Retail Prices and the Real Exchange Rate*

Alberto Cavallo
MIT, HBS, and NBER

Roberto Rigobon
MIT and NBER

May 22, 2017

Preliminary Draft

Please do not cite or distribute

Abstract

We use a dataset containing daily prices for thousands of matched retail products in nine countries to study tradable-goods real exchange rates. Prices were collected from the websites of large multi-channel retailers and then carefully matched into narrowly-defined product categories across countries, providing relative price *levels* data that collectively represents the bulk of expenditures on food, fuel, and consumer electronics in each country. Using bilateral results with the US, we show that relative prices in local currencies co-move closely with nominal exchange rates, in sharp contrast to the consensus view in the literature. In particular, exchange-rate passthrough into *relative* prices is approximately 75%, compared to just 30% with CPI data for the same countries and time periods. We decompose the difference and show that about 8 percentage points are explained by the exclusion of non-tradable sub-categories, 26 percentage points come from the use of closely-matched products across locations, and 11 percentage points are accounted by the price levels from new and disappearing goods, which are not captured by standard price indices. These results suggest that the retail prices for tradable goods can adjust quickly to nominal exchange rate movements and vice-versa, and have important implications for a vast literature that tries to characterize both the level and behavior of real exchange rates over time.

JEL-Codes: E3, F3, F4.
Keywords: Real Exchange Rate, Law of One Price, Purchasing Power Parity, Exchange Rate Passthrough, Online Prices.

*We thank Brent Neiman for his extensive help and guidance putting together this paper. We also thank Jesse Schreger and seminar participants at MIT Sloan and Boston University for helpful comments and suggestions. Yuriy Gorodnichenko gave an excellent discussion of preliminary results related to this paper.

1 Introduction

The relative price of tradable goods across countries is at the core of many issues in international economics, from comparisons of output and poverty levels to the dynamics of the real exchange rate (RER) and the speed by which shocks are transmitted across borders. A well-established fact in the literature is that retail prices in local currencies are not strongly correlated with the nominal exchange rate (NER), or equivalently, that the retail RER co-moves closely with the NER. This makes the RER both volatile and persistent, a fact labeled the "PPP puzzle" by Rogoff (1996). Initial hyphotesis that this reflected the price of non-tradable goods were challenged by the results in Engel (1999), who showed that tradable RERs constructed with Consumer Price Index (CPI) data are just as volatile and persistent as non-tradable RERs. Furthermore, when measured using import and export price indices, RERs are far less volatile and persistent (Gopinath et al. (2011a)), implying that tradable good prices adjust at-the-dock but not at the retail level.

Many possible explanations for these phenomena have been proposed in the literature, including the ideas that retail markets are segmented by high transportation and distribution costs (Burstein et al. (2003)), that local markups differ over time and space (Atkeson and Burstein (2008)), that existing measures of retail relative prices may reflect biases stemming from sectoral aggregation (Imbs et al. (2005)), temporal aggregation (Taylor (2001)), cross-country differences in goods sampled, and the disregard of entering and exiting goods (Nakamura and Steinsson (2012)).

While progress has been steady, the attempts to reach a consensus have been hampered by formidable empirical challenges. Most of the literature relies on CPIs, which are not designed for international price comparisons. In particular, CPIs preclude any price *level* comparisons, they are are constructed by separate national statistical offices (NSOs) using different methodologies and baskets of goods, and by construction they fail to capture the difference in prices between entering and exiting good varieties, which could be an important margin of price adjustment to nominal exchange rate fluctuations. Academic attempts to use micro data from closely-matched goods are often limited in coverage, with relatively few goods, retailers, and/or countries (See Crucini and Shintani (2008), Gopinath et al. (2011b), and Cavallo et al. (2014)). On the other extreme, the World Bank's International Comparisons Program (ICP) collects and aggregates matched-product prices covering a broad set of goods to compare output and poverty rates across countries, but its current data collection methods, coordinating the work of dozens of NSOs, means the whole process can take five years or longer.[1] Despite many similar efforts, it is still

[1]In addition to the low frequency, each round of ICP data collection uses different product lists and has a

hard to find a large set of highly similar products that span the traded consumption bundle and are sold simultaneously in a large cross-section of countries, with prices that are sampled consistently and at relatively high frequency. Taylor (2001) nicely summarizes these challenges and the literature's failure to meet them:

> To meet the desired standard we would be hoping that hundreds of price inspectors would leave a hundred or more capital cities on the final day of each month, scour every market in all representative locations, for all products, and come back at the end of a very long day, with a synchronized set of observations from Seoul to Santiago, from Vancouver to Vanuatu. We cannot pretend that this happens.

This paper re-examines the behavior of retail relative prices and RERs using a new dataset constructed, in part, with web-scraping programs that substitute for Taylor's "hundreds of price inspectors." After controlling for non-traded subcategories, using a basket of closely-match goods across countries, and accounting for products that enter and dissapear, we find that retail relative prices for tradable goods co-move closely with nominal exchange rates. This makes international relative prices less puzzling and provides support for simple tradable-goods assumptions in the theoretical literature.

Our data contain daily prices for thousands of varieties that collectively represent the bulk of expenditures on food, fuel, and consumer electronics in nine countries: Argentina, Australia, Brazil, China, Germany, Japan, the United Kingdom, South Africa, and the United States. We take the price information from over 50 thousand individual varieties matched to about 350 narrow product categories and use it to build daily relative price and real exchange rate statistics for eight countries relative to the United States.[2]

We enhance the existing data in the literature in two ways. First, relative to CPIs, we greatly improve the cross-country matching of products and methods. We measure prices in common physical units such as grams, items, or liters, so that we can reasonably compare the price levels of products sold in slightly different package sizes or configurations that, otherwise, are largely the same. Second, relative to ICP data, we dramatically increase the frequency of observations by using daily prices (though most of our findings are reported on a monthly basis for comparison

different methodology, leaving their panel dimension difficult to interpret. Similar issues arise with data collected for cost-of-living indices compiled by the Economist Intelligence Unit (EIU) and other consulting firms, which are designed for cross-section comparisons rather than time-series changes. The Center for International Price Research at Vanderbilt aggregates data sources, but the data are limited to european countries (collected by EuroStats) or within-country datasets. At a more narrow level, the Economist magazine's Big Mac index compares international prices of a single iconic product that is available in comparable form throughout most of the world. For the topics at hand, we can question the usefulness of a focus on any single product, no matter how tasty.

[2]Some countries and sectors extend further back, but I generally have stable and consistent coverage starting in 2013.

with CPI data). This also improves our ability to link relative prices to high-frequency exchange rate movements and minimizes the risk of biases emerging from temporal aggregation.

We first validate our measurements by showing that the relative price levels are comparable to those collected physically by the World Bank's ICP in 2011 (even though the comparison is complicated by temporal aggregation and differences in the products sampled). Consistent with a large body of earlier work, we find that price levels for individual goods differ meaningfully across countries. Differences tend to be largest for food and smallest for electronics, but in all cases there seem to be formidable limits to price level equalization across countries.

In contrast with most results in the literature, we find a strong negative co-movement of relative prices in local currencies and the NER. A simple graphical analysis suggests that, in most countries, depreciations are matched by increases in relative prices, and appreciations are accompanied by decreases in relative prices. In some cases relative prices tend to move first, with the NER adjusting to compensate. The RER, while volatile, tends to revert back to stable levels in a matter of months for most countries.

This co-movement of relative prices and nominal exchange rates can be quantified in different ways, but we focus on estimating passthrough rates (from the NER into *relative* prices) due to their simplicity and ease of comparison across datasets.[3] Our main specification, including all sectors and countries, provides an estimate long-run passthrough close to 75 percent. By contrast, when we use CPI data obtained from NSOs for the same countries and time period, the passthrough falls to 30 percent.

What is causing such a large discrepancy with CPI data? To answer this question, we decomposed the difference into stages, starting from equivalent measures using an all-sectors CPI, and gradually adjusting categories, the formulas used, the source of data, and other measurement characteristics.

First, using just CPI data we show that approximately 8 percentage points are explained by the exclusion of non-tradable subcategories and the use of expenditure weights from both countries. In particular, about 4 percentage points can be explained by our focus on food, fuel, and electronics. This is consistent with Engel (1999) and related papers that use goods' CPIs. However, a better proxy for tradable goods can be obtained in some countries by using dissagregated CPIs that better match the categories in our data, which implicitely exclude non-tradable services. Doing so raises passthough by an additional 7 percentage points. This, in turn, is partially compensated

[3]See the Appendix for results using correlations and Granger causality, autoregressive models, tests of stationarity, and Vector Error Correction Models. Our times series are too short to use some of these methods effectively, but the general results are consistent with those in the passthrough regressions: relative prices adjust more in our data than in CPIs.

by a reduction of 3 percentage points when we use a bilateral formula that takes into account expenditure weights from both countries.[4]

Second, passthrough increases about 26 percentage points when we switch from CPIs to the matched-product data. This does not appear to be caused by the online nature of the data, differences between branded or unbranded goods, or differences in the degree of nominal rigidity (see also Cavallo (2017)) and Cavallo (2016)). Instead, when we use data from the same source and categories, but without matching goods across countries, our results dissapear. So why does the product matching matter so much?

One reason is that these are *relative* price regressions. Using the relative price for a matched good can help control for factors that affect prices in both countries and are also correlated with NERs. One example is the currency of invoicing. Recent work by Gopinath et al. (2010) and Gopinath (2015) has shown that local prices adjust differently depending on the currency of invoicing in trade. In particular, passthrough is higher when goods are invoiced in a foreign currency. This implies that calculating relative prices with goods that are not well matched across countries may introduce a bias if the goods are invoiced in different currencies. For example, if we were to use a Lenovo laptop invoiced in pounds in the UK and compare it to an HP laptop invoiced in dollars in the US, the estimated passthrough would be lower than if we compared two lenovo laptops or two HP laptops that have the same invoicing currency in both locations. More generally, as we show analytically in Section 5.3, using prices from unrelated goods in the base country is similar to running the passthrough regression with an omitted variable. The fact that the local price of the exact same good in the base country is also correlated with the NER (in the opposite direction) tends to attenuate the estimated coefficient in the passthrough regression. In principle, the better the matching of products used across countries, the closer we are to controlling for the right base-country price, and therefore the smaller the ommited variable bias.

Another important reason is that the matching process can implicitely help us identify goods that are actually tradable (and therefore connected to the NER). The categories used to match CPI data are simply too broad, with names such as "Games, toys, and hobbies", or "Information Processing Equipment", so many products that go into the CPI calculation may not necessarily be goods that are actually traded or even tradable. We find evidence that this effect is important in our data by running single-country regressions (without relative prices). We find that passthrough is still relatively high, at 56%, which is similar to results that use at-the-dock prices to capture

[4] We use a bilateral Fisher formula to have symmetry in the measurement of the bilateral real exchange rate. In Section 5.2 we show that it has minor effects on measured passthrough.

"purely" tradable goods (Gopinath et al. (2011a)))

Finally, the measurement of RERs in levels allows us to account for the price differences between new and dissapearing varieties, potentially an important marging of adjustment to changes in the NER. We decompose the relative price series into a matched-model component that ignores entry and exit of products, and an additional extensive margin term that captures whether entering and exiting products are more or less expensive than continuing goods and, therefore, alter the average level of prices. We find that an additional 11 percentage points of the difference in our passthrough rates is explained by this extensive margin term. This effect is strongest in electronics, where much of the adjustment to NER shocks appears to take place by replacing old varieties with new ones at different prices. Traditional price indices tend to ignore it because they are constructed by using only price changes from goods available in two consecutive time periods.[5]

Our results challenge the common view that the primary source of RER movements are changes in the price of tradables (Engel (1999), Chari et al. (2002), Betts and Kehoe (2006)). Our data makes retail RERs behave more consistently with prices at-the-dock, and cast doubt on alternative explanations for slow adjustments in CPI data, such as the existence of a large non-tradable distribution costs or pricing-to-market.

At the same time, our findings are consistent with other results in the literature, such as the view that CPI data does not accurately reflect the price of pure-traded goods adjusting after large devaluations (Burstein et al. (2005)), the micro-price evidence linking tradability to PPP deviations (Crucini et al. (2005) and Crucini and Shintani (2008)), the importance of "product-replacement bias" to measure passthrough (Nakamura and Steinsson (2012)), other results showing higher passthrough in online price-comparison websites (Gorodnichenko and Talavera (2014), and the fact that passthrough is high for non-US countries when trade is invoiced in US dollars (Gopinath (2015))[6].

The paper proceeds as follows. Section 2 discusses the literature in greater detail. Section 3 describes the data, the product matching process, and other measurement topics. Section 4 discusses the measurement of real exchange rates in levels, compares the results with 2011 ICP data, and provides a simple graphical analysis of relative prices and nominal exchange rates in each

[5]Unless different goods are deemed to be comparable by the data collector, and some kind of adjustment for quality is made to impute at least part of the price difference between the old and new good as a price change. This, in turn, can cause problems when comparing relative prices with CPIs from countries with different quality adjustment methodologies.

[6]Using import and export prices, Gopinath (2015) finds high passthrough outside the US and low passthrough in the US. This implies high *relative* passthrough at-the-dock for countries other than the US. We find something similar is taking place at the retail level.

country. Section 5 provides the quantitative estimates of relative passthrough and decomposes the difference with results using CPI data. Section 6 concludes.

2 Related Literature

A large literature studies international price differences, absolute and relative law of one price (LOP) deviations, incomplete exchange rate passthrough, and real exchange rate behavior. These topics are intimately related, but data limitations have typically implied that they are studied in isolation.

Papers focusing on LOP deviations historically focused on a narrow set of goods, such as the analysis of prices of Big Macs and their ingredients in Parsley and Wei (2007), prices of *The Economist* magazine in Ghosh and Wolf (1994), or prices from a single retailer as in Haskel and Wolf (2001). Some more recent analyses span a much broader set of data but, typically, are only available at annual or lower frequency and with imperfect product matching. Important work, such as Crucini et al. (2005) and Crucini and Shintani (2008), use data from Eurostat and the Economist Intelligence Unit to highlight the role of non-tradables in generating cross-country price dispersion and to question the extent of persistence in good-level real exchange rates. Others assess international relative prices using retail scanner data, which offer a large set of well-measured relative prices but often constrain analyses to a single retailer and, typically, to comparisons of the United States and Canada. For example, Gopinath et al. (2011a) and Burstein and Jaimovich (2009) use data for a large retailer with presence throughout North America and provide evidence of a large border effect and extensive pricing to market at the wholesale level.

Closely related is the large set of papers documenting the extent of exchange rate passthrough. In recent years, this literature has been typically characterized by analyses of imports or exports of a particular country, using micro data which offer confidence that prices over time correspond to a fixed product or set of products, and which often allow the researcher to condition on price changes. Gopinath and Rigobon (2008), Gopinath et al. (2010), and Neiman (2010), among others, provide early estimates of incomplete passthrough using the micro data underlying the United States import price index constructed by the Bureau of Labor Statistics (BLS). Other work, including Berman et al. (2012), Fitzgerald and Haller (2014) and Amiti et al. (2014), document passthrough and pricing to market using micro-level export or transaction-level trade data for other countries. Our paper provides evidence in support of the hypothesis in Nakamura and Steinsson (2012) that product entry, exit, and substitution result in downward biases in passthrough regressions run on matched-model import price indices. Our work in Cavallo et al. (2014), like that of Gagnon et al. (2014), casted doubts on the quantitative salience of this explanation, while this paper's results suggest the explanation is important when taking into account goods sold by non-global retailers.

Finally, empirical work on the persistence of real exchange rates per se more commonly abstract from micro data and work with sector or aggregate indices obtained from NSOs or similar institutions. Engel (1999) uses real exchange rates constructed from consumer and producer price indices to attribute the bulk of real exchange rate variation at both short and long horizons to the relative price of traded goods. Imbs et al. (2005) offers a view that standard estimates of real exchange rate persistence appear implausible because of aggregation bias. Chen and Engel (2004) argues that sectoral heterogeneity is insufficient for aggregation bias to quantitatively reconcile the real exchange rate's puzzling persistence and documents a number of problems with the sectoral Eurostat data used by Imbs et al.. Carvalho and Nechio (2011) clarifies the source of discrepancy in these results and theoretically demonstrates that a plausibly calibrated multi-sector model of sticky prices can replicate standard measures of the real exchange rate's behavior.

Our focus on scraped online prices as a datasource offers the ability to more easily jointly interact with all three of the above strands of the literature. Like the papers focusing on the cross section of LOP deviations, we can match a large number of products with high degrees of precision. Like the passthrough literature, we can interpret our results conditional on an understanding of price stickiness in the data and can associate, at high frequency, exchange rate movements and relative price movements of the same goods. And like the real exchange rate literature, our data allow us to study multiple countries and with broad enough product coverage to reasonably approximate the behavior of traded goods price indices.

Given these helpful features of online pricing data, a number of recent papers have similarly focused on online prices to study LOP deviations and the behavior of the real exchange rate. Boivin et al. (2012) finds evidence consistent with low passthrough of exchange rates into prices of online book retailers in Canada and the United States. Cavallo et al. (2014) uses scraped data to demonstrate that LOP violations in prices from Apple, IKEA, H&M, and Zara are dramatically larger outside than inside common-currency areas such as the euro zone, and Cavallo et al. (2015) uses a subset of the same data to demonstrate that LOP violations with France, Germany, and Italy rapidly collapsed to zero when Latvia adopted the euro as its currency. In closely related work, Gorodnichenko and Talavera (2014) studies prices from online markets in Canada and the United States. The authors document high degrees of flexibility in online prices and passthrough rates that exceed typical estimates. Gorodnichenko et al. (2015), similarly, examines online prices in the United Kingdom and the United States from a large shopping platform and also emphasizes that online prices appear less sticky. Cavallo (2017) studies prices from large multinational retailers, rather than online markets or shopping platforms, and finds fewer differences when comparing online and offline prices.

International relative prices are also used in the large literature that focuses on the comparisons of national accounts and poverty levels across countries. From the work to create, improve, and maintain the Penn World Tables (e.g. Heston and Summers (88,96), Nuxoll (94), Feenstra, Inklaar and Timmer (13)) to the papers that study how to measure purchasing power parities across countries (e.g. Diewert (99), Hill(99), Deaton and Heston (10)), their implications for international comparison of poverty rates (e.g. Deaton (06, 10)), and the attempts to improve and expand the data collection of prices for similar goods across countries as part of the World Bank's ICP (e.g. World Bank (14), Inklaar and Rao (16), and Deaton and Aten (16)). Our approach to measure real exchange rates relies, heavily, from the methods used by this international comparisons literature; in particular, for the way we classify and match the data across countries. Our analysis in section 5.2 is also closely related to the discussion on "interpolation bias" in that literature (See Inklaar and Rao (16), and Deaton and Aten (16)). An important difference, however, is the fact that for simplicity -and comparison to the international finance literature- we only focus on bilateral real exchange rates, while more advanced multilateral methods are employed in this literature.

