

Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 1040

January 2012
(revised: December 2012)

Missing Import Price Changes and Low Exchange Rate Pass-Through

Etienne Gagnon, Benjamin R. Mandel, and Robert J. Vigfusson

NOTE: International Finance Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to International Finance Discussion Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors. Recent IFDPs are available on the Web at www.federalreserve.gov/pubs/ifdp/. This paper can be downloaded without charge from Social Science Research Network electronic library at <http://www.ssrn.com/>.

Missing Import Price Changes and Low Exchange Rate Pass-Through

Etienne Gagnon, Benjamin R. Mandel, and Robert J. Vigfusson^{*†}

Abstract

A large body of empirical work has found that exchange rate movements have only modest effects on inflation. However, the response of an import price index to exchange rate movements may be underestimated because some import price changes are missed when constructing the index. We investigate downward biases that arise when items experiencing a price change are especially likely to exit or to enter the index. We show that, in theoretical pricing models, entry and exit have different implications for the timing and size of these biases. Using Bureau of Labor Statistics (BLS) microdata, we derive empirical bounds on the magnitude of these biases and construct alternative price indexes that are less subject to selection effects. Our analysis suggests that the biases induced by selective exits and entries are modest over typical forecast horizons. As such, the empirical evidence continues to support the conclusion that exchange rate pass-through to U.S. import prices is low.

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

Keywords: exchange rate pass-through, import prices, item replacement

^{*}This research was conducted with restricted access to Bureau of Labor Statistics (BLS) data. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, the Federal Reserve Bank of New York, or any other person associated with the Federal Reserve System or the BLS. We thank Carolyn Evans, Hakan Yilmazkuday, Nicholas Li, Emi Nakamura, Jón Steinsson, John Rogers, Kei-Mu Yi, and seminar participants at the Federal Reserve Board, Federal Reserve Bank of New York, Bank of Canada, Banque de France, George Washington University, BLS Price and Index Number Research group, and conference on Microeconomic Sources of Real Exchange Rate Behavior at Vanderbilt University. Kelsey Ayres, Benjamin Hopkins, Gregory Howard, and Robert Sockin provided superb research assistance.

[†]E. Gagnon and R. Vigfusson are at the Federal Reserve Board. B. Mandel is at the Federal Reserve Bank of New York. Comments and suggestions can be directed to etienne.gagnon@frb.gov, benjamin.mandel@ny.frb.org, robert.j.vigfusson@frb.gov.

In conducting monetary policy, central bankers are interested in how much exchange rate movements affect the prices of imported goods (“exchange rate pass-through”) as fluctuations in these prices can, in turn, affect domestic prices and output. Commonly, exchange rate pass-through is measured by regressing changes in published import price indexes on changes in trade-weighted exchange rate indexes along with other explanatory variables. Using these regressions, researchers have estimated low rates of exchange rate pass-through for the United States. Recent estimates (for example, Campa and Goldberg [2005], and Marazzi and Sheets [2007]) suggest that, following a 10 percent depreciation of the dollar, U.S. import prices increase about 1 percentage point in the contemporaneous quarter and an additional 2 percentage points over the next year, with little if any subsequent increases. These low estimates have led several authors to formulate theories which generate incomplete pass-through.¹

In a recent paper, Nakamura and Steinsson (2012) argue that actual long-run pass-through is substantially higher than these standard estimates because published price indexes are missing price changes associated with item replacement. In their words, "*if the prices of new products entering the index have already adjusted to exchange rate movements ... then the response of these prices to movements in exchange rates ... will be 'lost in transit' (i.e., neither picked up by an observed price change of the exiting nor entering product).*" They derive an adjustment factor for standard long-run pass-through estimates under specific assumptions regarding the nature of item substitutions and the way firms set prices. They conclude that measured pass-through is substantially underestimated, implying that exchange rate pass-through to U.S. import prices is much more complete in the long run than was previously thought.

Our paper reassesses this conclusion by developing a general framework for understanding how sample turnover—in conjunction with pricing assumptions—impacts measured pass-through, and by proposing and implementing new ways of gauging the biases in standard estimates. Throughout our discussion, we emphasize the implications of sample exits and sample entries over the first couple years following an exchange rate shock, as such a short-term horizon is most commonly employed by policymakers and is especially useful for discriminating among theoretical models. Our conceptual framework captures the possibility that an item whose price is about to change is more likely than others to leave the index (resulting in a “selective exit”) and that an item whose price recently changed is more likely than others to be added the sample (resulting in a “selective entry”). In either cases, items tend to release their price pressure outside of the period spent in the sample, thus lowering

¹These papers include Atkeson and Burstein (2008), Gopinath, Itskhoki, and Rigobon (2010), Gopinath and Itskhoki (2010), and Gust, Leduc, and Vigfusson (2010).

the measured response of the price index. We show how the effects of item replacement on measured pass-through crucially depend on the nature of *both* exits and entries of items. In turn, the magnitude of these effects can be sensitive to pricing assumptions and the time horizon of interest. By contrast, the presence of real rigidities is typically of limited consequences. In presenting these findings, we substantially expand the analysis of Nakamura and Steinsson (2012), who are primarily concerned with the long-run effects of item substitutions and do not explicitly model sample exits.

Our theoretical work explores the implications of selective exits and selective entries under two popular pricing mechanisms, Calvo and menu costs, which together span a wide range of pricing behavior relevant for understanding pass-through dynamics. The most consequential situation is when price changes are missed both when items exit the sample and when they enter it. In such a case of selective exit and selective entry, standard pass-through estimates omit a roughly constant fraction of the actual price response to an exchange rate movement. This finding applies to all time horizons and both pricing models. When only exits are subject to a selection bias, items tend to leave the sample just before releasing their price pressure. The estimation bias as a share of the cumulative price response is largest initially, but then diminishes over time as items that have yet to respond to the exchange rate movement are brought into the sample and eventually release their price pressure. This compensating effect is largest in the Calvo model due to its relatively slow transmission of shocks. When only entries are subject to a selection bias, measured pass-through is initially unbiased, but a downward bias gradually appears as items entering the sample prove insensitive to past exchange rate movements. The bias grows largest in the Calvo model because of its relatively slow pass-through of shocks, which means that item substitutions are likely to take place before items have time to release their price pressure.

