Copyright © 1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014 by the American Economic Association.

Permission to make digital or hard copies of part or all of American Economic Association publications for personal or classroom use is granted without fee provided that copies are not distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation, including the name of the author. Copyrights for components of this work owned by others than AEA must be honored. Abstracting with credit is permitted.

The author has the right to republish, post on servers, redistribute to lists and use any component of this work in other works. For others to do so requires prior specific permission and/or a fee. Permissions may be requested from the American Economic Association Administrative Office.

American Economic Review 2012, 102(7): 3277–3316
http://dx.doi.org/10.1257/aer.102.7.3277

Lost in Transit: Product Replacement Bias and Pricing to Market[†]

By Emi Nakamura and Jón Steinsson*

In the microdata underlying US trade price indexes, 40 percent of products are replaced before a single price change is observed and 70 percent are replaced after two price changes or fewer. A price index that focuses on price changes for identical items may, therefore, miss an important component of price adjustment occurring at the time of product replacements. We provide a model of this "product replacement bias" and quantify its importance using US data. Accounting for product replacement bias, long-run exchange rate "pass-through" is substantially higher than conventional estimates suggest, and the terms of trade are substantially more volatile. (*JEL* F14, F31)

Despite large swings in the US dollar nominal exchange rate, US import and export prices appear remarkably stable in US dollar terms. Conventional measures of aggregate exchange rate "pass-through" imply that a 1 percent depreciation of the dollar leads to roughly a 0.2–0.4 percent long-run increase in nonoil import prices, and roughly a 0.1 percent long-run fall in export prices in US dollar terms (Campa and Goldberg 2005; Marazzi and Sheets 2007; Gopinath, Itskhoki, and Rigobon 2010). As a consequence, the ratio of export to import prices—the terms of trade—are much less volatile than the exchange rate. Low pass-through of exchange rates into aggregate price indexes persists at long horizons, despite highly persistent exchange rate movements, implying that it cannot be explained as a mechanical consequence of temporarily rigid prices alone.[1]

We argue, however, that conventional measures of exchange rate pass-through based on aggregate price indexes are seriously biased—both in the short run and in the long run—due to two pervasive features of the underlying micro-data: highly

*Nakamura: Graduate School of Business, Columbia University, 3022 Broadway, New York, NY 10027 (enakamura@columbia.edu); Steinsson: Department of Economics, Columbia University, 420 W. 118 St., New York, NY 10027 (jsteinsson@columbia.edu). We thank Ryan Chahrour, Vickram Mohan, Dmitriy Sergeyev, and especially Patrick Sun for excellent research assistance. We thank Ariel Pakes and John Rogers for inspiring conversations that encouraged us to work on this problem. In working on this project, we have benefited greatly from conversations with, among others, Bill Alterman, Mark Bils, Ariel Burstein, Menzie Chinn, Charles Engel, Tim Erickson, Doireann Fitzgerald, Etienne Gagnon, Penny Goldberg, Gita Gopinath, John Greenlees, Oleg Itskhoki, Patrick Kehoe, Timothy Kehoe, Narayana Kocherlakota, Greg Kurtzon, Benjamin Mandel, Rob McClelland, Alice Nakamura, Jaromir Nosal, Marshall Reinsdorf, Rozi Ulics, Robert Vigfusson, and Pierre Yared as well as three anonymous referees, and seminar participants at various institutions. This research has been supported by the National Science Foundation grant SES 0922011.

[†] To view additional materials, visit the article page at http://dx.doi.org/10.1257/aer.102.7.3277.

[1] Exchange rate movements are likely endogenous to other macroeconomic shocks that affect prices directly. Our focus is on measuring the relationship between exchange rates and prices as opposed to providing evidence for a particular causal interpretation of this relationship.

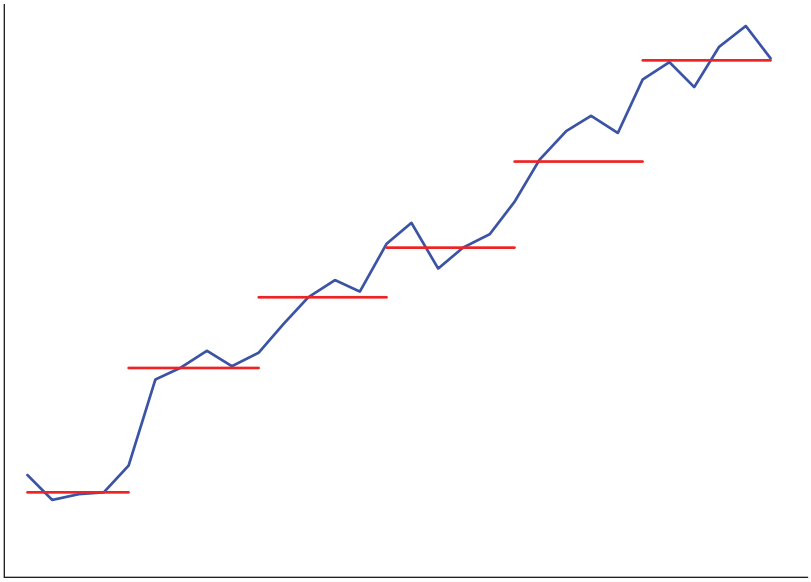

FIGURE 1. PRODUCT REPLACEMENT AND THE COMOVEMENT OF PRICES AND EXCHANGE RATES

rigid prices and frequent product replacements. In constructing price indexes, price changes that occur at the time of product replacements tend to be dropped. This causes a "product replacement bias" that leads aggregate import and export price indexes to be too smooth. Adjusting for product replacement bias raises aggregate pass-through and makes it more consistent with earlier work on pass-through for narrowly defined product categories (Goldberg and Knetter 1997). The coarser product definitions used in these earlier studies imply less product turnover and are thus less susceptible to product replacement bias.

To understand how product replacement bias arises, it is useful to consider an extreme example. Consider an economy 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 depreciating. 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 depreciates.

Now consider the problem of a statistical agency in this setting. Such an agency would ideally like to measure the change in quality adjusted prices. In practice, however, most price indexes (including the US import and export price indexes) are close approximations of a "matched-model index," in which all price changes used to construct the index are for identical items and product replacements are "linked into" the index. This means 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.[2] A matched model index will remain constant

[2] An alternative would be to make hedonic adjustments for quality change. For most products, however, it is extremely costly and difficult to accurately measure quality change (Abraham, Greenlees, and

throughout in our example since prices only change at the time of product replacements. Estimates of exchange rate pass-through using this price index will yield zero pass-through irrespective of what the true degree of pass-through is.[3]

While this is obviously an extreme example, it captures important features of the actual data underlying the US import and export price indexes. Price rigidity and frequent product turnover imply that about 40 percent of expenditure-weighted price series in these data have no price changes and roughly 70 percent have 2 price changes or less.[4] Even products that do have price changes while they are in the index typically exit the index after a prolonged spell of price rigidity. If the prices of new products entering the index have already adjusted to exchange rate movements over this interval (as in our simple example above), the response of these prices to movements in exchange rates over the last price spell of the exiting product will be "lost in transit" (i.e., neither picked up by an observed price change of the exiting nor the entering product). In this case, the price index will never fully reflect the true comovement of prices and exchange rates, even in the long run.

We develop a model of this "product replacement bias" and show how it depends on observable features of price data. We estimate our model using Bureau of Labor Statistics (BLS) micro-data on import and export prices. Our "corrected" measure implies that pass-through of changes in the trade-weighted US exchange rate into US import prices for the period 1982–2007 was 0.64—substantially higher than conventional measures of aggregate pass-through. For nonagricultural US exports, our correction implies that long-run pass-through was 0.79 for this period (in foreign currency terms)—a significantly lower number than conventional estimates yield. We also calculate a corrected series for the US terms of trade. The volatility of changes in this series is 75 percent higher than that of the official series. These estimates line up well with the implications of leading general equilibrium models such as Corsetti and Dedola (2005); Atkeson and Burstein (2008); and Drozd and Nosal (2012), which imply long-run pass-through between 0.7 and 0.9.[5]

We focus on how import prices respond to the trade-weighted exchange rate. If we were instead to study how import prices respond to changes in the bilateral exchange rates of the country of origin, our pass-through estimates would be dramatically lower. Gopinath and Itskhoki (2010b) find pass-through of bilateral exchange rates of roughly 0.1, suggesting that even our "adjusted" statistics would be less than 0.2.[6] The much larger response of prices to the trade-weighted exchange rate suggests

Moulton 1998). Feenstra (1994) and Broda and Weinstein (2006) develop and apply an alternative to hedonic methods. These papers apply a structural approach to back out quality-adjusted prices from data on prices and quantities.

[3] If the price comparisons that are dropped in this way are a representative sample of all price comparisons, the matched model index will provide an unbiased estimate of the true index. If the dropped price comparisons are in some way special, however, as in this example, product replacement bias will arise.

[4] Similar figures are reported in Gopinath and Rigobon (2008) and Gopinath, Itskhoki, and Rigobon (2010).

[5] Drozd and Nosal (2010) provide a detailed discussion of pass-through in a number of leading general equilibrium models. Most of these models have difficulty matching the low estimated pass-through for nonoil imports (though they can match the much higher pass-through estimates that arise when oil is included). The model of Erceg, Guerrieri, and Gust (2006) is designed to match the low nonoil pass-through estimates.

[6] This estimate of bilateral exchange rate pass-through is substantially lower than in Gopinath, Itskhoki, and Rigobon (2010) because the former paper uses the sample of all US trading partners, whereas the latter uses a sample of high-income Organisation for Economic Co-Operation and Development (OECD) countries selected as those with a substantial fraction of nondollar priced imports. Furthermore, the sample periods in Gopinath and Itskhoki (2010b) and Gopinath, Itskhoki, and Rigobon (2010) start in the mid-1990s—bilateral import price indexes cannot be constructed going further back—while our sample period starts in 1982.

that analyzing the price response to the bilateral exchange rate misses an important channel of adjustment. For example, the prices of British imports appear to respond not only to the sterling-dollar exchange rate, but also to the euro-dollar exchange rate, presumably due to the role of intermediate inputs and strategic complementarities in pricing versus competing European importers.

Our model of product replacement bias helps explain the wide range of pass-through estimates in the existing literature. Early work on pricing to market focused on industry studies of average prices for narrowly defined product categories (e.g., Knetter 1989; Gagnon and Knetter 1995). This literature typically found long-run pass-through of about 0.5 (Goldberg and Knetter 1997). An important drawback of this literature was that exchange rate changes potentially resulted in changes in both prices and average product quality within these categories. Recent work has addressed this concern by analyzing pass-through for exactly identical items over time. Gopinath, Itskhoki, and Rigobon (2010) (hereafter, GIR) present an extremely low aggregate pass-through estimate based on this approach. They, however, caution that this estimate should be interpreted with care given the large number of price series in the data that have no price changes. Building on this idea, we show that analyzing price responses for identical items has the downside that it is more likely to be sensitive to product replacement bias. The more one disaggregates the data, the shorter the horizon over which a "product" can be followed. This makes it increasingly likely that one's measure of the response to a shock over a product's lifetime reflects the short-run as opposed to the long-run response, and that the remainder of the response will be "lost in transit." The growing richness of product-level data—which identify not only individual products at an extremely disaggregated level but also individual firms—raises the importance of dealing with this measurement problem.

Two key inputs into our adjustment for product replacement bias are the frequency of price change and the frequency of product replacement. Product replacement bias disappears entirely if *either* prices are completely flexible or products last forever in the index. The magnitude of the bias also depends crucially on the degree of heterogeneity in the frequency of price change. Our empirical model allows for a highly flexible distribution of cross-sectional heterogeneity in price rigidity (at the level of individual products). This flexible specification is crucial in accounting for the large number of products with no observed price changes. Incorporating this heterogeneity amplifies product replacement bias substantially relative to a case with no variation in price rigidity across firms.

Another key input into our adjustment for product replacement bias is the degree of "overreaction" of the first observed price change of each product to past exchange rate changes. The simple example depicted in Figure 1 assumes that products enter the dataset with prices that have already adjusted to past exchange rate changes. This is not necessarily the case in practice—i.e., products may enter the index with "stale" prices. In this case, when these products do adjust their prices, they will "overreact" to historical exchange rate movements, potentially making up for the price adjustment that would otherwise have been "lost in transit." We provide direct empirical evidence on the magnitude of such "overreaction" by comparing the responsiveness of first versus second price changes to historical exchange rate movements.

Our estimates indicate that such overreaction is minimal. One potential reason for this finding is that contracts are likely to be renegotiated when firms start buying or selling a product (Carlton 1986). This is often when products enter the BLS dataset. In addition, the nature of the BLS repricing procedure makes products more likely to enter the dataset with disproportionately "fresh" prices. While initial prices are collected using a detailed interview, subsequent prices are collected using a "repricing form" asking firms to confirm a previously reported price (and providing them with their previous price). BLS internal studies suggest that the repricing procedure sometimes yields spurious rigidity for continuing products, a problem that does not arise for products newly initiated to the dataset. Finally, firms may choose to adjust their price more when they introduce new products because they perceive this as being less likely to antagonize their customers (Rotemberg 2005; Nakamura and Steinsson 2011).

Our results confirm and reinforce alternative micro-based estimates of long-run pass-through proposed by GIR. In addition to conventional aggregate pass-through, GIR propose a micro-based "lifelong" pass-through measure. They show that this approach yields an estimate of long-run pass-through of 0.49, almost twice as high as their estimates based on aggregate statistics. Lifelong pass-through is less sensitive to product replacement bias than aggregate pass-through for reasonable parameter values. Indeed, our adjusted long-run pass-through measure reduces to GIR's lifelong pass-through measure under certain conditions.

