

Lost in Transit:
Product Replacement Bias and Pricing to Market

Emi Nakamura and Jón Steinsson*

Columbia University

Preliminary and Incomplete

December 1, 2008

Abstract

Product replacement is frequent in the micro-data that underlie U.S. import and export price indices. Also, prices change infrequently in these data. In constructing price indices, price adjustments that occur at the time of product replacements tend to be dropped. If price adjustments disproportionately occur at the time of product replacements then price adjustments are disproportionately unobserved. We show that this “product replacement bias” has distorted the measured long-run relationship between import and export prices and the exchange rate by a factor of between 1.7 and 2.1. Accounting for this bias, we find that the price of non-oil U.S. imports (relative to domestic consumption) responds by roughly 0.7% for each 1% change in the U.S. real exchange rate, while the price of U.S. exports (relative to foreign consumption) responds by roughly 0.8%. This contrasts with conventional pass-through estimates of 0.2-0.4% for non-oil import prices and 0.9% for export prices. Thus, the degree of pricing to market for U.S. imports and exports is more symmetric and the degree of pricing to market for U.S. imports more moderate than conventional measures suggest. Adjusting for product replacement bias also substantially raises the volatility of the terms of trade. These results improve the fit of the data to standard models.

Keywords: Product Replacement Bias, Pricing to Market, New Goods.

JEL Classification: F31, F41, E30, E01, C81

*We thank Ryan Chahrour and Vickram Mohan for excellent research assistance. We thank Ariel Pakes and John Rogers for inspiring conversations that encouraged us to work on this problem. Our thinking about the ideas and empirical results presented in this paper has benefited greatly from conversations with among others Bill Alterman, Menzie Chinn, Charles Engel, Patrick Kehoe, Timothy Kehoe, Narayana Kocherlakota, Greg Kurtzon, Rob McClelland, Rozi Ulics and seminar participants at various institutions.

1 Introduction

The post-Bretton Woods period of flexible exchange rates has been characterized by large swings in the U.S. real exchange rate. Simple models suggest that these large swings in relative prices of U.S. products should lead to large amounts of expenditure switching between U.S. and foreign products. In practice, expenditure switching has been modest. This exchange rate disconnect is one of the major puzzles of open economy macroeconomics.

A leading potential explanation for this puzzle has been that exporters “price to market”. Exporters are said to “price to market” if they adjust their markups to stabilize the local currency prices of their products (Krugman, 1987). Pricing to market thus implies that import and export prices change less than 1% for every 1% change in the exchange rate. Krugman (1993) and Engel (2002), among others, emphasize that the extent of pricing to market has profound implications for the degree of expenditure switching associated with exchange rate changes. If firms stabilize their prices in local currency terms, expenditure switching associated with exchange rate changes can play little role in resolving external imbalances.

Conventional estimates suggest that the prices of non-oil imports into the U.S. relative to all domestic prices change by only 0.2-0.4% for each 1% change in the U.S. real exchange rate.¹ On the other hand, conventional measures suggest that the prices of U.S. exports relative to all foreign prices change by roughly 0.9% for each 1% change in the U.S. real exchange rate. In other words, these estimates suggest that pricing to market by foreign importers into the U.S. market is substantial, while pricing to market by U.S. exporters in foreign markets is much less substantial.

Simple models such as Backus, Kehoe and Kydland (1992), Stockman and Tesar (1995) and Obstfeld and Rogoff (1995) do not generate pricing to market. The large amount of measured pricing to market in U.S. imports has motivated a great deal of work on models designed to generate this kind of behavior by firms. One strand of the literature has focused on price rigidity as a source of pricing to market (Betts and Devereux, 1996, 2000; Chari, Kehoe and McGrattan, 2002). However, the apparent magnitude and persistence of pricing to market has shifted the focus of the

¹Campa and Goldberg (2005) report an estimate of roughly 0.4 for long run exchange rate pass-through into U.S. import prices for the period 1975 to 2003. Marazzi and Sheets (2007) report that long-run pass-through excluding oil was around 0.5 for the period from the early 1970’s to the late 1990’s, but has fallen to around 0.2 over the period from 1997 to 2004. Similar estimates are obtained by Gopinath, Itskhoki, and Rigobon (2008). These papers run pass-through regressions in nominal terms. Knetter (1989) and Marston (1990) run regressions in real terms that are closer to the specifications we adopt in this paper. Burstein, Eichenbaum, and Rebelo (2005) report much higher pass-through for developing countries during large devaluations.

literature towards mechanisms that can generate long run deviations in relative prices (Dornbusch, 1987; Goldberg and Verboven, 2001; Corsetti and Dedola, 2005; Atkeson and Burstein, 2008; Gust, Leduc and Vigfusson, 2006; Gopinath, Itskhoki and Rigobon, 2007).² Matching the measured amount of long run pricing to market for U.S. imports along the lines suggested in these papers requires a great deal of strategic complementarity in price setting. This leads to important changes in the dynamic implications of the model, often generating longer delays in the responsiveness of prices to exchange rate fluctuations than is consistent with empirical estimates. Standard models of pricing to market also do not explain why pricing to market appears to be so much greater for U.S. imports than for U.S. exports.

We argue that conventional measures of pricing to market are seriously biased due to measurement issues associated with the frequent replacement of products in the modern economy. The U.S. import and export price indices are based on micro price-data collected from U.S. firms. In these data, roughly 4.5% of products are replaced each month. In addition, reported prices change infrequently in these data. The frequency of price change of the median product is only about 8% per month (Gopinath and Rigobon, 2008). This implies that roughly 45% of price series in these data have no price changes and more than 70% have 2 price changes or less.³

A natural question is: What happens to prices when products are replaced? This is a difficult problem since the quality, size and specifications of the new product might not be the same as those of the old product. The ideal solution to this problem is to use hedonic methods to estimate the quality change associated with product replacements (Court, 1939; Griliches, 1961; Pakes, 2003). For most products, however, it is extremely costly and difficult to accurately measure quality change (Abraham et al., 1998). Hedonic adjustments are, therefore, used in only a tiny fraction of cases.⁴ In practice, a large fraction of product replacements are “linked-into” the index; meaning that the price comparison between the first observation of the new product and the last observation of the old product is dropped when changes in the index are calculated.⁵ Indices constructed by linking

²Other notable models of pricing to market include Alessandria (2004), Ravn, Schmitt-Grohe and Uribe (2007), Drozd and Nosal (2008), Hellerstein (2008) and Nakamura and Zerom (2008).

³Gopinath, Itskhoki, and Rigobon (2008) discuss the possibility that these features of the data might affect measured exchange rate pass-through.

⁴Feenstra (1994) and Broda and Weinstein (2006) develop and apply an alternative to hedonic methods. These papers make strong enough functional form assumptions about demand that quality adjusted prices can be backed out from prices and quantities.

⁵In some cases, the IPP deems a change in a product to be sufficiently small that the concurrent change in price is used in the index with no adjustment for a change in quality.

in new products are referred to as “matched model indices” since all price comparisons on which such indices are based are for identical items.

The literature on the “new goods problem” in index number theory points out that if new products are systematically introduced at a lower quality adjusted price than the old products they replace, then a matched model price index will yield an upward biased estimate of inflation and a matched model quantity index will yield a downward biased estimate of output growth.⁶ Our focus is not on this bias in the average level of measured inflation caused by product replacement. We argue that the combination of price rigidity and product replacement implies that matched model price indices yield biased estimates of the comovement between prices and other variables such as exchange rates.⁷

The implicit assumption embodied in the practice of linking new products into the price index is that the frequency and size of price changes at the time of product replacements are the same on average as the frequency and size of price changes for continuing products. We argue that this is a poor assumption since firms must set new prices for newly introduced products, while continuing products have a low probability of price change. In other words, it is likely that price changes occur disproportionately at the time of product replacement. In this case, linking new products into the price index disproportionately drops price change observations and overweights observations with no price change. Since it is the price change observations that reveal the comovement between prices and exchange rates, a price index with this property will yield biased estimates of this comovement. We refer to this bias as the “product replacement bias.”

It is helpful to consider an extreme example in which the price of each product remains fixed for the entire life of the product and all price adjustment occurs at the time of product replacements. Figure 1 depicts this type of setting. In the figure, the exchange rate is rising. Assume for simplicity that the product replacements involve no quality change. Agents living in this economy can observe the quality of each product. It is therefore obvious to them that prices are rising as the exchange rate rises. A price index based only on price comparisons for identical items will however remain

⁶The Boskin commission argued that the new goods bias led to an upward bias in the consumer price index of 0.6% per year (Boskin et al., 1996). Similarly, Bils and Klenow (2001), Hausman (2003), Nordhaus (1998), Pakes (2003) and Bils (2005), have emphasized the importance of the new goods bias in distorting measured inflation and economic growth. Goldberg et al. (2008) show that new imported varieties contributed substantially to effective price declines for Indian firms after a trade liberalization. In contrast, Moulton and Moses (1997), Abraham et al. (1998), Triplett (1997) and Hobijn (2003) argue that the Boskin Commission overestimated the new goods bias.

⁷The same bias applies to matched model quantity indices. We will, however, focus on the effects on price indices.

constant throughout in this example since prices only change at the time of product replacements. Estimates of the comovement of prices and exchange rates using this price index will yield no comovement irrespective of what the true degree of comovement is.

If the prices in figure 1 are U.S. dollar prices of products imported into the U.S., the figure depicts a case in which these prices display rigidity in the local currency. They are thus said to be local currency priced (LCP) products. The vast majority of U.S. imports are LCP products. A small fraction, however, display rigidity in the currency of the producing country. These goods are said to be producer currency priced (PCP) products. For U.S. exports, this pattern is reversed. The vast majority of U.S. exports are PCP products.

The example depicted in figure 1 shows that product replacement bias causes a *downward* bias in the measured comovement between prices and the exchange rate for LCP products when measured in the local currency. For PCP products, however, product replacement bias causes an *upward* bias in the measured comovement between prices and the exchange rate when measured in the local currency. Consider a U.S. export into the Euro area that has a completely rigid price in U.S. dollars. The Euro price of this product will vary exactly one-for-one with the U.S. dollar-Euro exchange rate. Now consider a sequence of such PCP products that only change their producer currency price at the time of product replacements. A Euro area import price index based on a collection of such product lines would display one-for-one comovement even if the true long-run comovement (taking the price changes at the time product replacements into account) were zero.

The quantitative importance of product replacement bias depends on the relative frequency of price changes and product replacements in the micro data used to construct import and export price indices. We use BLS micro data on import and export prices over the period 1994-2007 to estimate the magnitude of product replacement bias. We carry out this estimation separately for LCP imports, PCP imports, LCP exports and PCP exports. In each of these four cases, we find that conventional estimates of comovement are biased by a factor of between 1.7 and 2.1 toward zero pass-through in the currency in which prices are rigid. Adjusting for product replacement bias, we find that the price of non-oil imports into the U.S. relative to the price of all domestic goods changes by roughly 0.7% for each 1% change in the U.S. real exchange rate. In other words, “exchange rate pass-through” for non-oil U.S. import prices is 0.7 and the extent of pricing to market is 0.3. We find that the price of U.S. exports relative to all foreign prices changes by roughly 0.8% for each 1%

change in the U.S. real exchange rate. In other words, exchange rate pass-through for U.S. export prices is 0.8 and the extent of pricing to market is about 0.2.

Taken together, these estimates bring the data better in line with standard models. First, adjusting for product replacement bias yields estimates of pricing to market that are much more symmetric between U.S. imports and exports. Second, our adjusted estimates imply a more modest amount of pricing to market for U.S. imports. The large amount of pricing to market implied by conventional measures is difficult to match theoretically. It necessitates large markups and a degree of strategic complementarity in price setting that strains existing general equilibrium models along other observable dimensions. For instance, the preferred calibration of Corsetti and Dedola (2005) implies a roughly 0.9% change in import and export prices for each 1% change exchange rates in the long run, while the corresponding number for the preferred calibration of Atkeson and Burstein (2008) is roughly 0.75%. Adjusting for product replacement bias also brings estimates of pricing to market for U.S. imports better in line with estimates based on large devaluation in less developed countries (Burstein, Eichenbaum and Rebelo, 2005).