Item turnover is a potentially important source of bias in standard pass-through regressions. We find that there is about one item replacement for every 5 price changes in the U.S. import price index. And as reported by Gopinath and Rigobon (2008) and Nakamura and Steinsson (2012), between 30 and 40 percent of items leave the sample without ever displaying a price change. The potential for economically important biases is thus material, provided that many sample exits and entries are subject to a selection effect. However, based on a range of empirical exercises, we conclude that the biases induced by selective exits and entries, although a concern and worthy of continued research, do not materially alter the literature's view that pass-through to U.S. import prices is low over typical forecast horizons.

To inform our judgement, we first use BLS microdata to derive empirical bounds on the price level response to an exchange rate shock. To do so, we calibrate our Calvo and menu-cost models to match key features of exchange rate movements, sample turnover, and

individual price adjustments, and we then recover the model-implied correction factors. Our empirical analysis focuses on imported finished goods, for which the biases are most likely to be severe. Our baseline empirical estimate of pass-through for these prices after two years is 0.24. After correcting for the most severe case of selective exit and selective entry consistent with the data, this estimate rises only to 0.28.

We next use BLS microdata to construct alternative price indexes in which the inclusion of newly sampled items in the index is delayed. We formally prove that this delaying procedure can substantially mitigate the selective entry bias over typical forecast horizons in a Calvo model. The intuition is that, when entries are selective, the statistical agency is adding items to the index that are too insensitive to past exchange rate movements; delaying their entry in the index gives these items time to be affected by exchange rate shocks, making them more representative of the population and therefore reducing the selection effect. As such, theory predicts that if selective entries were a key source of bias, then these alternative price indexes should imply higher pass-through rates. However, when we estimate pass-through using these alternative price indexes, we find no evidence of bias reduction.

Given the close connection between our work and that of Nakamura and Steinsson (2012), we have attempted throughout to make the similarities and differences between the two papers clear with respect to both modelling assumptions and empirical conclusions. Section 3.5 compares the theoretical long-run correction factors derived by Nakamura and Steinsson (2012) to the ones obtained in our environment at various horizons. Turning to the empirical results, an important difference between our findings and those of Nakamura and Steinsson (2012) reflects our judgment over the incidence of selection biases in item turnover and our treatment of heterogeneity in the observed frequency of price changes across items. As we discuss in section 4.3, correction factors can differ greatly depending on, among other reasons, the assumed mass of observations at very low frequencies of price changes.

The remainder of the paper is structured as follows. Section 1 describes the sample of items used by the BLS to compute import price inflation and provides an overview of item exits and entries. Section 2 introduces the baseline Calvo and menu-cost models that we use to illustrate the nature of the various biases and to gauge their quantitative importance. Section 3 presents the possible biases associated with selection effects in sample exit and entry. Section 4 explores the empirical relevance of these biases by computing bounds on standard pass-through estimates and by constructing an alternative price index that mitigates these biases. Section 5 concludes.

1 Nature and occurrence of item exits and entries

Our study focuses on changes to the BLS import price sample, not on changes to the population (universe) of imported items. For clarity, we reserve the terms “exit” and “entry” for changes in the composition of the sample.² Throughout the presentation of the data and subsequent model-based analysis, we are concerned with the possibility that micro price changes tend to take place just after items exit the sample or shortly before items are added to the sample, so that part of the price response to shocks is censored. We define a “selective exit” as the subtraction of an item from the sample that is triggered by its price being about to change, and a “selective entry” as a systematic addition to the sample of an item that recently experienced a price change. By contrast, a “random exit” and a “random entry” are, respectively, the subtraction from and the addition to the sample of an item without regards to its pricing characteristics.

With the above terminology in mind, the remainder of this section provides background information about the construction of the import price indexes used in standard pass-through regressions, emphasizing the nature and occurrence of sample exits and entries, their treatment by the BLS, the potential for selection biases, and their relationship to micro price adjustments.

1.1 The International Price Program

Given identical data to Gopinath and Rigobon (2008) and Nakamura and Steinsson (2012), we rely on their work to convey the details of the BLS’ International Price Program (IPP) protocol and sample, as well as on the BLS Handbook of Methods. In brief, import prices are collected through a monthly survey of U.S. establishments. The sample consists of rolling groups of items, each item having a sampling duration of about three years, on average. The IPP chooses its firms and items based on a proportional-to-size sampling frame with some degree of oversampling of smaller firms and items.³ Respondents must provide prices for actual transactions taking place as closely as possible to the first day of the month. In total, we observe the price of approximately 13,000 imported items per month from September

²Our analysis thus excludes at the onset the well-known “new good bias” and “quality change bias,” which arise due to difficulties in imputing past reservation prices for items never observed before by consumers. For an introduction to the economics of new goods, see the volume edited by Breshnahan and Gordon (1996). Gordon (2006) and the conference summaries of the Ottawa Group provide overviews of subsequent research on these biases.

³For instance, if there are two items sampled at a firm, one of which has a 90 percent sales share and the other a 10 percent sales share, allocating weights uniformly would over-weight the smaller item. When constructing its aggregate price indexes, BLS corrects for this phenomenon with item-level probability weights.

1993 to July 2007. For the purpose of computing our sample statistics, and consistent with previous studies, we carry forward the last reported price to fill in missing values, effectively overwriting IPP price imputations and firm estimates of prices in non-traded periods.⁴ We also restrict our sample to U.S. dollar transactions, which account for about 90 percent of all observations.⁵

1.2 Nature of exits and entries

BLS price collectors take note when an item exits the sample and assign the retiring item one of the following codes: (1) regular phaseout, (2) accelerated phaseout, (3) sample dropped, (4) refusal, (5) firm out of business, (6) out of scope, not replaced, and (7) out of scope, replaced. Codes (1) through (3) indicate that item exit is driven primarily by the phaseout schedule of the IPP sampling protocol. Codes (4) and (5) describe situations in which price collection is impossible because the survey respondent refuses to respond or ceases to operate, even though the exiting items may continue to be traded in the universe. Codes (6) and (7), which we collectively refer to as “out of scope,” are those instances in which price quotes become unavailable because the item ceases to be traded by importers.⁶

The purpose of item phaseouts is to keep the sample representative of the universe of items; the BLS resamples approximately half of its disaggregated product categories every two years and typically plans to retire items five years after their entry. An item may retire early if it is insufficiently traded. We see such exits, given their planned nature, as unlikely to be selective. Contrary to phaseouts, refusals and importers going out of business are not foreseen events. Nevertheless, we view the risk that such exits systematically mask individual price adjustments as relatively modest, as there are several factors unrelated to micro price adjustments that could trigger them. Exits associated with items becoming out of scope likely present the greatest risk of masking price adjustments. For example, an importer could cease to order an item when faced with a price increase eating away its profit margins. The item could also exit because the foreign producer is adjusting the item’s effective price

⁴For a given item, reporting firms typically do not provide a transaction price every month. The BLS imputes an item’s missing price by either carrying forward the last reported price or by adjusting it by the average price change for the same firm and product category.