It is important to emphasize that neither of these estimates of long-run pass-through is inconsistent with the low estimate of pass-through "conditional on adjustment" reported in GIR (0.24). Pass-through "conditional on adjustment" measures short-run rather than long-run pass-through since a substantial fraction of pass-through is delayed beyond the first price change. The literature on pass-through has long recognized that there is a large difference between short-run and long-run pass-through (see, e.g., Gagnon and Knetter 1995; Campa and Goldberg 2005). Standard macroeconomic models generate such delays in pass-through (even beyond the duration of rigid prices) by incorporating staggered price adjustment and strategic complementarity in pricing.[7]

Product replacement bias causes a downward bias in pass-through when prices are reported in the buyer's currency (i.e., local currency priced (LCP)) but an upward bias in pass-through when prices are reported in the producer's currency priced (PCP).[8] This helps explain several empirical findings. First, GIR document a large difference in long-run aggregate pass-through for dollar prices (LCP) versus nondollar-priced (PCP) US imports, but a much smaller difference based on their alternative micro-based "lifelong" measure of long-run pass-through—which is less sensitive to product replacement bias. Second, conventional estimates of US import price pass-through are much lower than those of US price export pass-through (in foreign currency terms). Most US imports are LCP (and thus downward-biased),

[7] These two features together imply that any given firm will not fully adjust to past shocks when it changes its price because other prices are rigid and it wishes to avoid setting its price too far out of line with its competitors.

[8] To see why, notice that in the extreme example above, if we had assumed the price was reported in producer currency terms (instead of local currency terms), the price in local currency would appear to move one-for-one with the exchange rate irrespective of true pass-through.

while most US exports are PCP (and thus upward-biased). Third, measured exchange rate pass-through for imports is lower in the United States, where most imports are LCP, than in developing countries, where imports are more often PCP (see, e.g., Burstein, Eichenbaum, and Rebelo 2005).

While we have focused on the implications of product replacement bias for the effects of exchange rate movements, it may be important in other contexts as well. Diewert and Nakamura (2010) and Houseman et al. (2011) suggest that problems in measuring price changes at the time of sourcing changes from high-cost to low-cost countries (which often coincide with product replacements) may lead to mismeasurement of the effects of offshoring on the US economy. Product replacement bias may contribute to these problems (Houseman 2007; Mandel 2007, 2009). Product replacement bias may also lead measured *consumer* price inflation to respond less to aggregate shocks than true consumer price inflation. During the dramatic rise of inflation in the United States in the late 1970s, the BLS noticed that prices began rising rapidly in almost all sectors except apparel, which continued to show inflation rates near zero. A likely cause was that frequent product turnover accompanied by clearance sales was leading to substantial product replacement bias. In response, the BLS changed their procedure for linking in products in apparel (Reinsdorf, Liegey, and Stewart 1996).[9] Product replacement bias may limit the extent to which the "extensive margin" forces emphasized by Auer and Chaney (2009); Rodriguez-Lopez (2008); and Ghironi and Melitz (2005) show up in the data, in practice. Finally, mismeasurement of import and export price indexes also affects measured trade volumes and trade price elasticities.[10]

Product replacement bias is related to, but different from, "new goods" and "quality change" bias.[11] Product replacement bias addresses a bias in the *responsiveness* of inflation to aggregate shocks in the presence of price rigidity, whereas the new goods and quality change literature addresses a bias in the average level of inflation associated with systematic declines in the quality-adjusted price at the time of new product introductions.

Our analysis builds on several recent papers studying the micro-level US import and export price data beyond those we have already mentioned. Rogers (2006) argues that there is a negative relationship across industries between the frequency of product substitutions and exchange rate pass-through. Berger et al. (2012) study the relationship across products between product substitutions and distribution wedges. Neiman (2010) and Clausing (2001) provide additional evidence on the nature and reasons for price rigidity in the trade price data.

[9] Another indication that product replacement bias may affect the consumer price index is that price changes are disproportionately large at the time when new products are linked into price indexes (Armknecht and Weyback 1989; Liegey 1993; Reinsdorf, Liegey, and Stewart 1996; Triplett 1997; Greenlees and McClelland 2011).

[10] 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. This will cause estimates of trade price elasticities to be biased away from one.

[11] Important papers on "new goods" and "quality change" bias include Court (1939); Griliches (1961); Nordhaus (1998); Bils and Klenow (2001); Hausman (2003); Pakes (2003); Boskin et al. (1996); Bils (2009); Moulton and Moses (1997); Abraham, Greenlees, and Moulton (1998); Triplett (1997); and Hobijn (2002). Erickson and Pakes (2011) develop an experimental hedonic price index for televisions that accounts, among other things, for price rigidity. Goldberg et al. (2010) show that new imported varieties contributed substantially to effective price declines for Indian firms after a trade liberalization. Reinsdorf (1993) studies the related idea of "outlet substitution bias."

The paper proceeds as follows. Section I describes the BLS microdata underlying the US import and export price indexes that we use in our empirical analysis. Section II presents measures of pass-through for US imports and exports for the period 1982–2007 based on conventional methods using aggregate data. Section III derives expressions for product replacement bias as a function of the frequency of price change and the frequency of product replacement. Section IV presents estimates of the frequency of price change, the frequency of product replacement, the degree of overreaction of first observed price changes, and uses these to calculate our estimate of product replacement bias. Section V analyzes how product replacement bias relates to alternative estimates of exchange rate pass-through in the empirical literature. Section VI concludes.

I. Data Description

We use three main sources of data. First, we use the microdata underlying the US import and export price indexes. These data are collected by the International Prices Program (IPP) of the BLS. Second, we use aggregate US import and export price indexes 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 Federal Reserve Board (FRB) and the International Monetary Fund (IMF). We describe these data in turn.

The US import and export price indexes were introduced in the early 1980s to provide a more accurate alternative to unit value indexes. The micro data we use cover the time period 1994–2004 in the case of frequency statistics and 1994–2007 in the case of pass-through regressions.[12] We exclude intrafirm prices. The total number of product-months in our sample for which IPP attempts to record a price is roughly 1.6 million or about 150,000 per year. This dataset has previously been used by Clausing (2001); Gopinath and Rigobon (2008); Gopinath, Itskhoki, and Rigobon (2010); Gopinath and Itskhoki (2010 a, b); Berger et al. (2012); and Neiman (2010). Below, we provide a brief description of how these data are collected. See the IPP *Data Collection Manual* for a much more detailed description (US Department of Labor 2005).

The IPP data are collected using voluntary surveys. To initiate a product into the dataset, IPP collects a detailed item description and a particular set of transaction terms. Transaction terms may 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 during initialization is required to be a transaction price (rather than a list price), unless a discount structure can also be provided to adjust the list price. After initialization, subsequent prices are collected using a repricing form with prefilled information about the last reported price and product characteristics. One concern about the repricing form is that it creates a differential reporting friction for reporting a new price. We return to this issue in Section IV.[13]

[12] The version of the data we used had gaps in 2005 (missing months) as well as an incomplete year in 2007. For all of the pass-through calculations (Tables 5 and 7) we use all available time periods, but for the frequency calculations, we restrict attention to the data before 2005 to avoid biases associated with seasonality.

[13] To update the discount structure, the firm must cross out the existing discount information and pencil in new information. In practice, a discount structure is rarely reported and almost never changed.

Transaction prices are missing in about 40 percent of the product-months in the IPP dataset. This arises because of a combination of infrequent trade, "off-cycle" price collection carried out by the BLS to reduce reporting burden, and failure of reporters to return the repricing forms.[14] In the vast majority of cases, prices are missing for short periods (one to two months). During these periods, IPP imputes prices using a variety of methods (see Feenstra and Diewert 2000 for a detailed discussion). In many cases, IPP simply "pulls forward" the last available price through missing periods. Another standard procedure is "cell mean imputation" whereby IPP imputes prices over the intervening months using the average inflation rate in the category. In all cases, when the price is again observed, the price series reverts back to the actual observed price. Therefore, the choice of imputation method has no effect on the trajectory of prices beyond the horizon over which prices are imputed and thus no impact on long-run pass-through. We discuss this in more detail in online Appendix E.

The IPP accepts reported prices in any currency, but in practice about 92 percent of import prices and 97 percent of export prices are reported in US dollars.[15] 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.

We make use of detailed time-varying BLS product weights. Within product groups, the IPP sampling procedure is to sample product-firm pairs in proportion to their dollar sales from census data. This implies that the effects of product-firm size are accounted for by the sampling procedure itself. We account for the fact that different product groups have different fractions of market- and nonmarket-based items by allocating weight to market versus nonmarket goods within a classification group in proportion to the number of "net-price" observations in each category. We normalize the total weight allocated to a year to be equal across years. Since 1997, the BLS has also used additional weights to account for "sampling bias"—random deviations between the theoretical and actual sampling probabilities that arise in a finite sample and are uncorrelated with product-firm size. BLS studies confirm that these additional weights have little impact on the BLS index.[16]

The second set of data we use is from the US NIPA. We use the import price deflator for imported goods excluding oil. We use the export price deflator for exported goods excluding agricultural products. Finally, we make use of trade-weighted and bilateral, monthly, and daily exchange rates downloaded from the FRB's website,[17] as well as monthly bilateral exchange rates from the International Financial Statistics (IFS) database of the IMF.

[14] In the repricing process, the reporter is allowed to report an estimated or "list" price if there was no transaction or a transaction price is not available. We drop these prices from our analysis.

[15] This is a substantially higher fraction of dollar priced goods than in GIR. The main reason for the difference is that GIR condition explicitly on countries with a substantial fraction of nondollar goods.

[16] For estimates of the "lifelong" pass-through regression and the pass-through "conditional on adjustment" regression presented in Section V, we reweight the data in such a way as to avoid "overweighting" product categories with a high frequency of price change. Specifically, we reweight the observations such that the total weight within a given HS2 category in the new sample (including only price changes) is the same as in the original sample of all observations.

[17] http://www.federalreserve.gov/releases/h10/Hist/.

II. Prices and Exchange Rates: Evidence

The conventional approach to measuring long-run exchange rate pass-through for imports is to run the following regression:

$$\Delta(p_t^m - p_t) = \alpha + \sum_{k=0}^{6} \beta_k \Delta q_{t-k} + \epsilon_t, \tag{1}$$

where p_t^m denotes the log of the dollar price of US imports, p_t denotes the log of the dollar price of US production, and q_t denotes the log of the trade-weighted US real exchange rate. Long-run pass-through is then defined as the sum of the coefficients, $B = \sum_{k=0}^{6} \beta_k$. If $B < 1$, long-run pass-through is said to be incomplete.[18] Recent papers that use this type of regression to estimate long-run pass-through include Campa and Goldberg (2005); Marazzi and Sheets (2007); and GIR. We follow GIR in referring to estimates of pass-through from this type of regression as "aggregate" pass-through estimates.

A concern with the aggregate pass-through specification is that it is misspecified if relative import prices and the real exchange rate are cointegrated. To allow for cointegration we also consider the following vector error correction model (VECM):

$$\Delta \mathbf{y}_t = \Pi(\mathbf{A}\mathbf{y}_{t-1} + \boldsymbol{\alpha} + \gamma t) + \sum_{k=1}^{n-1} \boldsymbol{\Gamma}_k \Delta \mathbf{y}_{t-k} + \boldsymbol{\delta} + \boldsymbol{\epsilon}_t, \tag{2}$$

where $\mathbf{y}_t$ is the vector $(p_t^m - p_t, q_t)$, and A is the vector of coefficients in the cointegrating relationship given by $[1 \; -\beta]$. In this case, long-run pass-through is given by the parameter β. We find strong evidence of a cointegrating relationship between real import prices and the real exchange rate.[19]

To measure long-run pass-through for US exports, we use an aggregate pass-through regression and VECM analogous to equations (1) and (2).[20] To get a pass-through measure for exports that is comparable to the measure we use for imports, we adopt the viewpoint of foreign consumers. Our aggregate pass-through regression for exports thus regresses the foreign currency price of US exports relative to the foreign currency price of foreign production on current and past values of the real exchange rate.

We use the NIPA price deflator for nonoil goods imports and nonagricultural goods exports. Our sample period is from 1982 through 2007. We begin our

[18] It is common to interpret incomplete pass-through as evidence of "pricing to market." There are a number of other potential reasons for empirical estimates of $\beta < 1$, however, as emphasized, for example, by Goldberg and Knetter (1997). Knetter (1989) and Marston (1990) pioneered empirical estimation of pricing-to-market using micro-data, which allows one to control for some other reasons for $\beta < 1$ in ways that are not possible using aggregate data. One potentially important reason why $\beta < 1$ is that imports into the United States contain components produced in the United States and exported. In addition, general equilibrium effects can lead to incomplete pass-through even when there is no pricing-to-market (see, e.g., Corsetti, Dedola, and Leduc 2008, and Bouakez and Rebei 2008).

[19] We reject the 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 $n = 2$.

[20] For exports, we again reject the hypothesis of no cointegrating equations. 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.

Table 1—Comovement of Prices and Exchange Rates

	Imports	Exports
Measured:		
Aggregate	0.43	0.85
	(0.05)	(0.05)
VECM	0.41	0.87
	(0.05)	(0.06)
Adjusting for product replacement bias:		
Benchmark	0.64	0.79
Forced substitutions including refusals	0.69	0.76
All substitutions	0.74	0.74

Notes: The top panel presents alternative measures of the long-run relationship between the trade-weighted real exchange rate and aggregate import or export price indexes (standard errors in parentheses). "Aggregate" reports the sum of the coefficients on lagged exchange rate changes with six quarterly lags (equation (1)). "VECM" reports the estimated coefficient on exchange rates in the cointegrating relationship between prices and exchange rates (equation (2)). The bottom panel presents the estimated relationship based on the aggregate pass-through regression adjusted for "product replacement bias" according to the methods discussed in the paper, under alternative assumptions about heterogeneity in pricing behavior across products and the frequency of product replacement.

sample in 1982 because this is when the import and export price indexes were introduced in the United States. The exchange rate variable we use is the FRB's trade-weighted real exchange rate index for major currencies.[21] We use consumer prices as our proxies for the prices of overall US and foreign production.

Results for the four pass-through regressions discussed above are presented in Table 1. We find that the aggregate pass-through regression and the VECM yield similar estimates of long-run pass-through. For imports, the aggregate pass-through equation yields an estimate of 0.43, while the VECM yields an estimate of 0.41. These estimates are broadly in line with the existing literature on exchange rate pass-through. For example, Campa and Goldberg (2005) estimate long-run pass-through for US imports to be 0.42 for the period 1975 to 2003. For export prices, the aggregate pass-through equation yields an estimate of long-run pass-through 0.85, while the VECM yields an estimate of 0.87.