Finally, our results imply that conventional measures of import and export prices understate the volatility of the U.S. terms of trade. Measured in U.S. dollar terms, conventional indices suggest that non-oil import prices rise by 0.2-0.4% and export prices fall by 0.1% when the U.S. exchange rate depreciates by 1%. Taken together this implies a 0.1-0.3% improvement of the non-oil terms of trade. Adjusting for product replacement bias, we find that non-oil import prices rise by 0.7% and export prices fall by 0.2% when the U.S. exchange rate depreciates by 1%. These estimates imply a 0.5% improvement in the non-oil terms of trade. If movements in the terms of trade are dominated by movements in the exchange rate, our results suggest that conventional measures of the volatility of the terms of trade may be downward biased by a factor of between 1.7 and 5. As Drozd and Nosal (2008) emphasize, standard models without pricing to market imply that the terms of trade should be more volatile than the real exchange rate. Conventional measures suggest that the terms of trade is much less volatile than the real exchange rate. Our results suggest that these conventional measures greatly overstate this discrepancy between standard models and the data.

The closest antecedent to the ideas presented in this paper is the literature on downward bias in measured inflation due to linking of new products into price indices (Armknrecht and Weyback,

1989; Liegey, 1993; Reinsdorf, Liegey and Stewart, 1996; Triplett, 1997). Our work builds on a number of earlier papers that use the BLS microdata on import and export prices. Clausing (2001) and Gopinath and Rigobon (2008) document basic facts about the frequency of price changes and substitutions in these data. Gopinath, Itskhoki, and Rigobon (2008) argue that there is a substantial difference in conventionally measured long-run exchange rate pass-through for LCP versus PCP imports. Gopinath and Itskhoki (2008) argue that the degree of exchange rate pass-through is related to product's frequency of price change. Neiman (2008) documents a number of related facts regarding price rigidity and exchange rate pass-through for intrafirm prices. Fitzgerald and Haller (2008) study pricing-to-market using analogous data from Ireland. Our work is motivated in part by Erickson and Pakes (2008), who develop an experimental hedonic price index for televisions that accounts, among other things, for price rigidity. Finally, our findings regarding the importance of product replacement in generating low measured exchange rate pass-through are consistent with ongoing research by Jon Faust and John Rogers that argues that there is a negative relationship across industries between the frequency of product substitutions and measures of exchange rate pass-through (Rogers, 2006).

The paper proceeds as follows. Section 2 describes the BLS micro data underlying the U.S. import and export price indices that we use in our empirical analysis. Section 3 presents measures of pricing to market for U.S. imports and exports for the period 1982-2007 based on conventional methods. Section 4 derives expressions for product replacement bias as a function of the frequency of price change and the frequency of product replacement. Section 5 presents estimates of the frequency of price change and the frequency of product replacement and our quantitative estimates of product replacement bias. Section 6 discusses alternative measures of pricing to market and presents an example that allows for timing error. Section 7 discusses several additional implications of product replacement bias. Section 8 concludes.

2 Data Description

We use three sets of data. First, we use the microdata underlying the U.S. import and export price indices. These data are collected by the International Prices Program (IPP) of the Bureau of Labor Statistics (BLS). Second, we use aggregate U.S. import and export price indices produced by the Bureau of Economic Analysis (BEA) as a part of the National Income and Product Accounts

(NIPA). Third, we use exchange rate data from the the Federal Reserve Board and the International Monetary Fund (IMF). We describe these data in turn.

The U.S. import and export price indices were introduced in the early 1980’s to provide a more accurate alternative to unit value indices. The micro data we use cover the time period 1994-2007. The IPP is charged with collecting both interfirm and intrafirm prices for international trade. Because of difficulties in interpreting intrafirm transactions, we exclude intrafirm prices in our baseline analysis.⁸ For our sample of countries, excluding intrafirm prices, the total number of product-months for which IPP attempts to record a price is roughly 1.5 million or about 100,000 per year. However, price imputation is a major phenomenon in the import and export price data. According to our estimates, of the total number of product-months in the database, reported prices are not available in about 40% of cases.

The IPP data are collected using voluntary surveys filled out by a designated “reporter” at each firm. To initiate a product into the dataset, IPP collects a detailed item description and a particular set of transaction terms. Item descriptions include the physical characteristics and specifications of an item, while transaction terms include the number or type of units priced, the country of destination or origin, the port of exit or entry, the discount structure, and in some cases the duty applied to the product. The price provided by the reporter during the initialization must be a transaction price—rather than an “estimated” price—unlike in subsequent periods when the price may be estimated or imputed, as we discuss below. At initiation, reporters may also sometimes provide future prices if these are considered to be known, e.g., because of an existing contract.⁹

After the product is initiated, price information is collected using a repricing form. The repricing forms include pre-filled information collected during the initialization process, such as the characteristics of the product, the terms of the transaction, and the discount structure. The repricing form first asks whether the price has changed relative to the previous month and then asks the respondent to report a new price if the price did change. One concern about this procedure is that it may introduce a bias toward firms reporting that there has been no change in prices. To evaluate sensitivity of the price data to the method used to collect prices, Gopinath and Rigobon (2008) compared the behavior of prices during the anthrax scare of 2001, when the IPP data were

⁸See Neiman (2008) and Clausing (2001) for a discussion of the intrafirm trade data.

⁹In a small number of cases, IPP collects price data from secondary sources. Some products for which data are collected in this way are crude petroleum, grain, traded services, and various automobile and automobile part prices. In general, these data are not included in the database we study.

collected by phone survey rather than mail, to the behavior of prices during other time periods. They find no significant differences in the frequency of price change over the period when prices were collected by phone. It is important to note, however, that even in the phone survey, reporters were typically provided with the price they had previously reported—perhaps inducing a similar bias to the mail survey.

Until 2001, repricing forms were returned to the BLS by mail. They were generally received starting in the second week of the month. The IPP has collected repricing forms by FAX and electronic methods since that time, leading to a faster turnaround time—about half of the repricing forms are now returned by the end of the first week of the month. The IPP receives some pricing surveys with a significant delay. If pricing information is submitted after the indices for a particular month are finalized at the end of that month, this information is used to correct previously imputed prices in the revised version of the indices.

The reporter is asked to provide actual transaction prices whenever possible. In cases where the reporter can only provide estimated or list prices because there was no transaction, or because the transaction price is not available, the IPP accepts estimated or list prices. An estimated price is the reporter’s estimate of “the price that would have been charged during the reporting process.” List prices are only accepted if the reporter is also able to provide a complete discount structure that allows actual net transaction prices to be calculated. In principle, the reported discount schedule may be updated by the reporter as necessary. In practice, however, the discount structure is rarely updated (see Section 4.1.1 of the IPP Data Collection Manual).

The IPP indices are constructed monthly. However, since many goods are imported or exported in large shipments rather infrequently, many price quotes are collected “off-cycle” at a less than monthly frequency. This is particularly the case for seasonal items. The frequency with which data are collected may also be reduced on an ad hoc basis to limit the firm’s reporting burden. If prices are collected off-cycle, the months in which price data is collected are selected as those months when the product is normally traded, when price changes are expected to occur, in which price contracts are expected to be renegotiated, or when quality changes are likely to occur. In the intervening months, the IPP computer system either generates an imputed price or pulls forward the last month’s price. The price imputations are calculated by taking the average price increase within a product category and applying this change to the previous month’s price. When the

price is again observed, the price reverts back to actual observed price. This implies that long-run inflation for prices imputed in this way equals long-run inflation for a price series in which prices are simply pulled forward through missing spells (Feenstra and Diewert, 2000).

To make it easy for firms to report data, the IPP accepts reported prices in any currency. In practice, about 93% of import prices and 98% of export prices are reported in U.S. dollars. In situations where the item is covered by a written contract, the IPP also records whether the contract is contingent on the exchange rate, and incorporates this information when imputing prices. The IPP accepts prices according to any pricing basis. The IPP comments that while FOB prices for imports and CIF prices for exports may be viewed as preferable from the perspective of analyzing foreign trade, “IPP must be practical given the many other price bases companies actually use for international trade.” The IPP therefore requests that reporters at firms provide prices according to the price basis on which it is “most likely that the item can be repriced over time.”¹⁰

In principle it might be desirable to define a product sufficiently narrowly that products with any observable differences in characteristics are viewed as entirely different products. For import and export prices, this approach would lead to difficulties in repricing some types of products since the products imported by a firm are found to differ slightly from one shipment to the next. To deal with this problem, the IPP takes a more pragmatic approach. In cases where there has been a substantial change in quality, the IPP discontinues the former item and initiates a new item with a new description. The product is also replaced if it is not possible to quantify the magnitude of the quality change, or if it is discontinued. To avoid sacrificing repricability, however, the IPP may deem some product characteristics—such as color, fabric content, or country of origin—to be non-price determining. For machinery that is hard to track over time, the IPP may decide to price either a base machine or a machine with a “typical” set of options. Also, in some cases, the IPP attempts to adjust observed price changes for changes in quality using estimates of the associated changes in costs. Our analysis of substitutions focuses only on substitutions for which it is not possible to do this type of quality adjustment. This pragmatic approach is likely to lead to some spurious price changes that reflect changes in the product, rather than true price changes.¹¹

¹⁰The only exception to this policy is that the IPP does not accept prices that include U.S. duty unless sufficient information can be obtained to remove the duty charges in constructing the index. (See Section 4.2 of the IPP Data Collection Manual).

¹¹Substitutions in the IPP dataset are conceptually the same as non-comparable substitutions in the CPI data. The CPI data also contains information about the frequency of comparable substitutions. In contrast, the IPP does not record the frequency of non-comparable substitutions.

The IPP requests that reporters provide a price for the transaction that occurs as close as possible to the first day of the month. In practice, however, importers and exporters often go for long periods of time without purchasing from overseas. As a consequence, reporters often provide prices for other days in the month. In some cases, average monthly prices are provided. Furthermore, the prices requested by the IPP are the prices of products *received* by the firm as close as possible to the reference date. Production lags and delivery lags may therefore imply that this price will refer to a product ordered at a substantially earlier point in time.

It should be clear from the discussion above that the raw price data collected from firms, referred to as the “reported price”, must be manipulated in a variety of ways to construct the “net prices” that are eventually used in the import and export price indices. To arrive at the net price, prices are imputed to fill in any missing data. Prices are also adjusted for taxes and discounts reported in the discount structure, and converted to U.S. dollars. In our analysis, we start by dropping all imputed and estimated prices. We then use a simple approximation to the IPP procedure for imputing missing prices: we fill forward the last observed reported price through all missing periods, up to the last non-estimated price. To avoid introducing spurious price changes associated with numerical issues in converting prices quoted in foreign currencies into dollars, we use the “reported price” rather than the net price in our baseline analysis.

The sampling weights used in the IPP indices were, until 1997, specified at the level of detailed product groupings, called “classification groups”. Within a classification group, products were sampled with a probability equal to their trade weight, eliminating the need for more disaggregated weights. Since 1997, the IPP has introduced product-level weights below the level of individual classification groups. The IPP reports that a staff study in 1997 found that the change in methodology to the new item weights did not have a large effect on the index. The historical product-level weights are unfortunately not available. The most detailed weights available are classification group weights.

For the purpose of calculating the frequency of price change and substitutions and constructing aggregate indices, we construct item-level weights using a procedure analogous to the one that the IPP actually used until 1997—by dividing the classification group weight by the total number of “net price” observations used in the BLS index. In this regard, our analysis differs from earlier analyses by Gopinath and Rigobon (2008) and Gopinath, Itskhoki, and Rigobon (2008), which do

not incorporate sampling weights. However, for the regression analysis in our robustness analysis in section 6, we found that using weights at this detailed a level introduced a substantial amount of sampling error in our results. We therefore opted to calculate item-level weights as the average weight within a given HS2 code over the entire sample period and use these weights in our regression analysis.¹² This average is taken over all of the price observations used in our analysis including the “filled-in” observations. In section 6, we also carry out some regressions that condition on a price change having occurred. To avoid overweighting product categories with a high frequency of price, we calculate average item-level weights at the HS2-code level for these regressions by dividing the classification group weight by the total number of price change observations in that HS2 code.

Many of the products for which prices are sampled in the IPP data are intermediate products sold to other firms. An important question in interpreting the evidence on price rigidity for imported goods is whether the observed rigid prices are “allocative” (Barro, 1977). Since manufacturers and retailers interact repeatedly, the observed price in a particular month may not actually determine purchasing decisions. Rather, this price may be an “installment payment” on a “running tab” that adjusts continuously but is unobservable. This issue is less likely to influence the long run relationship between prices and exchange rates than it is to affect the short term dynamics of this relationship.

We carry out most of our analysis at an aggregate level. However, in some cases, we also report statistics for the following sectoral groupings based on HS2 codes: animal products (01-05); vegetable products (06-15), foodstuffs (16-24), mineral products (25-27), chemicals and allied industries (28-38), plastics and rubbers (29-40), raw hides, skins, leather and furs (41-43), wood and wood products (44-49), textiles (50-63), footwear and headgear (64-67), stone and glass (78-71), metals (72-83), machinery and electrical (84-85), transportation (86-89) and miscellaneous items (90-97).