3 Data Description

Our data were constructed in two steps. First, prices for individual goods were scraped off the websites of large retailers by an academic project at MIT called The Billion Prices Project (BPP) and a private firm called PriceStats.[7] Second, goods that could be matched across countries were organized and classified as varieties of a particular product. These product definitions are sufficiently narrow, in our view, as to be appropriately described as nearly the same product, with the exception that they are often available in different sized containers or bundles or with only minor variations in characteristics. This matching was performed by PriceStats using a machine-learning classifier combined with a manual verification of each good and its characteristics. Each product was also categorized into the appropriate 1-, 2-, and 3-digit (when possible) code using the

[7] Alberto Cavallo is a co-founder of both the BPP and PriceStats.

United Nations Classification of Individual Consumption According to Purpose (COICOP). All products belong to three sectors: "Food and non-alcoholic beverages" (COICOP 100), "Fuels and lubricans" (COICOP 722), and "Recreation and Culture" (COICOP 900), which we subsequently refer to as "Electronics" because it predominantly includes consumer electronics. We refer to these three sectors below as "1-digit" sectors.

Daily prices for the varieties are translated into per-unit prices and aggregated to form a price for each product in each country on each day. For countries other than Japan and the United States, internet prices are quoted inclusive of taxes (in some cases this is required by law). For Japan, we increase the scraped food and electronics prices by 5 percent prior to April 2014 and by 8 percent subsequently, reflecting a change in their VAT rates. For the United States, we increase food prices by 1.0 percent and electronics prices by 5.1 percent as those are the unweighted average of state sales tax rates in 2014 (food is frequently an exempt category). Our data on fuel prices are always collected inclusive of local tax rates. These prices, together with daily information on the nominal exchange rate, allow us to measure international prices, daily, at product-specific or at aggregated levels.

3.1 Scraping Online Prices

PriceStats scrapes millions of daily prices from hundreds of retailers' sites in over 50 countries to use in high-frequency inflation indices. The focus is, exclusively, on prices from multi-channel retailers, that sell both offline and online. Online-only stores, therefore, are excluded, as are brick and mortar stores that do not sell over the internet. Other than fuel, as discussed below, the data are always collected directly from the retailer's website, never from data aggregators, price comparison sites, or any other third party website, which may, potentially, alter the prices or the sample of products shown. This ensures that the data reflect the actual posted prices that an online buyer of these retailers would have seen when purchasing online on the date the price was collected.

Conceptually, web scraping is simple to understand. Every day, a software downloads the webpage where the product information is shown, scans the underlying code, uses a set of customized rules to identify the relevant information, and stores it in a database. More details on the web-scraping procedure and the BPP can be found in Cavallo and Rigobon (2016).

3.2 Choice of Countries, Stores, and Products

From this large set of pricing data, a subset of countries, stores, and products are used. The scope of the matching of cross-border products, essential for this paper, was dictated by PriceStats and the countries of interest to its clients. Our data on matched products therefore include: Argentina, Australia, Brazil, China, Japan, South Africa, the United Kingdom, and the United States. These countries include five of the seven largest economies, two important "commodity currency" countries in Australia and South Africa, and Argentina, a high-inflation country which was also the first country studied as part of the BPP.

Within each of these countries, pricing data are obtained from large retailers that sell both online and offline. All countries include, at least, the largest three such retailers by sales, though most countries include many more. None of the retailers have exited our sample, so there is no extensive margin in terms of stores.

A distinction is made between "products" and "varieties" in our data. A "product" is actually a narrowly defined category, such as "Regular Heinz Ketchup, 1 gram". Products are chosen based on their availability across multiple countries and, when combined, their representativeness of the bundle of consumption products. These products were chosen to ensure full coverage of all 3-digit COICOP subcategories in the three sectors (food, fuel, and electronics), for which online prices are available. Products are generally not available in all eight countries, but are always found in the United States and, typically, at least, several other countries. The median product is found in six of the eight countries in our data. "Varieties" are individual goods, such as a particular "Heinz Tomato Squeeze 20oz bottle" sold by a given retailer, whose prices are collected from the web and make up the raw data. The varieties present in the raw data of each country are then linked to the products chosen for the cross-country comparison. The price for each of these products is then calculated in each day and in each applicable country by aggregating over the prices observed for a number of varieties of that product.

One difficulty in making cross-country price comparisons of what would, otherwise, be equivalent products is that they are often priced based on slightly varying units, particularly given heterogeneous standards for measurement of weight and volume. For example, apples are sometimes priced in kilograms and sometimes in pounds. Milk can be sold in liters and in half-gallons. Hot dogs arrive in six packs and in eight packs. We calculate the unit price for each variety of each product by dividing the raw price by the number of grams, items, or liters of the product. For each product, we then aggregate across all varieties by taking the geometric mean of these unit prices, though we've verified that alternate moments, such as the median or arithmetic mean, do not change our results.[8]

Types of products for which particular brands enjoy significant global market share are typically divided into products made by those brands and products lacking a global brand. For example, ketchup products are classified according to three types: (i) regular, (ii) low sodium (e.g. no salt), and (iii) other (including flavored). Further, ketchup products manufactured by Heinz are distinguished from those made by other manufacturers. In total, we therefore consider six different ketchup products, whose prices for each day in each country reflect an aggregation over prices of a number of varieties. For example, by the end of our data, the prices of 26 varieties of non-Heinz regular ketchup in the United States and 18 varieties in China were used to create the prices in those countries for the product "non-Heinz regular ketchup". Some products, where branding can make an enormous difference, are only compared for a particular brand. For example, our data include prices for Logitech web cameras, which are separated into two products based on the definition of the cameras, but do not include web cameras of any other brands. Soy sauce products, by contrast, are not distinguished by brand. Table 1 lists a subset of the products that we scrape (the table shows 80 products, roughly one-quarter of the total number that we study).

Table 2 lists, for each of the eight countries in our data, the number of products and median number of varieties per product in each sector. We consider only those prices scraped during 2014 to better highlight meaningful cross-country differences rather than those emerging purely due to different initial dates in our data. As can be seen in the bottom rows corresponding to the United States, our data in 2014 contain 172 different food products and 76 different electronics products. Most countries have more than 100 of the 172 food products and all countries have at least 53 of the 76 electronics products. All countries have at least two of the four total fuel products.

Most of the largest grocery retailers in these countries show the prices for all their products online. In some cases, the prices posted online serve merely to communicate what customers must pay to purchase items in the store. But more commonly, the prices are actually transaction prices. In particular, each of the top five grocery retailers in the United States and in the United Kingdom sell online, while four of the largest five in Australia, three of the largest five in Argentina, China, and Japan, and two of the largest five in Brazil and South Africa also do so.[9] Online sales of food take a variety of forms. Sometimes, the sales are made online but require the customer to pick up the items at the store. In other cases, online food orders can be delivered, either for a flat fee

[8]This strategy reflects our view that the benefits of being able to study a broader set of the consumption bundle with prices averaged across outlets in each country outweighs the cost brought by ignoring price variation that may be due to quantity discounts.

[9]We obtained market shares from Euromonitor International.

or for a fixed percentage delivery charge.

Fuel prices, unlike food and electronics prices, are typically obtained from government websites or other pages unaffiliated with any particular retailer. For example, the prices for fuel in Brazil are obtained from the "National Petroleum Agency," which publishes weekly surveys of fuel prices that are then averaged based on reported fuel sales. Our source for Japan conducts similar weekly surveys. As such, these countries typically use only one variety to represent each fuel product's price. Significantly more varieties are used in China and the United States, where we use data from private organizations which report provincial (for China) and state or city (for the United States) prices. Cross-country differences in our price collection methodology are clearly much larger for fuel than for food or electronics. The ability of fuel prices to alter our overall inference is quite limited though, as we report most of our results excluding fuel items.

The last column of Table 2 lists the median number of varieties used to obtain the daily unit price in each of these countries and categories. The typical food product's price is obtained from aggregating over a large number of varieties, with the median exceeding 12 in five of our eight countries. The median number of varieties used to calculate unit prices of electronics products is always, at least, eight.

3.3 Comparing Online and Offline Prices

Online prices are interesting in their own right as internet purchases constitute a rapidly growing share of consumption spending (Euromonitor, 2014). We claim, however, that the prices in our data, both in terms of their levels and their dynamic properties, are representative of offline prices for equivalent products. The evidence is presented in another paper, Cavallo (2017), where the results of a large-scale, simultaneous, online-offline data collection effort are shown. Online and offline prices for the same products were collected for more than 50 of the largest retailers in 10 countries, including most of those used for this paper. On average, over 70% of the prices were found to be identical. Online and offline price changes, while not simultaneous, have similar frequencies and average sizes. The remaining price differences are driven by location-specific sales, lack of synchronization, and measurement error.

While the retailers in this paper and Cavallo (2017) are not always the same, these results provide strong evidence that large multi-channel retailers tend to have identical online and offline prices in most countries. Furthermore, we can measure the degree of nominal rigidity in our matched-product data and compare it to results in papers that use CPI data. We do this in Table 6 for the US by comparing the frequency of price changes in the three sectors with estimates in

Nakamura and Steinsson (2008). Our prices tend to be *stickier*, particularly for food. This is consistent with the results in Cavallo (2016), whichs shows that measurement bias can lead to an increase the frequency of price changes in both scanner and CPI datasets.

3.4 Product Turnover, Quantities, and other Measurement Issues

The use of individual varieties to compute the daily average price of a given product in each country is affected by goods entering and exiting the sample. The scraping methodology ensures, subject to some errors, that the raw data contain the prices of varieties from the day they are first sold until the day they disappear from the store. However, whether each individual product variety is used to compute the average price for a given product on a specific date depends on whether that variety has been "matched" to the product. As varieties disappear and are replaced with new ones in the raw data, this requires a frequent "matching" of varieties to products to ensure that the average price is being measured with all varieties that fit the description over time.

The varieties in our data were first linked to products in 2013, and then continued to be updated every month since then. Only varieties being sold in 2013 were initially included, so, while historical data from those varieties are used in our price averages before that year, there are varieties in the raw data that disappeared before 2013 and are not impacting our average prices until that time. The number of varieties and products rises gradually from 2010 to 2013, and remains stable afterwards in all countries. To control for these compositional effects, we restrict the sample to start only after a minimum number of products is available for sale in each category (food, fuel, and electronics).

Sale prices that affect all potential buyers are included in our data and statistics (though prices payed with coupons or other personalized discounts are not captured). Sales can greatly affect the prices per product on any given day, and, while they may introduce some high-frequency noise, they can also be an important margin of adjustment for real exchange rates.

Our data lack information about quantities, so we rely as much as possible on expenditure weights for subcategories provided by the national statistical agency in each country. In most cases, expenditure weights are publicly available at COICOP's level 3. Below that level, our averages assign the same weight to every variety. Unfortunately, there is no way of knowing whether a particular variety is sold more than another one, or how this changes over time. This means, for example, that if people switch from an expensive variety to a cheaper one after a devaluation of the currency, we are not able to observe any change. This problem affects not only our data, but also the CPI datasets, which do not contain information for expenditure weights for highly-dissagregated categories. However, an advantage of our data, in the context of this particular example, is that, if the expensive variety eventually disappears from the store or is replaced by a new cheaper version, this would immediately affect the average price measured for that product. Furthermore, if an adjustment via quantities also takes place, it would likely reinforce the high-passthrough results we discuss below.

4 Measuring Real Exchange Rates in Levels

4.1 Local Price Levels

We assume homothetic preferences in all countries over consumption of the same set of "products" $i \in \Omega$. In our empirical work, we will consider "Red Apples" and "Samsung 32 inch Basic LED Televisions" as examples of products i.

Consumption of each product i is itself an aggregation of consumption of a number of "varieties" $j \in \Omega_{i,t}$. In our empirical work, we will consider varieties j of the product "Red Apples" to include Gala, Pink Lady, and Red Delicious. Varieties of the product "Samsung 32 inch Basic LED Televisions" would include variations in color or other minor differences (major differences, such as high definition or "smart TV" features, would constitute a different product). One might think about preferences over these different varieties j of product i as a constant elasticity of demand aggregator, as in the nested CES demand structure in Atkeson and Burstein (2008). Alternatively, one might prefer thinking of this structure as reflecting the aggregation of ideal-type logit consumers within each product type, as described in Anderson et al. (1987).

Varieties j of product i might be packaged or measured in different units. In the case of "Red Apples", for example, prices in the United States are generally quoted per pound while in most other countries they are quoted per kilogram. Or sometimes product varieties include a package of six in one country, but a package of eight in another. For each product i, we define a consistent unit to be used across all countries to measure a unit price $p^y_{ij,t}$, so varieties j of product i can be appropriately compared across time and countries (and outlets therein).

We assume that all varieties $j \in \Omega_{i,t}$ have equal steady-state expenditure shares and that the number of varieties is constant across countries and remains fixed over time, $|\Omega_{i,t}| = |\Omega_i|$. Further, we assume that our price scraping technology captures a subset $N^y_{i,t} \in \Omega_{i,t}$ of the total varieties of product i in each country y at date t . We, therefore, approximate the log ideal price index of product i in country y at time t (up to a constant that is identical across time and countries) as the geometric mean of the unit prices of all scraped varieties. This gives:

$$\ln p_{i,t}^{y} = \frac{1}{N_{i,t}^{y}} \sum_{j \in N_{i,t}^{y}} \ln p_{ij,t}^{y}.$$

The number of true varieties is assumed equal across countries, even if the number of scraped varieties is not, so our measure does not allow for differences in the scale of $N_{i,t}^{y}$ to influence the price level. We are also implicitly assuming that the selection of the subset $N_{i,t}^{y} \in \Omega_{i,t}$ is orthogonal to price levels.

4.1.1 Dispersion in Dollar Prices

The average product prices $\ln p_{i,t}^{y}$ can be converted to US dollars to illustrate the failure of the law of one price across countries. This is done in Table 3, where we measure the cross-country dispersion in USD prices using the mean coefficient of variation for all products in different COICOP level 3 categories.

Electronics tend to have the lowest dispersion of dollar prices across countries, while food categories have the largest. In principle, this could be explained by higher tradability and opportunities for arbitrage in electronics. It can also be related to how narrow the product definitions tend to be within electronics and food, and how "similar" the goods really are across countries.

More generally, price dispersion in dollars is relatively high across countries. This is consistent with the results in the literature that show how the law of one price tends to fail for individual goods, particularly outside currency unions.[10]

4.2 Real Exchange Rate Levels

We now measure the bilateral real exchange rate with the United States for each country. We define e_t^{zy} to be the number of units of country z's currency per unit of country y's currency, so that an increase in e_t^{zy} is an appreciation of country y's currency. The good-level real exchange rate $q_{i,t}^{yz}$ is defined as the difference between prices in countries y and z after being translated into

[10] See Cavallo et al. (2014) for evidence of the law of one price holding within currency unions.

a common currency, so we have:

$$\begin{aligned}\ln\left(q_{i,t}^{yz}\right) &= \ln\left(p_{i,t}^{y}\right) - \ln\left(p_{i,t}^{z}\right) + \ln\left(e_t^{zy}\right) \\ &= \ln\left(rp_{i,t}^{yz}\right) + \ln\left(e_t^{zy}\right), \end{aligned} \tag{1}$$

where $rp_{i,t}^{yz}$ stands for the relative price of product i at time t, though measured in different local currencies of countries y and z. The real exchange rate $q_{i,t}^{yz}$ equals one when the LOP holds exactly.

To reach conclusions about pricing behavior at more general levels, we must aggregate across these product-level measures. We start by generating real exchange rate levels at the 3-digit COICOP level by taking the equally weighted geometric mean of product-level real exchange rates $q_{i,t}^{yz}$ for all products i within that 3-digit category. Examples of 3-digit COICOPs include "Bread and cereals" (111), "Fuels and lubricants for personal transport equipment" (722), and "Equipment for the reception, recording, and reproduction of sound and picture" (911). We then use fixed weights derived from CPI expenditures shares at this 3-digit level in 2014 to aggregate up to our sectors "Food", "Fuel", and "Electronics", or similarly for aggregations beyond that.

If the expenditure weights at level-3 were identical across countries, we could simply take the equally weighted average of product-level real exchange rates to reach more aggregated real exchange rate levels. But, in fact, the weights differ. To aggregate past the 3-digit level, therefore, we create Fisher (1922) indices, which equal the geometric mean of the Paasche and Laspeyres measures of the bilateral real exchange rate. The Paasche aggregates 3-digit real exchange rates using the weights of one of the countries in a bilateral pair, while the Laspeyres aggregates using the weights of the other country. All real exchange rate levels reported below for aggregated sectors use the Fisher measure. This ensures that the real exchange rate we measure is symmetric, meaning that the result is the same no matter which of the two countries is used as the base.

4.2.1 A comparison of 2011 Levels with ICP

The World Bank's ICP is the only public source that generates international price comparisons for such a wide variety of goods and countries, and with a focus on matching goods which are highly similar. In order to compare our results with theirs, we obtained from them the product-level micro data underlying their most recent release, which was in 2011. The ICP covers all countries in our data other than Argentina, and, to align with the time period they cover, we generate an average of our daily real exchange rates during 2011, where available.

We cannot exactly match our products to those used by the ICP but start by generating unweighted geometric means of the real exchange rate levels for all products within each 3-digit COICOP. The ICP data typically include from 6-12 products per 3-digit COICOP, with each product price calculated from a sampling of 5-10 varieties, so the level of aggregation this implies in the ICP data is somewhat similar to that in our data.

Figure 1(b) plots along the x-axis the real exchange rate with respect to the United States for each country and 3-digit COICOP in the ICP data, while the y-axis gives the value for that category in our data. Each data point is labeled to show the country, and hollow circles are used to identify 3-digit categories within food, solid squares are used to identify the category within Fuel, and hollow triangles are used to identify the categories within Electronics. Starting with fuel, we see that our results align quite well with the ICP measures. One reason fuel relative prices lie closest to the 45 degree line may be that concerns over product mismatch are clearly least important in that sector. Food and electronics cluster more or less evenly on both sides of the 45 degree line, but, aside from the two large outliers for South Africa at the bottom-right of the plot, we characterize our data as broadly consistent with ICP's in terms of relative price levels. A robust regression (which places less weight on outliers that it endogenously identifies) projecting the log real exchange rates in our data at the 3-digit level on those in ICP has a highly significant coefficient of 0.67 even after excluding the fuel categories.