Figures 2 and 3 display the stability of the relationship documented above. Figure 2 plots the relative dollar price of US imports $p_t^m - p_t$ and the fitted values based on the cointegrating relationship, $\hat{\beta} q_t$. The two series are normalized to have the same means and detrended. Figure 3 plots analogous series for the case of exports. Over this time period, these relationships—both for imports and exports—have been quite stable aside from an apparent one-time upward level shift in the price of exports after 2003.

[21] We focus on pass-through for the major currency exchange rate series rather than the broader index for two reasons. First, the weights in the import and export price index are often three to five 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 currency exchange rate index may better reflect the exchange rates for nonoil importers. Our adjustment factor, however, applies to pass-through regressions using any exchange rate measure. We report results for both the major currency and broad exchange rate indexes in Table 6.

FIGURE 2. US IMPORT PRICES AND THE REAL EXCHANGE RATE

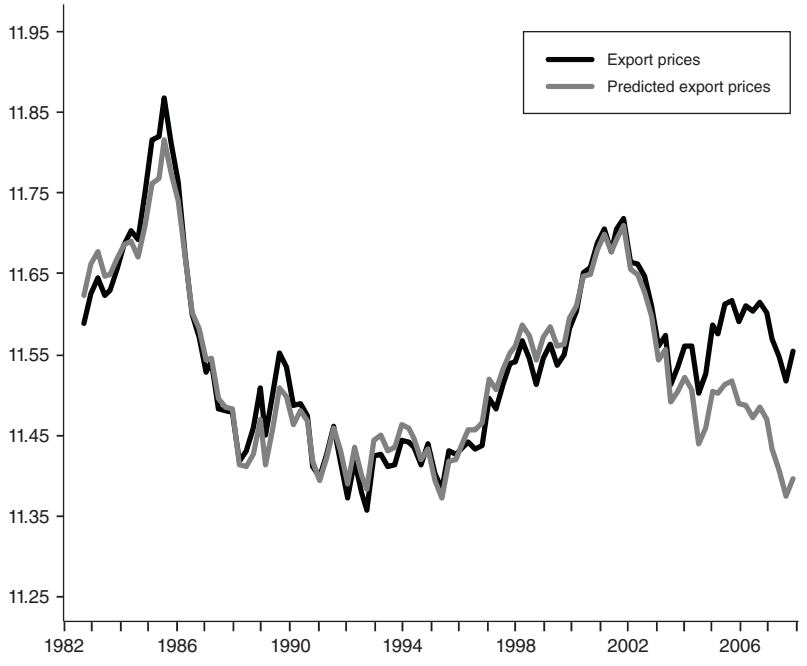

FIGURE 3. US EXPORT PRICES AND THE REAL EXCHANGE RATE

Table 2—Comovement of Prices and Exchange Rates over Subsamples

Period	Aggregate	VECM
1982–2008	0.43	0.41
	(0.05)	(0.05)
1994–2008	0.32	0.46
	(0.08)	(0.08)

Notes: The table presents alternative measures of the long-run relationship between trade-weighted exchange rate series and aggregate import or export price indexes for the time periods 1982–2008 and 1994–2008 (standard errors in parentheses). "Aggregate" reports the sum of the coefficients on lagged exchange rate changes with six quarterly lags (equation (1)). "VECM" reports the estimated coefficient on exchange rates in the estimated cointegrating relationship between prices and exchange rates (equation (2)).

Several researchers have argued that exchange rate pass-through into US imports has fallen in recent years (e.g., Olivei 2002; Marazzi and Sheets 2007). Table 2 reports pass-through estimates for US imports from both the aggregate pass-through regression and the VECM for two subsamples: 1982–2008 and 1994–2008. Aggregate 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. The results for the VECM, however, suggest that this apparent fall in pass-through might partly be due to model misspecification. For the VECM, the long-run pass-though estimate is slightly higher for the recent sample than it is for the longer sample—0.46 versus 0.41. Marazzi et al. (2005) and Hellerstein, Daly, and Marsh (2006) show that pass-through estimates for this period are sensitive to whether commodity prices are included in the regression as a separate regressor.

III. Prices and Exchange Rates: Theory

Consider an economy in which consumers purchase and consume a continuum of products, some of which are domestically produced and some imported. In each period, a fraction of both home and imported products are discontinued and an equal number of home and imported products are introduced. Some of the new products are new versions of the older products. Other new products are unrelated to the products that are discontinued that period.

Let C_{jit} denote the number of units of product j produced in region i and consumed at time t and P_{jit} denote the price per unit of these products. Let γ_{jit} denote the quality of each of these units measured in terms of utility. Furthermore, define $\hat{C}_{jit} = \gamma_{jit} C_{jit}$ to be the effective consumption of product j at time t and let $\hat{P}_{jit} = \gamma_{jit}^{-1} P_{jit}$ be its effective price. Notice that we allow product quality for the product indexed j to change over time. Below we equate changes in quality with product turnover.

Products from region i enter the consumer's utility function through the consumption aggregator

$$C_{it} = C(\hat{C}_{1it}, \ldots, \hat{C}_{jit}, \ldots, \hat{C}_{J_i it}).$$

The only two restrictions that we place on the function C is that it be homothetic and that C_{jit} and γ_{jit} enter through $\hat{C}_{jit}$. These restrictions are enough to guarantee the

existence of a price index—a function that gives the minimum cost of an additional unit of utility—written in terms of effective prices:

$$P_{it} = P(\hat{P}_{1it}, \ldots, \hat{P}_{jit}, \ldots, \hat{P}_{J_i it}). \tag{3}$$

A leading example of a consumption aggregator and corresponding price index that satisfies these assumptions is the "Dixit-Stiglitz" constant elasticity of substitution (CES) index:

$$C_{it} = \left[\int_{J_i} \hat{C}_{jit}^{\frac{\theta-1}{\theta}} \, dF(j) \right]^{\frac{\theta}{\theta-1}} \text{ and } P_{it} = \left[\int_{J_i} \hat{P}_{jit}^{1-\theta} \, dF(j) \right]^{\frac{1}{1-\theta}}.$$

Other examples include the nested CES preferences (e.g., Atkeson and Burstein 2008) and translog preferences (e.g., Kimball 1995; Bergin and Feenstra 2001).

On the firm side of the economy, firm j in region i produces products according to the production function

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

Quality enters the production function multiplicatively. This implies that to raise its output by a factor two is equally costly whether the firm chooses to do this by raising the number of units is produces by a factor of two or the quality of each unit by a factor of two.

Most models in macroeconomics and international economics abstract from product turnover. These models can be easily mapped into our framework, however, simply by viewing the quantities in those models as effective quantities and the prices as effective prices. The only restriction imposed by our framework is the requirement of a homothetic utility function. The additional complications we introduce in our framework beyond what is included in standard models arise from the fact that we wish to consider a situation where product quality is not observed, and the government therefore must calculate price indexes with raw prices and quantities as opposed to direct observation of effective prices and quantities.

Prices in the economy are sticky in local currency. Furthermore, the degree of price rigidity varies across firms. Let k index all firms that adjust their prices with probability $f_k(s)$ when the economy is in state s. Let P_{kt}^* denote the optimal price set by firms in product group k that adjust in period t. We assume that P_{kt}^* is independent of the time of the last price change of the product as well as the time the product was introduced into the economy. This means we rule out cases where the optimal degree of pass-through depends on the time since a product was introduced.[22] We make no assumption about the price firms set when they introduce products. Any

[22] For example, while it would not affect our results if new products were systematically introduced at low introductory prices (since this is uncorrelated with the exchange rate), we rule out the possibility that optimal exchange rate pass-through is systematically lower for new products than continuing products. Similarly, our bias calculation does not include biases that arise from a systematic tendency of new products to enter the index with high (or low) markups depending on the exchange rate (see Auer and Chaney 2009 and Rodriguez-Lopez 2008). Such features could either exacerbate or reduce the biases associated with product replacement bias.

difference between this price and P^*_{kt} is explicitly accounted for in our bias adjustment (by the α introduced below). We make no assumption about the dynamics of P^*_{kt} relative to the exchange rate. If strategic complementarity is important in price setting, movements in P^*_{kt} may lag movements in the exchange rate significantly.

Let $z(s)$ denote the fraction of products that are discontinued in state s. As noted above, we assume that an equal number of new products are introduced as are discontinued in each period. We therefore refer to $z(s)$ as the rate of product replacement. We model a product replacement as a change in γ_{jit}. For simplicity, we assume that each time a product is replaced, a new γ_{jit} is drawn from a distribution $\Gamma_t(s)$ and this level of quality remains constant for product j until it undergoes another replacement. In other words,

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

The distribution of product quality, $\Gamma(s)$, has no impact on our results since firm profits and consumer welfare depend only on quality-adjusted prices and since inflation depends on price relatives $(\hat{P}_{jt} / \hat{P}_{j,t-1})$ for which quality drops out.

A. *Price Measurement*

The government seeks to measure the changes in the price of imports over time. In practice, the government does not observe the functional form of C. We assume that the government is willing to make do with a first-order Taylor-series approximation to the logarithm of the ideal price index taken around the steady-state expenditure shares for each good. Written in terms of the change in the price index, this yields

$$\Delta p_{it} = \int_{J_i} \Delta \hat{p}_{jit} \, dF(j), \tag{4}$$

where lowercase variables denote logarithms of uppercase variables and $F(j)$ denotes the steady-state expenditure share of product j.[23] Notice that this price index depends on the quality-adjusted prices $\hat{} \, p_{jit}$. We will refer to it as measuring the "true" change in prices.

A major practical problem facing the government as it seeks to construct a price index is the fact that product quality γ_{ijt} is difficult to measure. The ideal solution would be to use hedonic methods to estimate product quality. Such methods are rarely used in practice, however, because they are extremely costly and difficult to apply in most cases. In practice, price comparisons that involve a change in quality—i.e., product replacements—are usually dropped from the index. Indexes constructed in this way are referred to as "matched model indexes." It is also prohibitively costly to measure the prices of all imports. The government therefore collects

[23] This price index is closely related to the fixed-weight price indexes used by statistical agencies.

a sample of import prices each period to construct a price index. In our notation, a matched model index based on a sample of products is

$$\Delta p_{it}^{mm} = \int_{N_i} \Delta p_{jit}\, dF(j), \tag{5}$$

where N_i denotes the sample from J_i for which raw price changes can be constructed. The raw price change cannot be constructed in the period in which the product is introduced into the government's sample since the difference in quality between the new product and the old product exiting the sample is unknown.

B. *Product Replacement Bias: A Factor Calculation*

Consider the following regression of the change in the true import price index on a vector of current, lagged (and possibly future) changes in the exchange rate denoted by $\mathbf{\Lambda}_t$,

$$\Delta p_{it} = \alpha + \mathbf{B}\mathbf{\Lambda}_t + \epsilon_t, \tag{6}$$

where $\mathbf{B}$ is a vector of coefficients. The vector $\mathbf{\Lambda}_t$ is meant to include all exchange rate changes that are correlated with Δp_{it}. In practice, this includes the current and a number of lagged exchange rate changes. Our primary interest is long-run pass-through or $\sum_n B_n$, where B_n denotes the nth element of $\mathbf{B}$. Given equation (4), it is straightforward to show (see online Appendix A1 for details) that the vector of regression coefficients for this regression, $\mathbf{B}$, may be "decomposed" as

$$\mathbf{B} = \int_S \int_{J_i} \mathbf{B}_k(s)\, dF_s(k) ds, \tag{7}$$

where $F_s(k)$ denotes the distribution of the frequency of price change in state $s \in S$ and $\mathbf{B}_k(s)$ is calculated as follows. Consider constructing a subprice index, $\Delta p_{ikt}(s)$, that consists of all price change observations for products in product group k and state s and regressing this subprice index on $\mathbf{\Lambda}_t$; $\mathbf{B}_k(s)$ denotes the resulting vector of coefficients. Equation (7) allows us to analyze each product group k and each state of the world s separately and then take an average over products and states.[24] An analogous decomposition is possible if the dependent variable in regression (6) is the change in the price index collected by the government, Δp_{it}^{mm}. Denote the regression coefficients from such regressions analogously by $\mathbf{B}^{mm}$ and $\mathbf{B}_k^{mm}(s)$.

In the absence of product replacements, changes in the optimal price—the price $p_{ikt}^*(s)$ that firms adjust to when they change their prices—over a particular span of time are always eventually incorporated into the government's price index and estimates of the pass-through regression (6) through the product's next price change. Thus, in this case, all movements in the optimal price are eventually "accounted for."

[24] Here we must assume that for each state of the world we have data from enough time periods that $\mathbf{B}_k(s)$ is identified. This implies that the state space for the frequency of price change and the frequency of substitutions must be somewhat "coarser" than the state space for the aggregate variable $\mathbf{\Lambda}_t$.

The presence of product replacements complicates matters by potentially leaving some time periods "unaccounted for" even in the long run.

To build intuition, consider a simple case in which the government's dataset consists of a sample of "product lines." Product replacement arises in this case because one model of a product is replaced by a new model. Suppose for simplicity that the new model is introduced with a new price equal to the optimal price $p^*_{ikt}(s)$. In this case, changes in the optimal price that occur following the last observed price change of the old model get incorporated into the price series by the price change that coincides with the introduction of the new model into the dataset. In the case of a matched model index, however, this price change is dropped (due to the difficulty of measuring quality changes). This implies that changes in the optimal price that occurred after the old model's last measured price change but before the product exits the dataset are never "accounted for" and thus never incorporated into the government's price index or the pass-through regression (6). The fact that product replacements leave a fraction of time periods in the governments' dataset "unaccounted for" in this way is what gives rise to product replacement bias.