The second set of data we use is from the U.S. NIPA. We use the import price deflator for imported goods excluding non-oil. And we use the export price deflator for exported goods excluding agricultural products. We use data for the period 1982 through 2007 since the IPP was introduced in 1982 and the price deflators before that time are based on unit value data. Finally, we make use of trade-weighted monthly and daily exchange rates downloaded from the Federal Reserve Board’s

¹²There are roughly 95 HS2 code groups, while there are roughly 6900 classification groups.

website at <http://www.federalreserve.gov/releases/h10/Hist/> as well as monthly bilateral exchange rates from the International Financial Statistics (IFS) database of the IMF.

3 Prices and Exchange Rates: Evidence

Measuring the degree of pricing to market for U.S. imports involves comparing the price of imports into the U.S. relative to all domestic prices and the U.S. real exchange rate.¹³ A simple way to measure the degree of pricing to market for U.S. imports is to run the following cointegrating regression:

$$p_t^m - p_t = \alpha + \gamma t - \beta q_t + \epsilon_t, \quad (1)$$

where p_t^m denotes the log of the dollar price of U.S. imports, p_t denote the log of the dollar price of U.S. production and q_t denotes the log of the trade weighted U.S. real exchange rate. Let e_t denote the log of the trade weighted U.S. nominal exchange rate (foreign currency price of U.S. dollars) and p_t^* denote the log of the foreign currency price of foreign production. The real exchange rate is defined as $q_t = e_t + p_t - p_t^*$. The estimated degree of “exchange rate pass-through” is given by β and the estimated degree of pricing-to-market is given by $1 - \beta$. It is because we are focusing on aggregate price indices that we use the terms “incomplete pass-through” and “pricing to market” interchangeably. Goldberg and Knetter (1997) emphasize that these concepts are distinct when applied to micro-data.

To capture potential dynamics in the relationship between movements in import prices and the real exchange rate, we consider the following vector error correction model (VECM),

$$\Delta y_t = \Pi(Ay_{t-1} + \alpha + \gamma t) + \sum_{k=1}^{n-1} \Gamma_k \Delta y_{t-k} + \delta + \epsilon_t, \quad (2)$$

where y_t is the vector $(p_t^m - p_t, q_t)$, and A is a vector of coefficients in the cointegrating relationship given by $[1 \ \beta]$. The parameter β is therefore a measure of the long-run responsiveness of import prices to exchange rates. We find strong evidence of a cointegrating relationship between import prices and the real exchange rate. We reject the null hypothesis of no cointegrating equations using the Johansen trace statistic method (Johansen, 1995). The Schwarz Bayesian Information Criterion selects one lag in the vector error correction model, so we set $n = 2$.

¹³Knetter (1989) and Marston (1990) pioneered empirical estimation of pricing to market.

The third specification we consider is the following “dynamic adjustment” specification:

$$\Delta(p_t^m - p_t) = \alpha - \sum_{k=0}^6 \beta_k \Delta q_{t-k} + \epsilon_t. \quad (3)$$

This is the type of regression that is typically run in the exchange rate pass-through literature (e.g., Campa and Goldberg, 2005; Marazzi and Sheets, 2007). Long-run pass-through is then defined as the sum of the coefficients, $\hat{B} = \sum_{k=0}^6 \beta_k$. This specification is misspecified if relative import prices and the real exchange rate are cointegrated.

For U.S. exports, measuring the degree of pricing to market involves comparing price of U.S. exports relative to all foreign prices and the U.S. real exchange rate. The regression equations we use to make this comparison for U.S. exports are analogous to equations (1)-(3). First we consider:

$$p_t^{x*} - p_t^* = \alpha + \gamma t + \beta q_t + \epsilon_t, \quad (4)$$

where p_t^{x*} denotes the foreign currency price of U.S. exports. We construct $p_t^{x*} - p_t^*$ as $p_t^x + q_t - p_t$, where p_t^x is the dollar price of U.S. exports. Second, we consider a VECM analogous to equation (2) but with $y_t = (p_t^{x*} - p_t^*, q_t)$. Again, we reject the null hypothesis of no cointegrating equations using the Johansen trace statistic method. The Schwarz Bayesian Information Criterion selects two lags in the vector error correction model, so we set $n = 3$. We allow for a structural break in the cointegrating relationship for exports in 2004 by adding a dummy variable that is equal to one in the first quarter of 2004 and thereafter into the vector y_t . This dummy variable accounts for an apparent level shift in the cointegrating relationship between export prices and the real exchange rate after 2004, as we discuss below.

Finally, we consider the following “dynamic adjustment” specification for exports,

$$\Delta(p_t^{x*} - p_t^*) = \alpha + \sum_{k=0}^6 \beta_k \Delta q_{t-k} + \epsilon_t. \quad (5)$$

Results for these six regressions are presented in table 1. We use the NIPA price deflator for non-oil goods imports and non-agricultural goods exports. Our sample period is from 1982 through 2007. We begin our sample in 1982 because this is when the import and export price indices were introduced in the U.S. The exchange rate variables we use are Federal Reserve’s trade weighted index for major currencies.¹⁴ We use consumer prices as our proxies for p_t and p_t^* .

¹⁴These indices are similar to the index used in Campa and Goldberg (2005). The major currency exchange rate series seems more appropriate than broader index for two reasons. First, the weights in the import and export price

We find that the dynamic adjustment regression and the VECM yield similar conclusions regarding the long-run relationship between trade prices and real exchange rates. For imports, the VECM yields an estimate of β of 0.41 while the dynamic adjustment equation yields an estimate of 0.43 for the sum of the coefficients \hat{B} . The levels regression yields a somewhat lower estimate of 0.36, presumably because the a substantial fraction of the adjustment of trade prices is lagged by 2-3 quarters relative to the real exchange rate adjustment. These estimates are broadly in line with the existing literature on exchange rate pass-through. For example, Campa and Goldberg (2005) estimate long run exchange rate pass-through for U.S. imports to be 0.42 for the period 1975 to 2003. For export prices, all three models yield estimates within the tight range of 0.85 to 0.87.¹⁵

Figures 2 and 3 display the stability of the relationship documented above. Figure 2 plots the detrended relative dollar price of U.S. imports $p_t^m - p_t$ and the fitted values based on the cointegrating relationship. To be specific, the “import price” is $p_t^m - p_t - \hat{\gamma}t$, while the “predicted import price” is given by the fitted values from the cointegrating equation $\hat{\beta}q_t$, where the parameter estimates are based on the estimates of the VECM. The two series are normalized to have the same means. Figure 3 plots analogous series for the case of exports. Specifically, Figure 3 plots the detrended relative foreign price of U.S. exports $p_t^{x*} - p_t^* - \hat{\gamma}t$ and the fitted values based on the cointegrating relationship $\hat{\beta}q_t$. Over this time period, these relationships—both for imports and exports—have been quite stable. The only significant instability in these relationships is an upward level shift in the price of exports after 2003 relative to their predicted value. We capture this feature of the data by allowing for a level shift in the constant term in the cointegrating relationship in the first quarter of 2004.

Several researchers have argued that exchange rate pass-through into U.S. imports has fallen in recent years (e.g., Olivei, 2002; Marazzi and Sheets, 2007). Table 2 reports pass-through estimates for U.S. imports from both the VECM and the dynamic adjustment regression for two subsamples: 1982-2008 and 1994-2008. For the dynamic adjustment regression, pass-through is indeed estimated to be quite a bit lower in the recent subsample—0.32 compared to 0.43 over the longer sample period.

index are often 3-5 years out of date. This implies that the growing role of countries outside the group of major currencies is captured only with a substantial lag and is therefore small for the majority of our sample period. Second, the major exchange rate index is potentially more similar to an index of non-oil import prices. However, we also estimated these regressions with the Federal Reserve’s broad exchange rate indices and we discuss the results for these alternative regressions below.

¹⁵If we instead use the Federal Reserve’s broad exchange rate index, the VECM yields 0.50 for imports and 0.84 for exports, while the dynamic adjustment regression yields 0.52 for imports and 0.84 for exports.

The results for the VECM, however, suggest that this apparent fall in pass-through might be due to model misspecification. For the VECM, the long-run pass-through estimate is slightly higher for the recent sample than it is for the longer sample—0.46 versus 0.41.

4 Prices and Exchange Rates: Theory

Consider an economy in which consumers purchase and consume products from a continuum of product lines. At each point in time, one product from each product line exists. Let C_{jit} denote the number of units of the current product from product line j produced in country i and consumed at time t and let γ_{jit} denote the quality of each of these units measured in terms of utility. Products from country i enter the consumer's utility function through the following consumption aggregator

$$C_{it} = \left[\int_{N_i} (\gamma_{jit} C_{jit})^{\frac{\theta-1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}}.$$

Let P_{jit} denote the price per unit of the current product from product line j in country i at time t . The price index that gives the minimum cost of an additional unit of utility is then given by

$$P_{it} = \left[\int_{N_i} \left(\frac{P_{jit}}{\gamma_{jit}} \right)^{1-\theta} dj \right]^{\frac{1}{1-\theta}}. \quad (6)$$

On the firm side of the economy, firm j in country i produces goods according to the following production function

$$C_{jit} = \gamma_{jit}^{-1} F(K_{jit}, L_{jit}).$$

The function F is homogeneous of degree one in capital, K_{jit} , and labor, L_{jit} . Quality enters the production function multiplicatively. This implies that to raise the quality of its products by a factor ξ and produce the same number of units the firm must employ ξ times as much of each factor input. In other words, it costs the firm twice as much to produce goods that are twice as good in utility terms.

Now let $\hat{C}_{jit} = \gamma_{jit} C_{jit}$ denote effective consumption of product line j at time t and let $\hat{P}_{jit} = \gamma_{jit}^{-1} P_{jit}$ denote the corresponding effective price. Using these concepts, it is possible to rewrite both the consumer problem and the firm problem entirely without reference to γ_{jit} . In particular, the consumption aggregator and the price index become

$$C_{it} = \left[\int_{N_i} \hat{C}_{jit}^{\frac{\theta-1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}} \quad \text{and} \quad P_{it} = \left[\int_{N_i} \hat{P}_{jit}^{1-\theta} dj \right]^{\frac{1}{1-\theta}}.$$

And the production function becomes

$$\hat{C}_{jit} = F(K_{jit}, L_{jit}).$$

This implies that the equilibrium allocations generated by models with this set of assumptions about product quality are the same as those of any number of standard models in the macroeconomics and international economics literature. The equilibrium allocations generated by standard models simply refer to effective consumption and effective prices in the corresponding model with changing product quality. Changing product quality, however, complicates the empirical evaluation of standard models if product quality is not observed. In this case, the models generate data on effective prices and quantities, while the real world data consists of raw prices and quantities. Below, we consider how the behavior of observed data on prices and quantities differs from the behavior of effective prices and quantities when product quality is unobserved.

Suppose that product replacement occurs in each product line with probability z . For simplicity, assume that each time a product is replaced a new γ_{jit} is drawn from a stationary distribution Γ and that this level of quality remains constant for the product until it undergoes another replacement. In other words,

$$\gamma_{jit} \begin{cases} \sim \Gamma & \text{with probability } z \\ = \gamma_{jit-1} & \text{otherwise} \end{cases}$$

Suppose that when a firm introduces a new product for product line j , it faces no pricing friction and sets the optimal price from its perspective. In other words, the prices of newly introduced products are fully flexible. For continuing products, however, prices are sticky in the local currency. Assume that firm j adjusts the price of continuing products with probability f_j , where the frequency of price change across different product lines has some distribution Φ .¹⁶

Let lower case variables denote logarithms of upper case variables. A first order Taylor-series approximation of equation (6) yields

$$p_{it} = \int_{N_i} \hat{p}_{jit} dj. \tag{7}$$

¹⁶We consider the simple case where product replacement is not associated systematically with variations in firm productivity or firm markups. This is consistent with the fact that over the sample period we analyze, product replacement does not vary systematically with the real exchange rate. However, if movements in the real exchange rate are systematically associated with movements in markups for new products, then the types of biases we consider associated with product replacement may be further exacerbated. See Auer and Chaney (forthcoming) and Rodriguez Lopez (2008).

First differencing this equation yields

$$\Delta p_{it} = \int_{N_i} \Delta \hat{p}_{jit} dj. \quad (8)$$

Assume for simplicity that the price index in equation (7) is the price index that the BLS seeks to construct when it calculates the U.S. import and export price indices.¹⁷ A major complication in constructing this type of price index is the fact that product quality is unobserved. The ideal solution would be to use hedonic methods to estimate product quality. However, such methods are almost never used because they are extremely costly and difficult to apply in most cases. In practice, price comparisons that involve a change in quality are usually dropped from the index. Indices constructed in this way are referred to as “matched model indices” since all price comparisons on which such indices are based are for identical items.