Figure 1(a) similarly compares average 1-digit real exchange rates in our dataset with those from the ICP. Food and fuel look are closely related, particularly given that there are necessarily differences due to the time aggregation in the ICP's annual survey. We note that our results differ more meaningfully with the ICP for electronics products, and these are, perhaps, the products where the relative quality of our matching procedure might be expected to be best.

While Figures 1(b) and 1(a) show the similarity of our results to those in the ICP, one might wonder if the same results could have been obtained without worrying about matching goods at such a disaggregated level. For instance, if we simply scraped thousands of prices for products from supermarket web pages and compared their average prices, would this have been sufficient to generate useful data for real exchange rates analysis? To test this out, we repeat the exact analyses above but capture prices that are classified up to a 3-digit COICOP category, the lowest level of aggregation for which CPI expenditure shares exist in all our countries. For example, consider a supermarket web page listing dairy products. Our baseline methodology would use only those prices for our particular matched products, such as various types of unbranded milk (whole, skim, etc.) or various types of Philadelphia cream cheese (regular, low fat, etc.). In this alternate methodology, we simply record and take the average of a random set of prices in the dairy section of the web page, ignoring heterogeneity in product sizes or qualities, but making sure we end up with roughly the same number of varieties from each retailer and in each 3-digit category as in our baseline matched sample. We cannot do this for China and Japan before 2013 because the raw data contained non-latin characters that could not be processed by the machine learning algorithm that classifies the data.

Figure 1(c) plots the results. For any given ICP real exchange rate value on the x-axis, there are, at least, two y-axis values for any given country: one showing the results in our data with exact matches (the solid circles) and another one showing results that ignore product matching (the hollow squares). The exact-matched products track the ICP's values (and, therefore, cluster around the 45 degree line) dramatically more closely than do the unmatched items, which often have average values several orders of magnitude larger or smaller than what is found in ICP. To credibly study international relative prices of comparable goods, effort must be made to standardize on units and quality beyond simply averaging across a large number of goods.

4.2.2 Relative Prices and Nominal Exchange Rates

The top row in Figures 2 to 4 show the daily bilateral real exchange rates with the United States over time for each country. The bottom row shows its components, the nominal exchange rate (a fall is a depreciation) and the relative price (a fall means local prices are falling relative to the United States). These figures only show the aggregate numbers for the three sectors (food, fuel, and electronics), but similar plots excluding fuel and for each individual sector can be found in the Appendix.

Consider first the graphs on the left of Figure 2, which correspond to Australia. Between 2010 and late 2016, the real exchange rate fluctuates around a level of approximately 1.3, which means that this basket of goods is, on average, 30% more expensive than in the United States. While the real exchange rate is quite volatile, there is a tendency for it to mean-revert back to this level over time. The reason for this can be seen in the bottom graph. When the nominal exchange rate was appreciating in 2010-2011, relative prices were falling, which kept the real exchange rate fluctuating around a stable level. Then, around the same time that the currency started to depreciate in mid-2013, relative prices started to rise.

The co-movement between relative prices (in local currencies) and the nominal exchange rate, appears to be even stronger in countries like Brazil and South Africa. Furthermore, in these cases, the level to which the real exchange rate seems to mean revert is closer to 1.

Argentina is a special case because it had both an official and black-market exchange rate during this time period. Figure 3(b) shows results using the official exchange rate, while Figure 3(c) shows the results with the black-market rate. In both cases, the nominal exchange rate tends to depreciate when prices rise, but they do so at different times and with different patterns. The official exchange rate lags behind relative price increases and tends to adjust via strong devaluations that take the relative exchange rate back to one (January 2014 and January 2016). The black-market exchange rate, on the other hand, depreciated faster than relative prices until 2015 , which made these goods in Argentina about 40% cheaper than in the United States for a considerable amount of time. When the foreign exchange market was liberalized in december 2015, both the black-market and the official rate adjusted to take the real exchange rate close to one.

Similar patterns are visible in the other countries, although there is heterogeneity in these behaviors. In China, for example, the real exchange rate rose steadily until 2013. Since then, the nominal exchange rate depreciation has been compensating a steady increase in relative prices. In Japan, prices have been stable while the nominal exchange rate depreciated. In the United Kingdom, it was the nominal exchange rate that was stable until the 2016's Brexit.

While these plots provide suggestive evidence that relative prices and nominal exchange rates are cointegrated, a more formal way to quantify these behaviors is needed to understand what drives them and be able to compare with results using CPI data.

There are many alternative ways to quantify these behaviors. We focus on measuring the rate of passthrough between nominal exchange rates and relative price level for two main reasons. First, it is a simple and transparent way to document the relationship between these two variables. Our data allow us to run a simple regression in levels. There are no complicated methodological assumptions or model specifications that can affect the results. We do not even have to impose a particular lag structure, as would be required for a passthrough regression in changes. Second, there are literally hundreds of papers with passthrough estimates, and surveys of the literature such as Goldberg and Knetter (1997) and Gopinath et al. (2011a) that we can use to compare our results.[11]

[11] In the Appendix, we show alternative results that compare simple correlations, granger causality, autoregressive models and stationarity tests, and vector error-correction models. Many of these methods, however, suffer from inherent limitations that are magnified by the relatively short time period covered by our data.

5 Passthrough Estimates and Decomposition

We can characterize the joint behavior of relative prices and the nominal exchange rate econometrically. To do this, we estimate the cointegration relationship between log relative prices and the log nominal exchange rate as:

$$\ln\left(rp_t^{yz}\right) = \alpha^{yz} + \beta \ln\left(e_t^{zy}\right) + \epsilon_t^{yz}, \tag{2}$$

where t_0 denotes the earliest observation for country pair yz in our data and where $-\beta$ constitutes our estimate of long-run exchange rate passthrough. We always use the United States as a base country in each bilateral pair, $z = USA$. In our main results, we use the overall relative price term in our data, calculated using a Fisher index and including the price levels at product introduction and exit. Later on, we consider a number of different measures for the rp_{t+1}^{yz} term on the left hand side of equation (2), including a term that replicates standard matched-model indices in our data, as well as versions of this relative price constructed using CPI data obtained from NSOs.

Our data are daily but, to be able to properly compare with monthly data on consumer prices collected by NSOs, we preserve only the observation corresponding to the last day of each month, effectively making our dataset monthly.

Our main results are presented in Table 4 for all countries and sectors, and are consistent with the qualitative analysis in the previous section. In each case, we report the β coefficient in equation 2 with the standard error below in parenthesis. Other than China and the United Kingdom, where the nominal exchange rates have barely moved, nearly all standard errors, presented in parentheses beneath the corresponding point estimates, are less than 5 percentage points.

The passthrough rate using all bilaterals and sectors is approximately 75%. Fuel has the highest passthrough rate of 96%, followed by food with 74% and electronics with 55%. Excluding fuel tends to decrease passthrough rates in all countries, with the only exception being Argentina, where fuel prices are strongly regulated by the government. These numbers are significantly higher than most estimates in the literature. For example, Burstein and Gopinath (2013) estimate the long-run passthrough for the tradable CPI in 8 countries to range between 11% and 36%. Why are our results so different?

To answer this question, we created alternative indices using CPI data for the same categories, countries, and time periods available in our online dataset, and, starting with an all-items CPI, we gradually changed the sectors, formulas, and data used. This made it possible to decompose the difference into several stages, as shown numerically in Table 5 and graphically in Figures 5 to 12.

5.1 Decomposing the Difference

In this section, we show that CPI-based passthrough measures are lower than passthrough measures calculated using the matched-model version of our data. And passthrough measures based on the matched-model version of our data are lower than passthrough using our full data inclusive of product entry and exit.

We start in column (1) in Table 5, which pools together all bilateral pairs with the United States and all sectors. Row (1) shows the passthrough estimate using an all-items CPI. At about 30%, this estimate is representative of the results in the literature. The difference with our benchmark estimate, now shown in row (6), is about 45 percentage points.

5.1.1 Sectors and Subsectors

One difference in our work, relative to other papers, is that we focus on just 3 tradable sectors. Row (2) shows the passthrough rate with CPI data when using the same sectors for which we have online data. At this stage, we are using the 1-digit CPIs as provided by the NSOs and aggregating them using their respective official expenditure weights. Passthrough rises by only 4 percentage points.

Under COICOP, these sectors are described as "Food and Non-Alchoholic Beverages", "Transport", and "Recreation and Culture". As the names suggest, these categories contain services and many non-tradable subcategories. So in row (3), we go deeper to try to match the exact same subcategories where we have online data. Fortunately, most of the countries in our sample provide CPIs and expenditure weights at the 3-digit COICOP level (eg. "Bread and Cereals"). The only exceptions are Argentina and China. Using these indices and aggregating them with expenditure weights and arithmetic means that replicate standard CPI methods, we are able to construct a 3-sector CPI that is more directly comparable to our data. This approach implicitely excludes many services and non-tradable subcategories in "Transport" and "Recreation and Culture", which are -by construction- not present in our online matched database. In these cases, passthrough rates rise 7 percentage points, to about 41%. This is, basically, the best proxy for a "tradable" CPI that we can construct for these sectors.

5.2 Formula Effects and "Extrapolation Bias"

Another difference in our methods relative to other passthrough papers is that we use of a Fisher index to measure relative price across countries. This is a common method in the international comparisons literature, and it allows us to take into account expenditure weights from both countries and make our real exchange rate symmetric (independent of which country is chosen as the base).

This could, in principle, explain some of our differences in passthrough rates since using CPIs to measure relative price changes across countries can lead to an "extrapolation bias" in the measurement of relative price changes. This term comes from the international comparisons literature, where annual PPP exchange rates are usually calculated by relying on CPIs to extrapolate relative prices until a new ICP round of data collection takes place.[12]

An extrapolation bias occurs when we use CPIs to substitute for the relative price levels in each country. To compute CPIs, national statistical offices mainly use the weighted geometric average of the price relative of each good. Expenditure shares in each good are given by s_i^y and s_i^z, and they are independetly used to construct the price index in each country.

$$\triangle \ln p_t^y = \sum_i^N s_i^y (\ln p_{i,t}^y - \ln p_{i,t-1}^y) \tag{3}$$

$$\triangle \ln p_t^z = \sum_i^N s_i^z (\ln p_{i,t}^z - \ln p_{i,t-1}^z) \tag{4}$$

The Fisher index of the relative price is given by

$$rp_t^{yz} = \left[\frac{\sum_i^N p_{i,t}^y q_{i,t}^z}{\sum_i^N p_{i,t}^z q_{i,t}^z} * \frac{\sum_i^N p_{i,t}^y q_{i,t}^y}{\sum_i^N p_{i,t}^z q_{i,t}^y} \right]$$

where the first term is a Laspeyres index using the quantities of country y, and the second is a Paasche index using the quantities of country z.

This can be re-written in terms of expenditure shares s_i in the following way (see proof in the Appendix):

$$rp_t^{yz} = \left[\sum_i^N \frac{p_{i,t}^y}{p_{i,t}^a} s_i^z * 1/ \sum_i^N \frac{p_{i,t}^z}{p_{i,t}^y} s_i^y \right]^{1/2}$$

[12] See Deaton (2012) and Inklaar and Rao (2016).

By taking logs, we get:

$$\ln rp_t^{yz} = \frac{1}{2}\left[\sum_i^N s_i^z(\ln p_{i,t}^y - \ln p_{i,t}^z) - \sum_i^N s_i^y(\ln p_{i,t}^z - \ln p_{i,t}^y)\right] \quad (5)$$

If we calculate the $\triangle \ln rp_t^{yz}$ and use the CPI definitions above, we get that the change in relative prices is (see proof in the Appendix):

$$\triangle \ln rp_t^{yz} = \triangle \ln p_t^y - \triangle \ln p_t^z + \frac{1}{2}\left[\sum_i^N (s_i^z - s_i^y)(\ln \frac{p_{i,t}^y}{p_{i,t-1}^y} + \ln \frac{p_{i,t}^z}{p_{i,t-1}^z})\right] \quad (6)$$

The change in relative prices is, therefore, equal to the change in CPIs plus an additional term. This term dissapears only if the expenditure shares or the inflation rates for all sectors are indentical. If not, it can be positive or negative depending on the correlation between relative inflation rates and expenditure shares.

In row (4) of Table 5, we show the passthrough rate when we construct a relative price measure using CPIs and a Fisher index. Relative to a simple ratio of CPIs, the passthrough rate *falls* by about 4 percentage points on average.

To understand the sign of the bias, consider a depreciation of the currency in country y. Our result suggests that the increase in relative prices in y is *smaller* than what is measured with the simple ratio of CPIs. So the term in the equation above has to be negative, and the bias from using CPIs is positive. One way this could happen is that goods with higher share of expenditures (in relative terms) are also the ones with higher relative passthrough rates. That is, if $(\ln \frac{p_{i,t}^y}{p_{i,t-1}^z} + \ln \frac{p_{i,t}^z}{p_{i,t-1}^z})$ is highest when $s_i^y > s_i^z$.

5.3 Matched Online Data and Relative Prices

Continuing with the decomposition, in Row (5) we show the passthrough rate when using our online data to construct a relative price measure that incorporates price changes of continuing goods. At this stage we are not yet incorporating the effect on the average price of a product caused by the price level differences of new and disappearing varieties. Nevertheless, the increase in the passthrough rate is the largest, at 26 percentage points.

A key characteristic of our data is the fact that we have carefully matched products across countries. This can be seen in row (), where we show the passthrough rate when using data from the same online sources but without any matching across countries. These are just products that are classified into the same 3-digit categories. To keep the sample consistent, we use the same number of products that we have in our main dataset for each country and 3-digit category. With un-matched data, there is a much lower passthrough rate.

But why does better product matching matter so much for calculating the relative price passthrough rate? The reason is that when measuring relative passthroughs, we also care about not only about local prices, but also the price in the base country, which is also correlated with the nominal exchange rate. If the base country price is ignored, this amounts to having a standard omitted variable bias. The outcome is the same if a base country price is included, but the "product" selected is completely different, so that it's price is uncorrelated with the local price. Intuitively, the better the quality of the matching of the products across countries, the lower the bias.

To see this more formally, note that the model that we want to estimate is:

$$\ln\left(p_t^y\right) - \ln\left(p_t^z\right) = \alpha + \beta \ln\left(e_t^{zy}\right) + \mu_t \tag{7}$$

or equivalently:

$$\ln\left(p_t^y\right) = \alpha + \beta \ln\left(e_t^{zy}\right) + \ln\left(p_t^z\right) + \mu_t \tag{8}$$

If products are not well matched across countries, the price that we observe in country z is just a proxy for the true price of that good, which can be modeled as:

$$\ln\left(p_t^{proxy}\right) = \nu + \gamma \ln\left(p_t^z\right) \tag{9}$$

This means that the regression being estimated is:

$$\ln\left(p_t^y\right) - \gamma \ln\left(p_t^z\right) = (\alpha + \nu) + \beta \ln\left(e_t^{zy}\right) + \mu_t \tag{10}$$

Our estimate for the β is therefore:

$$
\begin{aligned}
\hat{\beta} &= \frac{Cov\left(\ln\left(e_t^{zy}\right), \ln\left(p_t^{y}\right) - \gamma \ln\left(p_t^{z}\right)\right)}{Var\left(\ln\left(e_t^{zy}\right)\right)} \\
&= \frac{Cov\left(\ln\left(e_t^{zy}\right), \alpha + \beta \ln\left(e_t^{zy}\right) + \ln\left(p_t^{z}\right) + \mu_t - \gamma \ln\left(p_t^{z}\right)\right)}{Var\left(\ln\left(e_t^{zy}\right)\right)} \\
&= \frac{\beta Var\left(\ln\left(e_t^{zy}\right)\right) + (1-\gamma)\, Cov\left(\ln\left(e_t^{zy}\right), \ln\left(p_t^{z}\right)\right)}{Var\left(\ln\left(e_t^{zy}\right)\right)} \\
&= \beta + (1-\gamma) \frac{Cov\left(\ln\left(e_t^{zy}\right), \ln\left(p_t^{z}\right)\right)}{Var\left(\ln\left(e_t^{zy}\right)\right)}
\end{aligned}
\tag{11}
$$

The second term in this equation is equal to the bias in our estimate of the true β. Given our definition of the nominal exchange rate, the term $Cov\left(\ln\left(e_t^{zy}\right), \ln\left(p_t^{z}\right)\right)$ is expected to be positive (e.g. an increase in $\ln\left(e_t^{zy}\right)$ is a depreciation of the currency of country z, which should lead to an increase in $\ln\left(p_t^{z}\right)$.) The sign and magnitude of the bias, therefore, depends entirely on the coefficient γ.

If the products are closely matched, then γ will be close to 1, and the bias will tend to disappear. If the proxy is a different good altogether, then γ will tend to zero and this would be a standard omitted variable bias (equivalent to estimating equation 8 without the term $\ln\left(p_t^{z}\right)$). Given that the CPI data in different countries are collected for goods that are in similar categories, we can expect γ to be somewhere between 0 and 1, making the bias positive and the estimated CPI passthrough rate greater (less negative) than what it really is.

Note that our goal here is to estimate regression 7, so it is really the combination of better product matching and the use of relative prices at the product level that helps us control for the potential omitted variable bias. In some cases there may be interest in estimating the rate of passthrough on local prices, not relative ones. But even in this case, as we discuss in the Appendix, the use of matched relative prices can help reduce the bias caused by "global shocks" that can affect the exchange rate and prices in both locations at the same time.

Finally, another likely reason for why the matching matters might relate to our sampling procedure. It is possible that, by trying to find "matcheable" products, we are identifying goods that are more tradable or somehow more sensitive to movements in the nominal exchange rate. There is some evidence that prices for the matched goods are more flexible than other goods from the same retailers and categories. This can be seen in Table 7, where we compare the frequency of changes across samples in all countries. The matched products are between 19% and 82% more flexible than other goods found in the same online databases. So online prices in general are not more flexible than CPI prices, but the matched products in the online data tend to be more flexible than other online products in the same database.

5.4 Product Exit and Introductions

We now follow the argument presented in Nakamura and Steinsson (2012) and decompose our relative price series into a matched-model component that ignores entry and exit of products (discussed in the previous section) and an additional extensive margin term that captures whether entering and exiting products are more or less expensive than continuing goods and, therefore, alter the average level of prices.