In this product line example, it is simple to calculate the fraction of time "accounted for" in the matched model index. All time periods belonging to price spells that end with a price change that does not coincide with a product replacement are accounted for. All time periods that belong to price spells that end with a product replacement are not. The fraction of time accounted for in the matched model index is thus equal to the fraction of price spells that do not end with a product replacement adjusted for the fact that the time periods in which product replacements occur are dropped from the index. This fraction is

$$\frac{f_k(s)}{f_k(s) + z(s) - f_k(s)z(s)}. \tag{8}$$

Since this is the fraction of time periods accounted for in the government's matched model price index, this is also the fraction of exchange rate movements that are captured in the pass-through regression (6) when the government's matched model price index is used as the dependent variable. True pass-through in product group k and state s is then understated because of product replacement bias by the factor in equation (8).[25] Notice that the assumption that gives rise to product replacement bias in this simple case is the assumption that periods of product replacement are "special" in that all new models are introduced with a "fresh" price—i.e., product replacements represent a "free" opportunity to change prices. This implies that a disproportionate amount of price adjustment occurs at these points (because the frequency of price change is one as opposed to $f_k(s)$) and, thus, a disproportionate amount of adjustment is "lost in transit" by a matched model index.

[25] An important question is whether this logic is substantially changed in a menu cost model. In such models, the "selection effect" leads the aggregate price level to respond particularly quickly to aggregate shocks. New products may be less subject to this selection effect than other price changes. The selection effect only affects the speed of price adjustment to an aggregate shock, however; it does not affect long-run pass-through. In both the Calvo model and the menu cost model, all prices eventually adjust to the shock. A menu cost model generates a larger initial response of prices to an exchange rate movement, but a smaller subsequent price responses, with the eventual effect on prices being the same as in a Calvo model. This point is illustrated in detail in recent work by Bils, Klenow, and Malin (2012).

Integrating over k and s and assuming for simplicity that long-run pass-through is the same for all products yields an overall bias factor of

$$\int\int \frac{f_k(s)}{f_k(s) + z(s) - f_k(s)z(s)}\, dF_s(k)ds. \tag{9}$$

Since the function $f_k(s)/\big(f_k(s) + z(s) - f_k(s)z(s)\big)$ is concave in $f_k(s)$, product replacement bias is greater the greater is the amount of heterogeneity in the frequency of measured price changes across products.

In practice, not all product replacements in the dataset are product upgrades for a well-defined "product line." Rather, some products are discontinued without being replaced and other unrelated products are introduced into the dataset. This gives rise to the complication that the timing of product introductions into the dataset may not always coincide with the timing of price adjustments. In other words, products may enter the dataset with "stale" prices. This implies that it is unclear what periods the first observed price change for any given product accounts for, since we don't actually observe when the previous price change for that product occurred.

To derive the extent of product replacement bias for this more general case, we must introduce some additional notation. Let $\mathbf{B}_k^{ch}(s)$ denote the coefficients from regression (6) with the dependent variable constructed as the average price change for products that are changing their price but excluding those that are changing the price for the first time since they entered the dataset. Let $(1 + \alpha_k(s))\sum_n B_{nk}^{ch}(s)$ denote the sum of the regression coefficients when regression (6) is run with the dependent variable constructed as the average size of the price changes of those products in the dataset that are changing their price for the first time since they entered the dataset. This notation implies that $\alpha_k(s)$ is a factor governing the extent of "overreaction" of the first versus subsequent observed price changes to $\mathbf{\Lambda}_t$.[26] In Section IV, we will present an empirical estimate of $\alpha_k(s)$.

Given this notation, we can state the main result of this section as follows.

PROPOSITION 1: *The relationship between measured long-run pass-through,* $\sum_n B_{nk}^{mm}(s)$, *and true long-run pass-through,* $\sum_n B_{nk}(s)$, *in equation* (6) *is*

$$\sum_n B_{nk}^{mm}(s) = \frac{f_k(s)}{f_k(s) + z(s) - f_k(s)z(s)}\Big[\Phi_k(s)(1 + \alpha_k(s))\sum_n B_{nk}(s) + (1 - \Phi_k(s))\sum_n B_{nk}(s)\Big], \tag{10}$$

where $\Phi_k(s)$ *denotes the fraction of all price changes that are the first observed price change for a product.*

[26] If all products enter the dataset with a fresh price, $\alpha_k(s) = 0$. If, however, new products entering the dataset are drawn randomly from the pool of all existing products, the average duration of completed first price spells will be $2/\hat{f}_k(s) - 1$, while the average duration of subsequent completed spells will be $1/\hat{f}_k(s)$ (where we define $\hat{f}_k(s) = f_k(s) + z(s) - f_k(s)z(s)$). If we assume for simplicity that price rigidity is the only source of delay in pass-through, this implies that $\alpha = 1 - \hat{f}$. See Lancaster (1992) for a thorough discussion of such "length biased" or "stock" sampling from a population.

Notice that the term in front of the square brackets on the right-hand side of equation (10) is the same factor as in the simple product line case. The bias expression also incorporates an adjustment for overreaction of the first observed price change for each product, however. Intuitively, we allow for the fact that the first observed price change may react more strongly to past exchange rate changes because products may enter the dataset with stale prices. Notice that if $\alpha_k(s) = 0$ we are back to the simple case of the product-line model.

PROOF OF PROPOSITION 1:

The derivation of equation (10) proceeds in several steps. First, it is useful to consider the set of price spells in product group k and state s that are uncensored; i.e., that are neither the first nor the last observed price spell for the product in question. Let $\Delta p^u_{ikt}(s)$ denote the change in a price index constructed from a large sample of such price spells (including the price change at the end of each spell but not the price change at the beginning of the spell). Let $\mathbf{B}^u_k(s)$ denote the vector of regression coefficients for regression (6) with $\Delta p^u_{ikt}(s)$ as the dependent variable. The first step of the proof is to show that the dynamics of true pass-through and pass-through for the sample of uncensored spells are the same:

$$B_{nk}(s) = B^u_{nk}(s). \tag{11}$$

We establish this in online Appendix A2. The intuition for this result is that the price changes in the sample used to construct $\Delta p^u_{ikt}(s)$ account for movements in the optimal price of products in this sample over exactly the time period for which these products are included in this sample.

Second, notice that the frequency of price change in the sample of uncensored spells used to construct $\Delta p^u_{ikt}(s)$ is equal to the hazard that such price spells end each period, which is $f_k(s) + z(s) - f_k(s)z(s)$ (see online Appendix A2 for a derivation). Intuitively, the frequency of price change in this sample is higher than the frequency of price change in the full government sample because long price spells are more likely to be censored than short spells. Given this, the properties of the ordinary least squares (OLS) estimator imply that the pass-through coefficients for the sample of uncensored spells is

$$B^u_{nk}(s) = \big(f_k(s) + z(s) - f_k(s)z(s)\big)B^{ch}_{nk}(s). \tag{12}$$

Intuitively, $B^u_{nk}(s)$ is a scaled-down version of $B^{ch}_{nk}(s)$ since the observations with no price change contribute nothing to $B^u_{nk}(s)$. The scaling factor is the fraction of observations that have a price change in the sample that is used to construct $B^u_{nk}(s)$.

Third, combining equations (11)–(12), implies that true long-run pass-through is

$$B_{nk}(s) = \big(f_k(s) + z(s) - f_k(s)z(s)\big)B^{ch}_{nk}(s). \tag{13}$$

Fourth, we derive a relationship between pass-through for price change observations—$B^{ch}_{nk}(s)$—and measured pass-through for a matched model index—$B^{mm}_{nk}(s)$. To do this we must take account of the potential overreaction of the first observed

price change represented by $\alpha_k(s)$. In the full sample, the frequency of observed price changes is $f_k(s)$. Using this fact and the same logic as we used to derive equation (12), we find that measured long-run pass-through is

$$(14)\quad \sum_n B_{nk}^{mm}(s) = f_k(s)\left[\Phi_k(s)(1+\alpha_k(s))\sum_n B_{nk}^{ch}(s) + (1-\Phi_k(s))\sum_n B_{nk}^{ch}(s)\right],$$

where $\Phi_k(s)$ the fraction of price changes that are a first observed price change for a product (see online Appendix A3 for a derivation).

Finally, combining equations (13) and (14) yields equation (10).

To arrive at a factor that applies to pass-through for the entire price index, we must integrate over sectors k and states of the world s. This yields

$$(15)\quad \sum_n B_n^{mm} = \int\int \frac{f_k(s)}{f_k(s)+z(s)-f_k(s)z(s)}\left[\Phi_k(s)(1+\alpha_k(s))\sum_n B_{nk}(s) + (1-\Phi_k(s))\sum_n B_{nk}(s)\right] dF_s(k)ds.$$

Our model accommodates heterogeneity in the degree of true pass-through across sectors. Only heterogeneity that generates a correlation between true pass-through and the adjustment factor $f_k(s)/(f_k(s)+z(s)-f_k(s)z(s))$ or the degree of overreaction $\alpha_k(s)$ will affect the degree of product replacement bias. In particular, heterogeneity in true pass-through that is correlated with the frequency of price change—emphasized by Gopinath and Itskhoki (2010a)—is potentially important in determining the extent of product replacement bias. Empirical estimates of this heterogeneity are one input into our factor calculation in Section IV.

The factor calculation above focuses on regressions in which the dependent variable is Δp_t^m rather than $\Delta(p_t^m - p_t)$. In practice, there is a slight positive correlation between prices and the real exchange rate, which implies that equation (15) yields a slight underestimate of product replacement bias. For the case of $\alpha_k(s) = 0$, our derivations for product replacement bias also extend easily to the VECM specification.[27] In Section IV, we show that $\alpha_k(s) = 0$ is the empirically relevant case.

The discussion above considers the case of products' prices in the buyer's currency (LCP) and shows that for these products, product replacement bias causes a *downward* bias in measured pass-through. Product replacement bias causes an *upward* bias in measured pass-through for producer currency priced (PCP) products, however. To see this, consider a US export into the euro area that is priced in US dollars (i.e., PCP) and the price of which is only adjusted at the time of product replacements. A matched model euro-area price index based on a collection of such products would display one-for-one comovement with the exchange rate regardless of the true relationship between prices and exchange rates.

Notice that the lifelong regression considered by GIR coincides with our adjusted estimate of aggregate pass-through in the special case when firms' optimal prices are a function only of the current exchange rate, and there is no overreaction of the

[27] With $\alpha_k(s) = 0$, the measured index is simply missing a random fraction of price changes within each product group k and state s. The regression coefficients in the cointegrating relationship for each product group k and state s will therefore be biased downward by this fraction.

first observed price change; i.e., $\alpha_k(s) = 0$. (See online Appendix A4 for details.) If adjustment to exchange rate changes is delayed, however—e.g., due to strategic complementarities in price setting—lifelong pass-through will be downward-biased, while our estimator will yield an unbiased estimate of long-run pass-through. To see this, it is helpful to consider an extreme example. Suppose firms' desired prices are related to the exchange rate from two years prior but that products last only for two years in the dataset. In this case, the lifelong regression will yield pass-through of zero regardless of true long-run pass-through since observed price changes are responses to exchange rate movements before a product was introduced, which are not included in the lifelong regression. In practice, this bias is likely to be modest relative to the bias in aggregate pass-through. We explore this issue further in Section V.

A potential alternative approach to solving the problem of product replacement bias in our model economy is to calculate the price index simply as the weighted average of all prices including both new and continuing products. In practice, an import price index calculated from average prices in this way is extremely noisy. The simple average of prices for imports and exports routinely fluctuates by 10–20 percent per month. This reflects the fact that such an index is comparing the prices of entirely different products—say, last year's wool jacket versus this year's down coat—and products are highly heterogeneous as well as being measured in highly variable units. The massive amount of sampling error in this type of index generates sufficiently large standard errors in the estimated relationship between an average price index and the exchange rate that almost nothing can be concluded about the nature of exchange rate pass-through.[28] A second problem with analyzing average prices is that an appreciation of the exchange rate may lead consumers to systematically switch toward higher-quality products. This could bias upward the estimated relationship between prices (per unit quality) and the exchange rate (see, e.g., Ghironi and Melitz 2005).

Many products in the IPP data are intermediate products. An important question in interpreting the evidence on price rigidity for imported products is thus whether the observed rigid prices are "allocative" (Barro 1977). We do not address this issue, since our focus is on documenting rather than interpreting the observed relationship between prices and exchange rates. This phenomenon seems 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, however.

C. *Sample Rotation and Reporting Errors*

The discussion above implicitly makes two assumptions that we can relax easily. First, we can allow the frequency of product replacement in the government's dataset to differ from that in the economy as a whole. This may arise if the government

[28] In constructing unit value indexes, large price changes are often excluded as outliers or trimmed to deal with this concern. Alterman (1991) estimates that the US unit value indexes, produced in 1985, were calculated from only 56 percent of the value of imports and 46 percent of the value of exports. Observations are also dropped in unit value indexes because of lack of availability of data for the previous period—for example, because the product was not traded in the previous period. These practices may cause price changes associated with product replacements to be "lost in transit," leading to product replacement bias.

rotates products out of the dataset once they have been in the dataset for a certain amount of time and more generally replaces products periodically to ensure that new products are adequately represented in the dataset.

As we discuss in detail in Section IVB, an additional potential measurement problem is "satisficing" behavior by firms that are contacted about their prices; i.e., the tendency of firms to report no change in price even if a price change or product substitution has occurred because reporting "no change" requires less effort by the reporter. We can allow for this type of reporting error by assuming that firms make accurate price reports with probability $g(s)$ and report no-change irrespective of the accuracy of that report with probability $1 - g(s)$.[29]

Sample rotation and satisficing behavior by firms imply that the measured frequency of price change and product replacements in the government's dataset will be different from their true values. Let $\tilde{f}_k(s)$ denote the measured frequency of price change in the government's data for product group k in state s and $\tilde{z}(s)$ the measured frequency of product replacement in the government's dataset. Expressions for these rates in terms of $f_k(s)$, $z(s)$, and $g(s)$ are derived in online Appendix A5.