4.1 Product Replacement Bias: A Factor Calculation

Consider the following regression,

$$\Delta p_t = \alpha + B\Lambda_t + \epsilon_t, \quad (9)$$

where B is a vector of coefficients and Λ_t is a vector of aggregate variables. Given equation (8), it is straight-forward to show that the vector of regression coefficients for this regression, B , may be “decomposed” as follows

$$B = \int_{N_i} B_j dj, \quad (10)$$

where B_j denotes the regression coefficient from estimating equation (9) with micro-price change observations from product line j , $\Delta \hat{p}_{jit}$, as the dependent variable.

Suppose that data on γ_{jit} was available and that it was therefore possible to estimate equation (9) using *all* observations on changes in effective relative prices for product line j as the dependent variable. Our assumptions about price rigidity and product replacements imply that the effective price of product line j changes with probability $f_j + z - f_j z$. We can divide the observations on price changes and exchange rates for this regression into two groups based on whether $\Delta \hat{p}_{jit} = 0$ or $\Delta \hat{p}_{jit} \neq 0$. Let B_j denote the regression coefficients using all the available observations from product line j . Let B_j^{ch} denote the coefficients based on estimating equation (9) using only the observations for which $\Delta \hat{p}_{jit} \neq 0$. Since we have assumed that the timing of price changes is unrelated to

¹⁷This price index is a special case of the Tornqvist index with fixed and equal weights.

movements in the exchange rate, the regressors for the observations with $\Delta\hat{p}_{jit} = 0$ and those with $\Delta\hat{p}_{jit} \neq 0$ have the same distribution.¹⁸ In this case, it is easy to show that $B_j = (f_j + z - f_j z)B_j^{ch}$ asymptotically.

Now, consider the case where data on γ_{jit} is not available. In this case, the effective price change is not observed when a product replacement occurs. Consider estimating equation (9) using data from the remaining time periods and let B_j^{mm} denote the vector of coefficients from this regression. Price changes occur in a fraction f_j of these time periods. Since the change in the effective price has the same distribution in these periods as in the periods in which product replacements occur in our model, we have that $B_j^{mm} = f_j B_j^{ch}$ asymptotically. We thus have that,

$$B_j^{mm} = f_j B_j^{ch} = \frac{f_j}{f_j + z - f_j z} B_j. \quad (11)$$

In other words, B_j^{mm} yields an estimate of the true regression coefficients (defined as B_j) that is downward biased by a factor $f_j/(f_j + z - f_j z)$.

Integrating over j yields

$$B^{mm} = \int \frac{f_j}{f_j + z - f_j z} B_j dj.$$

This implies that the matched model index yields an estimate of B that is biased by a factor

$$\frac{B}{B^{mm}} = \frac{\int B_j dj}{\int \frac{f_j}{f_j + z - f_j z} B_j dj}. \quad (12)$$

In some cases we may be interested in the sum of the regression coefficients. One such instance is the dynamic adjustment regression—equation (3). Define $B(\text{sum})$ to be the sum of the elements of B . Since $B(\text{sum})$ is a linear function of the individual regression coefficients, we have,

$$\frac{B(\text{sum})}{B^{mm}(\text{sum})} = \frac{\int B_j(\text{sum}) dj}{\int \frac{f_j}{f_j + z - f_j z} B_j(\text{sum}) dj}. \quad (13)$$

Under the further simplifying assumption that $B_j(\text{sum})$ is the same for goods from different product lines j , we have that

$$\frac{B(\text{sum})}{B^{mm}(\text{sum})} = \left[\int \frac{f_j}{f_j + z - f_j z} dj \right]^{-1}. \quad (14)$$

Since the function $f_j/(f_j + z - f_j z)$ is concave, product replacement bias will be greater the greater is the amount of heterogeneity in the frequency of price change across products.

¹⁸In particular, we need to assume that $E(\Lambda_i \Lambda'_i)$ is the same for the regressors in both groups, where i denotes a particular observation.

4.2 Product Replacement Bias and Pricing to Market

We now apply the product replacement bias factor adjustment derived above to pricing to market regressions. Consider first the dynamic adjustment regression—equation (3). The factor calculation does not in general apply directly to this regression since the dependent variable in this regression is the change in the *relative* price of imports, $\Delta(p_t^m - p_t)$. However, the factor calculation does provide a reasonable approximation when the covariance of p_t with the real exchange rate is small enough that equation (3) yields similar results with Δp_t^m as with $\Delta(p_t^m - p_t)$ as the dependent variable. This is true in practice.¹⁹

Deriving the extent of product replacement bias for the VECM is somewhat more involved. Under the assumptions that p_t is uncorrelated with the real exchange rate and that the real exchange rate and the relative price of imports are integrated of order one, we can show that the bias in β in the VECM—equation (2)—is given by equation (12). This derivation is presented in appendix B.

The discussion given above is for LCP products. For PCP products, matched model indices yield an *upward* biased measure of exchange rate pass-through. For concreteness, consider a product imported from the Euro zone that has a sticky price in Euros. Product replacement bias then causes a downward bias to pass-through in Euro terms. This turns into an upward bias when prices are converted into U.S. dollars. Consider the extreme case where prices are perfectly rigid over the course of product lifetimes, and all pass-through occurs at the time of product replacements. Measured pass-through in Euro terms is then zero. And stable prices in Euro terms are equivalent to perfect pass-through in dollar terms. Therefore, in this extreme, estimated long-run pass-through would be one even if true long-run pass-through were small.

4.3 Discussion

The assumption that prices are more responsive to the economic environment for newly introduced products than for continuing products is a crucial feature of our model. Matched model price indices drop the price changes that occur for newly introduced products. If prices are more likely to change at these points than at other points, dropping these observations leads to a selection bias in the measurement of the relationship between prices and aggregate variables. Matched model

¹⁹If p_t is positively correlated with the real exchange rate (prices rise when the real exchange rate depreciates), equation (13) understates product replacement bias. If they are negatively correlated, equation (13) overstates product replacement bias.

price indices will then underestimate the responsiveness of prices to aggregate variables since they disproportionately drop price change observations.

One way to view matched model price indices is as implicitly imputing a price change for the effective price of newly introduced products that is equal to the average change in the price of all continuing products. In our model this is a poor assumption. The effective price change for newly introduced products has the same distribution as the effective price change of all continuing products *that change their price* in that period in our model. This implies that the change in the effective price of newly introduced products is much larger on average than the average price change of all continuing products.

A more general model would allow for differences in the pricing behavior of firms when products are introduced versus when the price of continuing products changes. In such a model, the extent of product replacement bias would be given by a generalized factor:

$$\frac{f_j}{f_j + \alpha z - \alpha z f_j},$$

where α represents the difference in the pricing behavior of firms at the time of product introductions versus when they change the price of continuing products. One example of a more general model would be a model in which firms have both a menu cost and a quadratic cost of changing prices for continuing products but no such quadratic cost when setting the first price of new products. This type of pricing behavior could arise because of implicit contracts or consumer antagonism (Nakamura and Steinsson, 2008; Rotemberg, 2005) In this case, α is greater than one and the product replacement bias is larger than in our baseline model. On the other hand, if prices change less at the time of product replacements or if some products are introduced without the firm reoptimizing their prices, α is less than one and product replacement is smaller than in our baseline model.

It is well known that prices may respond much more quickly to monetary shocks in menu cost models than in the Calvo model (Caplin and Spulber, 1987; Golosov and Lucas, 2007). An important question, therefore, is the extent to which our results on product replacement bias depend on our assumption that price adjustments and product replacements occur with a constant probability. In fact, our results may be generalized in at least two ways. First, they hold whenever the price changes at the time of product replacements are a random subsample of all price changes. For example, consider a menu cost model in which firms flip a coin each time they hit their

Some firms band about whether to change the specifications of their product as well as reoptimize their price. Second, consider a model in which the timing of price adjustments for a given product are determined by a menu cost but the timing of product replacements is random (Nakamura and Steinsson, 2007). When this type of model is calibrated to match micro-data on price changes and product replacement it is found that the timing of price changes is largely determined by idiosyncratic shocks and is roughly constant over time. In the limit when the frequency of price change is constant in such a model, our factor argument goes through. The only difference caused by the state dependence of price change is that pass-through will be faster.

The assumptions that the frequency of price change and the frequency of product replacement are constant are therefore not crucial. They are merely simplifying assumptions. The key assumption is that the distribution of regressors in equations (1)-(5) is the same for product replacements as for other price changes. This assumption holds in a larger class of models than the Calvo-type model considered here. In fact, we find no statistically or economically significant relationship between the frequency of product replacements and the magnitude of exchange rate movements in our data.²⁰ This finding is not, however, definitive given the short time period we analyze. In contrast, Burstein, Eichenbaum, and Rebelo (2005) document clear evidence of a rise in the number of products that ceased to be imported into Argentina at the time of Argentina’s 2000-2002 financial crisis and devaluation. These results suggest that there may be a relationship between product introduction or destruction and exchange rates for larger exchange rate movements. Such a relationship would constitute an additional source of product replacement bias. Investigating the nature of this bias is beyond the scope of this paper.

The observed price rigidity in our micro price data may reflect problems with the way in which the data is reported to the BLS—barriers to *reporting* price changes—rather than truly rigid prices. From the perspective of our analysis, it does not matter whether the observed price rigidity is “real” or simply reflects frictions in reporting price movements to the BLS. As we describe in section 2, once a product has been initiated by the BLS, firms fill out “repricing forms” that first ask whether the price has changed relative to the previous month and then ask the respondent to report a new price if the price did change. This reporting procedure may, therefore, introduce a cost for firms of

²⁰We regressed the frequency of product replacements for dollar-priced imports and exports on the absolute magnitude of log movements in the trade-weighted exchange rate for the years 1995-2006 (the period of years when a sufficiently large sample of price microdata is available). The resulting coefficient is -0.0089 (0.093) for imports and 0.120 (0.100), where we report standard errors in parentheses.

reporting price changes for continuing products that leads reported price to be sticky even if actual prices are not.

In our model economy, one could, in principle, calculate a price index simply as the weighted average of all prices including both new and continuing products. Since this procedure makes use of all price comparisons, it avoids the selection bias problem associated with dropping price changes associated with new product introductions. Yet, it implicitly means comparing the prices of entirely different products—say, last year’s wool jacket versus this year’s down coat. In a world of highly heterogeneous products, these price comparisons introduce a large amount of sampling error into the price index. In the IPP data, we find that an import price index calculated from average prices in this way is extremely noisy. To match this feature of the data in our model, we assume that the distribution Γ is sufficiently dispersed that an index based on average prices yields an unacceptably large amount of sampling error. In other words, there is no easy fix for product replacement bias based on calculating an index of average prices.

If movements in the exchange rate were exogenous, our estimate of exchange rate pass-through would represent a causal link from the exchange rate to prices. In this case, $B < 1$ could indicate that as the U.S. dollar appreciates versus foreign currencies, firms adjust their prices to stabilize it in local currency terms. In equilibrium, however, exchange rates vary as a consequence of demand and supply shocks that also affect firm’s costs and market shares directly. For instance, if exchange rate movements arise largely from monetary shocks then the covariance between exchange rates and import prices may be low even if the price response to an exogenous exchange rate movement is much larger (Corsetti, Dedola and Leduc, 2008; Bouakez and Rebei, 2008). Our analysis of product replacement bias does not rely on a particular model of equilibrium exchange rate determination. Our objective is simply to improve on the measurement of the equilibrium relationship between prices and exchange rates.

5 Product Replacement Bias: Measurement

We show in section 4 that product replacement bias is most severe when the frequency of product replacements is large relative to the frequency of price change. Table 3 reports our estimates of the weighted fraction of products that have less than or equal to 0, 1, 2, 3 and so on price changes. For LCP imports, 41% of products have no price changes, while 68% have two or less price changes. For

PCP exports, 51% of products have no price changes, while, 78% have two price changes or less. These statistics motivate the idea that product replacement bias may be an important phenomenon in import and export price data.²¹

The small number of price changes per product reflects substantial price rigidity in the microdata on import and export prices collected by the BLS. Table 4 reports the weighted mean and median frequency of price change for imports and exports, separately for LCP and PCP goods. These statistics parallel those reported in Clausing (2001) and Gopinath and Rigobon (2008), though our analysis differs somewhat from theirs in that we study a longer time period, and make use of product-level weights.²²

Table 4 shows that both imports and exports exhibit substantial price rigidity. Most U.S. imports are local currency priced (93%). For these goods, the mean frequency of price change is 15.7% and the median is 9.8%. The fact that the median frequency of price change is much lower than the mean reflects the long right tail in the distribution of the frequency of price change across goods—while most products have a frequency of price change less than 10%, a significant fraction have a frequency of price change over 50%. This skewness is similar to the skewness in the distribution of the frequency of price change for consumer and producer prices (Nakamura and Steinsson, 2008). Most U.S. exports are producer currency priced (98%). For these goods, the mean frequency of price change is 17.8% while the median is 7.2%. Table 4 also reports the weighted mean frequency of product substitutions for imports and exports. This fraction varies from 4.0-5.5% across the four classes of goods considered.