⁵Gopinath, Itskhoki, and Rigobon (2010) estimate nearly complete pass-through rates (0.95) to non-dollar import prices after a single price adjustment, compared to only 0.25 for dollar import prices. While excluding non-dollar transactions nudges our estimation results toward finding low pass-through rates, it has the benefit of centering our analysis on transactions for which pass-through is most likely to be underestimated. As such, the correction factors derived in the paper may slightly overstate the extent of the bias when applied to estimates obtained using the full sample of transactions.

⁶In some instances, the firm can provide an alternative item that meets BLS sampling needs (called “replaced”), though that new item would still be recorded as a separate entry. About a third of out-of-scope exits are flagged as “replaced.”

through a change in its characteristics. Other situations leading to out-of-scope items may be unrelated to micro price adjustments. For example, the importer may be curtailing the range of products on offer to streamline its inventory management.

It is worthwhile to note that exits are not generally accompanied by the simultaneous entry of a newly sampled item. When an item suddenly becomes out of scope, BLS price analysts ask the reporting firm whether it can provide another item that meets the same sampling criteria. When possible, the BLS may link the price of the entering and exiting items through a one-time quality adjustment, in which case the estimated change in the effective price is recorded. However, this practice is relatively infrequent. In other instances, the firm may provide an alternative item meeting the BLS sampling needs, though that item is recorded as a separate entry. More often, when no item with similar characteristics is available in the same establishment, or when the planned phaseout date is within the next 18 months, the BLS simply waits until the next biennial sample redrawing. The lag between an unplanned exit and the subsequent item entry can thus be fairly long. Even in the case of planned phaseouts, BLS protocol does not necessitate synchronizing exit and entry.⁷

Notwithstanding the fact that exits and entries are staggered, the size of the IPP sample has been roughly constant since 1993 as the gross number of exits has typically been matched by a corresponding number of entries. The BLS uses probability sampling techniques to select establishments within broad strata of items, and then to select product categories within each stratum-establishment combination. A BLS field agent next conducts an interview with the establishment to select specific items. Probability sampling may be used at that stage. In general, special efforts are made to ensure that selected items are traded regularly, which implies that higher-volume items with established price histories are more likely to be selected.

In principle, the BLS's decision to sample a given item from within the universe should be unrelated to the timing of that item's price changes. Indeed, our reading of the BLS methodology is that the risk of selective entry is somewhat low, especially for those items entering the sample through planned sample redrawing. The risk of selective entries is arguably larger for items entering the sample concurrently with or shortly after an unplanned exit when no quality adjustment is made. Our assessment that the risk of a selective entry is somewhat low, especially if occurring in isolation from a selective exit, stands in contrast with Nakamura and Steinsson's (2012) conclusion that selective entries are an empirically important phenomenon.

⁷For instance, during biennial sample redrawings, some disaggregate product categories may be retired from further sampling but their items may remain in the index until their planned phaseout. Where outgoing and incoming product groups are dissimilar, the benefit to overlapping their items is unclear.

1.3 Accounting for exit and entry

For a given month t , let $exit(t)$ and $entry(t)$ be the number of items exiting and entering the sample, respectively. These items cannot be used in the computation of inflation at month t because their price in either month $t - 1$ or t is missing. For items whose price is available in both month $t - 1$ and t , let $change(t)$ and $no_change(t)$ be the number of observations with a price change and no price change, respectively. We define the exit rate as

$$exit_rate(t) = \frac{exit(t)}{entry(t-1) + change(t-1) + no_change(t-1)}.$$

The denominator in the above expression is the number of items whose price was collected in month $t - 1$. The exit rate thus measures the fraction of items present in the sample at the end of month $t - 1$ that leave in the next month. Analogously, the entry rate is measured as

$$entry_rate(t) = \frac{entry(t)}{entry(t-1) + change(t-1) + no_change(t-1)}.$$

Table 1 shows summary statistics about exits and entries over the period October 1995 to April 2005.⁸ For industry groupings, we use the Bureau of Economic Analysis three-digit Enduse classification to bring descriptions of the microdata closer to the groups of goods commonly used in aggregate pass-through regressions (for instance, Bergin and Feenstra [2009] and Marazzi, *et al.* [2005]). In aggregating up from unique items in a given month to industry-level statistics, we weight each measure by its importance to overall U.S. import purchases.⁹ We aggregate the measures defined above in two stages: first, by computing unweighted statistics for each Enduse category in each month. Then, we aggregate across categories and time periods using the 2006 import sales value of each Enduse category.¹⁰

The rates of item exits and entries are both approximately 3 percent, indicating that the *average* size of the IPP sample remained about the same over the sample period. However,

⁸Incomplete reporting for item discontinuation reasons in the database made available to outside researchers by the IPP truncates our sample at its beginning and end. October 1995 is the first month for which the discontinuation reason field is populated, while the months following April 2005 contained incomplete information about exits at the time of our data extraction.

⁹Doing so assigns the average item frequencies for sampled items and products to those not sampled within the same industry.

¹⁰An alternative weighting scheme would be to use the BLS product weights, which are akin to annual import values at the Harmonized System 10-digit (HS10) level, spread evenly across items within each HS10 product. The Enduse weights for a given month would be the sum total of the individual item weights across items and HS10 products within that Enduse. However, due to incomplete weight data for petroleum (Enduse 100), that method tends to under-weight those high-frequency products in the aggregate statistics. Otherwise, at the end-use level, the measures are quite similar.

Also, ignoring the BLS probability weights for items and firms within each HS10 product, as we do, does not drastically change the summary statistics. Probability-weighted and unweighted statistics are available upon request.

the steadiness of the overall sample size hides a degree of heterogeneity in exit and entry rates at the Enduse product level. For instance, computers and semiconductors (Enduse 213) had an entry rate of 5.0 percent, nearly twice that of agricultural machinery and equipment (Enduse 212). Certain categories (like computers) expanded over the course of the sample as evidenced by higher entry rates relative to exit rates. Differences in net entries likely reflect the changing trade intensity of certain categories over the course of the sample. Exiting items coded as out of scope, which we see as presenting the highest risk of selective exits, accounted for half (1.5 percentage points) of the total exit rate. Enduse categories with a relatively high share of out-of-scope exits include computers and semiconductors, home entertainment equipment, as well as trucks and buses.