In our model, extended to include sample rotation and reporting errors, we can follow exactly the same steps as above to derive a product replacement bias factor. The only difference is that the frequency of price change and the frequency of substitutions that enter the factor become the measured frequency of price change, $\tilde{f}_k(s)$, and the measured frequency of product replacement $\tilde{z}(s)$.[30]

IV. Prices and Exchange Rates: Measurement

Before the introduction of the IPP, import and export price indexes were based on unit value data for highly disaggregated categories. This practice was criticized by, among others, the Stigler Commission (Stigler et al. 1961) because it did not control for changes in quality and composition within these categories. It has since become an important part of the BLS mandate to track the prices of exactly identical items over time to avoid mistaking quality changes for price changes. The IPP therefore takes great care in the way it defines a product. The definition of a product in the IPP data includes not only a unique product identifier such as a bar code, but also other "price determining characteristics" identified by the BLS such as the terms of the transaction, size of the shipment, and in some cases even the identity of the seller. We adopt the product definitions in the IPP. A product, as we use the term, is therefore often a contract between a particular buyer and seller. Carlton (1986) shows that defining products in this way is important when analyzing price rigidity for producer prices.

[29] We would like to thank Virgiliu Midrigan for suggesting that we incorporate this feature into our model.

[30] For tractability, here we assume that both the true price index—equation (4)—and the matched model price index—equation (5)—are subject to reporting error. In other words, we define the "true" price index as an index of the quality-adjusted prices firms would report if they were all in the government's dataset and had the same tendency to erroneously report no-change as the firms that are actually in the government's dataset. The difference between the two indexes is then entirely due to product replacement bias. Our results on long-run pass-through are not sensitive to this assumption, however.

Table 3—Price Changes per Quote Line

	Imports		Exports	
Number of price changes	LCP	PCP	LCP	PCP
0	44.3	43.8	16.3	39.2
1 or less	59.3	65.2	27.0	57.6
2 or less	68.7	77.7	30.0	68.2
3 or less	74.8	84.4	32.3	75.6
4 or less	79.2	90.7	34.7	80.8
5 or less	82.1	93.7	35.3	83.8
10 or less	88.4	98.1	77.0	90.1
15 or less	91.6	99.2	84.3	93.0

Notes: The table presents the fraction of products in the BLS microdata on import and export price data with less than or equal to a given number of price changes over the entire time span for which they are in the dataset. These statistics are for market-based products for 1994–2004 and are reported for both local currency priced (LCP) and producer currency priced (PCP) products. The statistics are weighted percentiles, using as weights the cumulative product-level weights over each product's lifetime.

A. *The Frequencies of Price Change and Product Replacement*

We show in Section III 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, 44 percent of products have no price changes, while 69 percent have 2 or fewer price changes. For PCP exports, 39 percent of products have no price changes, while 68 percent have 2 or fewer price changes. These statistics motivate the idea that product replacement bias may be a quantitatively important phenomenon in import and export price data.[31]

The small number of price changes per product reflects substantial price rigidity and frequent product replacement in the microdata on import and export prices collected by the BLS. Table 4 reports statistics on the frequency of price change and product replacement. We report these statistics separately for imports and exports as well as for LCP and PCP products.[32] Most US imports are local currency priced (92 percent), while most US exports are producer currency priced (97 percent). Our discussion therefore focuses on these two categories.

[31] Our estimate of the fraction of price spells with no price change is somewhat higher than the estimate of Gopinath, Itskhoki, and Rigobon (2010). Most of the difference arises because our estimate is for the entire dataset, while theirs is for a subset of high-income OECD countries. Another difference is that our estimates incorporate product-level weights.

[32] We 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 I. This is the procedure used by the BLS in many but not all cases when prices are missing. See Section I for a discussion of BLS imputation procedures. The key object when it comes to the size of product replacement bias is the amount of time over which exchange rate movements may be "unaccounted for" by subsequent price changes because they occur at the end of a price series. This is unaffected by price changes that are subsequently reversed, such as those associated with the BLS cell mean imputation procedure. We discuss this issue in greater detail in online Appendix E. All of the statistics we report are calculated as weighted averages using the item-level weights described in Section I.

TABLE 4—THE DISTRIBUTION OF PRICE CHANGES AND SUBSTITUTIONS

	Imports		Exports	
	LCP	PCP	LCP	PCP
Fraction of imports/exports	0.922	0.078	0.032	0.968
Mean frequency of price change	0.151	0.061	0.572	0.130
Median frequency of price change	0.066	0.033	0.573	0.060
Mean frequency of substitutions				
Forced	0.025	0.016	0.062	0.020
Forced including refusals	0.037	0.026	0.064	0.032
All	0.049	0.046	0.067	0.046
Distribution of the frequency of price change				
a	0.44	0.82	0.36	0.50
	(0.01)	(0.06)	(0.06)	(0.01)
b	3.50	20.72	3.52	4.59
	(0.05)	(1.74)	(0.87)	(0.10)

Notes: The top panel reports summary statistics for the mean and median frequency of price change and product substitution calculated using IPP microdata on import and export prices. The sample period is 1994–2004. Statistics are reported for both local currency priced (LCP) and producer currency priced (PCP) products. The weighted means and medians are calculated using the item-level weights described in the paper. The lower panel reports our estimates of "a" and "b," which are the parameters in the estimated distribution of the frequency of price change, assumed to be Beta(a,b). This distribution is estimated using the BLS microdata on imports and exports.

Product replacements occur for a number of reasons in the IPP data. Consequently, we report three different measures for the frequency of substitution. About half of all product replacements occur either because the firm no longer sells the product in question or because the firm itself has gone out of business. We refer to these product replacements as forced substitutions. The frequency of forced substitution is 2.5 percent for LCP imports and 2.0 percent for PCP exports.[33] In the case of roughly one-quarter of product replacements, the firm refuses to provide a new price quote without giving a reason. Some of these cases may also involve the product being discontinued.[34] The frequency of forced substitutions, including refusals, is 3.7 percent for LCP imports and 3.2 percent for PCP exports. The remaining 25 percent of product substitutions in the IPP dataset are due to product phase-out

[33] It is interesting that the frequency of product substitutions is slightly higher for exports than for imports because the export price data is gathered from sellers while the import price data is gathered from buyers. If sellers have list prices that apply to a large number of customers for each product, one might expect substitutions to be more frequent in data gathered from buyers. Each time the buyer switched products, a product substitution would occur in import price data. But in the export price data the BLS would continue to sample the product as long as there was another buyer buying at the same price. That exports have more substitutions thus suggests that many substitutions occur because of product replacements by sellers or contract renegotiation, as opposed to only product switching by buyers.

[34] It is difficult to estimate the fraction of substitutions that involve a version change or upgrade. There are at least two reasons why this variable in the dataset 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 have different HS10 codes as do 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 agent. 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.

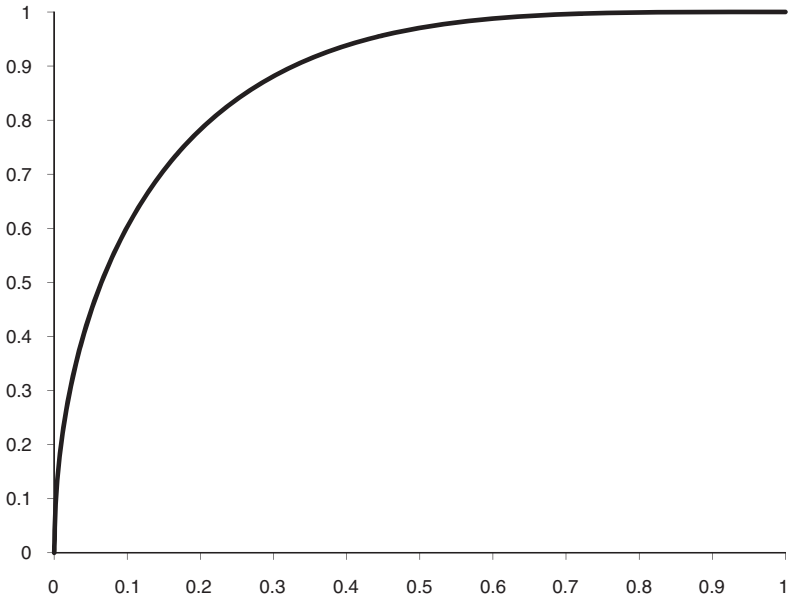

FIGURE 4. CUMULATIVE PROBABILITY DISTRIBUTION OF BETA (0.50, 3.65)

by the BLS. The overall frequency of substitutions is 4.9 percent for LCP imports and 4.6 percent for PCP exports. For our baseline results on product replacement bias, we use the more conservative measure of forced substitutions.[35]

Table 4 also shows that both imports and exports exhibit substantial price rigidity. For LCP imports, the mean monthly frequency of price change is 15.1 percent and the median is 6.6 percent. For PCP exports, the mean monthly frequency of price change is 13.0 percent while the median is 6.0 percent. These statistics parallel those reported in 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.

Heterogeneity in the frequency of price change is an important determinant of the quantitative impact of product replacement bias, as we discuss in Section IIIB. The frequency of price change varies widely across different sectors for imported products, from over 40 percent for Animal Products and Vegetables to less than 10 percent for such categories as Footware, Textiles, and Machinery. There is also a great deal of heterogeneity within each sector. Since the adjustment factor for product replacement bias is highly concave in the frequency of price change, ignoring intrasector 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. Suppose also that $f_j \sim \text{Beta}(a,b)$. We

[35] In 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 these cases, the IPP does not record a product substitution. Also, if a product tends to differ from one shipment to the next, it is often considered "out of scope" by the IPP since the IPP seeks to select products that can be repriced consistently. The IPP index is, therefore, likely to have somewhat less product turnover than the universe of products.

denote the product's lifetime by n_j. The total number of price changes x_j for a product are then distributed according to a binomial distribution, i.e., $x_j \sim \text{Bin}(n_j, f_j)$.

In online Appendix C, 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 $a = 0.44$ and $b = 3.50$. 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 Beta(0.44,3.50). For PCP exports, the estimated parameters are $a = 0.50$ and $b = 4.59$.

The average frequency of product replacement varies less across industry groups than the frequency of price change. The frequency of all product replacements in most industry groups is between 3 and 6 percent.[36] In our baseline results, we assume a homogeneous frequency of product replacement across goods. We have also considered a sectoral model in which the frequency of substitutions is allowed to vary across sectors and the distribution of the frequency of price change across products within a sector is a different beta distribution for each sector.[37] This model yields very similar results.

B. *Do First Price Changes React More Strongly to Past Exchange Rates?*

In Section III, we show that an important determinant of product replacement bias is the extent to which the first observed price change "overreacts" to historical exchange rate movements relative to subsequent price changes. Simple manipulation of equation (10) (see online Appendix D for details) implies that a conservative measure of product replacement bias is

$$\sum_n B_n^{mm} = \int \frac{f_k}{f_k + z - f_k z}\, dF(k) \sum_n B_n + \Psi, \tag{16}$$

where

$$\Psi = z \frac{\bar{d}}{\bar{d}_2} \left[(1+\alpha) \sum_n B_{2n}^{ch} - \sum_n B_{2n}^{ch} \right], \tag{17}$$

$\bar{d}$ denotes the average duration of all price spells, $\bar{d}_2$ denotes the average duration of all price spells of products with two or more price changes, and $(1+\alpha)\sum_n B_{2n}^{ch}$ and $\sum_n B_{2n}^{ch}$ denote the sum of the pass-through coefficients from regression (6) with the first observed price change and second observed price change, respectively, as the dependent variables and with the sample restricted to products with two or more price changes.

Equations (16) and (17) clearly indicate that if the degree of overreaction of the first observed price change is sufficiently strong, this can completely eliminate any

[36] Product replacement bias is especially important for a number of durable goods categories such as autos, furniture, and computers. For autos, the frequency of price change in the IPP is 6.8 percent, while the frequency of substitution is 5.1 percent. For furniture, these frequencies are 8.2 percent and 4.4 percent, respectively, and for computers they are 13.7 percent and 5.8 percent, respectively.

[37] This version of the model is discussed in greater detail in online Appendix B.

product replacement bias. There are several reasons, however, why such overreaction may not occur. First, a large fraction of product substitutions occur because one product is discontinued, to be replaced by a new product. Firms tend to negotiate new prices when they sign new contracts with their customers and this is also the time when many products are initiated into the dataset (Carlton 1986).

Second, "satisficing" behavior by firms may result in spurious rigidity for continuing products, while newly introduced products are more likely to be recorded correctly. For continuing products, the BLS collects prices using a "repricing form" that 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. The easiest response is to simply check the box indicating "no change" in price.[38] In contrast, the prices of products that are newly initiated into the BLS dataset are collected using a detailed personal interview and are therefore less likely to be spuriously stale.

The BLS has had a longstanding concern about this issue. In 1988, the BLS carried out a study to investigate it, known as the "Quality through Correspondence" initiative. In this study, the BLS contacted a sample of firms who had reported "no change" in prices for 24 months or more and asked them to either confirm that their prices were unchanged or provide updated information. They found that the vast majority of firms either reported an updated price or reported that the product had been discontinued. Given the success of this initiative, the BLS implemented a second "Quality through Correspondence" initiative on a broader scale in 1999, again targeted at firms who had reported "no change" for 24 months or more. During the initiative, the frequency of price change and discontinuation for the targeted firms rose by 50–100 percent relative to surrounding periods. Most recently, the BLS carried out a study to analyze how firms' reporting behavior changes over time. They found that when reporters are first initiated into the dataset, they tend to report many price changes, but that fewer price changes are reported over time, suggesting that seasoned reporters tend to become fatigued with the reporting process.

Because of these concerns, the BLS has at some points tried to systematically contact reporters who have reported no change in prices for more than 12 months on the repricing form, as noted by Gopinath and Rigobon (2008). Unfortunately, funding limitations have meant that this policy has not been implemented consistently.[39] Other researchers have considered alternative ways of investigating how important underreporting of price changes might be. In particular, Gopinath and Rigobon (2008) argue that these biases are likely to be small since the frequency of price change was essentially unchanged when the BLS was forced to switch to phone surveys during the 2001 anthrax attacks. Since reporters were provided with their previous prices over the phone, however, the phone surveys may have exhibited a similar bias toward "no change" as the regular repricing forms.[40] Our model of product replacement bias in Section IIIC explicitly allows for satisficing behavior by firms.

We measure the extent to which the first observed price change for each product reacts more strongly to past exchange rates than subsequent price changes—i.e., we

[38] Liu et al. (2009) document the severity of this kind of bias in a different context.