Equation (12) indicates that the degree of heterogeneity in the frequency of price change is an important determinant of the quantitative impact of product replacement bias. A simple way to assess the degree of heterogeneity in the frequency of price change is to present results on this statistic for different industry groups. Table 5 reports the mean frequency of price change and the mean frequency of product replacement across different industry groups of imports and exports.

²¹Our estimate of the fraction of price spells with no price change is somewhat higher than the estimate of Gopinath, Itskhoki, and Rigobon (2008). Their estimate for LCP imports is around 30%, while our estimate is 41%. Most of the difference arises because our estimate is for the entire dataset, while theirs is for a subset of OECD countries. Another difference is that our estimates incorporate product-level weights.

²²We calculate the frequency of price change by constructing an indicator variable for whether a price change occurred and taking the mean of this variable. We calculate the frequency of product substitutions as the total number of product substitutions observed in the data, divided by the total number of periods that the price series are observed. The series we use for this are constructed by “filling in” the previously observed price through the large number of missing spells in the import price data as we discuss in section 2. All of the statistics we report are calculated as weighted averages using the item-level weights described in section 2.

The frequency of price change varies widely—from 64% for LCP mineral product imports, a group that includes gasoline, to less than 5% for several categories. In contrast, the average frequency of product replacement varies much less across industry groups. Most industry groups have frequencies of product replacement of 3-6%. In what follows, we will assume for simplicity that the frequency of product replacement is the same for all goods. We will, however, take explicit account of the large amount of heterogeneity in the frequency of price change.

While a great deal of heterogeneity in the frequency of price change is embodied in differences across sectors, there is also much heterogeneity within each sector. Ignoring intra-sector heterogeneity would seriously bias our estimate of product replacement bias. We therefore estimate a flexible distribution for the overall heterogeneity in the frequency of price change across products. Suppose that a product j has a constant hazard of adjusting its price, f_j , in each month. And suppose that $f_j \sim \text{Beta}(a, b)$. We denote the product’s lifetime by n_j . Given that the product has a constant hazard of adjusting its price, the total number of price changes x_j are distributed according to a binomial distribution, i.e., $x_j \sim \text{Bin}(n_j, f_j)$.

In appendix A, we derive a simple expression for the log-likelihood function in this setting. We estimate this model by maximum likelihood for four groups of products: LCP imports, PCP imports, LCP exports and PCP exports. Table 4 reports our estimates for the parameters of the beta distribution. For LCP imports, the estimated parameters are $a = 0.50$ and $b = 3.65$. These parameters imply a very large amount of heterogeneity in the frequency of price change across products. Figure 4 plots the cumulative distribution function of the distribution $\text{Beta}(0.50, 3.65)$. The mean of this distribution is 0.121 while the standard deviation is 0.144. For PCP exports, the estimated parameters are $a = 0.53$ and $b = 4.60$. The mean of this distribution is 0.103 while the standard deviation is 0.123. We have estimated the amount of heterogeneity in the frequency of price change within sectors using this same method. We find the within-sector heterogeneity to be almost as large as the overall heterogeneity for most sectors.

For both LCP imports and PCP exports, the mean of the beta distribution we estimate is somewhat lower than our estimate of the weighted mean frequency of price change discussed above. These discrepancies likely arises because the beta distribution is an imperfect approximation to the true distribution of the frequency of price change across products. The median of the beta distribution and is much closer to the weighted median we estimate from the data. To assess the

robustness of our results, we did all our calculations both using the aggregate beta distribution and using a weighted mixture of the sectoral beta distributions. The results were very similar for both cases. We present the results based on the aggregate beta distribution for simplicity.

Defining what is a new product in the IPP data is not always completely straightforward. We use the product definitions of the BLS. This implies that the definition of a product includes not only the specifications of the product itself but also other price determining factors such as the shipment size, delivery terms and possibly the identity of the seller or buyer. Carlton (1986) provides evidence that defining a product in this way is important when studying the flexibility of prices.

In the model presented in section 4, all products belong to product lines. Each time a product is discontinued, it is immediately replaced with another product. All products are observed from the time they come into being. Things are not quite this simple in the real world. While most product substitutions in our data occur because the BLS staff are not able to reprice an item that they have been following and a large fraction of these “forced substitutions” occur because of a version change or upgrade, some product substitutions occur for other reasons.²³ This is, however, not the key determinant of product replacement bias. Its key determinant is the extent to which the new products that enter the data set have prices that have been changed more recently on average than the products that exit—irrespective of whether the entering products and the exiting products are related.

New products in the model presented in section 4 have completely fresh prices. In the data, a large fraction of product substitutions occur because of actual product replacements, where an entirely new product or a new version of an older product—presumably with an entirely fresh price—replaces an old product. Also, the way in which the BLS gathers price data suggests that reporting frictions may be more important for continuing products than new products, as we discuss in section 2. To the extent this is true, the reported price of the newly introduced product will on

²³It is difficult to estimate the fraction of substitutions that involve a version change or upgrade. The dataset contains a flag indicating whether a product substitution is due to such a version change or upgrade. However, there are at least two reasons why this flag is unreliable. First, for most of the time period we study, to qualify as a version change or upgrade, the replacement product must fall into the same HS10 category. Since these categories are extremely disaggregated, it often happens that the replacement product falls in a different HS10 code. For example, male cows and female cows are different HS10 as are VHS players and DVD players. Second, BLS economists have indicated to us that many product discontinuations are followed by reinitiations of similar products by a BLS field agents. This may happen because firms find it easier to simply discontinue a product than to report the details of a replacement product to the BLS.

average be less stale than the reported price of the product it replaces. Furthermore, the BLS may choose to replace products that they rotate out of the dataset with new products. In fact, this is the stated purpose of the product rotation scheme.

In figure 5, we present direct evidence on the degree to which newly introduced prices have been recently changed. For products with exactly two price changes, we run two regressions. First, we regress the size of the second price change on 18 monthly exchange rate changes before the beginning of the price spell that ends with the price change in question. Second, we regress the size of the first price change on 18 monthly exchange rate changes before the introduction of the product in question. If the introduction of the product coincides with a price change, the coefficients of these two regressions should be the same. If, however, the first price spell of the product reaches back before the introduction of the product into the data set, the coefficients in the regression for the first price change should be larger than in the other regression.

We find that the coefficients on the first two exchange rate changes are substantially larger in the regression on the first price change than the second price change. However, the coefficients further back are quite similar. Figure 5 plots two lines. For the second price change, the j th point on the line is the sum of the coefficients on the first j lagged exchange rate changes. For the first price change, the j th point is the sum of the coefficients on the 3rd, 4th, ..., $(j + 2)$ th exchange rate change (as though the first price spell had begun two periods before the product was introduced into the data set). These two lines are very similar. This supports the notion that for most products the price reported at their introduction was set less than two months before the product was introduced. We use this information to adjust our estimate of product replacement bias. Specifically, we assume that the true average duration of price spells—accounting for product replacement—is two months longer than what our estimates of the frequency of price change and frequency of product substitutions.

Given the estimates in table 4 we can use equations (12) and (13) and the adjustment discussed in the last paragraph to produce estimates of the factor by which pricing to market is mismeasured because of product replacement bias. These estimates are reported in table 6. We carry out these calculations for two sets of assumptions regarding heterogeneity in true pass-through across goods with different frequencies of price change. First, we assume that true long-run pass-through is constant within each of the four groupings considered in table 4. For this case, we can apply equation

(13) and find that the factor for both LCP imports and PCP exports is roughly 1.8. Second, we consider a case in which true pass-through is lower for products with a lower frequency of price change. Gopinath and Itskhoki (2008) argue that this pattern exists in the data. Specifically, their estimates suggest that true pass-through for LCP imports with a frequency of price change below about 25% per month may be only about 65% of true pass-through for LCP imports with a higher frequency of price change.²⁴ We apply this degree of heterogeneity to each of the four groups considered in table 4. Applying equation (12), the factors for LCP imports and PCP export are both between 1.7 and 1.8.

Using these factors we can recompute the comovement between prices and exchange rates adjusting for product replacement bias. The results of these calculations for the VECM are reported in the lower panel of table 1. Gopinath, Itskhoki, and Rigobon (2008) argue that there is a large difference in pass-through between LCP and PCP imports. We allow for this difference in our calculations. In particular, we adopt their estimate of 0.94 for measured pass-through of PCP imports.²⁵ For exports, virtually all products are PCP. So, any reasonable heterogeneity across LCP and PCP products makes virtually no difference. As in table 6, we present results on pass-through adjusting for product replacement bias for both the case of no heterogeneity of this kind and a case calibrated based on the results of Gopinath and Itskhoki (2008).

Adjusting for product replacement bias raises the comovement of U.S. import prices (relative to all U.S. prices) and the U.S. real exchange rate from 0.41 to between 0.65 and 0.68. And it lowers the comovement between U.S. export prices (relative to all foreign prices) and the U.S. real exchange rate from 0.87 to roughly 0.80. This improvement in measurement brings the data closer in line with standard models along two dimensions. First, our adjusted estimates imply a much more similar degree of pricing to market for U.S. imports and exports. The unadjusted numbers imply a great deal of asymmetry between imports and exports that is difficult to match in theoretical models. Second, adjusting for product replacement bias implies more modest pricing to market for U.S. imports. This improves the fit of the data to models such as Corsetti and Dedola (2005) and Atkeson and Burstein (2008). The authors' preferred calibration of the models in these

²⁴This difference in measured true pass-through could alternatively arise due to a modest number of spurious price changes in the micro data. See section 2 for a discussion of reasons why spurious price changes may exist in the micro data on imports and exports.

²⁵Given this assumption about the pass-through of PCP imports we back out an estimate for LCP imports that yields our estimate of measured pass-through of 0.41 for all imports. Reasonable variations in this set of assumptions have negligible effects on our results.

papers implies exchange rate pass-through of roughly 0.9 and 0.75, respectively. Other models can in principle generate more pricing to market (Gust, Leduc and Vigfusson, 2006, Gopinath, Itskhoki and Rigobon, 2007). But to do so they must introduce a great deal of strategic complementarity. The degree of strategic complementarity that is needed to match the unadjusted estimates of pricing to market for imports in these models implies extremely slow pass-through. This is inconsistent with the empirical finding that the bulk of pass-through occurs in less than one year (Campa and Goldberg, 2005; Marazzi and Sheets, 2007). Our adjusted estimates of pricing to market for U.S. imports and exports are also more in line with industry studies. For example, Nakamura and Zerom (2008) estimate the degree of pricing to market for ground coffee in the U.S. to be roughly 1/3.

Table 7, presents results for several alternative data and measurement assumptions. In our baseline results, we use the trade weighted “Major Currency” U.S. real exchange rate published by the Federal Reserve Board. Table 7 presents alternative results using the Fed’s “Broad” real exchange rate series. Measured pass-through for somewhat higher for imports and slightly lower for exports using the Broad real exchange rate. Using this measure and adjusting for product replacement bias yields pass-through of roughly 0.80 for imports and 0.75 for exports. Also, in our baseline results, we count all substitutions. Table 7 presents alternative results where we count only forced substitutions. This change has only a minor effect on the resulting estimate of pass-through.

6 Alternative Measures of Comovement

Not all measures of exchange rate pass-through are affected by product replacement bias. Using the micro price data that underlie the U.S. import and export price indices, it is possible to avoid the problems associated with product replacement under certain assumptions. Consider regressing changes in relative prices on the change in the real exchange rate since the time of the previous price change:

$$\Delta(p_{jit} - p_t) = \alpha - \beta_{\Delta} \Delta^* q_t + \epsilon_{jit}. \tag{15}$$

where Δ^* is a difference operator representing the difference between the current real exchange rate and the real exchange rate at the time of the previous price change of product j (or in the case of the first price change of product j , the introduction of product j). In a simple model of price adjustment in which firms’ optimal prices depend only on current exchange rates, the coefficient

β_Δ in this regression provides a measure of the long-run relationship between prices and exchange rates that is unaffected by product replacement bias.