By definition, a selective exit entails a price change concurrent with an item leaving the sample. This pattern suggests that the rate of selective exit should vary over time along with macroeconomic variables triggering price adjustments. Evidence of this phenomenon in scanner data is provided by Broda and Weinstein (2010) in their analysis of barcode creation and destruction over the business cycle. To see if exits of imported goods similarly respond to the exchange rate, the top panel of figure 1 presents the time series of the exit rate restricted to out-of-scope items along with an index of the broad nominal dollar.¹¹ This measure is very close to the “endogenous exit” measure reported in Berger *et al.* (2009), with the minor difference that we also exclude exits resulting from refusals and firms out of business. The series is flat at about 1 percent throughout most of the early periods with a transient peak at the beginning of 2000. Then, the out-of-scope rate rises by about 50 basis points in 2003 through 2005. These three prominent features of the time series (i.e., flatness or slight decline early, peak in 2000, and uptick in 2003-5) correspond inversely to the pattern of the broad nominal dollar index, shown in black. The intuition for this relationship is straightforward: as the dollar depreciates, the profitability and viability of a higher proportion of imported items is adversely affected, leading firms to pull the items before the end of their scheduled sample life. We view this evidence as suggestive that exits may, in fact, occur in tandem with price changes.

The occurrence of exits related to factors other than items falling out of scope, which we see as presenting a relatively low risk of selection bias, varies far less systematically with the exchange rate. Rather, the random exit series exhibits the fairly normal pattern of peaks every two years (i.e., the end of 1996, 1998, 2000, 2002, and 2004), which is in line with the biennial shuffling of IPP items. For the most part, the overall entry rate shows a similar pattern with peaks in the middle of the year in 1997, 1999, and so on. Of note, similarly to the out-of-scope exit rate, the rate of overall entry also ticks up towards the end of the

¹¹The exit rates shown in the figure are 12-month moving averages.

sample.

We also note that the timing of the changes in out-of-scope exit rates and, to a lesser extent, in entry rates, does not seem to account for the decline in measured exchange rate pass-through documented in the literature, which has roughly halved since the 1980's. The decrease in pass-through took place primarily in the 1990's, preceding the upticks in exit rates by quite a few years.

1.4 Micro price adjustments

As will be made clear in the next sections, the implications of selective exits and selective entries can be sensitive to underlying pricing frictions. It is convenient for our discussion to define the observed frequency of individual price changes as

$$frequency(t) = \frac{change(t)}{change(t) + no_change(t)}.$$

The overall weighted incidence of price changes for finished goods is estimated to be 6.2 percent. The analogous statistic for the entire IPP import sample (i.e., additionally including industrial supplies, foods, feeds and beverages) is 15.3 percent.¹² These levels are consistent with the weighted average of 14.1 percent in Nakamura and Steinsson (2012) and the median of 15 percent in Gopinath and Rigobon (2008). The average absolute (nonzero) price change is 6.7 percent for finished goods and 8.0 percent overall, in line with the mean overall estimate of 8.2 percent in Gopinath and Rigobon (2008). Here, again, there is significant dispersion across categories, with items belonging to computers and peripherals (Enduse 213) having an average price change of 9.6 percent, compared to 2.0 percent for passenger cars (Enduse 300).

1.5 Other data considerations

We conclude the data description by mentioning two additional elements important for the interpretation of the results. First, in any given month, prices are missing for about 40 percent of items in the sample, which could reflect the absence of a transaction or reporting issues. Second, nearly half of all observations in the BLS sample refer to items that are traded

¹²Some micro price studies include carried-forward prices in their count of usable observations (i.e., the denominator of the frequency formula), while others do not. We follow the former approach. One benefit is that the number of usable observations from one month to the next is directly determined by the number of entries and exits. If we instead excluded carried-forward prices, then our statistics would need to account for the fact that some quotes are inactive. Our decision makes price changes, exits, and entries somewhat less frequent than if carried-forward prices were excluded. The broad findings of the paper do not hinge on this methodological choice, however.

between affiliates or entities of the same company. Although the BLS prefers that these intra-company transfer prices be market-based or market-influenced, some have expressed concern over whether these prices play the same allocative role as market transactions. Excluding intra-company transfer prices from the sample has a negligible impact on our analysis because intra-firm and market transactions have roughly similar entry rates, exit rates, and frequency of price changes. See Neiman (2010) and Gopinath and Rigobon (2008) for a comparison of intra-firm and market transactions.

2 Pass-through and micro price adjustments: a baseline case

This section introduces the baseline Calvo and menu-cost models that we will use to illustrate the nature of the selection biases occurring during sample turnover. Although these models are only two of the many pricing mechanisms used in the literature, together they span a wide range of behavior that are relevant for understanding biases in measured pass-through.

2.1 Economic environment

We assume that the universe comprises an infinite number of imported items, each produced by a single firm. There are no changes over time to the composition of the universe or to item characteristics. The data-generating process for the change in the price (in logs) of item i at period t is

$$\Delta p_{it} = \begin{cases} 0 & \text{if } \mathcal{I}_{it}^f = 0 \\ u_{it} + \beta \Delta x_t + \varepsilon_{it} & \text{if } \mathcal{I}_{it}^f = 1 \end{cases} .$$

The occurrence of a price change is indicated by the variable \mathcal{I}_{it}^f . Given an opportunity (or decision) to change its price, a firm sets Δp_{it} equal to the sum of (a) the amount of price pressure inherited from previous periods, u_{it} , (b) a fraction β of the current change in the exchange rate, Δx_t , and (c) the contribution of a (mean-zero) idiosyncratic factor, ε_{it} . The price pressure carried to the beginning of the next period is

$$u_{it+1} = \begin{cases} u_{it} + \beta \Delta x_t + \varepsilon_{it} & \text{if } \mathcal{I}_{it}^f = 0 \\ 0 & \text{if } \mathcal{I}_{it}^f = 1 \end{cases} .$$

When the firm does not change its price, the shocks occurring in period t are simply added to the price pressure that had already accumulated. If the firm adjusts its price, then the

price pressure is fully released.¹³ The set up so far is quite general and not specific to import prices. One could, for example, interpret Δx_t as the contribution of aggregate shocks, such as wage inflation, to a firm's reset price. In what follows, we simply assume that Δx_t can be represented by an $AR(1)$ process,

$$\Delta x_t = \alpha + \rho \Delta x_{t-1} + \xi_t,$$

with Gaussian innovations, ξ_t .

