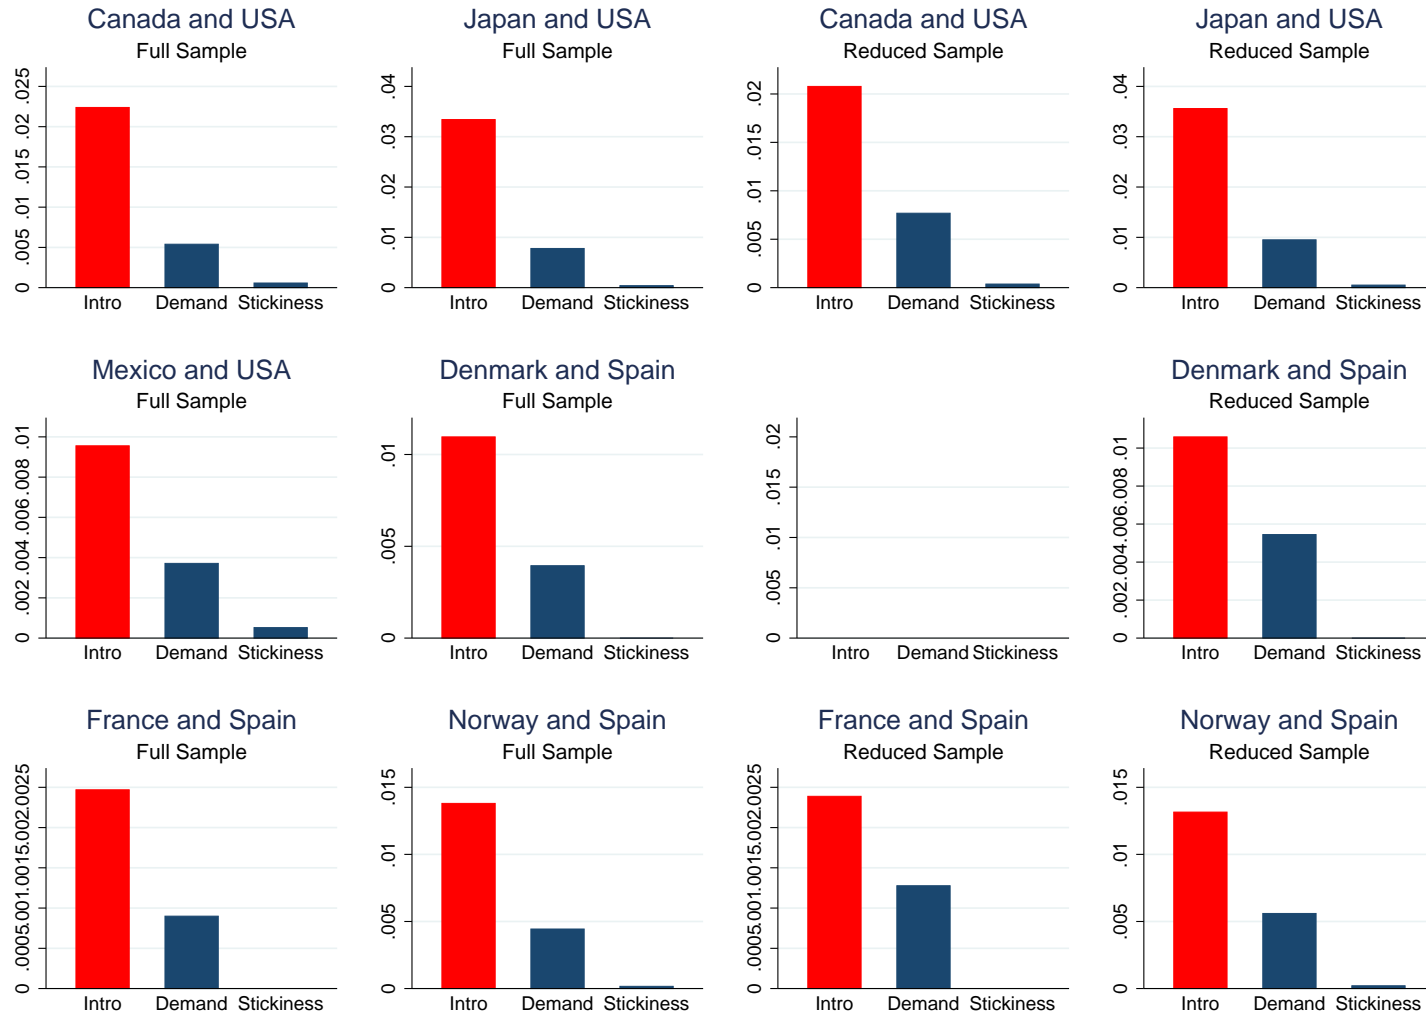


Figure A.1: Decomposing the Cross-Section of Good-Level RERs q_{ij} for Selected Bilateral Pairs

Notes: Figure plots the three terms from the cross-sectional decomposition (A.4). The decomposition in the left two columns (“Full Sample”) is calculated for each country pair at each date that contains at least 100 goods and then the results are averaged across all available dates. The decomposition in the right two columns (“Reduced Sample”) is equivalently calculated but on the subsample of goods that ultimately have at least one price change in one of the countries and that remain in the sample for at least one year. Weights are used to equalize the contribution of all stores within each country pair. We exclude the small number of observations where $|q_{ij}| > 0.75$.