Another way to measure long-run pass-through is to run the following “levels” regression:

$$p_{jit} - p_t = \alpha_j + \gamma t - \beta q_t + \epsilon_{jit}, \quad (16)$$

on the subsample of observations for which a price change or product introduction occurs. Again, in a model in which firms’ optimal prices depend only on current exchange rates, the regression coefficient β provides a measure of the long-run relationship between prices and exchange rates that is unaffected by product replacement bias.²⁶

Table 8 presents estimates of these two regressions based on U.S. import price micro-data for the period 1995-2007, the years in which a full set of data are available in the IPP database. We estimate these regressions for dollar-priced non-oil imports for a subset of high income OECD countries: Canada, Sweden, U.K., Netherlands, Belgium, France, Germany, Switzerland, Spain, Italy, Japan. Observations from these countries account for the majority of observations in the IPP data set. For each country, we match the price data to exchange rate and consumer price data for that country from the IFS database of the IMF. The real exchange rate variable q_{it} is the end-of-month exchange rate for the month previous to the month identified for the reported price. We estimate both regressions using weighted OLS. The weights in the regressions are constructed at the HS2-code level to avoid overweighting HS2 codes with particularly large numbers of price changes. This procedure is described in more detail in section 2.²⁷

For equation (15), we estimate $\beta_\Delta = 0.26$. For equation (16), we estimate $\beta = 0.54$. Both of these estimates are substantially lower than our adjusted estimate of roughly 0.7 for the long-run relationship between import prices and the exchange rate in section 5. Furthermore, our estimate of β_Δ for equation (15) is much lower than our estimate of β for equation (16).

One important difference between the three regressions that these three estimates are based on is their differential reliance on the assumption that the price change at time t depends only on the exchange rate movement since the last price change. In particular, equation (15) relies much more

²⁶This regression yields similar results as the “life-changes” regression run in Gopinath, Itskhoki, and Rigobon (2008).

²⁷We have also run these regression on only price change observations, i.e. excluding the first price change for each product because the previous price change is not observed for these observations. This yields similar estimates. The similarity of these estimates is consistent with our assumption that, for a given product, pure price changes and product replacements generate similar responses to exchange rate movements.

heavily on this assumption that the other equations (Griliches and Hausman, 1986) and the VECM model estimated in section 3 is more flexible regarding the relative timing of exchange rates and prices than equation (16). This patterns of regression coefficients, therefore, suggests that price adjustments may respond to exchange rate changes further in the past than the last price change.

There are a number of reasons why price adjustments in a particular time period may depend on exchange rate movements before the last price change, invalidating the use of β_Δ and β as measures of long-run comovement. One reason is the presence of strategic complementarities in price-setting (Gopinath, Itskhoki and Rigobon, 2007). Another reason is measurement error in the timing of price observations in the IPP data. In section 2, we discuss evidence on a number of potential sources of this type of timing error in the IPP data that arise from the inherent difficulty of collecting survey data on import prices.

First, the IPP instructs firms to report prices for products that they *receive* on a particular day rather than prices for products that they *order* on that day. For example, suppose a product is ordered in January at a price p_1 and delivered in March, but that the order price of new identical products in March is p_2 . In this case, the IPP requests that the firm report p_1 rather than p_2 as its March price. While p_2 may reflect movements in exchange rates up to March, p_1 is set in January, and therefore does not reflect unforecastable movements in exchange rates and costs between January and March. The estimates in Abel and Blanchard (1988) suggest that common delivery lags for domestic manufacturers are in the range of 2 to 3 quarters.

Second, the prices of many products are only reported intermittently. For example, an importing firm might report the price of a product only in January, April, July and October—either because it only receives the product in those months or simply to reduce its reporting burden. If a price change occurs between, say, April and July, it is not possible to say exactly when the price changed within this interval. Overall, about 40% of the monthly price observations in the IPP data are missing. Moreover, even in the months for which prices are reported, firms often do not report prices that correspond precisely to the “reference date” specified by the BLS.

In this section, we analyze whether the results of the regressions described above, as well as standard measures of pricing to market, can be explained by a model with timing error and product replacement bias. We consider a simple model in which prices are set as a function of only the current exchange rate when they change—ruling out long delays in adjustment due to strategic

complementarities. This simple model provides a concrete example of how product replacement bias affects various measures of the comovement between prices and exchange rates. However, it is important to note that the analysis in this section is much less general than in the previous sections. Models with strategic complementarities in price-setting, forecastable movements in exchange rates or other general equilibrium factors that cause prices to adjust slowly to exchange rates are not consistent with the simple data generating process considered here.

As in our previous analysis, firm j adjusts its price with probability f_j and $f_j \sim \text{Beta}(a, b)$. Product replacements occur with constant probability, z . The key new structural assumption is that conditional on adjusting its price the firm sets its price as follows:

$$p_{jit} - p_t = -\phi q_{it} + \eta_{jit}, \quad (17)$$

where η_{jit} is orthogonal to the exchange rate. We assume that η_{jit} follows the stochastic process,

$$\eta_{jit} = \mu + \rho \eta_{jit-1} + \epsilon_{jit}, \quad (18)$$

where $\epsilon_{jit} \sim N(0, \sigma_\epsilon^2)$.

We define true pass-through as in previous sections as the long-run adjustment of prices to an exchange rate shock. In this model, true pass-through is given by ϕ . As in our previous analysis, we allow for heterogeneity in ϕ across products that is correlated with the frequency of price change.

To capture the type of timing error discussed above, we assume that the price that the IPP records for a firm at time t is the price that firm charged at a different time $t + u_t$. Specifically, we assume that $p_t^r = p_{t+u_t}$, where p_t^r denotes the price recorded by the IPP for the firm at time t and u_t reflects random timing error. We assume that the timing error u_t has two components, $u_t = u_{1,t} + u_{2,t}$. The first component $u_{1,t}$ is distributed $u_{1,t} \sim \text{Unif}[-1, 0]$. This term reflects the simple fact that the reported price changes for a particular month occur randomly over the course of the preceding month, but are observed at discrete intervals.

The second term $u_{2,t}$ is distributed, $u_{2,t} = -k d_{jit}$, where d_{jit} is the time in months since the last price change, and $k \sim \text{Unif}[0, K]$. This term is motivated by the various forms of timing error discussed above, which lead recorded prices to be somewhat “stale”. On average, we assume that the reported prices are “stale” by a fraction k of the duration since the last reported price.²⁸

²⁸In the case of the first observed price, we set d_{jit} to the unconditional duration of prices for the product in question. We never allow the price for the current month to be older than the time of the last observed price change.

In our simulation exercise below, we use actual daily observations on the U.S.-German exchange rate over the time period 1995-2007. When we simulate the model, we use the fractional values generated by the timing error model described above to infer on which day of the month a price change occurs. For robustness, we have carried out analogous experiments using the Canadian, Japanese, and U.K. exchange rates. These experiments yield almost identical results. For simplicity, in our simulations we assume that the difference between home and foreign inflation is constant.

6.1 Estimation

We estimate and analyze this model for the subset of high income OECD countries discussed above. For this subset of countries, the weighted average frequency of product substitution is 4.67%. The estimated distribution of the frequency of price change is Beta(0.52, 3.35). As in our earlier analysis, we follow Gopinath and Itskhoki (2008) in assuming that products fall into two groups: products with a low frequency of price change (lower than 25% per month), have a pass-through parameter of ϕ , while products with a higher frequency or price change have a pass-through parameter of ϕ_{high} , where we assume that $\phi = 0.65\phi_{high}$. Below, we find that this assumption provides a good explanation for the observed heterogeneity in the data. In practice, our assumptions regarding the process for idiosyncratic shocks have no impact on our results. We set $\rho = 0.5$ based on previous estimates in Nakamura and Steinsson (2008a) for consumer price data, and we set σ_ϵ^2 to match the average size of price changes in the data.

The remaining parameters are the timing error parameter K and the parameter ϕ . We use a simulated method of moments procedure to estimate these parameters. The moments we use in this procedure are the coefficients of the regression equations (15)-(16), β_Δ and β , as well as the coefficient β_2 obtained from re-estimating equation (16) on only the subset of products that have exactly two price changes.

The difference between β_Δ and β is directly related to the amount of timing error in the model. In our simple model, if we set the timing error to zero $\beta_\Delta = \beta$ and both coefficients provide good estimates of long-run pass-through. As the amount of timing error is increased, β_Δ declines sharply since the timing of the price changes are offset from the exchange rate movements they are suppose correspond to. The value of β is also affected by the presence of timing error, but the bias is much less severe. Finally, the presence of timing error also differentially affects the measured value of

β in equation (16) for products with different frequencies of price change (and therefore different numbers of price changes). Including β_2 in the set of moments we seek to match incorporates this effect in our estimation procedure.

Our estimation procedure is to select the values of K and ϕ that minimize the sum of the squared deviations between β_Δ , β and β_2 in the simulated and actual data. In our baseline estimation procedure, we weight the three moments equally. Since we are able to come very close to exactly matching the actual moments in the data, the choice of a weighting matrix makes little difference to our results.

This estimation procedure yields $K = 0.96$ and $\phi = 0.48$. The estimated value of K implies that on average, delivery lags and other sources of timing error account for a delay in price reporting of about 48% of the average duration since the last price change or product replacement. Since the median duration since a price change or product replacement is 8.4 months in our model, this corresponds to an average reporting lag of about 4 months. This magnitude of timing error seems well within the limits of existing estimates of delivery lags (e.g., Abel and Blanchard, 1988). The estimated value of ϕ implies that the low frequency of price change products in our model have desired pass-through of 0.48, while the high frequency of price change products have desired pass-through of 0.73, yielding aggregate pass-through of 0.52.

6.2 Results

Table 8 compares results from the data with results from the estimated model discussed in this section. The first three rows of the table present results for the three moments used to estimate the model: β_Δ , β and β_2 . We match these statistics almost exactly. Notice that in both the model and the data, β_Δ yields a measure of pass-through that is much lower than β . The assumed amount of heterogeneity in desired pass-through is also quantitatively successful in matching the observed difference between β and β_2 in the data.

The fourth row of table 8 presents results for equation (16) for the subset of goods that have eight or more price changes. This subset of goods yields a substantially higher value of β than does the same regression run on all products. The model is able to capture this difference. The high value of β for this subset of goods largely reflects a combination of higher desired pass-through and diminished importance of timing error for this subset of goods.

The fifth row of table 8 presents results for the micro dynamic adjustment equation—equation (3) with individual price observations as the dependent variable. The covariates in this regression are the current change in exchange rates, one future change and 22 lagged changes. This regression yields an estimate of long-run pass-through of 0.26, while in the model the estimated value is 0.31. In both cases, the estimated long-run pass-through is only approximately half of the true value. This discrepancy arises as a consequence of product replacement bias.

The fact that measured pass-through from the dynamic adjustment equation for the LCP imports is 0.26—substantially lower than the aggregate estimate of 0.43 reported in section 3—reflects the fact that our analysis of the microdata covers a substantially shorter sample period and includes only LCP goods. We report in table 2 that for the 1994-2008 period, estimated pass-through based on the dynamic adjustment model is 0.32 (as opposed to 0.43 based on the longer sample discussed in section 3). Adjusting this lower aggregate estimate for the fraction of LCP vs. PCP goods, using the assumptions about PCP goods described in section 5, yields an estimate of LCP pass-through of 0.27, which is similar our estimate based on the microdata of 0.26. The remaining difference between the aggregate and micro estimates likely arises from differences in product weights used in our analysis as compared to the official IPP index (see section 2) as well as effective weighting differences arising from aggregation (Theil, 1954).

7 Discussion

Product replacement may cause apparent long-run failures of purchasing power parity. Consider a situation where the nominal U.S. exchange rate versus Canada depreciates by 20%. Suppose that U.S. consumer prices rise by 20%, implying that the real exchange rate between the U.S. and Canada remains constant. Consumer prices may, nevertheless, appear to increase by less than 20% due to product replacement bias, leading the real exchange rate to appear to have permanently changed.

It is useful to note that similar measurement problems may also apply to unit value indices. Import price indices were introduced in the U.S. as a more accurate alternative to unit value indices.²⁹ While we do not study unit value indices in this paper, many of the measurement

²⁹Alterman (1991) estimated that the U.S. unit value indices, produced in 1985, were calculated for only 56 percent of the value of imports and 46 percent of the value of exports.

concerns we explore here also apply to unit value indices. In particular, for standard unit value indices, price comparisons are often dropped because of lack of availability of data for the previous period (for example, because the product was not traded in the previous period), potentially leading to product replacement bias.