We are ultimately interested in the impact of exchange rate movements on import prices. To this end, we define aggregate price inflation (in the universe or in the IPP sample) as the average change in item prices,

$$\Delta p_t = \int \Delta p_{it} di.$$

To recover pass-through dynamics the models, we will estimate linear regressions of the form

$$\Delta p_t = a + \sum_{l=0}^L b_l \Delta x_{t-l} + r_t, \quad (1)$$

where r_t is an error term.

2.1.1 Calvo model

In the Calvo model, the decision to change the price is exogenous to the firm. \mathcal{I}_{it}^f equals 1 with constant probability f and 0 with probability $1 - f$. This assumption has strong implications for the dynamic responses of import prices to exchange rate movements. It is convenient to consider the case in which innovations to the exchange rate, Δx_t , are uncorrelated over time ($\rho = 0$), as it allows us to derive analytical expressions for the regression coefficients.

As appendix A shows (in a more general environment), the (plim) linear estimate of b_l for the baseline Calvo model is

$$b_l = f (1 - f)^l \beta. \quad (2)$$

Intuitively, for a movement in the exchange rate l periods earlier to impact an item's price today, the firm must be given the opportunity to adjust its price today (probability f) and no price change must have occurred in each of the previous l periods (probability $1 - f$ in each period). Otherwise, the current price would already reflect Δx_{t-l} . The Calvo model provides a textbook example of a geometric lag model in which the coefficient on the explanatory

¹³Our pricing rule abstracts from forward-looking concerns, which greatly simplifies the exposition. Monthly exchange rate innovations are only weakly correlated, so our pricing rule should nevertheless capture central features of micro price adjustments in response to exchange rate movements.

variable decays exponentially with the number of lags. Summing up the (plim) coefficients in the regression, we get

$$\sum_{l=0}^L b_l = \left(1 - (1 - f)^{L+1}\right) \beta, \quad (3)$$

which converges to β as $L \rightarrow \infty$. Thus, although the effects of an exchange rate shock are never passed-through fully to import prices, the econometrician can nevertheless approximate β (the “long-run” pass-through) with an arbitrary degree of precision.

2.1.2 Menu-cost model

In the menu-cost model, the decision to change the price is the result of a cost-benefit analysis performed by the firm. As shown by Sheshinski and Weiss (1977), it is optimal for the firm to keep its price unchanged if the deviation from the reset price, $u_{it} + \beta\Delta x_t + \varepsilon_{it}$, falls within a certain range. For simplicity, we assume that this range is symmetric and constant across time and firms. Hence, the price adjustment decision is

$$\mathcal{I}_{it}^f = \begin{cases} 0 & \text{if } |u_{it} + \beta\Delta x_t + \varepsilon_{it}| \leq K \\ 1 & \text{if } |u_{it} + \beta\Delta x_t + \varepsilon_{it}| > K \end{cases} .$$

Unfortunately, deriving analytical results is challenging for the menu-cost model unless one is willing to make stringent assumptions (see Danziger [1999] and Gertler and Leahy [2008] for examples). However, the assumptions required for tractability seem less suitable here. Therefore, we will proceed by simulations to illustrate our main points. Note that the decision to change the price now depends on the value of β : All else equal, the larger the pass-through coefficient, the more a shock to the exchange rate is likely to trigger a price adjustment. More generally, the more shocks are large and persistent (and thus associated with relatively large benefits of adjusting the price), the more likely is a firm to change the price immediately. The estimated coefficients in equation 1 are thus sensitive to the particular realization of the shocks in the menu-cost model.

2.2 Calibration of the baseline models

To illustrate pass-through in our baseline models for various degrees of price rigidity, we first set the mean, standard deviation, and autoregressive coefficient of exchange rate innovations to match the corresponding moments of the broad dollar index computed by the Federal Reserve from January 1995 to March 2010. The standard deviation of monthly (end-of-period) exchange rate movements was 1.5 percent over that period, with no apparent drift. Exchange rate movements were slightly autocorrelated ($\rho = 0.19$). We will report results for

$\beta = 0.3$, which is in-line with recent estimates in the literature (e.g., Marazzi *et al.* [2005], Gopinath, Itskhoki, and Rigobon [2010]), but somewhat lower than the consensus value for pass-through in the 1980s (e.g., Goldberg and Knetter [1997]).

The remaining parameters are calibrated to match salient features of individual import price adjustments. As shown by Gopinath and Itskhoki (2010), the median size of import price changes is insensitive to the frequency of price change changes, hovering between 6 and 7 percent. For the Calvo model, we set the probability of a price change equal to the frequency of interest then calibrate the variance of individual innovations to match a median size of price changes of 6.5 percent. For the menu-cost model, we choose the menu cost K and the standard deviation of ε_{it} to match both the observed median size and the average frequency of price changes of interest. We make the additional assumption that ε_{it} is normally distributed with mean zero.¹⁴ The larger is K , the less frequent and the larger are the individual price changes. Likewise, the larger is the standard deviation of ε_{it} , the more frequent and large are individual price changes.

2.3 Impulse response to an exchange rate movement

The choice of a particular pricing model can have important consequences for the dynamic response of import prices. The upper, middle, and bottom panels of figure 2 show the response of import price inflation to an exchange rate movement in our baseline Calvo and menu-cost models for (steady-state) frequencies of price changes of 5 percent, 20 percent, and 35 percent, respectively. Although the models share the same average frequency, average size of adjustments, and long-run pass-through coefficient, the transmission of an exchange rate shock is markedly faster in the menu-cost model than in the Calvo model.¹⁵ In the Calvo model, the frequency of price changes has a direct impact on the speed at which exchange rate disturbances are transmitted to the import price index. For a relatively low frequency (upper panel), the exchange rate movement is still diffusing by the end of the forecast horizon. For a frequency of 20 percent (middle panel), the shock is almost entirely passed-through by the end of the forecast period. Higher frequencies lead to even faster pass-through. The speed of pass-through is markedly higher in the menu-cost model than in the Calvo model at any given frequency. Under our low-frequency calibration, there is

¹⁴We do not posit leptokurtic idiosyncratic shocks as Midrigan (2011) or Gertler and Leahy (2008), an assumption that would make pass-through in the menu-cost model more similar to that in the Calvo model. While leptokurtic shocks help improve the fit of the distribution of individual price changes in the menu-cost model, they work against finding fast pass-through, which is a feature that we want to explore later in our discussion.