[39] We thank Rozi Ulichs and Will Adonizio for detailed correspondence on this issue.

[40] See the IPP Data Collection Manual (US Department of Labor 2005) for a discussion of the survey procedure during this episode.

TABLE 5—PRICE CHANGE FOR FIRST AND SECOND SPELL ON EXCHANGE RATE

	Imports		Exports	
	First price change	Second price change	First price change	Second price change
Panel A				
Exchange rate change:				
During price spell that is ending	0.23	0.26	0.13	0.11
	(0.03)	(0.03)	(0.03)	(0.03)
First quarter before price spell that is ending	0.19	0.18	0.03	0.14
	(0.05)	(0.04)	(0.07)	(0.06)
Second quarter before price spell that is ending	0.13	0.11	−0.06	−0.02
	(0.04)	(0.04)	(0.06)	(0.06)
Third quarter before price spell that is ending	−0.01	0.05	0.14	0.15
	(0.05)	(0.04)	(0.06)	(0.05)
Fourth quarter before price spell that is ending	0.07	0.09	−0.09	0.09
	(0.04)	(0.04)	(0.05)	(0.05)
Fifth quarter before price spell that is ending	0.05	0.04	0.19	0.08
	(0.04)	(0.04)	(0.06)	(0.06)
Sixth quarter before price spell that is ending	0.14	0.06	0.06	−0.01
	(0.05)	(0.04)	(0.07)	(0.06)
p-value for the null that coefficients for price spell and second quarter before price spell are equal	0.052	0.002	0.008	0.059
Panel B				
Exchange rate change:				
During the price spell that is ending and first quarter before this spell	0.21	0.24	0.09	0.12
	(0.02)	(0.02)	(0.03)	(0.03)
Second to fourth quarter before the price spell that is ending	0.06	0.09	−0.02	0.08
	(0.02)	(0.02)	(0.03)	(0.03)
p-value for the null that the two coefficients are equal	0.000	0.000	0.021	0.340
Average duration of spells in months	8.7	8.5	9.5	9.5
Observations	3356	3339	2212	2230

Notes: Panel A presents coefficients from OLS regressions of price changes on current and lagged exchange rate changes, including the exchange rate change during the price spell that is ending, and exchange rate changes over the first to sixth quarters before this price spell (standard errors in parentheses). These regressions are run for market-based transactions for high-income OECD countries. We drop observations where the absolute values of the log price change is larger than 0.5. The sample period is 1994–2007. The table also reports the *p*-value testing the null that the coefficients on the current price spell and the second quarter before the price spell are equal. Panel B reports the results of similar regressions where the regressors are (i) the exchange rate change during the current price spell and the quarter before the spell and (ii) the exchange rate change over the second to fourth quarters before the spell started. The *p*-value for the null that these coefficients are equal is also reported.

estimate Ψ—using the BLS micro-data. Our results on this are reported in Table 5. Panel A of Table 5, reports results for the following regression:

$$(18) \quad \Delta p_{jk}^{r} = \alpha + \beta_S \Delta q_{jk,S} + \beta_{1Q} \Delta q_{jk,1Q} + \cdots + \beta_{6Q} \Delta e_{jk,6Q} + \epsilon_{jk},$$

where Δp_{jk}^{r} denotes the log size of the kth price change for product j relative to the change in the price of domestic production over the kth price spell for product j, Δq_{jkS} denotes the log change in the real exchange rate over the course of the kth price spell for product j, and $\Delta q_{jk\#Q}$ denotes the log change in the real exchange rate over the

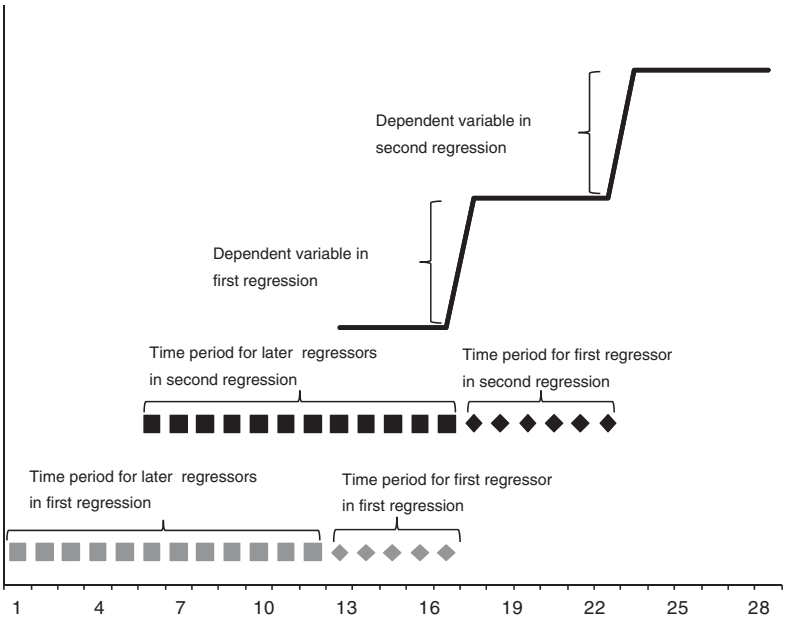

FIGURE 5. GRAPHICAL DEPICTION OF REGRESSIONS IN TABLE FIVE

Notes: This figure provides a graphical depiction of the regression equation presented in equation (17). The solid line denotes the price of a product that has two price changes. The first regression has as its dependent variable the first price change and as explanatory variables exchange rate movements over the first price spell as well as the six quarters preceding the product's introduction. The second regression has as its dependent variable the second price change, and as explanatory variables exchange rate movements over the course of the second price spell as well as exchange rate movements in the preceding six quarters. If the first price is "stale" then the first price change should respond more to exchange rate movements preceding the first price spell than the second price change does to exchange rate movements preceding the second price spell.

course of the #th quarter prior to the *k*th price spell for product *j*. We run this regression for the first and second observed price changes ($k = 1$ and 2) of all products that have two or more price changes. We run these regressions separately for import and exports for high-income OECD countries using bilateral real exchange rates.

Figure 5 provides a graphical illustration of these "first" and "second" price change regressions—i.e., equation (18) for $k = 1$ and 2. While we do not observe how much prices change when a new good is introduced into the dataset, we can observe how responsive subsequent price adjustments are to exchange rate movements that occurred before the good was introduced. If the first observed price change for each product is more strongly related to exchange rate movements that occur before that product's introduction into the dataset than the second observed price change is to exchange rate changes that occurred before the first observed price change, this would suggest that the initial prices of products entering our dataset were not newly reset. In fact, we find no evidence of this.

For imports, the pattern of coefficients is very similar for the first and second price changes. In both cases β_S is larger than 0.2. The coefficients then fall rapidly over the first two quarters before the price spell in question and are insignificant in most cases after that. There is no evidence that the first observed price change responds more to exchange rate changes before the first price spell than the second

observed price change responds to exchange rate changes before the second price spell. Importantly, the average number of months between the time the products in these regressions are introduced and the first observed price is almost exactly the same as the average time between the first and second observed price changes (8.5 and 8.7 months, respectively). The pattern is similar for exports.[41]

Given the pattern of coefficients reported in panel A, we have also run the following more parsimonious specification:

$$\Delta p^r_{jk} = \alpha + \beta_{S+1} \Delta q_{jk,S+1} + \beta_{2-4Q} \Delta q_{jk,2-4Q} + \epsilon_{jk}, \tag{19}$$

where $\Delta q_{jk,S+1}$ denotes the log change in the real exchange rate over the course of the kth price spell and one quarter before this price spell for product j and $\Delta q_{jk,2-4Q}$ denotes the log change in the real exchange rate over the course of the second, third, and fourth quarters prior to the kth price spell for product j. This regression again yields very similar results for the first and second observed price changes for both imports and exports. There is no evidence that the first observed price change responds more to past exchange rate movements than the second observed price change. If anything, the opposite is true.

We use these results to estimate Ψ (see equation 17). Since the second price change is estimated to be slightly more responsive to past exchange rates, we estimate $\Psi < 0$. Recall that if the first price change overreacts to past exchange rates, Ψ will be positive. Our results thus suggest that factors such as firms' preferences for raising their price when a product is introduced rather than at other times—which can make $\Psi < 0$—may be more important than factors that would lead to $\Psi > 0$. In what follows we set $\Psi = 0$, to be conservative, when we calculate our product replacement bias factor.

To gauge the robustness of our estimates of product replacement bias to sampling variation, we consider an "upper bound" estimate for Ψ. Specifically, we add two standard errors to the point estimates for the first price change regression and subtract two standard errors from the point estimates for the second price change regression, using our estimates for imports from Panel B in Table 5. We recalculate Ψ based on these alternative values. Using this estimate of Ψ, we report this "lower bound" estimate for true pass-through in Panel C of Table 6. Even for this very conservative measure, our analysis implies a large adjustment to measured pass-through.[42]

C. *Adjusting Pass-through for Product Replacement Bias*

Given the estimates in Table 4, we can use equation (16) with $\Psi = 0$ to produce estimates of the factor by which exchange rate pass-through is mismeasured because of product replacement bias. These estimates are reported in Table 6. We assume

[41] For exports, we report the pass-through from the domestic exporters' perspective since these are more easily compared to the results for imports. Pass-through from the foreign importers' perspective is one minus the pass-through reported in the Table 5.

[42] We thank an anonymous referee for encouraging us to do this exercise. We have also run a Monte Carlo in which we compare data from a model in which products enter with fresh prices with data from a model in which new additions to the dataset are drawn randomly from the population of products in the economy (and thus enter with stale prices on average). We find substantial differences between the regression on the second and first observed price changes in the latter dataset but not the former.

TABLE 6—ROBUSTNESS

	Major Country RER			Broad RER		
	Frequency of subsitutions			Frequency of subsitutions		
	Forced (benchm)	Forced (incl. ref)	All	Forced (benchm)	Forced (incl. ref)	All
Measured pass-through						
Imports		0.43			0.52	
		(0.05)			(0.06)	
Exports		0.85			0.84	
		(0.05)			(0.06)	
Panel A. With heterogeneity in pass-through						
Factors						
LCP imports	1.63	1.78	1.91	1.63	1.78	1.91
PCP imports	1.80	2.12	2.72	1.80	2.12	2.72
LCP exports	2.30	2.32	2.36	2.30	2.32	2.36
PCP exports	1.57	1.74	1.91	1.57	1.74	1.91
Adjusted pass-through						
Imports	0.64	0.69	0.74	0.80	0.86	0.92
Exports	0.79	0.76	0.74	0.77	0.75	0.73
Panel B. No heterogeneity in pass-through						
Factors						
LCP imports	1.71	1.88	2.03	1.71	1.88	2.03
PCP imports	1.80	2.12	2.72	1.80	2.12	2.72
CP exports	2.47	2.50	2.54	2.47	2.50	2.54
PCP exports	1.62	1.81	2.01	1.62	1.81	2.01
Adjusted pass-through						
Imports	0.67	0.73	0.78	0.83	0.90	0.97
Exports	0.77	0.76	0.73	0.77	0.74	0.71
Panel C. "Lower bound" based on sampling variation in Ψ						
Adjusted pass-through						
Imports	0.56	0.56	0.55	0.71	0.73	0.73

Notes: The table presents the estimated adjustment factor for product replacement bias under various alternative assumptions. We present results for three frequencies of substitutions: "Forced (benchm)" is our benchmark measure of forced substitutions, "Forced (incl. ref)" is forced substitutions including refusals, "All" is the frequency of product substitutions using all observed substitutions. Panel A and B present results for different assumptions regarding heterogeneity in pricing behavior across products. The results are presented for local currency priced (LCP) and producer currency priced (PCP) imports and exports. The first two columns present results for the "Major Country" real exchange rate (RER) while the last two columns present results for the "Broad" real exchange rate measure. Panel C presents "lower bound" estimates of true pass-through based on the sampling variation in our estimate of Ψ.

for empirical tractability that the frequency of price change and the frequency of substitutions are constant over time for each product.[43] We present estimates of the product replacement bias factor for the three different measures of the frequency of product replacement reported in Table 4 and two sets of assumptions regarding

[43] Empirical evidence suggests that these are reasonable assumptions for the particular application we study. We have 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 resulting coefficient is -0.023 (0.131) for imports and -0.078 (0.308) for exports, where we report standard errors in parentheses. For periods and countries for which exchange rate variation was more dramatic, there may be a stronger relationship between the frequency of product replacement and the real exchange rate. 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. In this case, it would be important to account for time variation in the frequency of price change and frequency of substitutions when calculating the adjustment factor.

heterogeneity in true pass-through across products with different frequencies of price change.

Our benchmark results are presented in the first data column of panel A of Table 6. In this case, we use the most conservative measure of product replacement. We also allow for variation in true pass-through across products with different frequencies of price change. Gopinath and Itskhoki (2010a) argue that this pattern exists in the data. Their estimates suggest that true pass-through for dollar-priced (LCP) imports with a frequency of price change below about 25 percent per month is only about 65 percent of true pass-through for dollar-priced (LCP) imports with a higher frequency of price change. Under these assumptions, the factor by which traditional estimates of pass-through are biased because of product replacement bias is 1.63 for LCP imports and 1.57 for PCP exports. The next two columns report factors for our other two measures of the frequency of product replacement.

Using these factors, we can adjust measured exchange rate pass-through for product replacement bias. For our benchmark measure of the frequency of product replacement, adjusting for product replacement bias raises exchange rate pass-through for US imports from 0.43 to 0.64 and lowers exchange rate pass-through for US exports from 0.85 to 0.79. The adjustment is even larger when we use our other measures of the frequency of substitutions. The adjustment based on the overall frequency of substitutions yields exchange rate pass-through for both US imports and exports of 0.74.[44]

Table 6 presents results for several additional cases. First, we report results for cases in which true long-run pass-through does not vary with the frequency of price change. This raises the product replacement bias factor by about 0.1. We also present results using the FRB's "Broad" real exchange rate series. Measured pass-through is somewhat higher for imports and slightly lower for exports using the Broad real exchange rate. Using this exchange rate series and our benchmark assumptions about the frequency of product replacement and heterogeneity in true pass-through yields a pass-though estimate of 0.80 for imports and 0.77 for exports.