We denote with $N_{i,t}^{y,I}$ the set of varieties of product i which are first sold (or just introduced) at date t, and denote with $N_{i,t}^{y,X}$ the set of varieties last sold (or about to exit) at date t. The number of sampled varieties in our data changes over time, but, since we model the magnitude of the true set of consumed varieties as stable, we do not allow for changes in the product's price index that are purely due to increases or decreases in the number of varieties, as in Feenstra (1994). Entry and exit only matters here through implications on the average per-variety price. We, therefore, write the time $t+1$ price of good i in y as:

$$\begin{aligned} \ln p_{i,t+1}^{y} &= \frac{1}{|N_{i,t+1}^{y}|}\left(\sum_{j\in N_{i,t+1}^{y,I}} \ln p_{ij,t+1}^{y} + \sum_{j\in N_{i,t+1}^{y}-N_{i,t+1}^{y,I}} \ln p_{ij,t+1}^{y}\right) \\ &= n_{i,t+1}^{y,I}\ln\left(\bar{p}_{i,t+1}^{y,I}\right) + (1-n_{i,t+1}^{y,I})\ln\left(\bar{p}_{i,t+1}^{y,I*}\right), \end{aligned} \tag{12}$$

where $n_{i,t+1}^{y,I} = |N_{i,t+1}^{y,I}|/|N_{i,t+1}^{y}|$ is the number of new varieties at time $t+1$ as a share of the total scraped varieties, and where $\bar{p}_{i,t+1}^{y,I}$ and $\bar{p}_{i,t+1}^{y,I*}$ are the geometric means at time $t+1$ of prices in country y of newly entering and preexisting varieties of product i, respectively.

Instead of focusing on the divide between newly entering and preexisting varieties, one can also disaggregate the price level for product i at any time t between varieties that will exit and varieties that will continue to the following period:

$$\begin{aligned} \ln p_{i,t}^{y} &= \frac{1}{|N_{i,t}^{y}|}\left(\sum_{j\in N_{i,t}^{y,X}} \ln p_{ij,t}^{y} + \sum_{j\in N_{i,t}^{y}-N_{i,t}^{y,X}} \ln p_{ij,t}^{y}\right) \\ &= n_{i,t}^{y,X}\ln\left(\bar{p}_{i,t}^{y,X}\right) + (1-n_{i,t}^{y,X})\ln\left(\bar{p}_{i,t}^{y,X*}\right), \end{aligned} \tag{13}$$

where $n_{i,t}^{y,X} = |N_{i,t}^{y,X}|/|N_{i,t}^{y}|$ is the number of exiting varieties as a share of the total scraped varieties, and where $\bar{p}_{i,t}^{y,X}$ and $\bar{p}_{i,t}^{y,X*}$ are the geometric means at time t in country y of prices of varieties that will subsequently exit and varieties that will continue to the following period, respectively.

Combining equations (12) and (13), we can, therefore, write our measure for the change in the price index for product i as:

$$\begin{aligned}\Delta \ln p_{i,t+1}^{y} &= \ln p_{i,t+1}^{y} - \ln p_{i,t}^{y} \\ &= \Delta \ln p_{i,t+1}^{y,MM} + n_{i,t+1}^{y,I} \ln \left(\frac{\bar{p}_{i,t+1}^{y,I}}{\bar{p}_{i,t+1}^{y,I*}} \right) - n_{i,t}^{y,X} \ln \left(\frac{\bar{p}_{i,t}^{y,X}}{\bar{p}_{i,t}^{y,X*}} \right), \end{aligned} \tag{14}$$

where $\Delta \ln p_{t+1}^{MM}$ is the change in a matched-model price index and simply equals the average log change in the price of all varieties that existed in both periods t and $t+1$:

$$\Delta \ln p_{i,t+1}^{y,MM} = \frac{1}{|N_{i,t}^{y} \cap N_{i,t+1}^{y}|} \sum_{j \in N_{i,t}^{y} \cap N_{i,t+1}^{y}} \Delta \ln p_{ij,t+1}^{y}. \tag{15}$$

Equation (14) captures that the price of a product in a country can increase for three reasons. First, continuing varieties may cost more today than previously. Second, varieties that are newly available may cost more than varieties that are not new. Third, those varieties that became unavailable this period used to cost less, on average, than varieties which did not become unavailable in this period.

As with the real exchange rate levels, to reach conclusions about real exchange rate dynamics at more general levels, we use weights from CPI expenditure surveys and create Fisher indices. NSOs generally approximate changes in the log ideal price index as the sum of log changes in item-level prices, using a measure of expenditure shares for the fixed weights. We can combine the definition of the real exchange rate (1) and the decomposition (14) to better understand the drivers of real exchange rate variation. We write:

$$\Delta \ln q_{i,t+1}^{yz} = \Delta \ln \left(rp_{i,t+1}^{yz,MM} \right) + \ln \left(rp_{i,t+1}^{yz,I} \right) - \ln \left(rp_{i,t+1}^{yz,X} \right) + \Delta \ln \left(e_{t+1}^{zy} \right), \tag{16}$$

where:

$$\Delta \ln \left(rp_{i,t+1}^{yz,MM} \right) = \Delta \ln \left(p_{i,t+1}^{y,MM} \right) - \Delta \ln \left(rp_{i,t+1}^{z,MM} \right),$$

and:

$$\ln\left(rp_{i,t+1}^{yz,k}\right) = \eta_{i,t+1}^{y,k} \ln\left(\frac{\bar{p}_{i,t+1}^{y,k}}{\bar{p}_{i,t+1}^{y,k*}}\right) - \eta_{i,t+1}^{z,k} \ln\left(\frac{\bar{p}_{i,t+1}^{z,k}}{\bar{p}_{i,t+1}^{z,k*}}\right),$$

for $k = I, X$. Finally, we can then write the change in the real exchange rate at some level of aggregation as:

$$\begin{aligned} \Delta \ln q_{t+1}^{yz} &= \Delta \ln\left(rp_{t+1}^{yz}\right) + \Delta \ln\left(e_{t+1}^{zy}\right) \\ &= \Delta \ln\left(rp_{t+1}^{yz,MM}\right) + \ln\left(rp_{t+1}^{yz,I}\right) - \ln\left(rp_{t+1}^{yz,X}\right) + \Delta \ln\left(e_{t+1}^{zy}\right), \end{aligned} \tag{17}$$

where $\Delta \ln\left(rp_{t+1}^{yz,MM}\right)$, $\ln\left(rp_{t+1}^{yz,I}\right)$, and $\ln\left(rp_{t+1}^{yz,X}\right)$ are Fisher indices of the equivalent objects at the product level (where they'd have an i subscript), and where $\Delta \ln\left(rp_{t+1}^{yz}\right) = \Delta \ln\left(rp_{t+1}^{yz,MM}\right) + \ln\left(rp_{t+1}^{yz,I}\right) - \ln\left(rp_{t+1}^{yz,X}\right)$. For now, we separate relative price terms from the nominal exchange rate – as opposed, for instance, to merging them to create a term $\Delta \ln q_{t+1}^{yz,MM}$ – to be able to distinguish cases when real exchange rates are stable due to stable prices and nominal exchange rates, as opposed to the case where relative price adjustment offsets movement in the nominal exchange rate.

Row (6) shows that taking product entry and exit into account increases passthrough rates an additional 11 percentage points. Estimates uniformly increase in all sectors, although the increase is small for fuel (where varieties seldom enter or exit) and is large in electronics (where variety turnover is very common).

6 Conclusions

We use a new dataset with carefully-matched products across countries to study the behavior of tradable real exchange rates over time. In contrast to previous results in the literature, we find that tradable good's relative prices co-move strongly with the nominal exchange rate. We quantify this phenomenon by estimating a rate of passthrough from nominal exchange rates to relative prices of 75%, over 45 percentage points higher than consensus estimates in the literature and our own results using CPIs for the same countries and time periods. We decompose this difference by gradually changing sectors, methods, and data, and study the potential reasons for the increase in passthrough at each stage.

Two main factors explain most of the difference. First, the better matching of products across countries allows us to control for an omitted variable bias that exists when prices in other countries are not included, or are proxied by products that are not well matched. The fact that we can run the passthrough regression in relative levels is also key for this result. Second, we find evidence that some of the adjustment in prices to nominal exchange rate occurs through the introduction and discontinuation of varieties of the same product over time. Traditional price indices do not account for the price level differences between old and new varieties, and therefore miss an important margin of relative price adjustment.

Our main result therefore suggest that retail prices for tradable goods can adjust quickly to nominal exchange rates, and vice versa. This supports a common assumption about the behavior of tradable goods in theoretical papers, and casts doubt on the importance of other factors, such as non-tradable distribution costs, as important limits to exchange rate passthrough in retail goods.

While using CPIs for international comparison can lead to measurement biases, we note that our findings do not imply that CPIs are being incorrectly measured. Our point is simply that the passthrough rate of relative prices in tradable goods is much higher than can be observed with CPI data. This can have important implications for Central Banks' policies and the way inflation data are interpreted. For example, by assuming low passthrough rates, a sudden change in inflation might be incorrectly interpreted as coming from other factors, such as a shock in the aggregate demand or supply.

Our work is one of the first attempts to build and characterize the behavior of high-frequency real exchange rates, so the potential uses for these data extend beyond these findings. For example, the measurement of price level differences across countries can help examine barriers to arbitrage and competitiveness in different sectors, improve measurement and cross-country comparisons of output and poverty estimates, and potentially improve forecasts of both inflation and nominal exchange rate levels. To facilitate this research, we publicly share all time series produced for this paper at The Billion Prices Project website (www.thebillionpricesproject.com).

References

Amiti, Mary, Oleg Itskhoki, and Jozef Konings, "Importers, Exporters, and Exchange Rate Disconnect," *American Economic Review*, 2014, *104* (7), 1942–78. bibtex: amiti2014importers.

Anderson, Simon P, Andr De Palma, and Jacques-Francois Thisse, "The CES is a discrete choice model?," *Economics Letters*, 1987, *24* (2), 139–140. bibtex: anderson1987ces.

Atkeson, Andrew and Ariel Burstein, "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 2008, *98* (5), 1998–2031. bibtex: AAAB2008.

Berman, Nicolas, Philippe Martin, and Thierry Mayer, "How do different exporters react to exchange rate changes?," *The Quarterly Journal of Economics*, 2012, *127* (1), 437–492. bibtex: berman2012different.

Betts, Caroline M. and Timothy J. Kehoe, "U.S. real exchange rate fluctuations and relative price fluctuations," *Journal of Monetary Economics*, October 2006, *53* (7), 1297–1326.

Boivin, C., R. Clarck, and N. Vincent, "Virtual Borders," *Journal of International Economics*, 2012, *86* (2). bibtex: JBRCNV2012.

Burstein, Ariel and Gita Gopinath, "International prices and exchange rates," Technical Report, National Bureau of Economic Research 2013.

_ **and Nir Jaimovich**, "Understanding movements in aggregate and product-level real exchange rates," *unpublished paper, UCLA and Stanford University*, 2009. bibtex: burstein2009understanding.

_ **, Joao Neves, and Sergio Rebelo**, "Distribution Costs and Real Exchange Rate Dynamics During Exchange-Rate-Based Stabilizations," *Journal of Monetary Economics*, 2003, *50* (6), 1189–1214. bibtex: ABJNSR2003.

_ **, Martin Eichenbaum, and Sergio Rebelo**, "Large Devaluations and Real Exchange Rates," *Journal of Political Economy*, 2005, *113* (4), 742–784. bibtex: ABMESR2005.

Carvalho, Carlos and Fernanda Nechio, "Aggregation and the PPP Puzzle in a Sticky-Price Model," *The American Economic Review*, 2011, *101* (6), 2391–2424. bibtex: carvalho2011aggregation.

Cavallo, Alberto, "Scraped Data and Sticky Prices," *Review of Economics and Statistics*, 2016, *Forthcoming.*

_ , "Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers," *American Economic Review*, 2017, *1* (107).

_ **and Roberto Rigobon**, "The Billion Prices Project: Using Online Data for Measurement and Research," *Journal of Economic Perspectives*, 2016, *30* (2), 151–78.

_ **, Brent Neiman, and Roberto Rigobon**, "Currency Unions, Product Introductions, and the Real Exchange Rate," *The Quarterly Journal of Economics*, 2014, *129* (2), 529–595. bibtex: ACBNRR2014.

_ **,** _ **, and** _ , "The Price Impact of Joining a Currency Union: Evidence from Latvia," *IMF Economic Review*, 2015, *63* (2), 281–297. bibtex: ACBNRR2015.

Chari, Varadarajan V., Patrick J. Kehoe, and Ellen R. McGrattan, "Can sticky price models generate volatile and persistent real exchange rates?," *The Review of Economic Studies*, 2002, *69* (3), 533–563.

Chen, Shiu-Sheng and Charles Engel, "Does" aggregation bias" explain the PPP puzzle?," Technical Report, National Bureau of Economic Research 2004. bibtex: chen2004does.

Crucini, M.J. and M. Shintani, "Persistence in law of one price deviations: Evidence from micro-data," *Journal of Monetary Economics*, 2008, *55* (3), 629–644. bibtex: MCMS2008 bibtex[owner=cbarger].

_ **, C.I. Telmer, and M. Zachariadis**, "Understanding European real exchange rates," *The American Economic Review*, 2005, *95* (3), 724–738. bibtex: MCCTMZ2005 bibtex[owner=cbarger].

Deaton, Angus, "Consumer price indexes, purchasing power parity exchange rates, and updating.," *ICP Techinical Advisory Panel Meetings*, 2012.

Engel, Charles, "Accounting for U.S. Real Exchange Rate Changes," *Journal of Political Economy*, 1999, *107* (3), 507–538. bibtex: CE1999.

Euromonitor, "Internet vs Store-Based Shopping: The Global Move Towards Omnichannel Retailing," 2014. bibtex: euromonitorinternet2014.

Feenstra, Robert, "New Product Varieties and the Measurement of International Prices," *American Economic Review*, May 1994, *84*, 157–177. bibtex: RF1994.

Fisher, Irving, *The making of index numbers: a study of their varieties, tests, and reliability*, Houghton Mifflin, 1922. bibtex: fisher1922making.

Fitzgerald, Doireann and Stefanie Haller, "Pricing-to-market: evidence from plant-level prices," *The Review of Economic Studies*, 2014, *81* (2), 761–786. bibtex: fitzgerald2014pricing.

Gagnon, Etienne, Benjamin R Mandel, and Robert J Vigfusson, "Missing Import Price Changes and Low Exchange Rate Pass-Through," *American Economic Journal: Macroeconomics*, 2014, *6* (2), 156–206. bibtex: gagnon2014missing.

Ghosh, Atish R and Holger C Wolf, "Pricing in international markets: lessons from The Economist," Technical Report, National Bureau of Economic Research 1994. bibtex: ghosh1994pricing.

Goldberg, Pinelopi K. and Michael M. Knetter, "Goods Prices and Exchange Rates: What Have We Learned?," *Journal of Economic Literature*, 1997, *35* (3), 1243–1272. bibtex: knetter1997goods.

Gopinath, Gita, "The international price system," Technical Report, National Bureau of Economic Research 2015.

_ and Roberto Rigobon, "Sticky Borders," *Quarterly Journal of Economics*, 2008, *123* (2), 531–575. bibtex: GGRR2008 bibtex[owner=cbarger].

_ , Oleg Itskhoki, and Roberto Rigobon, "Currency Choice and Exchange Rate Pass-through," *American Economic Review*, March 2010, *100*, 304–336. bibtex: GIR2010.

_ , Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li, "International Prices, Costs, and Markup Differences," *American Economic Review*, 2011, *101* (6), 2450–2486. bibtex: GGHL2011.

_ , _ , _ , and _ , "International Prices, Costs, and Markup Differences," *American Economic Review*, October 2011, *101*, 2450–2486. bibtex: GGHL2011b.

Gorodnichenko, Y. and O. Talavera, "Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-Border Integration," 2014. Working Paper bibtex: GT2014.

Gorodnichenko, Yuriy, Viacheslav Sheremirov, and Oleksandr Talavera, "Price Setting in Online Markets: Does IT Click?," 2015. Working Paper bibtex: GST2015.

Haskel, Jonathan and Holger Wolf, "The Law of One Price – A Case Study," *The Scandinavian Journal of Economics*, 2001, *103* (4), 545–558. bibtex: JHHW2001.

Imbs, Jean, Haroon Mumtaz, Morten Ravn, and Helene Rey, "PPP Strikes Back: Aggregation and the Real Exchange Rate," *Quarterly Journal of Econoimcs*, 2005, *120* (1), 1–43. bibtex: IMRR2005.

Inklaar, Robert and Prasada Rao, "International Comparison Program (ICP): A Move towards the Compilation of Purchasing Power Parities (PPPs) and Real Expenditures on an Annual Basis," *Working Paper*, 2016.

Nakamura, Emi and Jn Steinsson, "Lost in Transit: Product Replacement Bias and Pricing to Market," *The American Economic Review*, 2012, *102* (7), 3277. bibtex: ENJS2012.

_ **and Jon Steinsson**, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464. bibtex: NS2008.

Neiman, Brent, "Stickiness, Synchronization, and Passthrough in Intrafirm Trade Prices," *Journal of Monetary Economics*, 2010, *57* (3), 295–308. bibtex: BN2010b.

Parsley, David C and Shang-Jin Wei, "A prism into the PPP puzzles: The Micro-Foundations of big mac real exchange rates*," *The Economic Journal*, 2007, *117* (523), 1336–1356. bibtex: parsley2007prism.

Rogoff, Ken, "The Purchasing Power Parity Puzzle," *Journal of Economic Literature*, 1996, *34* (2), 647–668. bibtex: KR1996.

Taylor, Alan, "Potential Pitfalls for the Purchasing-Power-Parity Puzzle? Sampling and Specification Biases in Mean-Reversion Tests of the Law of One Price," *Econometrica*, 2001, *69* (2), 473–498. bibtex: AT2001.