¹⁵In practice, the frequency of price changes and the degree of exchange rate pass-through appear inter-related. Gopinath and Itskhoki (2010) present evidence that items with relatively low frequencies of price changes tend to be associated with relatively low rates of pass-through.

negligible amounts of import price inflation left one year after the shock, even for frequencies as low as 5 percent.

3 Selection effects in sample exits and entries

We now alter the baseline models by assuming that the econometrician has only access to a index computed using a large sample of prices from the universe. Whereas the composition of the universe is constant over time, the composition of this index varies as items exit or enter the sample. Prices are collected at the end of the period after nominal adjustments, sample exits, and sample entries have taken place. Entering and exiting items thus cannot be used to compute inflation because either their past or their current price is not observed.

Exits occur through two channels. First, items face a probability d of dropping out of the sample every period (the “random exit” channel). These exits are akin to the sample rotation performed by the BLS in that they do not depend on the amount of price pressure that has cumulated. Second, some exits are triggered by firms changing their prices (the “selective exit” channel). Conditional on its price being changed in the period, an item faces a probability e of exiting the sample. Such a situation could occur if, for example, a firm changed an items’ effective price by altering its characteristics and the price collector, rather than making the required hedonic adjustment, simply dropped the old item from the sample. Selective exits are not “endogenous” since the decision to exit is exogenous to firms. Nevertheless, they capture in a straightforward way the possibility that some exits partly censor the adjustment of the price index. In total, of all items present in the sample at the beginning of the period, a fraction $s_t = d + (1 - d)ef_t$ exits by the end of the period, where f_t is the true frequency of price changes among those items.

For convenience, we postulate that exiting items are replaced by an equal number of entering items, which is a rough approximation of the BLS’ practice over the past two decades. Entries also occur through two channels. A constant fraction $1 - n$ of entering items are drawn at random from the universe of items (the “random entry” channel). The distribution of deviations from the optimal price, u_{it} , is thus the same for these entering items as that of the entire universe, with some fraction of entering items having their price reset during the period. Another fraction n of entering items are sampled *systematically* from price trajectories with a price change in the current period (the “selective entry” channel). Their price already reflects current and past movements in the exchange rate and other disturbances (i.e., $u_{it} = 0$). Similarly to the selective exit channel, the selective entry channel has the feature that items tend to release their price pressure outside of the period spent in the sample.

As was the case earlier, it is convenient to first consider a Calvo model with *i.i.d.* innovations to the exchange rate. We show in the appendix that the (plim) coefficient on the l -th lag of the exchange rate is

$$b_l = \left(\frac{1-e}{1-fe} \right) \left((1-d)^l + \frac{s(1-n)}{d} \left(1 - (1-d)^l \right) \right) f(1-f)^l \beta. \quad (4)$$

The expression in front of β is the probability that a price change occurs in period t and the previous price change took place over l periods earlier. Relative to equation 2, the above expression has two new terms, $\frac{1-e}{1-fe}$ and $(1-d)^l + \frac{s(1-n)}{d} \left(1 - (1-d)^l \right)$, that capture the biases associated with selective exits and selective entries. To gain some intuition about these biases, it is useful to consider four canonical cases. Following our discussion of these canonical cases, we will then relate these cases with the theoretical work in Nakamura and Steinsson (2012).

3.1 All exits and entries are random

When all exits and entries are random (i.e., $s = d$ and $n = 0$), the (plim) coefficients in the Calvo model with *iid* exchange rate innovations are identical to equation 2. In short, standard pass-through regressions are unbiased even though, every period, some fraction $s = d$ (with $d < 1$) of items in the basket is replaced. Intuitively, usable items in the index have the same distribution of deviations from the optimum price as items in the universe. For the same reason, biases are absent in the menu-cost model or when exchange rate innovations are autocorrelated.

3.2 All exits and entries are selective

Consider now the case where both exits and entries are selective (i.e., $s = fe$ and $n = 1$). This situation could arise, for example, if a surveyed firm were to implement an effective price change by altering an item's characteristics and the statistical agency, noting this alteration, were to censor the effective price change by treating the old version as a sample exit and the new version as an unrelated sample entry. When a fraction e of all price changes is censored in this way, the statistical agency is missing a roughly constant share of the index response

to shocks.¹⁶ In the Calvo model with *iid* exchange rate innovations, we have

$$b_l = \left(\frac{1-e}{1-fe} \right) f(1-f)^l \beta. \quad (5)$$

All coefficients are downwardly biased by the same factor $(1-e)/(1-fe)$ relative to the true response. Note that $\hat{f} = \frac{(1-e)}{(1-fe)}f$ is the frequency of price changes observed by the econometrician, so that the estimated coefficients are downwardly biased by a factor \hat{f}/f . This bias can be large even when the exit rate is low because what crucially matters is the prevalence of exits among price changes rather than among all observations in the index.

The left and right panels of figure 3 show the cumulative response of the price index to an exchange rate movement in the Calvo and menu-cost models, respectively, as a share of true long-run pass-through. In addition to $n = 1$, the figure shows the special case $n = 0$ (no selective exit), which we will consider shortly. We tentatively assume that a quarter of all price changes are accompanied by an exit, a proportion roughly equal to the median across three-digit Enduse categories of the worse-case probability of selective exit (0.28) that we estimate later in section 4.2. We leave the other parameters unchanged relative to the baseline calibration described in section 2.2. In both models, the bias is important regardless of how frequently prices are updated. And as noted above, the censoring of price changes through exit reduces the frequency of price changes observed by the econometrician. For underlying frequencies of 5, 20 and 35 percent in the universe of items, the econometrician would observe frequencies of about 4, 16, and 29 percent, respectively, among usable observations in the index.