D. *Adjusting the Terms of Trade for Product Replacement Bias*

Our results indicate that the US import and export price indexes are too smooth. One consequence of this is that the US terms of trade are also too smooth. Above, we focus on results for long-run pass-through. In Section V, we estimate a simple model of the dynamics of pass-through. Using this estimated model, we can construct time series for the US import and export price indexes that are adjusted for product replacement bias.[45]

Adjusting for product replacement bias raises the standard deviation of the quarterly change in the log price index for nonoil imported goods from 1.1 percent to

[44] Gopinath, Itskhoki, and Rigobon (2010) argue that there is a large difference in pass-through between LCP and PCP imports. We allow for this difference in our calculations. We adopt their estimate of 0.94 for measured pass-through of PCP imports and use this in our calculation of aggregate pass-through adjusted for product replacement bias. Reasonable variations on this assumption have negligible effects on our results. For exports, virtually all products are PCP. So, any reasonable heterogeneity across LCP and PCP products makes virtually no difference.

[45] We first simulate data from the estimated model in Section V. We then construct a matched model index and the true price index from this data (we can construct the true index since we know the relative quality of different products in our simulation). We then run a regression of the true index on the current value and eight lags of the matched model index. Finally, we apply the resulting filter to the US import and export price indexes from 1982 to 2010.

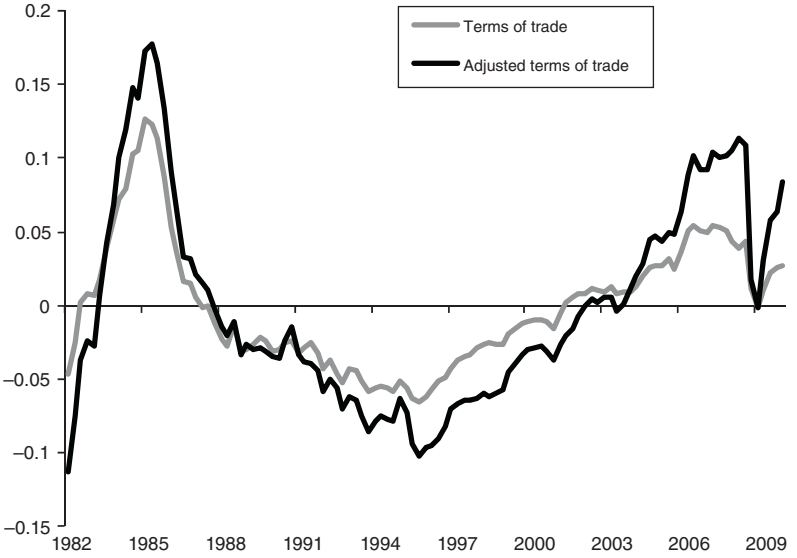

FIGURE 6. US TERMS OF TRADE ADJUSTED FOR PRODUCT REPLACEMENT BIAS

1.6 percent, while it raises this measure of volatility for nonagricultural exported goods from 1.1 percent to 1.9 percent. Figure 6 plots the US terms of trade adjusted for product replacement bias along with the same series without such an adjustment.[46] The standard deviation of the quarterly change in the terms of trade rises by 75 percent from 0.97 percent to 1.70 percent.

This adjustment for product replacement bias brings the data closer in line with standard models. Simple two-country real business cycle models imply that the terms of trade should be more volatile than the real exchange rate. The official nonoil terms of trade are, however, only about 30 percent as volatile as the real exchange rate. Adjusting for product replacement bias raises the volatility of the terms of trade to about half the volatility of the real exchange rate. New Keynesian models designed to match the volatility of the real exchange rate generate a more volatile terms of trade series than the official series (Corsetti, Dedola, and Leduc 2008). Adjusting the terms of trade for product replacement bias raises its volatility in the data to roughly match its volatility in these models.

V. Alternative Measures of Pass-through

In addition to analyzing "aggregate" pass-through—based on a regression similar to equation (1)—GIR propose two alternative pass-through measures: pass-through "conditional on adjustment" and "lifelong" pass-through.[47] In this section, we study the implications of product replacement bias for these alternative statistics. We focus

[46] For comparability to the rest of our analysis, the measure of the terms of trade we present here is the ratio of the price index for nonoil imported goods and nonagricultural exported goods.

[47] Lifelong pass-through is motivated in part by concerns regarding the effects of product substitutions on aggregate indexes similar to those we emphasize here.

TABLE 7—RECONCILING DIFFERING MEASURES OF PASS-THROUGH

	Real data	Simulated data
Pass-through conditional on price adjustment	0.24	0.24
	(0.07)	
Lifelong pass-through	0.51	0.52
	(0.12)	
Lifelong for frequency less than 0.25	0.41	0.41
	(0.12)	
Aggregate pass-through	0.33	0.31
	(0.08)	
True pass-through		0.60

Notes: The table reports regression coefficients from analogous regressions run using actual and simulated data. The "Real data" statistics are based on calculations using BLS micro-data on the prices of local currency priced imports. The sample period is 1994–2007. The "Simulated data" statistics are based on output from our simulation model using our benchmark measure of forced substitutions. A detailed discussion of the regressions is presented in Section 6. The last row presents the average long-run pass-through assumed in the model that gives rise to the simulated data. For the "Real data" results, standard errors are reported in parentheses. The standard errors for the first three regressions have been clustered by year. Clustering these standard errors by country yields somewhat smaller standard errors. For the aggregate pass-through regression, robust standard errors are reported.

on dollar-priced (LCP) imports and make two main points. First, low pass-through "conditional on adjustment" is entirely consistent with a large bias in aggregate pass-through and much higher pass-through in the long-run. Second, GIR's finding that lifelong pass-through is much higher than aggregate pass-through are corroborating evidence for our findings of large biases in aggregate pass-through associated with product replacement bias. We also use the analysis in this section—and analogous results for nondollar (PCP) imports as well as exports—to construct the adjusted measure of the US terms of trade presented in Section IVD.

GIR define pass-through "conditional on adjustment" as β_{Δ} in the regression,

$$\Delta(p_{jit} - p_t) = \alpha + \beta_{\Delta}\Delta^{*}q_t + \epsilon_{jit}, \tag{20}$$

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). They define "lifelong" pass-through as $\beta_{\tilde{\Delta}}$ in the regression,

$$\tilde{\Delta}(p_{jit} - p_t) = \alpha + \beta_{\tilde{\Delta}}\tilde{\Delta}q_t + \epsilon_{jit}, \tag{21}$$

where $\tilde{\Delta}$ is a difference operator representing the difference between the time the product is introduced into the dataset and the time of the last price change.

Table 7 reports the results of these regressions for our main specification and sample, which extends the sample used in GIR to more countries and a slightly longer time period. The results are very similar to their original findings.[48] Pass-through

[48] Our sample and specification differ from GIR's in the following ways. First, they focus on a sample of high-income OECD countries that have a substantial number of nondollar priced goods, while we include all countries. In addition, our analysis covers a slightly longer time period, and uses the trade-weighted as opposed to the bilateral

conditional on adjustment is 0.24, while lifelong pass-through is 0.51, roughly twice as high. Lifelong pass-through is lower for the sample of products with a low frequency of price change, consistent with the results reported in Gopinath and Itskhoki (2010b) and used to calibrate our model in Section IV. The aggregate measure of long-run pass-through based on the dynamic adjustment equation is 0.33, much lower than the estimate of long-run pass-through based on the "lifelong" approach.[49]

To be able to capture the large differences between the alternative pass-through measures documented in Table 7, we extend the model presented in Section III by specifying a process for the evolution of firms' desired prices. This yields a full quantitative model of pass-through that is nested within the framework presented in Section III and shares many features with the models analyzed in GIR and Gopinath and Itskhoki (2010a), but incorporates additional features that generate product replacement bias. We allow for delayed adjustment to exchange rate movements due to strategic complementarities. We parameterize these in a reduced-form way by assuming that desired prices p^*_{jit} are affected by a distributed lag of past real exchange rates

$$p^*_{jit} - p_t = \phi \Upsilon \sum_{s=0}^{23} \psi^s q_{it-s} + \eta_{jit}, \tag{22}$$

where $\Upsilon = (1 - \psi)/(1 - \psi^{24})$. The parameter ϕ governs the overall level of long-run pass-through, while ψ determines the degree to which pass-through is delayed.[50] As in our previous analysis, we allow for heterogeneity in ϕ across products that is correlated with the frequency of price change. The variable η_{jit} captures all influences on p^*_{jit} that are orthogonal to the real exchange rate. It follows the stochastic process,

$$\eta_{jit} = \mu + \rho \eta_{jit-1} + \epsilon_{jit}, \tag{23}$$

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

We also allow for measurement error in the timing of price observations. This is motivated by features of the data gathering procedure used by the IPP. 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 these prices will refer to products ordered at a substantially earlier point in time. Furthermore, while 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, importers

real exchange rate. While the choice of high-income OECD countries yields higher pass-through than when all countries are included, the use of bilateral exchange rates yields lower pass-through than the trade-weighted exchange rate. The two effects roughly cancel out, implying that the results in Table 7 is very similar to the results reported in GIR.

[49] This estimate is very similar to the estimate of long-run pass-through reported in Gopinath and Itskhoki (2010b) for the same sample and specification. It is considerably higher than the estimate reported in GIR. This difference arises from a combination of a different sample (GIR focus on a sample of high-income OECD countries selected as those with a substantial fraction of nondollar priced goods), a different exchange rate measure (GIR focus on bilateral exchange rates) and the fact that GIR do not use product-level weights in constructing their price index.

[50] Woodford (2003) discusses several sources of strategic complementarity that generate gradual delayed adjustment.

and exporters often go for long periods of time without importing or exporting. As a consequence, reporters often provide prices for other days in the month. In some cases, average monthly prices are provided. The large amount of price imputation done in the IPP and discussed in Section I is an additional source of timing error.

To allow for such timing error in reported prices, 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 $t + u_t$, where 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 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, $u_{2,t} = -xd_{jit}$, where d_{jit} is the time in months since the last price change, and $x \sim \text{Unif}[0,X]$. This term is motivated by the various forms of timing error discussed above, which lead recorded prices to be somewhat "stale."

We calibrate the distribution of the frequency of price change as well as the frequency of substitutions based on the results reported in Table 4. For the frequency of substitutions, we use the frequency of forced substitutions of 2.5 percent per month. When simulating the model, we use actual daily observations on the United States–German exchange rate over the time period 1995–2007. 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.[51] For simplicity, we assume that the difference between home and foreign inflation is constant. We set $\rho = 0.5$ based on previous estimates in Nakamura and Steinsson (2008), and we set σ_ϵ^2 to match the average size of price changes in the data (in practice, these assumptions have little impact).

The remaining parameters are the two delayed adjustment parameters—ϕ and ψ —and the timing error parameter X. 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 reported in Table 7. We select the values of (ϕ, ψ, X) that minimize the sum of the squared deviations between these moments in the simulated and actual data. 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.[52]

This estimation procedure yields $\phi = 0.86$, $\psi = 0.85$, and $X = 0.66$. The estimated value of X implies that, on average, delivery lags and other sources of timing error account for a delay in price reporting of about 33 percent of the average duration since the last price change or product replacement. This corresponds to an average reporting lag of about three months, which is consistent with existing estimates of delivery lags (e.g., Abel and Blanchard 1988). The estimated value of ψ implies that there is a substantial amount of delayed adjustment due to strategic complementarity. Prices respond to a distributed lag of past levels of the exchange rate with about 30 percent of the weight on values of the exchange rate from more than 6 months earlier. The estimated value of ϕ implies that the high frequency of price change products in our model have a desired pass-through of 0.86, while the

[51] For robustness, we have carried out analogous experiments using the Canadian, Japanese, and UK exchange rates. These experiments yield almost identical results.

[52] Each simulated dataset in this analysis consists of 3,000 price series.

low frequency of price change products have a desired pass-through of 0.57. This implies that aggregate pass-through is 0.60.

A. *Discussion*

The second column of Table 7 reports the fit of the model to the data using the parameters described above. The model is able to match the data for all the pass-through measures we consider. Incorporating product replacement bias allows us to explain the large observed difference between the lifelong and aggregate measures of long-run pass-through. Consistent with our analysis in Section III, aggregate pass-through is substantially downward-biased due to product replacement bias, but lifelong pass-through is much less biased.[53] Our estimates also imply that a substantial amount of delayed adjustment in firms' desired prices is required to explain the large difference between "pass-through conditional on adjustment" and longer-run measures of pass-through. Substantial amounts of timing error in the IPP data due to the delivery lags described above also contribute to this difference between short and long-run measures.

It is useful to note that delayed pass-through alone—due, for example, to strategic complementarities in pricing—cannot explain why lifelong pass-through is so much higher than long-run aggregate pass-through. For example, GIR present a model with strategic complementarities in pricing but no product replacement. Their model yields virtually the same estimates for long-run aggregate pass-through and lifelong pass-through. Gopinath and Itskhoki (2010a) present a model with heterogeneity in the frequency of price change and true pass-through as well as strategic complementarity and product replacement. This model also yields nearly identical estimates for long-run aggregate pass-through and lifelong pass-through.

The absence of product replacement bias in the model of Gopinath and Itskhoki (2010a) arises largely from two features of their model that our analysis in Section III shows are crucial in determining product replacement bias. First, they assume that products fall into 1 of 7 frequency of price change "bins," ranging from about 10 percent to 35 percent per month and do not allow for heterogeneity in the frequency of price change within these bins. To accurately assess the quantitative force of product replacement bias, however, it is important to allow for the large number of products in the dataset with no observed price changes—about 40 percent of products in the data have no price changes.[54] Second, Gopinath and Itskhoki (2010a) assume that products enter the dataset randomly, implying a substantially higher value of the "overreaction" parameter, α, than we estimate in the data. In other words, their model assumes that price changes at the time of product replacements are "accounted for" by subsequent observed price changes—something we do not find evidence for in our empirical analysis.

[53] The downward bias in the lifelong pass-through measure grows with the extent of delays in price adjustment. The downward bias is substantially larger in a Calvo model than in a menu cost model since the latter implies more rapid adjustment of prices to exchange rates. See the working paper version of GIR for a more detailed discussion of this issue (Gopinath, Itskhoki, and Rigobon 2007).