Food Products	Electronics Products
Basmati White Rice	LG Basic Blu-Ray Player
Jasmine White Rice	LG Specialized Blu-Ray Player
Wheat All-Purpose Flour	Samsung Blu-Ray Player
Barilla Spaguetti (including whole grain)	Samsung Specialized Blu-Ray Player
Non-Barilla Spaguetti (including whole grain)	Sony Blu-Ray Player
Kellogg's Breakfast Cereal (excluding gluten free)	Sony Specialized Blu-Ray Player
Kellogg's Granola Breakfast Cereal	Samsung 32 Inch LED TV (excluding HD, Smart, 3D)
Non-Kellogg's Breakfast Cereal (excluding gluten free)	Philips 32 Inch LED TV (excluding HD, Smart, 3D)
Non-Kellogg's Granola Breakfast Cereal	Panasonic 32 Inch LED TV (excluding HD, Smart, 3D)
Ground Beef	Sony 44-47 Inch LED TV (Full HD, Smart, or 3D)
Chicken Breast (whole)	Toshiba 44-47 Inch LED TV (Full HD, Smart, or 3D)
Honey-Baked Ham Cold Cut	Samsung 61-65 Inch LED TV
Smoked Ham Cold Cut	LG 61-65 Inch LED TV
Low-Sodium Ham Cold Cut	Apple Ipod Shuffle 2GB
Low Fat Hot Dogs	Apple Touch 32GB
Regular Hot Dogs	Sony In-Ear Earphones
Canned Tuna in Oil	Beats In-Ear Earphones
Canned Tuna in Water	Sennheiser Over-Ear Headphones
Philadelphia Regular Cream Cheese	Skullcandy Over-Ear Headphones
Philadelphia Fat Free or Low Fat Cream Cheese	Logitech Basic Webcam
Brown Eggs	Non-Logitech Basic Webcam
White Eggs	Apple 13 Inch Macbook
Nutella Chocolate Spread	Sony VAIO 14-16 inch Laptop (No Touchscreen)
Extra Virgin Olive Oil	Apple Ipad Air 32GB (excludes 3G)
Illy Coffee Beans (excluding decaf)	Apple Ipad Air 64GB
Illy Decaf Coffee Beans	Apple Ipad 4 16GB with 3G
Non-Illy Coffee Beans (excluding decaf)	Samsung 7inch Tablet
Non-Illy Decaf Coffee Beans	HP Color Laser Printer
Nesquik Chocolate Milk Mix	Xerox Color Laser Printer
Twinnings Earl Grey Tea Bags	Sandisk 4GB Memory Card
Twinnings Green Tea Bags	Sandisk 32GB Memory Card
Non-Twinnings Earl Grey Tea Bags	Sony 4GB Memory Card
Non-Twinnings Green Tea Bags	Sony 32GB Memory Card
Nestle Mineral Water	Sony Playstation 3 500GB
Dasani Mineral Water	Sony Playstation 3 500GB Super Slim
Tropicana Pulp Free Orange Juice	Sony Playstation 4
Tropicana Orange Juice With Pulp	Microsoft Xbox 360
Non-Tropicana Pulp Free Orange Juice	GoPro Full HD Camcorder
Non-Tropicana Orange Juice With Pulp	Nikon 20-24mpx Digital SLR Camera

Table 1: Sample Product Definitions

Notes: TBD.

Country	Sector	Products	Median Varieties per Product
Argentina (ARG)	Food and non-alcoholic beverages	100	24
Argentina (ARG)	Fuels and lubricants	2	1
Argentina (ARG)	Electronics	58	9
Australia (AUS)	Food and non-alcoholic beverages	119	12
Australia (AUS)	Fuels and lubricants	4	1
Australia (AUS)	Electronics	64	10
Brazil (BRA)	Food and non-alcoholic beverages	120	14
Brazil (BRA)	Fuels and lubricants	2	1
Brazil (BRA)	Electronics	70	14
China (CHN)	Food and non-alcoholic beverages	105	7
China (CHN)	Fuels and lubricants	4	32
China (CHN)	Electronics	70	17
Japan (JPN)	Food and non-alcoholic beverages	58	4
Japan (JPN)	Fuels and lubricants	2	1
Japan (JPN)	Electronics	53	11
South Africa (ZAF)	Food and non-alcoholic beverages	89	5
South Africa (ZAF)	Fuels and lubricants	3	1
South Africa (ZAF)	Electronics	54	8
United Kingdom (GBR)	Food and non-alcoholic beverages	134	17
United Kingdom (GBR)	Fuels and lubricants	3	2
United Kingdom (GBR)	Electronics	69	22
United States (USA)	Food and non-alcoholic beverages	172	28
United States (USA)	Fuels and lubricants	4	51
United States (USA)	Electronics	76	50

Table 2: Summary Statistics

Notes:

L3 COICOP Category	**Mean CV USD Price** (1)	**Countries** (2)
Fish and seafood	0.23	9
Games, toys and hobbies	0.26	8
Information processing equipment	0.27	8
Photographic and cinematographic equipment and optical instruments	0.28	8
Equipment for the reception, recording and reproduction of sound and picture	0.30	8
Fuels and lubricants for personal transport equipment	0.30	7
Recording media	0.34	8
Sugar, jam, honey, chocolate and confectionery	0.36	9
Milk, cheese and eggs	0.36	8
Coffee, tea and cocoa	0.37	7
Mineral waters, soft drinks, fruit and vegetable juices	0.39	5
Food products n.e.c.	0.39	8
Meat	0.41	8
Bread and cereals	0.41	8
Vegetables	0.44	7
Oils and fats	0.49	8
Fruits	0.53	7

Table 3: Which Categories have Greatest Dispersion in US Dollar Prices?

Notes: We first calculate the price in USD in each country. We then get the coefficient of variation across countries (by product and month). We then average for all months, and, finally, for all products.

Relative Price	**3 Sectors** (1)	**Ex-Fuel** (2)	**Food** (3)	**Fuel** (4)	**Electronics** (5)
(1) All Countries	-0.749	-0.721	-0.738	-0.955	-0.553
	(0.013)	(0.025)	(0.027)	(0.016)	(0.031)
(2) Argentina	-0.790	-0.987	-1.010	-0.914	-0.988
	(0.029)	(0.055)	(0.058)	(0.041)	(0.105)
(3) Australia	-0.655	-0.508	-0.577	-0.855	-0.164
	(0.027)	(0.044)	(0.052)	(0.031)	(0.065)
(4) Brazil	-0.852	-0.575	-0.592	-1.383	-0.392
	(0.042)	(0.052)	(0.053)	(0.057)	(0.062)
(5) China	-1.122	-0.921	-1.062	-1.690	-0.369
	(0.154)	(0.143)	(0.169)	(0.367)	(0.115)
(6) Germany	-0.776	-0.593	-0.580	-0.920	-0.435
	(0.061)	(0.096)	(0.100)	(0.058)	(0.095)
(7) Japan	-0.208	-0.170	-0.266	-0.660	0.106
	(0.037)	(0.066)	(0.075)	(0.046)	(0.090)
(8) South Africa	-0.780	-0.591	-0.508	-0.956	-0.843
	(0.020)	(0.058)	(0.065)	(0.020)	(0.060)
(9) UK	-0.582	-0.113	-0.069	-1.330	-0.219
	(0.097)	(0.113)	(0.149)	(0.108)	(0.055)

Table 4: Passthrough Estimates

Notes: All bilaterals calculated with respect to the United States. Results for benchmark series labelled “CN Overall” in other tables.

Price Measure	Relative Price Regressions					Price Regressions	
	3 Sectors	Ex-Fuel	Food	Fuel	Electronics	3 Sectors	3 Sectors -USA
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1) CPI All items	-0.296					-0.374	-0.208
	(0.007)					(0.007)	(0.032)
(2) CPI 1-Digit	-0.344	-0.269	-0.251	-0.452	-0.183	-0.361	0.040
	(0.008)	(0.013)	(0.014)	(0.010)	(0.023)	(0.008)	(0.041)
(3) CPI 3-Digit	-0.414	-0.299	-0.278	-0.743	-0.219	-0.357	0.355
	(0.011)	(0.015)	(0.015)	(0.021)	(0.031)	(0.010)	(0.063)
(4) CPI 3-Digit Fisher	-0.376	-0.268	-0.268	-0.701	-0.194	-0.344	0.431
	(0.010)	(0.015)	(0.014)	(0.019)	(0.028)	(0.010)	(0.062)
(5) PPP Matched Model	-0.638	-0.475	-0.513	-0.948	-0.117	-0.557	
	(0.013)	(0.024)	(0.022)	(0.016)	(0.040)	(0.019)	
(6) PPP Overall	-0.749	-0.721	-0.738	-0.955	-0.553		
	(0.013)	(0.025)	(0.027)	(0.016)	(0.031)		
(7) PPP Overall Branded		-0.662	-0.661		-0.586		
		(0.026)	(0.028)		(0.033)		
(8) PPP Overall Unbranded		-0.69	-0.736	-0.955	-0.348		
		(0.026)	(0.028)	(0.016)	(0.047)		

Table 5: Passthrough Decomposition - All countries

Notes: All bilaterals calculated with respect to the United States.

Category	Monthly Frequency US PPP Online Data	US CPI Data Nakamura & Steinsson (08)	Ratio PPP/CPI
Panel A: Weigthed Means by Sector			
3- Sectors (matched)	46.7	48.5	0.96
Food	25.0	32.3	0.77
Fuel	96.1	87.4	1.10
Electronics	20.8	17.9	1.17
Panel B: Weighted Means by Sub-sector			
Cereals	17.2	23.7	0.73
Flour	29.6	15.4	1.92
Bread	10.7	27.1	0.39
Beef and Veal	29.1	45.6	0.64
Poultry	29.6	34.2	0.87
Whole Milk	24.0	32.3	0.74
Butter	23.2	38.3	0.60
Vegetable Oils	33.5	24.7	1.35
Fruits	20.5	45.1	0.45
Vegetables	23.1	46.7	0.49
Sugar	24.2	18.1	1.34
Sauces & Condiments	23.1	20.3	1.14
Tea	24.3	26.4	0.92
Cocoa and chocolate	17.2	16.5	1.04
Mineral water	13.6	29.4	0.46
Soft Drinks	31.7	29.7	1.07
Fuel	96.1	87.4	1.10
Cameras	29.6	17.9	1.66
Computers	26.1	32.9	0.79
Games and Toys	14.9	11.5	1.30

Table 6: Stickiness - Comparison with CPI data from US

Notes: Coicop Level 4 and US ELI classification matched using coding from Berardi et al (2015). Weights from Nakamura and Steinsson (08) ELI table. Not all categories could be matched.

Category	Country	Monthly Frequency PPP Online Data	BPP Online Data	Ratio PPP/BPP
3 Sectors	Argentina	38.4	32.2	1.19
	Australia	41.1	33.9	1.21
	Brazil	52.4	38.5	1.36
	China	30.2	16.7	1.82
	Germany	30.8	22.3	1.38
	Japan	25.8	19.5	1.33
	South Africa	37.3	23.8	1.57
	UK	43.2	35.1	1.23
	USA	48.7	30.5	1.60

Table 7: Stickiness - PPP Online Data vs BPP Online Data

Notes: The PPP online data is a subset of the BPP online data. It includes only the goods that are matched across countries and used for computing the bilateral RERs.

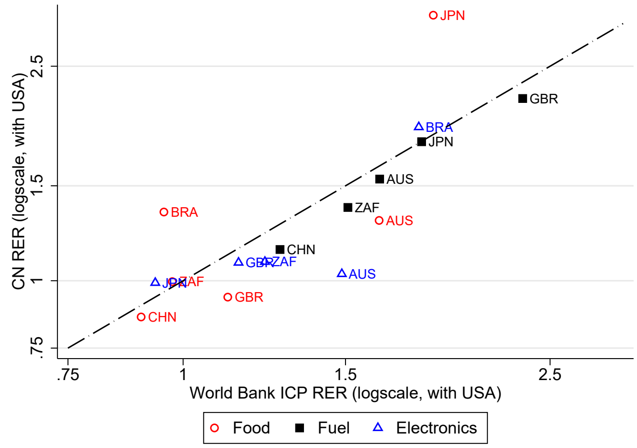

(a) 1-Digit

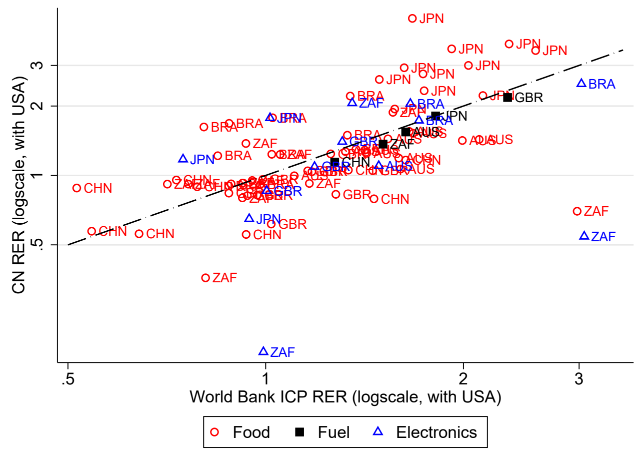

(b) 3-Digit

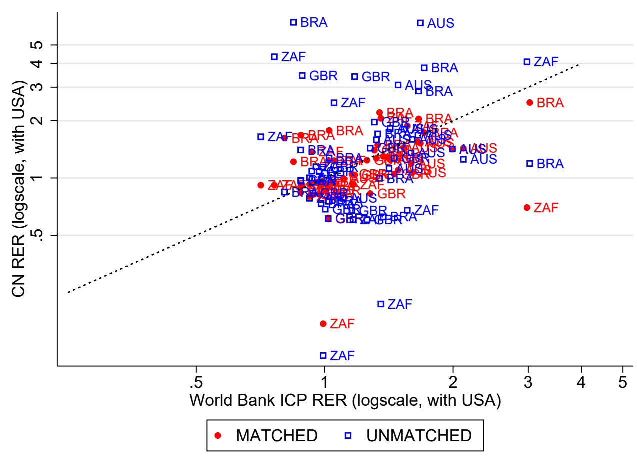

(c) 3-Digit - Matched vs Unmatched

Figure 1: Real Exchange Rate Levels Relative to the United States, ICP vs. CN in 2011

Notes: 2011 is the only year with ICP data in our 2010-2016 sample. ICP prices are collected by NSOs at some unpublished date during that year. We calculate the ICP real exchange rate using the average nominal exchange rate for that year. The CN real exchange rate is the average daily real exchange rate for the year 2011.

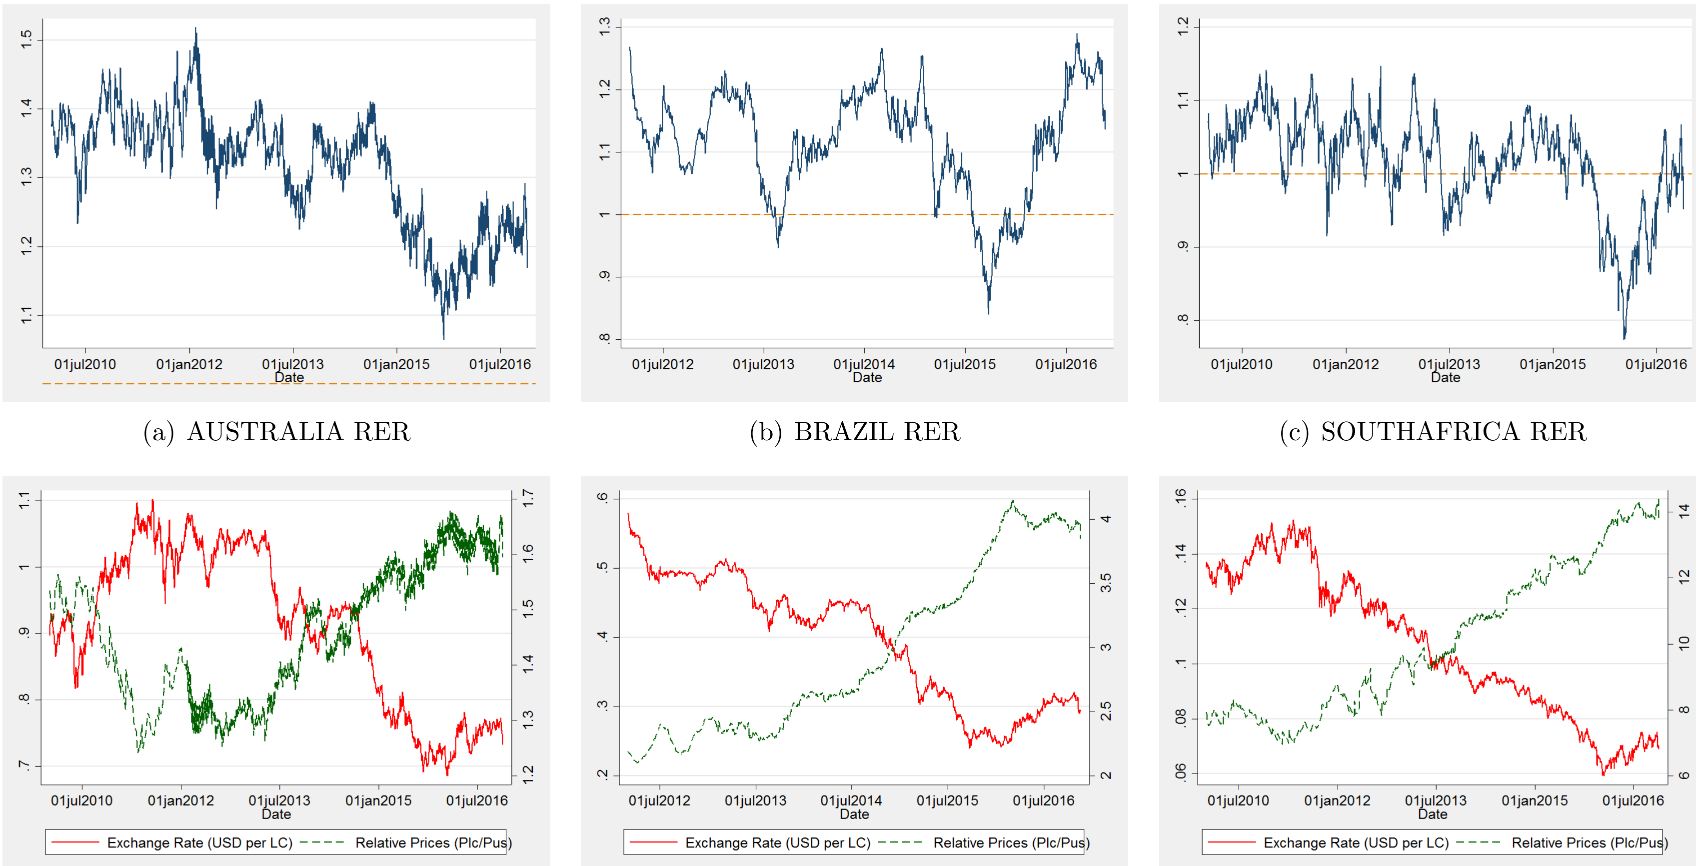

(a) AUSTRALIA RER

(b) BRAZIL RER

(c) SOUTHAFRICA RER

(d) AUSTRALIA RP & E

(e) BRAZIL RP & E

(f) SOUTHAFRICA RP & E

Figure 2: Real Exchange Rates, Relative Prices, and Nominal Exchange Rates - All Sectors

Notes: The solid (blue) line is the bilateral real exchange rate relative to the US in all sectors. It is computed as the relative price (Plc/Pus) multiplied by the nominal exchange rate (USD per local currency). The dashed (orange) line is drawn at the level where the RER is equal to one (the value predicted by absolute PPP). The solid (red) line is the nominal exchange rate expressed as local currency per unit of US dollars (an increase means the local currency depreciates). The dashed (green) line is the relative price expressed as the price in local currency over the price in the US. Relative prices are first calculated at the level of the product, and then aggregated with a geometric weighted average and a Fisher price index that uses the official CPI expenditure weights in both countries.