3.3 All exits are selective and all entries are random

It can be challenging for statistical agencies to know if exits are selective or random as they have to press respondents for information about the circumstances in which exits take

¹⁶There is an interesting analogy between this case of selective exit and selective entry and the well-known “quality-change bias” by which statistical agencies have difficulties accounting for changes in quality from one vintage to the next. In both cases, part or all of an item’s effective price adjustment is censored. The direction of the bias and the effects on measured pass-through differ between the two situations, however. Mismeasured changes in quality can result in either systematic upward or systematic downward biases in inflation, depending on whether the quality change is underestimated (e.g., ignoring improvements in a computer’s processing power) or overestimated (e.g., failing to account for the use of lower-quality components). In practice, the quality change bias is associated with a systematic overestimation of inflation (see BLS [1997b] and Bils [2010]). By contrast, in our canonical case with both selective exits and selective entries, only a fraction of the aggregate price adjustment is recorded, so that inflation is underestimated when it is positive and overestimated when it is negative. As a consequence, pass-through is underestimated, which needs not be the case if inflation is systematically mismeasured in a particular direction due to mismeasured quality.

place. Statistical agencies have more leeway to avoid selection biases in the entry of items in the basket since, in principle, they can design the sampling procedure to randomly select observations from the universe of items. The special case we now consider assumes that all exits are selective while all entries are random (i.e., $s = fe$ and $n = 0$). Although we model the sample exit decision as exogenous to the firm, this case captures the essence of “endogenous exits” problems: Item exits tend to be associated with unobserved price adjustments, so that the price index response to shocks is underestimated.¹⁷ When firms choose the timing of price changes, as in our menu-cost model, exits also tend to be associated with relatively large deviations of individual prices from their optimum.

Starting again with the Calvo model with uncorrelated exchange rate innovations, we have

$$b_l = \left(\frac{1 - e}{1 - fe} \right) (1 + lfe) f (1 - f)^l \beta. \quad (6)$$

The size of the bias now depends on the relative strength of two opposing forces. On the one hand, selective exits censor price adjustments, thus dampening the response of the price index to past exchange rate movements. This force is represented by the term $(1 - e) / (1 - fe)$, which we encountered earlier. On the other hand, exits also create opportunities to introduce items in the index whose price has not changed for some time. This possibility subsequently makes the price level more responsive to past exchange rate movements. This second force is captured by $1 + lfe$. For short lags, the downward bias is the predominant force. In particular, the initial response of the index, $b_0 = \hat{f}\beta$, is always downwardly biased. As we increase the number of lags, $(1 + lfe)$ grows linearly to any arbitrarily large number, so that individual coefficients become upwardly biased at sufficiently long lags. Nevertheless, the cumulative index response remains downwardly biased because the coefficients converge more rapidly to zero.¹⁸

As shown in the left panels of figure 3, assuming that exiting items are replaced by sampling at random from the population ($n = 0$) reduces the size of the bias noticeably over the forecast horizon in the Calvo model relative to the case in which entries are selective ($n = 1$). For frequencies of about 20 percent, the estimated two-year cumulative response is nearly the same as the true one. The randomization of entries mitigates the bias from

¹⁷Greenlees and McClelland (2011) offer evidence that exiting food items in the CPI have larger price changes (including zeros) on average than continuing items. They also argue that current CPI techniques may overestimate the extent of quality adjustments for these categories.

¹⁸When all exits are selective and all entries are random, the long-run response of the index is given by $\frac{1 - fe - (1 - f)e^2}{1 - fe} \beta < \beta$. The same conclusion applies to the long-run response in the general case, $\left(\frac{1 - e}{1 - fe} \right) \left(\frac{f + (d + (1 - d)fe)(1 - n)(1 - f)}{f + d - fd} \right) \beta$. This response is smaller than β and increasing in the fraction of entering items randomly sampled from the universe. Randomizing entries thus reduces the downward bias without eliminating it entirely.

selective exits because some of the entering items have not had a price change in a while, making them responsive to past exchange rate movements. As our figure illustrates, this counterbalancing effect can be quite large, offsetting much of the bias by the end of typical forecast horizons.

The gains from resampling at random are more modest in the menu-cost model (right panels) because pass-through is very rapid. As hinted in equation 6, the counterbalancing effect of random substitutions grows with the number of lags, l , but since coefficients are tiny after a small number of lags, the ultimate impact on cumulative pass-through is modest.

3.4 All exits are random and all entries are selective

We next turn our attention to the case in which all exits are random and all entries are selective ($s = d$ and $n = 1$). This case is similar in spirit to that described by Weber (2012), in which the statistical agency only observes new goods with a lag and hence misses price changes occurring shortly before an item's initiation in the sample. Under these assumptions, the (plim) coefficient on the $l - th$ lag of exchange innovations in our baseline Calvo model is

$$b_l = (1 - s)^l f (1 - f)^l \beta. \quad (7)$$

This expression has a very intuitive interpretation. For a movement in the exchange rate l periods ago to contribute to inflation in the current period, one must observe a price change in the current period (probability f) and no price change or substitution in the previous l periods (constant probability $1 - f$ and $1 - s$, respectively, each period). Price changes and substitutions from period $t - l$ to $t - 1$ result in posted prices that already reflect movements in the exchange rate at period $t - l$. Relative to equation 2, the above expression is downwardly biased by a factor $(1 - s)^l$.

A few comments are worth making. First, the nature of the downward bias is that items entering the basket *systematically* are less sensitive to past movements in the exchange rate than items in general. Including these entering items in pass-through regressions thus lowers estimated pass-through rates. Second, in the special case of $l = 0$, we have $b_0 = f\beta$; the estimated initial impact of an exchange rate movement on the price index is always unbiased. We also note that the share of the true coefficient correctly measured decays exponentially with the number of lags considered. The importance of the bias as a share of the cumulative response thus grows over time, with estimates of the short-run cumulative response being less biased than estimates of the long-run response.

Third, as stressed by Nakamura and Steinsson (2012), the long-run bias under Calvo pricing is most important for product categories with very low frequency of price changes.

The left panels of figure 4 illustrate the bias over the policy-relevant horizon by plotting the cumulative contribution of the coefficients. As seen in the figure, the bias increases in severity with the degree of price stickiness. Only two thirds of the actual cumulative pass-through is correctly estimated at a two-year horizon when the frequency of price changes is 5 percent, and almost one fifth is still missing when the frequency is 20 percent. For a frequency of 35 percent, the econometrician captures more than 95 percent of the true response over the forecast horizon. Under Calvo pricing, only $(1 - s)^l$ of the contribution of lag l to pass-through is correctly estimated. This term typically is decreasing at a slow rate because s is small in practice, meaning that the downward bias kicks in most strongly when much of the exchange rate response occurs at long lags. Under low frequencies of price changes, the coefficients associated with long lags in the Calvo model account for a substantial share of the long-run price response, so that the bias can become large over long horizons.