[54] Since the lowest frequency bin in Gopinath and Itskhoki (2010a) has a frequency of price change of roughly 10 percent (excluding substitutions) there is less than a 3 percent probability of observing 0 price changes over the course of a product's 35-month lifetime—substantially less than in the data. Gopinath and Itskhoki (2010a) exclude products with no price changes from their empirical analysis.

VI. Conclusion

This paper argues that the simultaneous presence of price rigidity and frequent product replacements lead aggregate price indexes to appear smoother than they actually are, biasing import price pass-through measures and the volatility of the terms of trade toward zero. We propose a model of this "product replacement bias." Our model yields an adjustment to aggregate indexes based on empirical measures of the frequency of price change and product replacements. The adjustment depends importantly on the nature of cross-sectional heterogeneity in the frequency of price change and true "pass-through." Our results illustrate the enormous dependence of such statistics on what is assumed about the response of prices at the time they enter the index.

It is important to recognize that our estimates of the bias rely on the assumption that the optimal price response of new products can be inferred from the behavior of other products. If the optimal price for a new product is much more responsive than for a continuing product—say because of adverse customer reactions to price changes during a product's lifetime—then product replacement bias could be much larger than what we have assumed; or, alternatively, if the optimal price is less responsive at the time of product replacements, product replacement bias could be substantially smaller. Our adjustment is designed to address a situation where characteristics data are unavailable, and therefore it is not possible to calculate the quality-adjusted prices using hedonic methods. Ideally, future research using more detailed data will allow for more direct estimates of product replacement bias based on comparisons of the quality-adjusted prices of entering and exiting items.

REFERENCES

Abel, Andrew B., and Olivier Jean Blanchard. 1988. "Investment and Sales: Some Empirical Evidence." In *Dynamic Econometric Modeling: Proceedings of the Third International Symposium in Economic Theory and Econometrics*, edited by William A. Barnett, Ernst R. Berndt, and Halbert White, 269–96. New York: Cambridge University Press.

Abraham, Katharine G., John S. Greenlees, and Brent R. Moulton. 1998. "Working to Improve the Consumer Price Index." *Journal of Economic Perspectives* 12 (1): 27–36.

Alterman, William. 1991. "Price Trends in U.S. Trade: New Data, New Insights." In *International Economic Transactions: Issues in Measurement and Empirical Research*, edited by Peter Hooper and J. David Richardson, 109–39. Chicago: University of Chicago Press.

Armknecht, Paul A., and Donald Weyback. 1989. "Adjustments for Quality Change in the U.S. Consumer Price Index." *Journal of Official Statistics* 5 (2): 107–23.

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

Auer, Raphael, and Thomas Chaney. 2009. "Exchange Rate Pass-Through in a Competitive Model of Pricing-to-Market." *Journal of Money, Credit, and Banking* 41 (1): 151–75.

Barro, Robert J. 1977. "Long-Term Contracting, Sticky Prices, and Monetary Policy." *Journal of Monetary Economics* 3 (3): 305–16.

Berger, David, Jon Faust, John H. Rogers, and Kai Steverson. 2012. "Border Prices and Retail Prices." *Journal of International Economics* 88 (1): 62–73.

Bergin, Paul R., and Robert C. Feenstra. 2001. "Pricing-to-Market, Staggered Contracts, and Real Exchange Rate Persistence." *Journal of International Economics* 54 (2): 333–59.

Bils, Mark. 2009. "Do Higher Prices for New Goods Reflect Quality Growth or Inflation?" *Quarterly Journal of Economics* 124 (2): 637–75.

Bils, Mark, and Peter J. Klenow. 2001. "Quantifying Quality Growth." *American Economic Review* 91 (4): 1006–30.

Bils, Mark, Peter J. Klenow, and Benjamin A. Malin. 2012. "Reset Price Inflation and the Impact of Monetary Policy Shocks." *American Economic Review 102* (6): 2798–825.

Board of Governors of the Federal Reserve System. 1982–2007. "Foreign Exchange Rates." http://www.federalreserve.gov/releases/h10/hist (accessed 2008).

Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale W. Jorgenson. 1996. "Toward a More Accurate Measure of the Cost of Living." Final Report to the Senate Finance Committee, December 4.

Bouakez, Hafedh, and Nooman Rebei. 2008. "Has Exchange Rate Pass-Through Really Declined? Evidence from Canada." *Journal of International Economics* 75 (2): 249–67.

Broda, Christian, and David E. Weinstein. 2006. "Globalization and the Gains from Variety." *Quarterly Journal of Economics* 121 (2): 541–85.

Bureau of Economic Analysis. 1982–2007. "National Income and Product Accounts." United States Department of Commerce (accessed 2008).

Bureau of Labor Statistics. 1994–2007. "International Prices Program Micro-Data 1994–2007." United States Department of Labor (accessed 2008–2011).

Burstein, Ariel, Martin Eichenbaum, and Sergio Rebelo. 2005. "Large Devaluations and the Real Exchange Rate." *Journal of Political Economy* 113 (4): 742–84.

Campa, Jose Manuel, and Linda S. Goldberg. 2005. "Exchange Rate Pass-Through into Import Prices." *Review of Economics and Statistics* 87 (4): 679–90.

Carlton, Dennis W. 1986. "The Rigidity of Prices." *American Economic Review* 76 (4): 637–58.

Clausing, Kimberly A. 2001. "The Behavior of Intrafirm Trade Prices in U.S. International Price Data." Unpublished.

Corsetti, Giancarlo, and Luca Dedola. 2005. "A Macroeconomic Model of International Price Discrimination." *Journal of International Economics* 67 (1): 129–55.

Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc. 2008. "High Exchange-Rate Volatility and Low Pass-Through." *Journal of Monetary Economics* 55 (6): 1113–28.

Court, Andrew. 1939. "Hedonic Price Indexes with Automobile Examples." In *The Dynamics of Automobile Demand*, 99–117. New York: General Motors Corporation.

Diewert, Erwin, and Alice O. Nakamura. 2010. "Bias Due to Input Source Substitutions: Can It Be Measured?" Unpublished.

Drozd, Lukasz A., and Jaromir B. Nosal. 2010. "Pricing-to-Market in Business Cycle Models." Unpublished.

Drozd, Lukasz A., and Jaromir B. Nosal. 2012. "Understanding International Prices: Customers as Capital." *American Economic Review* 102 (1): 364–95.

Erceg, Christopher J., Luca Guerrieri, and Christopher Gust. 2006. "SIGMA: A New Open Economy Model for Policy Analysis." *International Journal of Central Banking* 2 (1): 1–50.

Erickson, Tim, and Ariel Pakes. 2011. "An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods." *American Economic Review* 101 (5): 1707–38.

Feenstra, Robert C. 1994. "New Product Varieties and the Measurement of International Prices." *American Economic Review* 84 (1): 157–77.

Feenstra, Robert C., and Erwin W. Diewert. 2000. "Imputation and Price Indexes: Theory and Evidence from the International Price Program." Unpublished.

Gagnon, Joseph E., and Michael M. Knetter. 1995. "Markup Adjustment and Exchange Rate Fluctuations: Evidence from Panel Data on Automobile Exports." *Journal of International Money and Finance* 14 (2): 289–310.

Ghironi, Fabio, and Marc J. Melitz. 2005. "International Trade and Macroeconomic Dynamics with Heterogeneous Firms." *Quarterly Journal of Economics* 120 (3): 865–915.

Goldberg, Pinelopi Koujianou, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova. 2010. "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India." *Quarterly Journal of Economics* 125 (4): 1727–67.

Goldberg, Pinelopi Koujianou, and Michael M. Knetter. 1997. "Goods Prices and Exchange Rates: What Have We Learned?" *Journal of Economic Literature* 35 (3): 1243–72.

Gopinath, Gita, and Oleg Itskhoki. 2010a. "Frequency of Price Adjustment and Pass-Through." *Quarterly Journal of Economics* 125 (2): 675–727.

Gopinath, Gita, and Oleg Itskhoki. 2010b. "In Search of Real Rigidities." In *NBER Macroeconomics Annual*, edited by Daron Acemoglu and Michael Woodford. Chicago: University of Chicago Press.

Gopinath, Gita, Oleg Itskhoki, and Roberto Rigobon. 2007. "Currency Choice and Exchange Rate Pass-Through." National Bureau of Economic Research Working Paper 13432.

Gopinath, Gita, Oleg Itskhoki, and Roberto Rigobon. 2010. "Currency Choice and Exchange Rate Pass-Through." *American Economic Review* 100 (1): 304–36.

Gopinath, Gita, and Roberto Rigobon. 2008. "Sticky Borders." *Quarterly Journal of Economics* 123 (2): 531–75.

Greenlees, John S., and Robert McClelland. 2011. "Does Quality Adjustment Matter for Technologically Stable Products? An Application to the CPI for Food." *American Economic Review* 101 (3): 200–05.

Griliches, Zvi. 1961. "Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change." In *The Price Statistics of the Federal Government*, edited by the Price Statistics Review Committee, 173–96. Chicago: National Bureau of Economic Research.

Hausman, Jerry. 2003. "Sources of Bias and Solutions to Bias in the Consumer Price Index." *Journal of Economic Perspectives* 17 (1): 23–44.

Hellerstein, Rebecca, Deirdre Daly, and Christina Marsh. 2006. "Have U.S. Import Prices Become Less Responsive to Changes in the Dollar?" *Federal Reserve Bank of New York Current Issues in Economics and Finance* 12 (6): 1–7.

Hobijn, Bart. 2002. "On Both Sides of the Quality Bias in Price Indexes." Federal Reserve Bank of New York, Staff Report 157.

Houseman, Susan. 2007. "Outsourcing, Offshoring, and Productivity Measurement in United States Manufacturing." *International Labour Review* 146 (1-2): 61-80.

Houseman, Susan, Christopher Kurz, Paul Lengermann, and Benjamin Mandel. 2011. "Offshoring Bias in U.S. Manufacturing." *Journal of Economic Perspectives* 25 (2): 111–32.

International Monetary Fund. 2008. "International Financial Statistics." (accessed 2008).

Johansen, Soren. 1995. *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. New York: Oxford University Press.

Kimball, Miles S. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model." *Journal of Money, Credit, and Banking* 27 (4): 1241–77.

Knetter, Michael M. 1989. "Price Discrimination by U.S. and German Exporters." *American Economic Review* 79 (1): 198–210.

Lancaster, Tony. 1992. *The Econometric Analysis of Transition Data*. New York: Cambridge University Press.

Liegey, Paul R., Jr. 1993. "Adjusting Apparel Indexes in the Consumer Price Index for Quality Differences." In *Price Measurements and their Uses*, edited by Murray F. Foss, Marilyn E. Manser, and Allan H. Young, 209–26. Chicago: University of Chicago Press.

Liu, Jingchen, Xiao-Li Meng, Chih-nan Chen, and Margarita Alegria. 2009. "Statistics Can Lie But Can Also Correct for Lies: Reducing Response Bias in NLAAS via Baysian Imputation." Unpublished.

Mandel, Michael. 2007. "The Real Cost of Offshoring." *BusinessWeek*, June 18.

Mandel, Michael. 2009. "Growth: Why the Stats Are Misleading." *BusinessWeek*, June 3.

Marazzi, Mario, and Nathan 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 (6): 924–47.

Marazzi, Mario, Nathan Sheets, Robert J. Vigfusson, Jon Faust, Joseph Gagnon, Jaime Marquez, Robert F. Martin, Trevor Reeve, and John Rogers. 2005. "Exchange Rate Pass-Through to U.S. Import Prices: Some New Evidence." Board of Governors of the Federal Reserve System (U.S.) International Finance Discussion Paper 833.

Marston, Richard C. 1990. "Pricing to Market in Japanese Manufacturing." *Journal of International Economics* 29 (3–4): 217–36.

Moulton, Brent R., and Karin E. Moses. 1997. "Addressing the Quality Change Issue in the Consumer Price Index." *Brookings Papers on Economic Activity* 1997 (1): 305–49.

Nakamura, Emi, and Jón Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123 (4): 1415–64.

Nakamura, Emi, and Jón Steinsson. 2011. "Price Setting in Forward-Looking Customer Markets." *Journal of Monetary Economics* 58 (3): 220–33.

Nakamura, Emi, and Jón Steinsson. 2012. "Lost in Transit: Product Replacement Bias and Pricing to Market: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.102.7.3277.

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

Nordhaus, William D. 1998. "Quality Change in Price Indexes." *Journal of Economic Perspectives* 12 (1): 59–68.

Olivei, Giovanni P. 2002. "Exchange Rates and the Prices of Manufacturing Products Imported into the United States." *Federal Reserve Bank of Boston New England Economic Review* 2002: 3–18.

Pakes, Ariel. 2003. "A Reconsideration of Hedonic Price Indexes with an Application to PC's." *American Economic Review* 93 (5): 1578–96.

Reinsdorf, Marshall. 1993. "The Effect of Outlet Price Differentials on the U.S. Consumer Price Index." In *Price Measurements and Their Uses*, edited by Murray F. Foss, Marilyn E. Manser, and Allan H. Young, 227–54. Chicago: University of Chicago Press.

Reinsdorf, Marshall B., Paul Liegey, and Kenneth Stewart. 1996. "New Ways of Handling Quality Change in the US Consumer Price Index." Bureau of Labor Statistics Economic Working Paper 276.

Rodriguez-Lopez, Jose Antonio. 2008. "Prices and Exchange Rates: A Theory of Disconnection." *Review of Economic Studies* 78 (3): 1135–77.

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

Stigler, George J., Dorothy S. Brady, Edward F. Denison, Irving B. Kravis, Philip J. McCarthy, Albert Rees, Richard Ruggles, and Boris C. Swerling. 1961. *The Price Statistics of the Federal Government*. New York: National Bureau of Economic Research.

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

United States Department of Labor. 2005. *IPP Data Collection Manual*. Washington, DC: Government Printing Office.

Woodford, Mike. 2003. *Interest and Prices*. Princeton, NJ: Princeton University Press.