Mismeasuring import and export price indices also affects measured trade volumes and trade price elasticities. Holding fixed nominal quantities, if the increase in import prices in response to an exchange rate depreciation is underestimated, then the corresponding decline in import quantities will be underestimated as well. If the corresponding price and quantity series are used to estimate trade price elasticities, then the resulting elasticity estimates will be biased away from one. If the true trade price elasticity is less than or equal to one, then the estimated price elasticity will be biased toward zero.

Finally, product replacement bias may help to explain differences in measured pass-through across countries. Burstein, Eichenbaum, and Rebelo (2005) report much higher pass-through for developing countries during large devaluations than for developed countries. Many developing countries have a large fraction of imports and exports priced in a foreign currency. The high fraction of PCP (foreign currency priced) imports in developing countries implies that measures of import price pass-through are likely to be upward biased as a consequence of product replacement bias. In contrast, product replacement bias is likely to generate a downward bias in measures of import price pass-through for developed countries, in which most imports are priced in domestic currency.

8 Conclusion

This paper argues that the simultaneous presence of import price rigidity and frequent product replacements may lead to serious problems in measuring the relationship between exchange rates and prices. Since many imports and exports exhibit only a small number of price changes over their lifetimes, we argue that the comovement of import and export prices and exchange rates is seriously mismeasured. We refer to the resulting bias as “product replacement bias”.

We quantify the magnitude of product replacement bias using empirical estimates of price rigidity and the frequency of product replacements for U.S. import and export price data. We find that conventional estimates of the comovement between import prices and exchange rates are biased

by a factor of two. On the other hand, export prices may comove somewhat less with exchange rates than conventional estimates suggest.

Overall, adjusting for product replacement bias makes the behavior of import and export prices easier to reconcile with standard models. The degree of pricing to market for U.S. imports and exports is much more symmetric, and the degree of pricing to market for U.S. imports more moderate than conventional estimates suggest.

A Log-Likelihood in the Presence of Unobserved Heterogeneity in the Frequency of Price Change

We assume that product i has a constant hazard of adjusting, f_i , in each month, where $f_i \sim \text{Beta}(a, b)$. Let us denote the product's lifetime by n_i . These assumptions imply that the total number of price changes in a product's lifetime is distributed according to the binomial distribution, $x_i \sim \text{Bin}(n_i, f_i)$. We assume, furthermore, that f_i is distributed according to the beta distribution, $f_i \sim \text{Beta}(a, b)$.

Given this model, we can write the likelihood of observing a product with length n_i and the total number of price changes x_i as,

$$L = \prod_{i=1}^I \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} f_i^{a-1} (1-f_i)^{b-1} \binom{n_i}{x_i} f_i^{x_i} (1-f_i)^{n_i-x_i} \quad (19)$$

$$= \prod_{i=1}^I \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} f_i^{x_i+a-1} (1-f_i)^{n_i-x_i+b-1} \binom{n_i}{x_i} \quad (20)$$

We can integrate out the f_i 's to get,

$$L = \prod_{i=1}^I \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \binom{n_i}{x_i} \frac{\Gamma(a+x_i)\Gamma(b+n_i-x_i)}{\Gamma(a+b+n_i)}. \quad (21)$$

The log-likelihood function is, therefore,

$$\log L = n \log \Gamma(a+b) - n \log \Gamma(a) - n \log \Gamma(b)^n + \sum_{i=1}^I [\log n_i! - \log x_i! \quad (22)$$

$$- \log(n_i - x_i)! + \log \Gamma(a + x_i) + \log \Gamma(b + n_i - x_i) - \log \Gamma(a + b + n_i)]. \quad (23)$$

B Product Replacement Bias in the VECM

Consider estimating equation (9) on a price index of all product lines with a frequency of price change f_j —which we refer to as sector j —and with time dummies for each quarter as the aggregate variables z . The coefficient on each dummy variable gives the average change in prices in sector j for that quarter. Based on equation (8), it is straightforward to show that the coefficients B_j equal Δp_{it}^j , the measured change in the price index for sector j . Equation (11) implies that these measured changes, Δp_{it}^j , are all biased by the factor $f_j/(f_j + z - f_j z)$. Since the price index p_{it}^j is constructed by stringing together these changes, the true price index for sector j is equal to the measured price index for sector j multiplied by $f_j/(f_j + z - f_j z)$. This implies that for each sector of the economy, the estimated cointegrating vector in the VECM is biased by the factor $f_j/(f_j + z - f_j z)$.

To calculate the overall bias in the estimated cointegrating vector, assumption that the cointegrating vectors across different sectors of the economy are the same. Without loss of generality, assume that the estimated cointegrating vector for sector j is $[1 \ \beta_j^{mm}]$. Define $\beta^{mm} = \int_{N_i} \beta_j^{mm} dj$ and notice that

$$p_{it} - \beta^{mm} q_t = \int_{N_i} (p_t^m - p_t) - \beta_j^{mm} q_t dj.$$

This implies that $p_{it} - \beta^{mm} q_t$ is stationary. The aggregate cointegrating vector is therefore $[1 \ \beta^{mm}]$. The ratio of the true long run relationship between relative import prices and the real exchange rate, β , and the measured value β^{mm} from the VECM is thus

$$\frac{\beta}{\beta^{mm}} = \left[\int \frac{f_j}{f_j + z - f_j z} dj \right]^{-1}. \quad (24)$$

References

- ABEL, A. B., AND O. J. BLANCHARD (1988): "Investment and Sales: Some Empirical Evidence," in *From Dynamic Econometric Modeling: Proceedings of the Third International Symposium in Economic Theory and Econometrics*, ed. by W. A. Barnett, E. Berndt, and H. White, pp. 269–296, New York. Cambridge University Press.
- ABRAMHAN, K. G., J. S. GREENLEES, AND B. R. MOULTON (1998): "Working to Improve the Consumer Price Index," *Journal of Economic Perspectives*, 12(1), 27–36.
- ALESSANDRIA, G. (2004): "International Deviations From the Law of One Price: The Role of Search Frictions and Market Share," *International Economic Review*, 45(4), 1263–1291.
- ALTERMAN, W. (1991): "Price Trends in U.S. Trade: New Data, New Insights," in *International Economic Transactions*, ed. by P. Hooper, and J. D. Richardson, pp. 109–143, Chicago. University of Chicago Press.
- ARMKNECHT, P. A., AND DONALD WEYBACK (1989): "Adjustments for Quality Change in the U.S. Consumer Price Index," *Journal of Official Statistics*, 5(2), 107–123.
- ATKESON, A., AND A. BURSTEIN (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, Forthcoming.
- AUER, R., AND T. CHANEY (forthcoming): "Exchange Rate Pass-Through in a Competitive Model of Pricing-to-Market," *Journal of Money, Credit and Banking*.
- BACKUS, D. K., P. J. KEHOE, AND F. E. KYDLAND (1992): "International Real Business Cycles," *Journal of Political Economy*, 100(4), 745–775.
- BARRO, R. (1977): "Long-Term Contracting, Sticky Prices, and Monetary Policy," *Journal of Monetary Economics*, 3, 305–316.
- BETTS, C., AND M. B. DEVEREUX (1996): "The Exchange Rate in a Model of Pricing to Market," 40(3-5), 1007–1021.
- (2000): "Exchange Rate Dynamics in a Model of Pricing to Market," *Journal of International Economics*, 50(1), 215–244.
- BILS, M. (2005): "Measuring Growth from Better and Better Goods," Working Paper, University of Rochester.
- BILS, M., AND P. J. KLENOW (2001): "Quantifying Quality Growth," *American Economic Review*, 91(4), 1006–1030.
- BOSKIN, M. J., E. R. DULLBERGER, R. J. GORDON, Z. GRILICHES, AND D. W. JORGENSON (1996): "Toward a More Accurate Measure of the Cost of Living," Final Report to the Senate Finance Committee.
- BOUAKEZ, H., AND N. REBEI (2008): "Has Exchange Rate Pass-Through Really Declined? Evidence from Canada," *Journal of International Economics*, Forthcoming.
- BRODA, C., AND D. E. WEINSTEIN (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121(4), 541–585.

- BURSTEIN, A., M. ECHENBAUM, AND S. REBELO (2005): “Large Devaluations and the Real Exchange Rate,” *Journal of Political Economy*, 113(4), 742–784.
- CAMPA, J. M., AND L. S. GOLDBERG (2005): “Exchange Rate Pass-Through into Import Prices,” *The Review of Economics and Statistics*, 87(4), 679–690.
- CAPLIN, A., AND D. SPULBER (1987): “Menu Costs and the Neutrality of Money,” *Quarterly Journal of Economics*, 102(4), 703–725.
- CARLTON, D. W. (1986): “The Rigidity of Prices,” *American Economic Review*, 76(4), 637–658.
- CHARI, V., P. J. KEHOE, AND E. R. MCGRATTAN (2002): “Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates,” *Review of Economic Studies*, 69, 533–563.
- CLAUSING, K. A. (2001): “The Behavior of Intrafirm Trade Prices in U.S. International Price Data,” BLS Working Paper.
- CORSETTI, G., AND L. DEDOLA (2005): “A Macroeconomic Model of International Price Discrimination,” *Journal of International Economics*, 67(1), 129–155.
- CORSETTI, G., L. DEDOLA, AND S. LEDUC (2008): “High Exchange Rate Volatility and Low Pass-Through,” *Journal of Monetary Economics*, Forthcoming.
- COURT, A. (1939): “Hedonic Price Indexes with Automobile Examples,” in *The Dynamics of Automobile Demand*, pp. 99–117, New York. General Motors Corporation.
- DORNBUSCH, R. (1987): “Exchange rates and prices,” *American Economic Review*, 77(1), 93–106.
- DROZD, L., AND J. NOSAL (2008): “Understanding International Prices: Customers as Capital,” Working Paper, Columbia University.
- ENGEL, C. (2002): “Expenditure Switching and Exchange-Rate Policy,” in *NBER Macroeconomics Annual*, ed. by M. Gertler, and K. S. Rogoff, pp. 231–272, Cambridge, Massachusetts. MIT Press.
- ERICKSON, T., AND A. PAKES (2008): “An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods,” Working Paper.
- FEENSTRA, R. C. (1994): “New product Varieties and the Measurement of International Prices,” *American Economic Review*, 84(1), 157–177.
- FEENSTRA, R. C., AND E. W. DIEWERT (2000): “Imputation and Price Indexes: Theory and Evidence from the International Price Program,” Working Paper, UC Davis.
- FITZGERALD, D., AND S. HALLER (2008): “Exchange Rates and Producer Prices: Evidence from Micro-Data,” Working Paper, Stanford University.
- GOLDBERG, P., A. KHANDELWAL, N. PAVCNIK, AND P. TOPALOVA (2008): “Imported Intermediate Inputs and Domestic Product Growth: Evidence from India,” NBER Working Paper No. 14416.
- GOLDBERG, P. K., AND M. KNETTER (1997): “Goods Prices and Exchange Rates: What Have We Learned?,” *Journal of Economic Literature*, 35(3), 1243–1272.

- GOLDBERG, P. K., AND F. VERBOVEN (2001): “The Evolution of Price Dispersion in the European Car Market,” *Review of Economic Studies*, 68, 811–848.
- GOLOSOV, M., AND R. E. LUCAS (2007): “Menu Costs and Phillips Curves,” *Journal of Political Economy*, 115, 171–199.
- GOPINATH, G., AND O. ITSKHOKI (2008): “Frequency of Price Adjustment and Pass-through,” Working Paper.
- GOPINATH, G., O. ITSKHOKI, AND R. RIGOBON (2008): “Currency Choice and Exchange Rate Pass-Through,” Working Paper.
- GOPINATH, G., AND R. RIGOBON (2008): “Sticky Borders,” *Quarterly Journal of Economics*, 123, 531–575.
- GRILICHES, Z. (1961): “Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change,” in *Price Statistics of the Federal Government*, pp. 173–196, New York. National Bureau of Economic Research.
- GRILICHES, Z., AND J. A. HAUSMAN (1986): “Errors in Variables in Panel Data,” *Journal of Econometrics*, 10, 93–118.
- GUST, C., S. LEDUC, AND R. J. VIGFUSSON (2006): “Trade Integration, Competition, and the Decline in Exchange Rate Pass-Through,” International Financial Discussion Paper No. 864, Board of Governors of the Federal Reserve System.
- HAUSMAN, J. (2003): “Sources of Bias and Solutions to Bias in the CPI,” *Journal of Political Economy*, 17, 23–44.
- HELLERSTEIN, R. (2008): “How Bears the Cost of a Change in the Exchange Rate: Pass-Through Accounting for the Case of Beer,” *Journal of International Economics*, forthcoming.
- HOBIIJN, B. (2003): “One Both Sides of the Quality Bias in Price Indexes,” .
- JOHANSEN, S. (1995): *Likelihood-Based Inference in Cointegrated Vector Auto-Regressive Models*. Oxford University Press, New York.
- KNETTER, M. M. (1989): “Price Discrimination by U.S. and German Exporters,” *American Economic Review*, 79(1), 198–210.
- KRUGMAN, P. (1987): “Pricing to Market When the Exchange Rate Changes,” in *Real-Financial Linkages among Open Economies*, ed. by S. Arndt, and J. D. Richardson, pp. 49–70. MIT Press, Cambridge.
- KRUGMAN, P. R. (1993): *Exchange Rate Instability*. MIT Press, Cambridge, Massachusetts.
- LIEGEY, P. R. (1993): “Adjusting Apparel Indexes in the CPI for Quality Differences,” in *Price Measurement and Their Uses*, ed. by M. Foss, M. Manser, and A. Young, pp. 209–226, Chicago. University of Chicago Press.
- MARAZZI, M., AND N. SHEETS (2007): “Declining exchange rate pass-through to U.S. import prices: The potential role of global factors,” *Journal of International Money and Finance*, 26, 924–947.