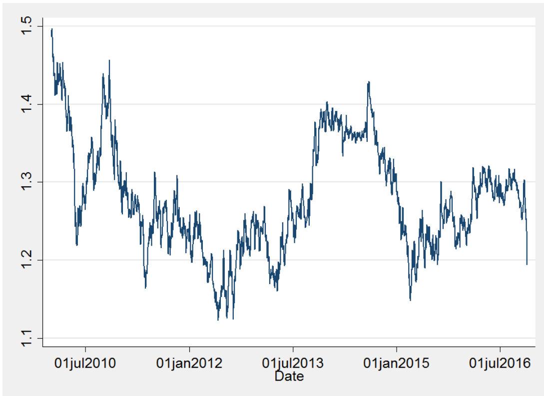

(a) GERMANY RER

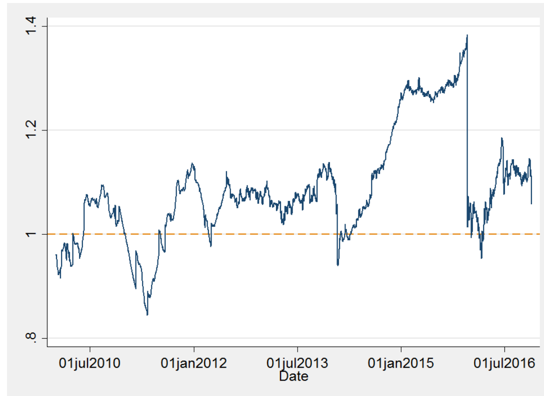

(b) ARGENTINA RER

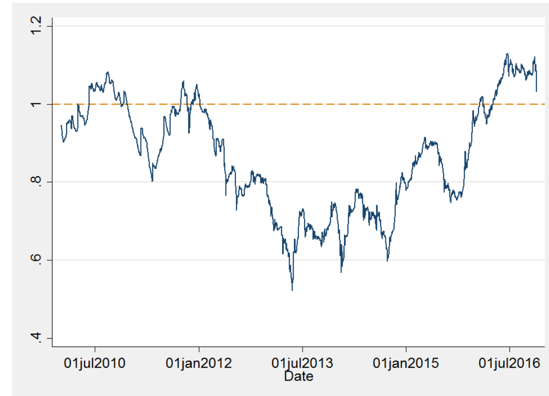

(c) ARGENTINA BLACK MARKET RER

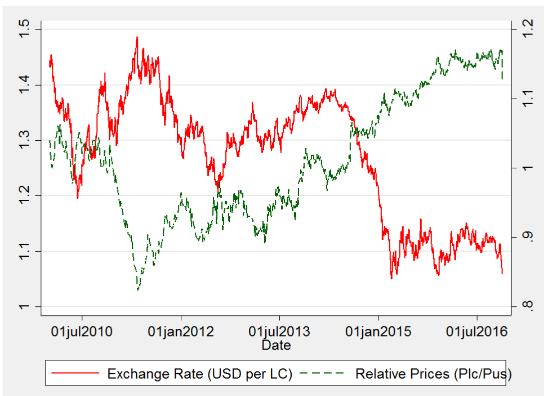

(d) GERMANY RP & E

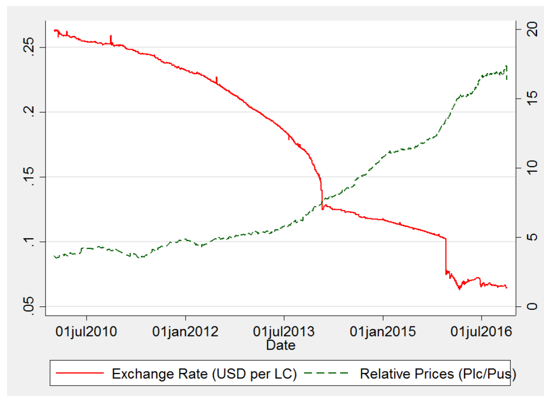

(e) ARGENTINA RP & E

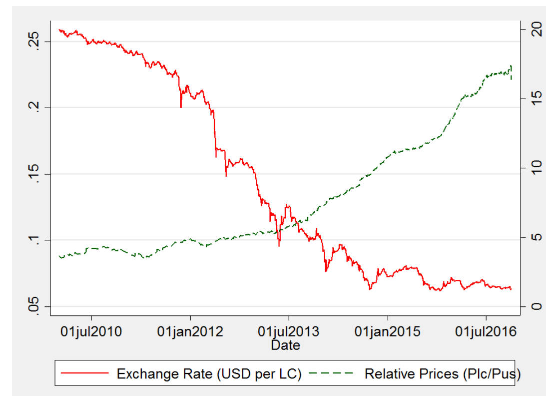

(f) ARGENTINA BLACK MARKET RP & E

Figure 3: Real Exchange Rates, Relative Prices, and Nominal Exchange Rates - All Sectors

Notes: The solid (blue) line is the bilateral real exchange rate relative to the US in all sectors. It is computed as the relative price (Plc/Pus) multiplied by the nominal exchange rate (USD per local currency). The dashed (orange) line is drawn at the level where the RER is equal to one (the value predicted by absolute PPP). The solid (red) line is the nominal exchange rate expressed as local currency per unit of US dollars (an increase means the local currency depreciates). The dashed (green) line is the relative price expressed as the price in local currency over the price in the US. Relative prices are first calculated at the level of the product, and then aggregated with a geometric weighted average and a Fisher price index that uses the official CPI expenditure weights in both countries.

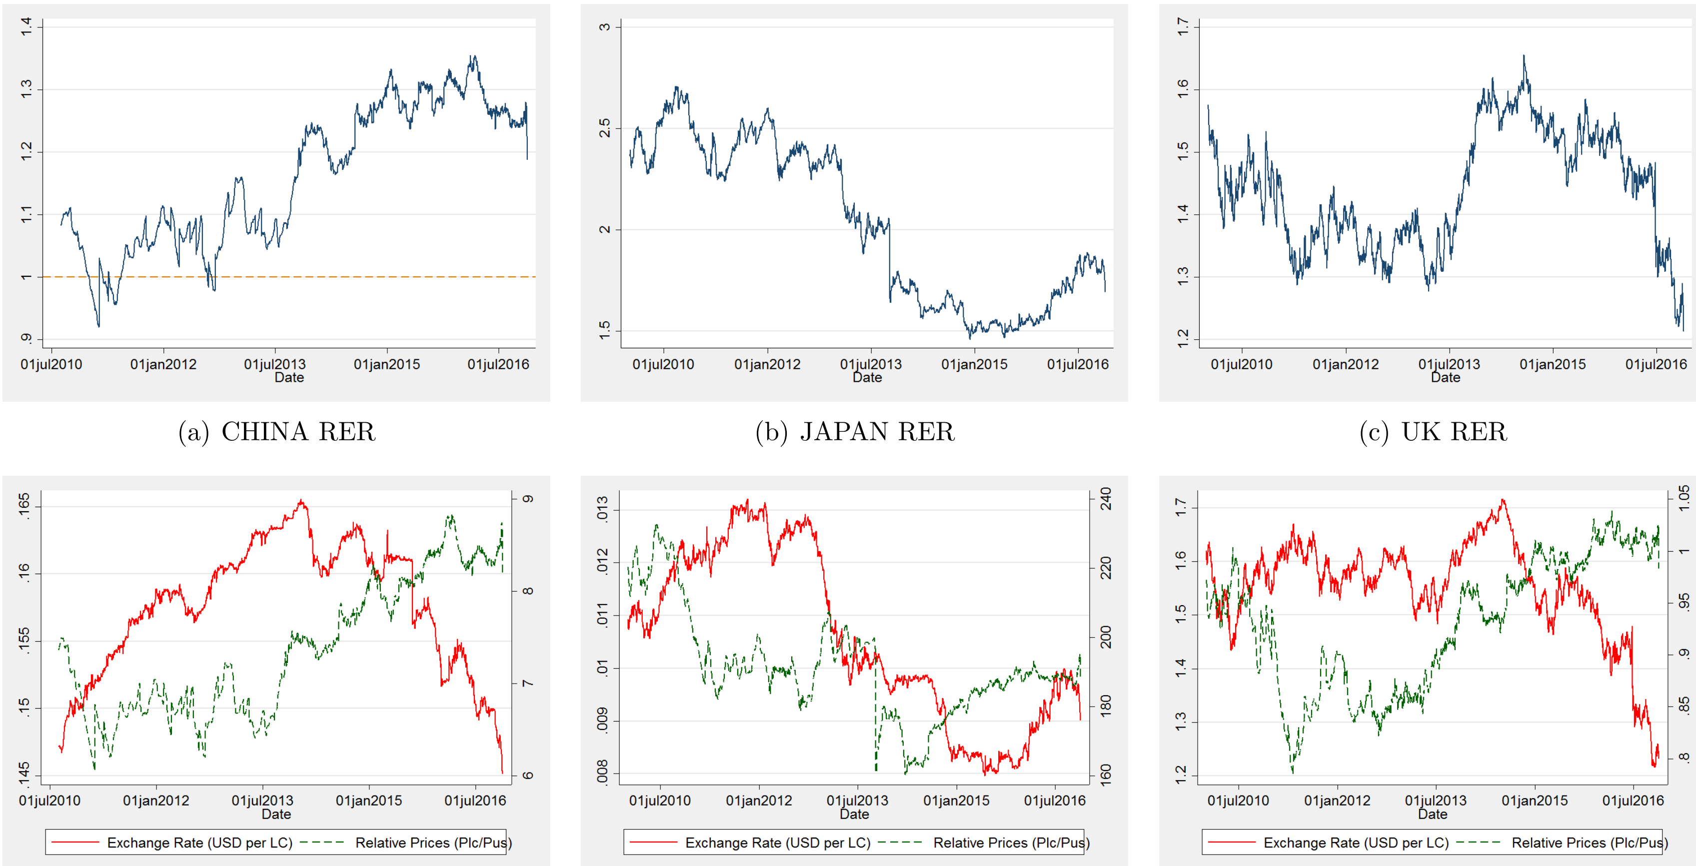

(a) CHINA RER (b) JAPAN RER (c) UK RER

(d) CHINA RP & E (e) JAPAN RP & E (f) UK RP & E

Figure 4: Real Exchange Rates, Relative Prices, and Nominal Exchange Rates - All Sectors

Notes: The solid (blue) line is the bilateral real exchange rate relative to the US in all sectors. It is computed as the relative price (Plc/Pus) multiplied by the nominal exchange rate (USD per local currency). The dashed (orange) line is drawn at the level where the RER is equal to one (the value predicted by absolute PPP). The solid (red) line is the nominal exchange rate expressed as local currency per unit of US dollars (an increase means the local currency depreciates). The dashed (green) line is the relative price expressed as the price in local currency over the price in the US. Relative prices are first calculated at the level of the product, and then aggregated with a geometric weighted average and a Fisher price index that uses the official CPI expenditure weights in both countries.

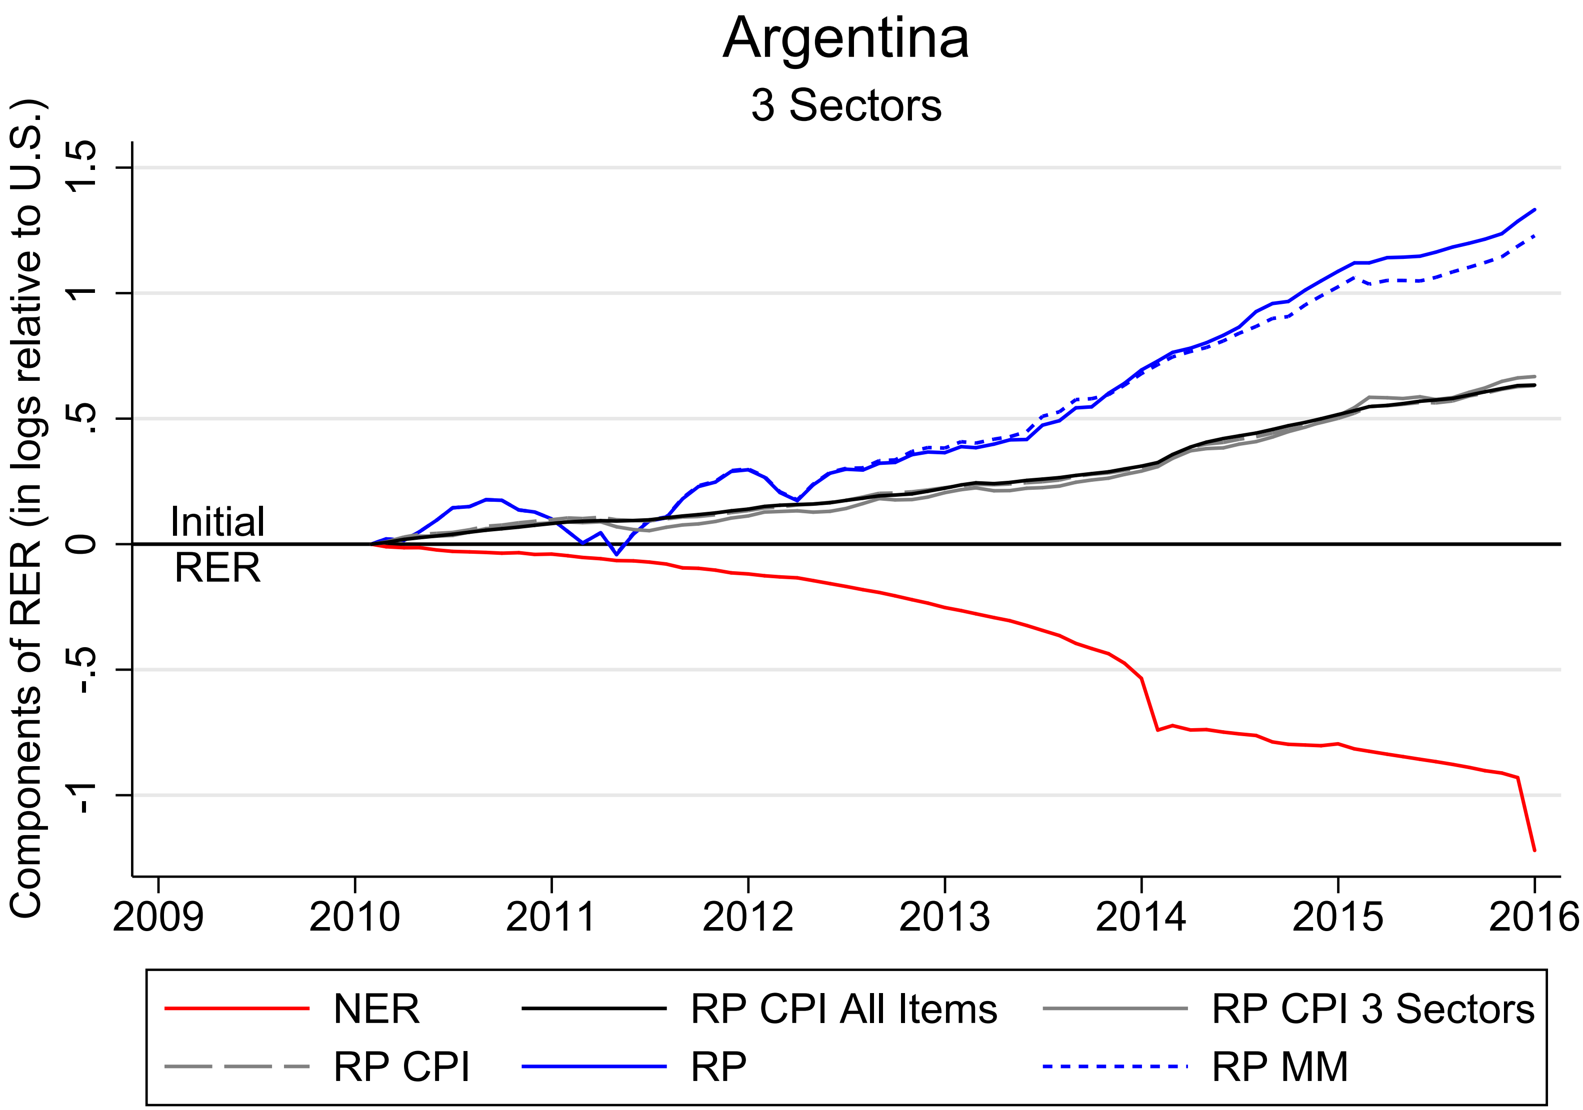

Figure 5: Components of Real Exchange Rate Relative to the United States

Notes: TBD

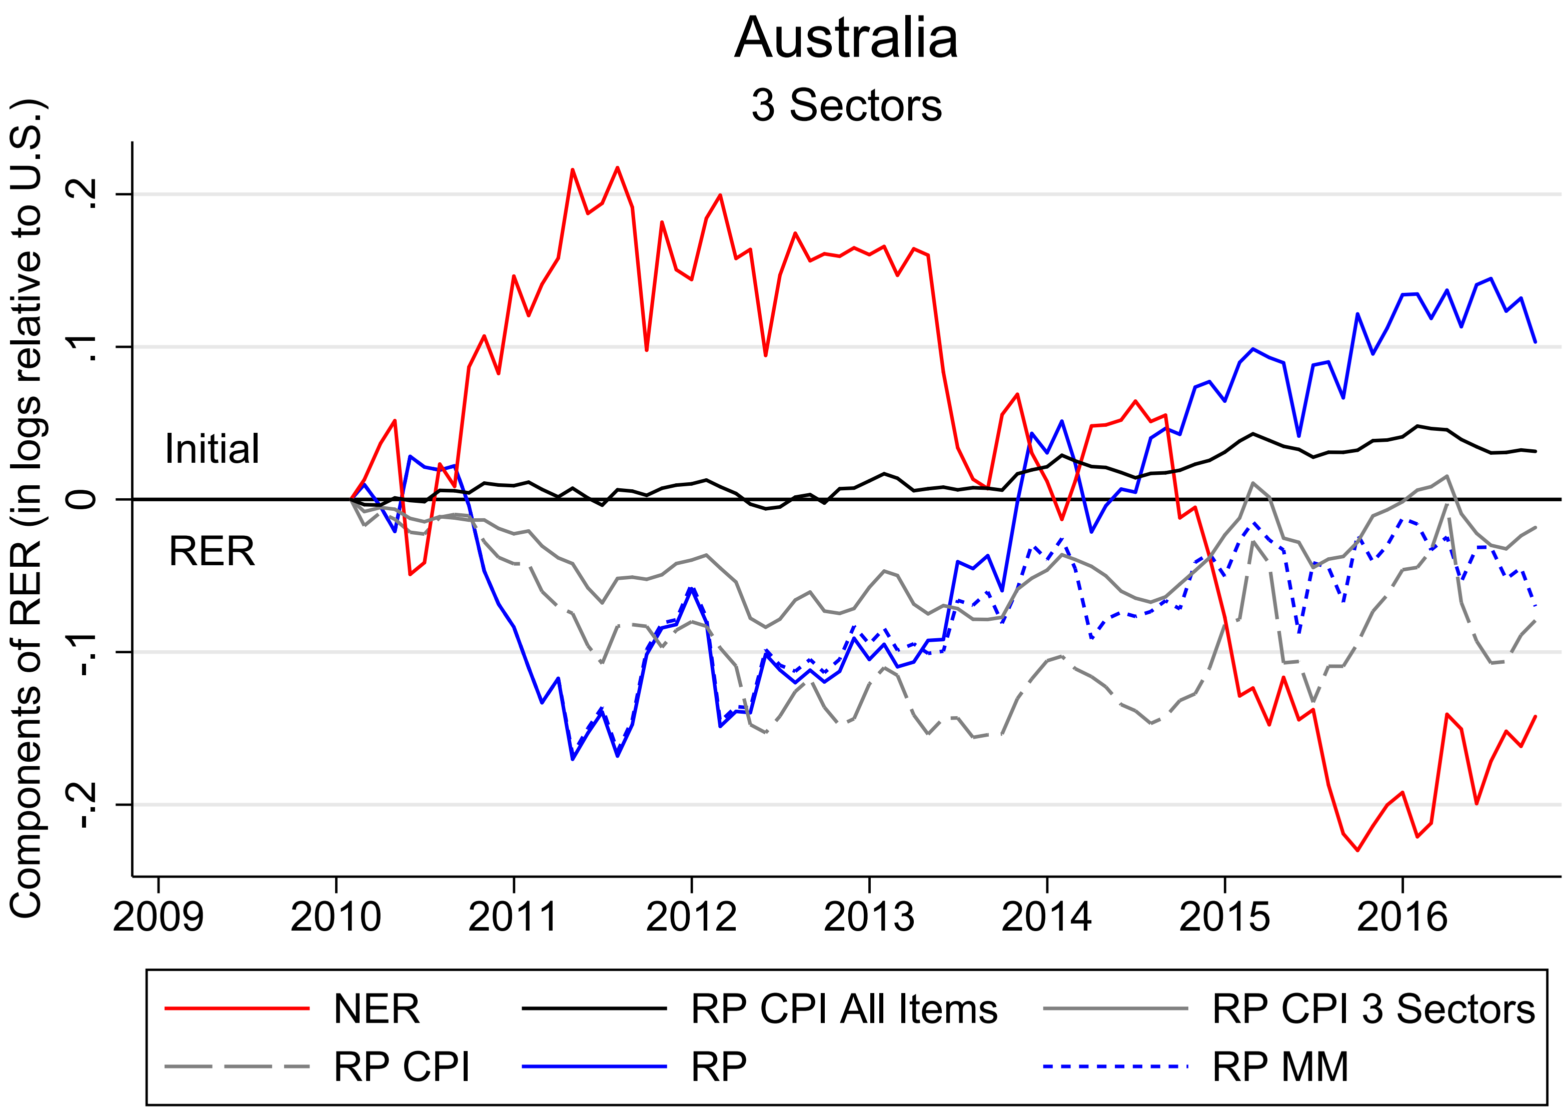

Figure 6: Components of Real Exchange Rate Relative to the United States

Notes: TBD

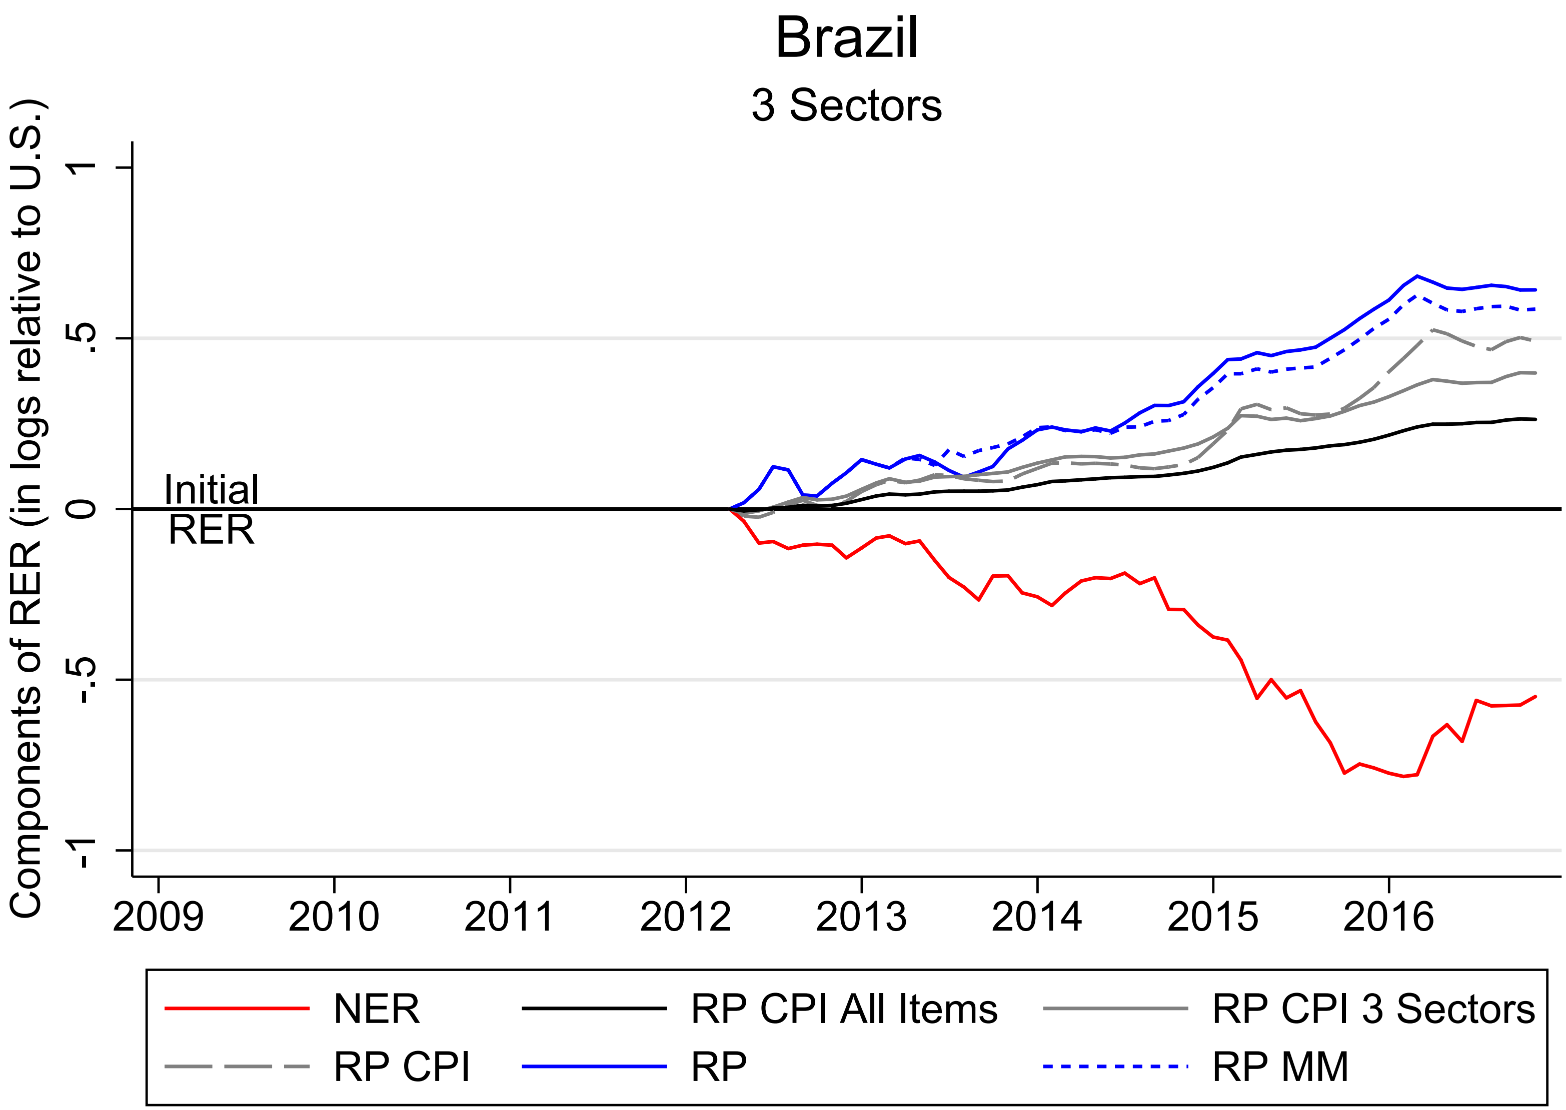

Figure 7: Components of Real Exchange Rate Relative to the United States

Notes: TBD

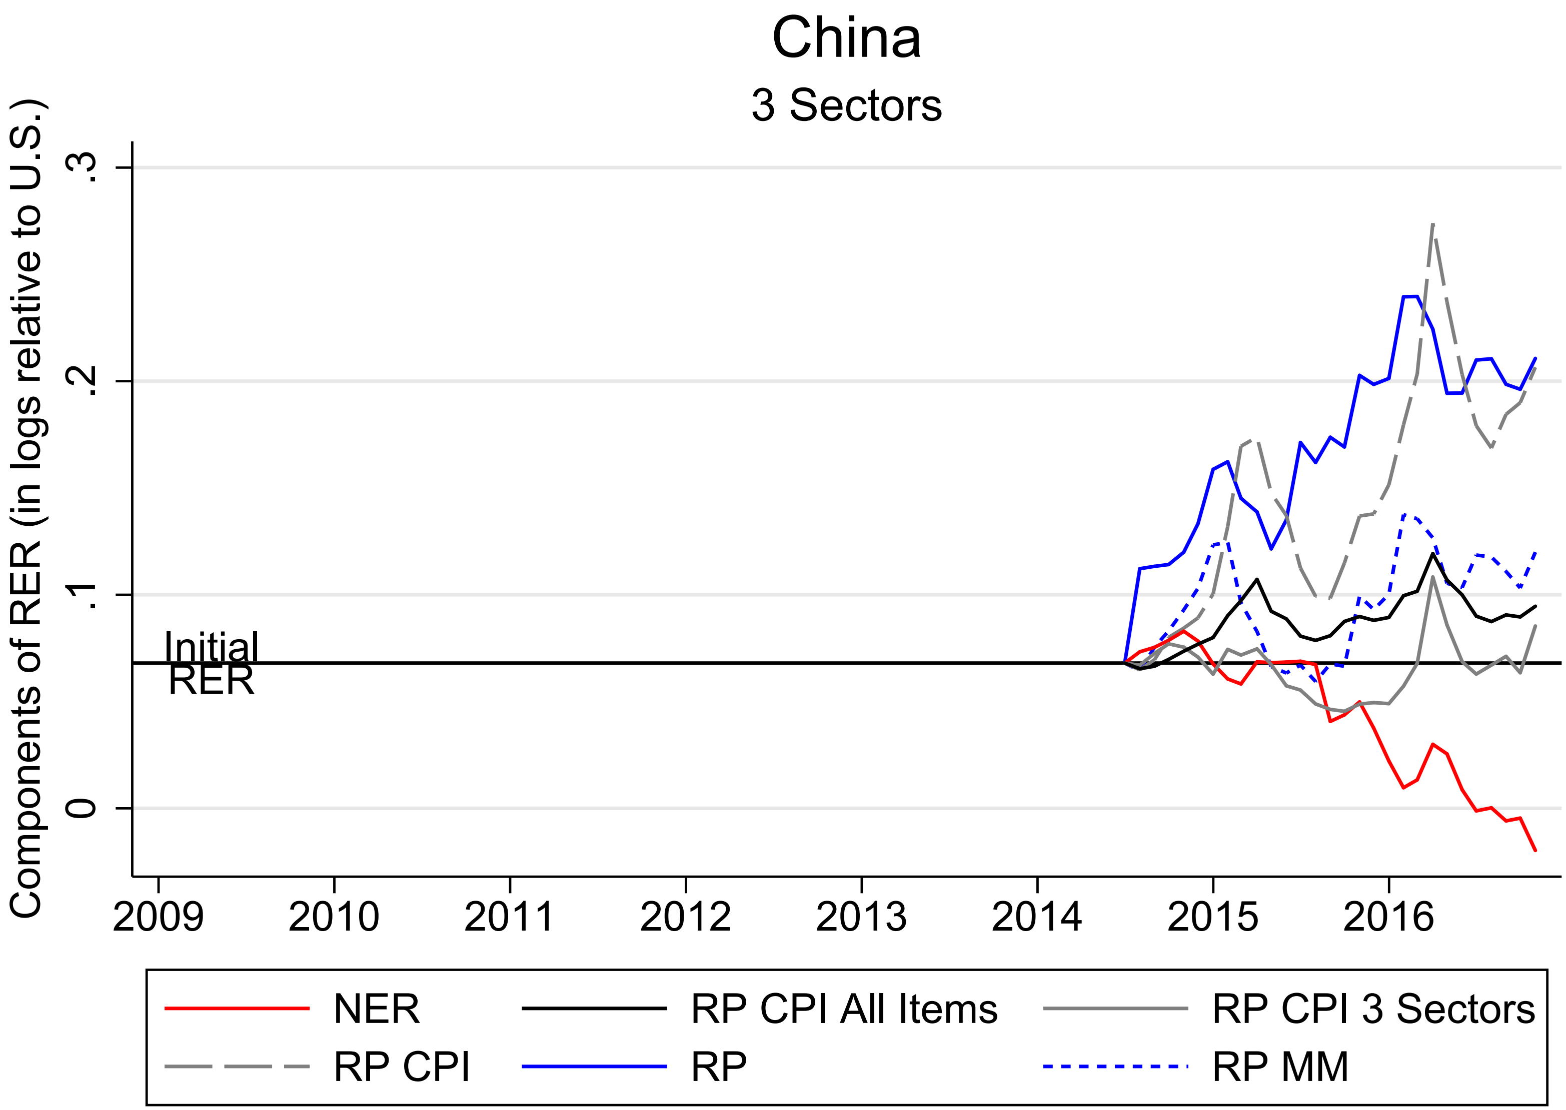

Figure 8: Components of Real Exchange Rate Relative to the United States

Notes: TBD