- MARSTON, R. C. (1990): “Pricing to Market in Japanese Manufacturing,” *Journal of International Economics*, 29, 217–236.
- MOULTON, B. R., AND K. E. MOSES (1997): “Addressing the Quality Change Issue in the Consumer Price Index,” *Brookings Papers on Economic Activity*, 1997(1), 305–349.
- NAKAMURA, E., AND J. STEINSSON (2007): “Monetary-Non-Neutrality in a Multi-Sector Menu Cost Model,” Working Paper, Columbia University.
- (2008a): “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, Forthcoming.
- (2008b): “Price Setting in Forward-Looking Customer Markets,” Working Paper, Columbia University.
- NAKAMURA, E., AND D. ZEROM (2008): “Accounting for Incomplete Pass-Through,” Working Paper, Columbia University.
- NEIMAN, B. (2008): “Multinationals, Intrafirm Trades and International Macro Dynamics,” Working Paper, Harvard University.
- NORDHAUS, W. D. (1998): “Quality Change in Price Indexes,” *Journal of Economic Perspectives*, 12(1), 59–68.
- OBSTFELD, M., AND K. ROGOFF (1995): “Exchange Rate Dynamics Redux,” *Journal of Political Economy*, 103(3), 624–660.
- OLIVEI, G. P. (2002): “Exchange Rates and the Prices of Manufacturing Products Imported into the United States,” *New England Economic Review*, 2002(1), 3–18.
- PAKES, A. (2003): “A Reconsideration of Hedonic Price Indexes with an Application to PC’s,” *American Economic Review*, 93, 1578–1596.
- RAVN, M., S. SCHMITT-GRHOE, AND M. URIBE (2007): “Incomplete Cost Pass-Through Under Deep Habits,” Working Paper, Columbia University.
- REINSDORF, M. B., P. LIEGEY, AND K. STEWART (1996): “New Ways of Handling Quality Change in the US Consumer Price Index,” Bureau of Labor Statistics Economic Working Paper No. 276.
- RODRIGUEZ LOPEZ, J. A. (2008): “Prices and Exchange Rates: A Theory of Disconnect,” Working Paper, University of California, Irvine.
- ROGERS, J. H. (2006): “Exchange Rate Pass-Through to U.S. Import Prices: A Look Under the Hood,” Staff Presentation, Board of Governors of the Federal Reserve System.
- ROTEMBERG, J. J. (2005): “Customer Anger at Price Increases, Changes in the Frequency of Price Adjustment and Monetary Policy,” *Journal of Monetary Economics*, 52(4), 829–852.
- STOCKMAN, A., AND L. TESAR (1995): “Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements,” *American Economic Review*, 85(1), 168–185.

THEIL, H. (1954): *Linear Aggregation of Economic Relations*. North Holland, Amsterdam.

TRIPLETT, J. E. (1997): "Measuring Consumption: The Post-1973 Slowdown and the Research Issues," *Federal Reserve Bank of St. Louis Review*, 79(3), 9–42.

TABLE I
Pricing to Market

	Imports	Exports
<u>Measured</u>		
VECM	0.41	0.87
Dynamic Adjustment	0.43	0.85
Levels	0.36	0.86
<u>Adjusting for Product Replacement Bias</u>		
No Heterogeneity in Comovement	0.68	0.79
With Heterogeneity in Comovement	0.65	0.80

Statistics based on aggregate import price index and trade-weighted exchange rate. Dynamic adjustment model has 6 lags. Adjusted statistics are based on the VECM model.

TABLE II
Pricing to Market over Subsamples

Period	VECM	Dynamic Adj.
1982-2008	0.41	0.43
1994-2008	0.46	0.32

Statistics based on calculations using aggregate import price index and trade-weighted exchange rate.

TABLE III
Number of Price Changes Per Product

Number of Price Changes	Imports		Exports	
	LCP	PCP	LCP	PCP
0	40.6	42.8	42.2	51.0
1 or less	57.2	65.1	61.6	68.8
2 or less	67.7	79.4	72.5	78.0
3 or less	74.6	86.0	79.3	84.6
4 or less	79.3	91.4	83.6	88.7
5 or less	82.3	94.4	86.3	94.2
10 or less	90.0	98.1	93.1	98.6
15 or less	93.1	99.4	95.4	98.9
20 or less	94.7	99.5	97.0	98.9

Statistics are weighted mean number of price changes per product, using as weights the cumulative product-level weights over each product's lifetime.

TABLE IV
The Distribution of Price Changes and Substitutions

	Imports		Exports	
	LCP	PCP	LCP	PCP
Fraction of Imports/Exports	0.927	0.073	0.024	0.976
Mean Frequency of Price Change	0.157	0.065	0.086	0.178
Median Frequency of Price Change	0.098	0.073	0.084	0.072
Mean Frequency of Substitutions	0.044	0.041	0.055	0.040
Distribution of the Frequency of Price Change				
a	0.50	1.07	0.763	0.53
b	3.65	19.12	7.762	4.60

Statistics are weighted means and medians calculated using IPP microdata on import and export prices. Weighted medians are calculated as medians across industry-level means. "a" and "b" are parameters in the estimated distribution of the frequency of price change, assumed to be Beta(a,b).

Table V
Sectoral Frequency of Price Change and Product Substitution

Industry Group	LCP Imports			PCP Exports		
	Freq. PC	Freq. Subs.	Weight	Freq. PC	Freq. Subs.	Weight
Animals & Animal Products	0.424	0.025	0.031	0.393	0.039	0.019
Vegetable Products	0.422	0.046	0.025	0.800	0.012	0.078
Foodstuffs	0.215	0.027	0.037	0.168	0.033	0.034
Mineral Products	0.163	0.076	0.010	0.174	0.032	0.019
Chemicals & Allied Industries	0.163	0.037	0.047	0.173	0.040	0.070
Plastics / Rubbers	0.170	0.042	0.033	0.153	0.036	0.043
Raw Hides, Skins, Leather, & Furs	0.056	0.039	0.019	0.385	0.030	0.006
Wood & Wood Products	0.306	0.036	0.064	0.224	0.035	0.049
Textiles	0.054	0.060	0.132	0.209	0.039	0.026
Footwear / Headgear	0.044	0.043	0.040	0.040	0.117	0.000
Stone / Glass	0.211	0.021	0.072	0.425	0.035	0.030
Metals	0.359	0.039	0.088	0.265	0.042	0.043
Machinery / Electrical	0.098	0.052	0.200	0.072	0.047	0.303
Transportation	0.089	0.045	0.107	0.062	0.036	0.179
Miscellaneous	0.048	0.041	0.094	0.048	0.061	0.101
Weighted Average	0.157	0.044	1.000	0.178	0.040	1.000
Weighted Median Across Industries	0.098	0.043	1.000	0.072	0.265	1.000

Statistics are based on IPP microdata on import and export prices. All statistics are weighted means using product level weights. "Freq. PC" is the frequency of price change, "Freq. Subs" is the frequency of substitutions and "Weight" is the total weight on the industry group.

TABLE VI
Product Replacement Bias

	LCP	PCP
No Heterogeneity in Comovement		
Imports	1.80	2.04
Exports	1.93	1.82
With Heterogeneity in Comovement		
Imports	1.70	2.03
Exports	1.88	1.75

TABLE VII
Pricing to Market and Product Replacement Bias: Robustness

	Major Country RER		Broad RER	
	All Subs	Forced Subs	All Subs	Forced Subs
Measured Pass-Through				
Imports		0.41		0.50
Exports		0.87		0.84
<i>A. No Heterogeneity in Comovement</i>				
Adjusted Pass-Through				
Imports	0.68	0.65	0.84	0.81
Exports	0.79	0.81	0.73	0.75
Factors				
LCP Imports	1.80	1.71	1.80	1.71
PCP Imports	2.04	1.72	2.04	1.72
LCP Exports	1.93	1.84	1.93	1.84
PCP Exports	1.82	1.70	1.82	1.70
<i>B. With Heterogeneity in Comovement</i>				
Adjusted Pass-Through				
Imports	0.65	0.63	0.80	0.77
Exports	0.80	0.81	0.74	0.76
Factors				
LCP Imports	1.70	1.63	1.70	1.63
PCP Imports	2.03	1.72	2.03	1.72
LCP Exports	1.88	1.80	1.88	1.80
PCP Exports	1.75	1.64	1.75	1.64

TABLE VIII
Results for Simulated Data with Timing Error

	LCP Imports	
	Real Data	Simulated Data
Change in Exchange Rate Since Last Price Change	0.26	0.25
Levels Conditional on a Price Change	0.54	0.54
Conditional on 2 Price Changes	0.24	0.25
Conditional on 8 or more Price Changes	0.67	0.70
Micro Dynamic Adjustment Equation	0.26	0.31

Statistics from "Real Data" based on calculations using IPP microdata on import prices denominated in dollars for high income OECD countries. Statistics from "Simulated Data" based on output from simulation model.




Figure II
U.S. Import Prices and the Real Exchange Rate




Figure III
U.S. Export Prices and the Real Exchange Rate




Figure IV
Cumulative Probability Distribution of Beta(0.50,3.65)

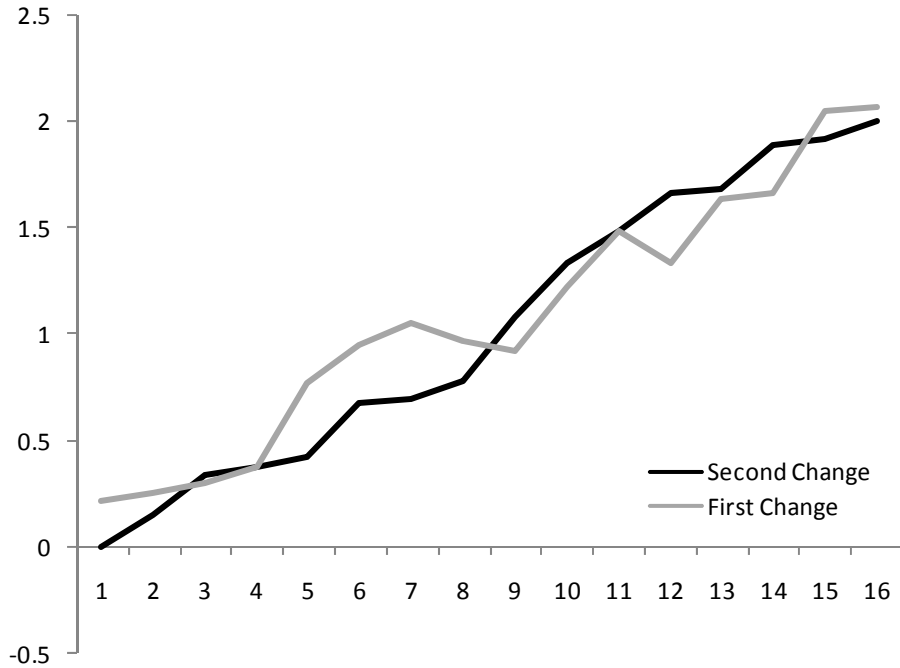


FIGURE V

Prices and Exchange Rates before Product Introduction

Note: For products with exactly two price changes, we run two regressions. First, we regress the size of the second price change on 18 monthly exchange rate changes before the beginning of the price spell that ends with the price change in question. Second, we regress the size of the first price change on 18 monthly exchange rate changes before the introduction of the product in question. The figure plots two lines. For the second price change, the j th point on the line is the sum of the coefficients on the first j lagged exchange rate changes. For the first price change, the j th point is the sum of the coefficients on the 3rd, 4th, ..., $(j+2)$ th exchange rate change (as though the first price spell had begun two periods before the product was introduced into the data set).