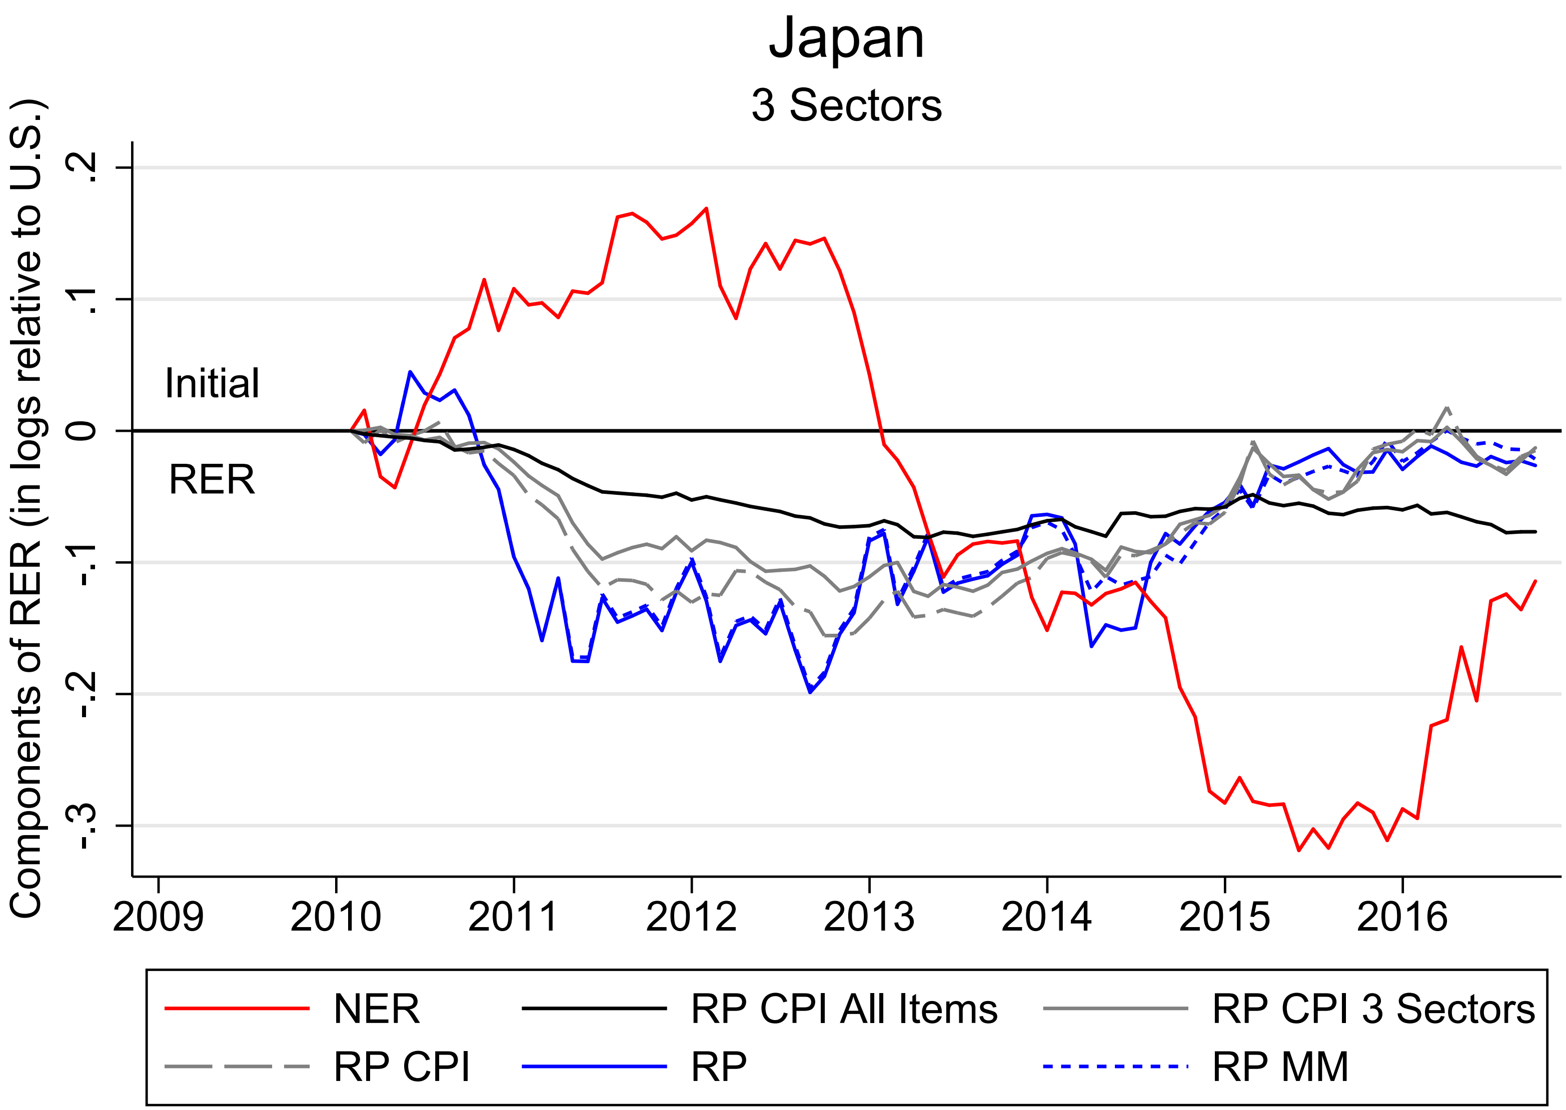

Figure 9: Components of Real Exchange Rate Relative to the United States

Notes: TBD

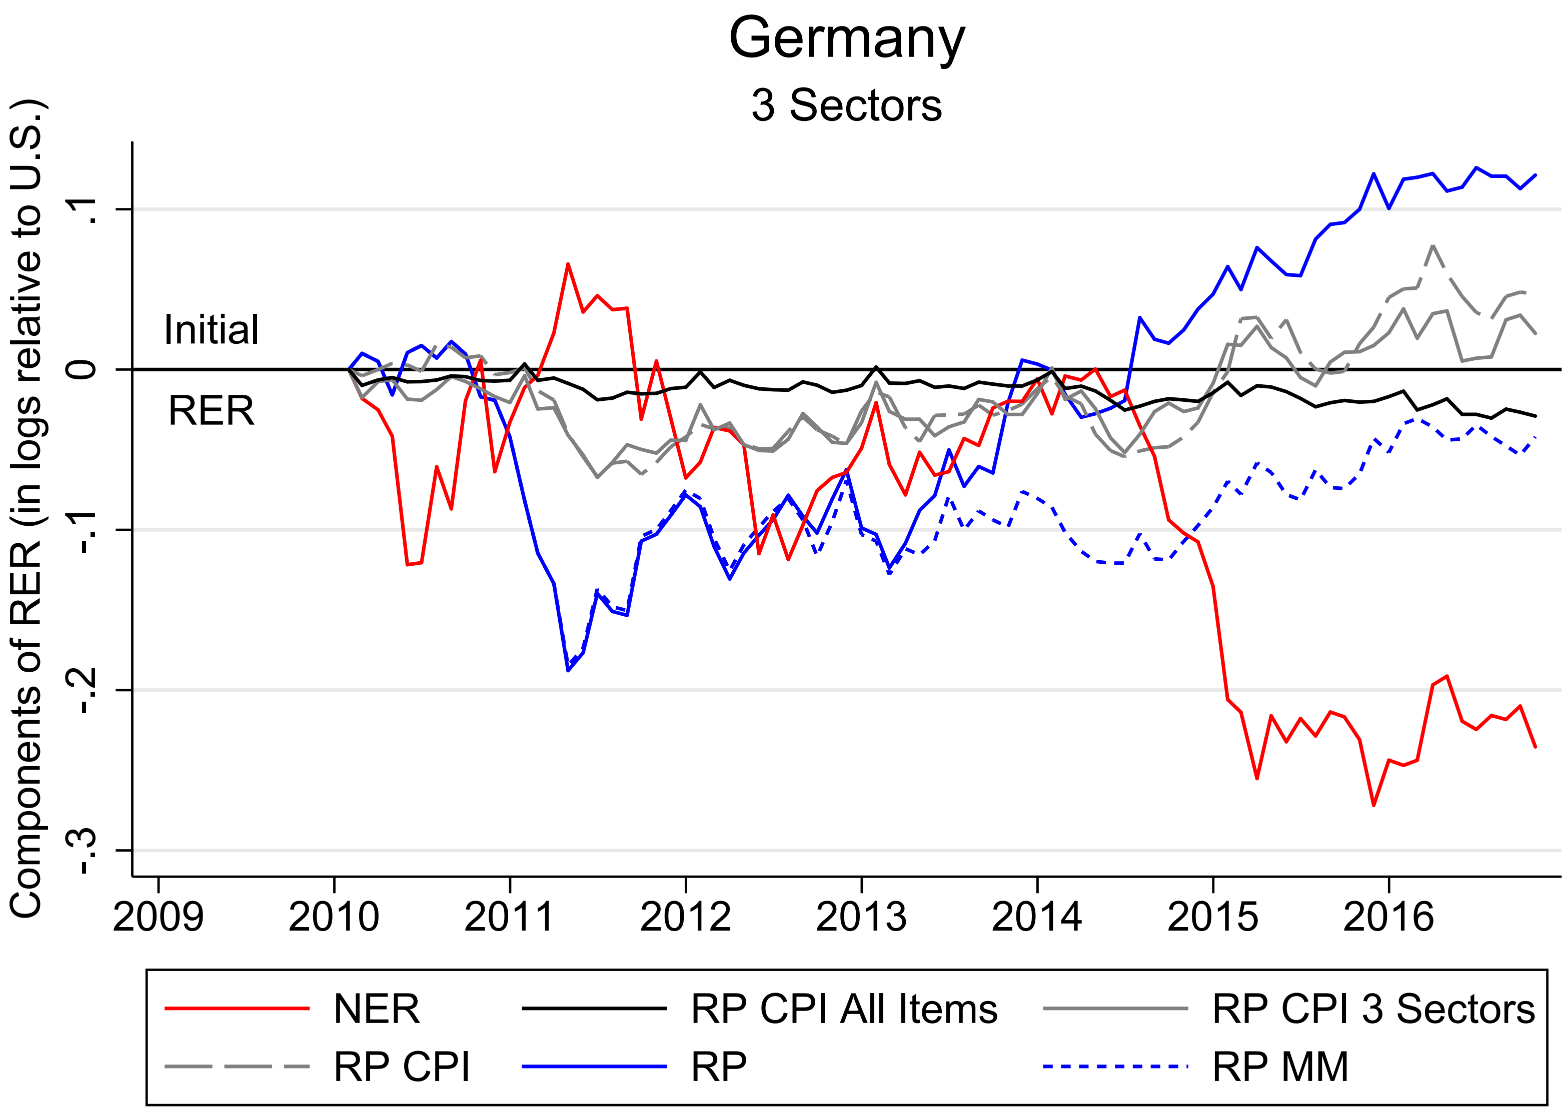

Figure 10: Components of Real Exchange Rate Relative to the United States

Notes: TBD

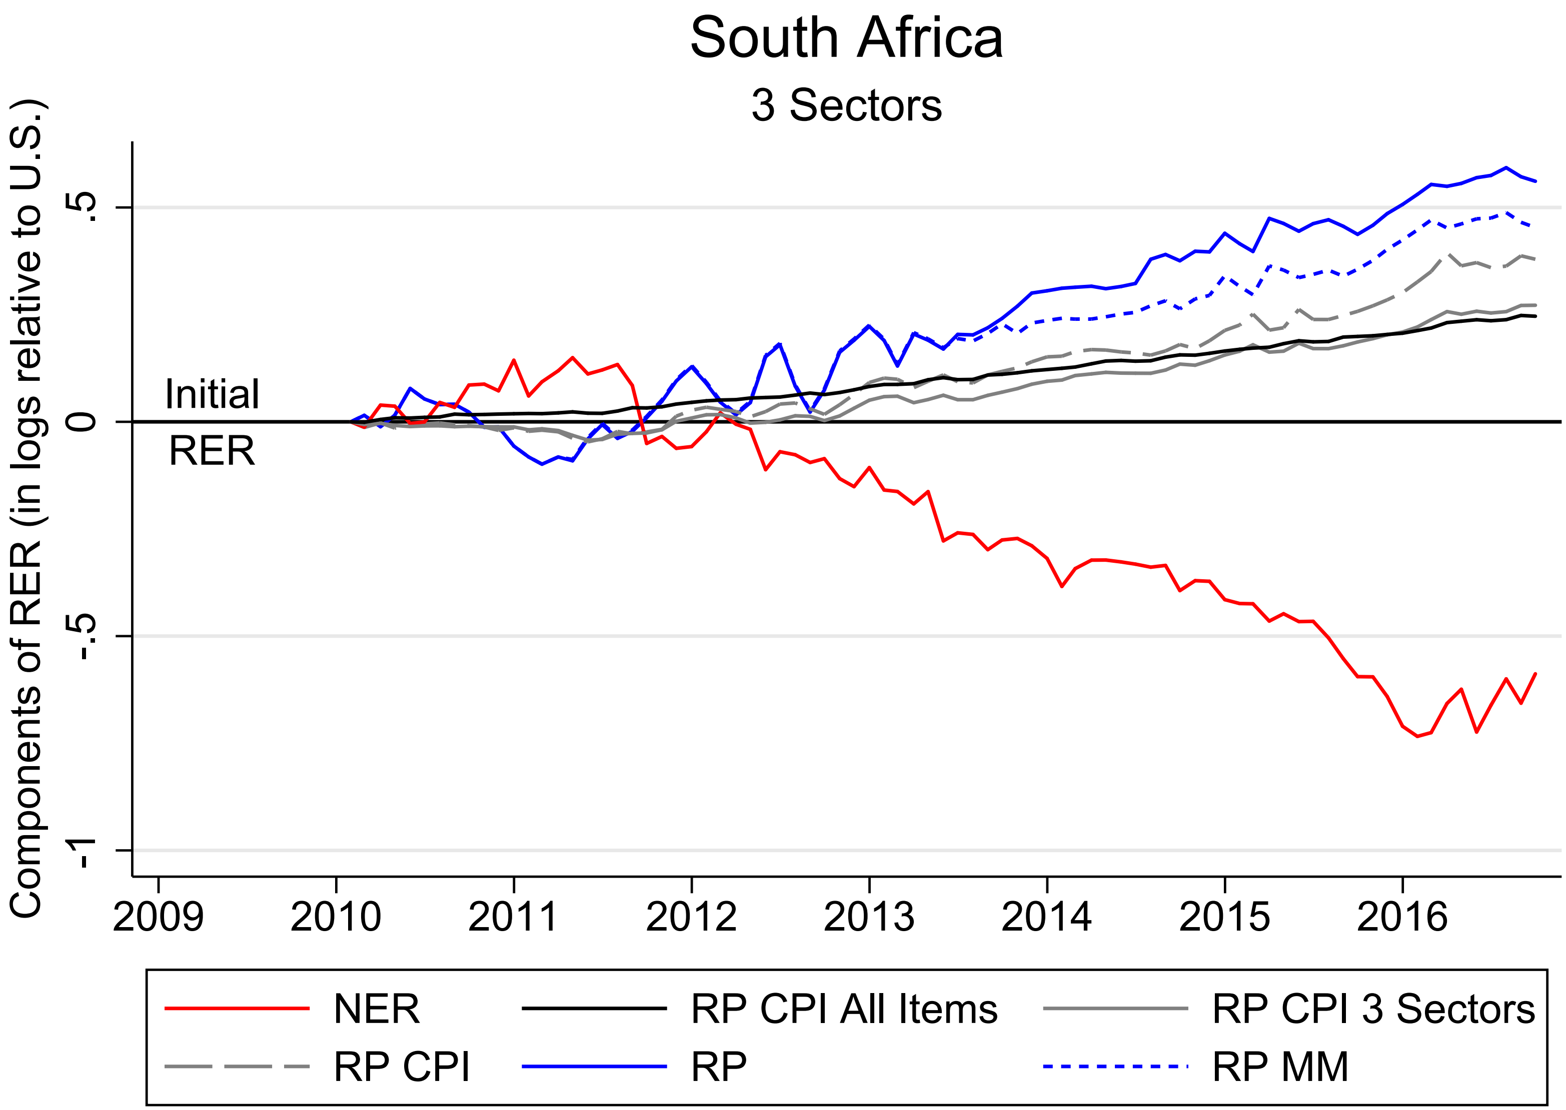

Figure 11: Components of Real Exchange Rate Relative to the United States

Notes: TBD

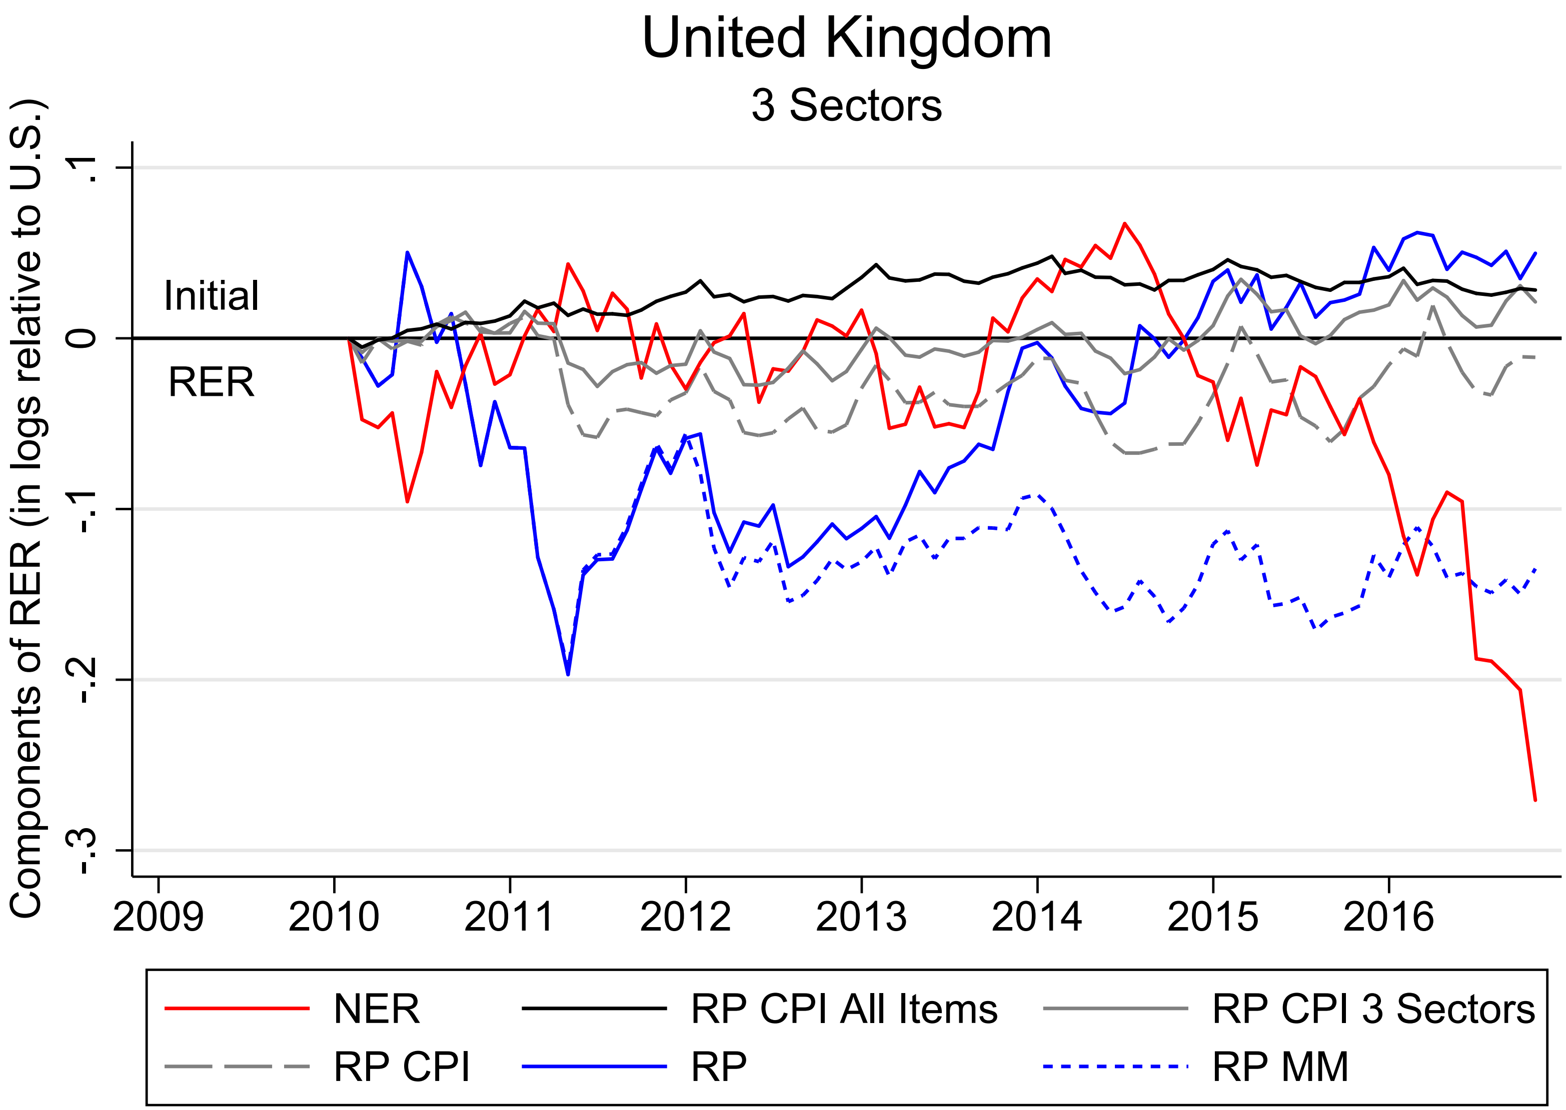

Figure 12: Components of Real Exchange Rate Relative to the United States

Notes: TBD