More generally, the size of the bias appears to be related to the *speed* at which the price index responds to an exchange rate shock. The right panels of figure 4 show the estimated cumulative contribution of the regression coefficients on the various lags of exchange rate movements (the dark-shaded bars), along with the bias left out by the econometrician (the light-shaded bars), under menu-cost pricing. The bias is much less severe than under Calvo pricing. Even for frequencies of price changes as low as 5 percent (upper-left panel), the econometrician captures almost 90 percent of the price index response at the two-year horizon. In the menu-cost model, most of the long-run pass-through occurs in the first few periods following a shock – even at low frequencies – so that the bias does not have time to cumulate to something large.

Finally, our figure depicts the worst-case assumption that all entries are selective ($n = 1$). As noted in section 3.1, there would be no bias if price collectors were replacing exiting items by observations randomly selected from the population ($n = 0$). In the more general case where all exits are random and a fraction n of entries are selective, for the Calvo model, (plim) coefficients are a linear combination of the true coefficient and the biased coefficient under our canonical case,

$$b_l = \left(1 - n + n(1 - s)^l\right) f(1 - f)^l.$$

Given this expression, we have that long-run pass-through is

$$\beta(1 - n) + n \frac{f}{f + s - fs} \beta.$$

Departing from the extreme case of $n = 1$ can substantially reduce the size of the bias. As

a rule of thumb, the bias reduction achieved by the end of our forecast horizon is roughly proportional to $1 - n$, so that, for example, randomizing half of the entries would roughly halve the area represented by the light bars.

3.5 Comparing our four canonical cases with Nakamura and Steinsson (2012)

Having laid out these four canonical cases, we now discuss how the selection biases in our environment relate to the product replacement bias explored by Nakamura and Steinsson (2012). Those authors derive correction factors that apply to pass-through over the infinite horizon under the assumption of Calvo pricing. Our paper extends their analysis in three ways: by emphasizing dynamics, by explicitly considering selection biases in sample exit (Nakamura and Steinsson do not spell out their assumptions about exit), and by exploring both Calvo and menu-cost pricing.

Nakamura and Steinsson (2012) proposition 1 relates observed pass-through to true pass-through using the following equation (expressed here using our notation),

$$\sum_{l=0}^{\infty} b_l = \frac{\tilde{f}}{\tilde{f} + s - \tilde{f}s} [\Phi(1 + \alpha)\beta + (1 - \Phi)\beta],$$

where Φ denotes the fraction of all price changes that are first observed price changes. The parameter α is a factor governing the extent of *overreaction* of the first versus subsequent observed price changes after an innovation in the exchange rate. This parameter is estimated rather than derived from the axioms of their theoretical model. To derive an expression for α requires being explicit about the nature of exits and the frequency of random entry. As such, we compare the model-derived cases in our paper to the model-derived cases in Nakamura and Steinsson (2012) (i.e., where α equals zero).

For the infinite-horizon case, there is a direct connection between Nakamura and Steinsson's (2012) long-run correction factors and those in our environment under Calvo pricing. When all entries are selective ($n = 1$), we can express measured long-run pass-through as $\tilde{f}/(\tilde{f} + s - \tilde{f}s)\beta$, where \tilde{f} is the observed frequency of price changes in the price index. This expression implies the same long-run correction factor, $(\tilde{f} + s - \tilde{f}s)/\tilde{f}$, derived by Nakamura and Steinsson (2012). Further, over the infinite horizon, that correction factor is independent of whether exits are random or selective.¹⁹ The infinite horizon case is

¹⁹The expression $\tilde{f}/(\tilde{f} + s - \tilde{f}s)\beta$ for long-run measured pass-through is perhaps most easily derived for our random exit-selective entry case by setting $\tilde{f} = f$ in equation 7 and summing over all lags. One obtains the same expression in our selective exit-selective entry case by using $s = fe$ and $\tilde{f} = (1 - e)f/(1 - fe)$ to

special, however. As the parsing of our canonical cases showed, the size of biases over finite horizons can vary greatly with the severity of selection biases in exits. For example, when all entries are selective, the initial response of the price index to a shock is unbiased when exits are random (with $b_0 = f\beta$) but biased downward when exits are selective (with $b_0 = (1 - f) / (1 - fe) f\beta$).

Another reason for being explicit about item exit assumptions is that the degree of selective exit affects the bias to measured pass-through at *all* horizons when some entries are random. Our framework has the advantage of allowing us to easily reconstitute the bias to short and long-run pass-through with Calvo pricing, by summing up the coefficients given by our general equation 4. Doing so shows that the size of the bias at *both* short and long horizons is a function of the severity of selection bias in sample exit (e) whenever $n < 1$.

As shown by our comparison of the menu-cost and Calvo models, the nature of the pricing frictions can be quite important in determining the size of the biases. In particular, models with otherwise identical true frequencies of price changes can differ greatly in the severity of the biases. For example, in our random exit-selective entry case, holding constant the observed frequency and exit rate, the menu-cost model had a much smaller bias by the end of our forecast horizon than the Calvo model. What mattered in that case was the speed of pass-through rather than the frequency of price changes per se. In sharp contrast, in our selective exit-selective entry case, the Calvo and menu-cost models with identical exit rates and true frequency of price changes have similar bias.

Summing up, our conceptual analysis shows the importance of understanding the nature of *both* exits and entries, as well as the relevant pricing frictions, in judging of the implications for measured pass-through of selection biases in sample turnover. Our section 4 on the empirical relevance of the biases addresses each of those aspects.

3.6 Robustness to the presence of real rigidities

Gopinath and Itskhoki (2011) find that movements in the exchange rate are passed-through to import prices over more than one price adjustment, consistent with the presence of real rigidities slowing the diffusion of the shock to reset prices. The model adopted so far abstracts from this possibility. However, we show in appendix B that the inclusion of real rigidities has a negligible impact on the theoretical correction factors one should apply to standard pass-through estimates. In particular, we prove that the correction factors for the initial and the long-run responses are independent of real rigidities in the Calvo model, and then show that the factors are overall insensitive to real rigidities at intermediate horizons.

eliminate e and f in equation 5. We thank Emi Nakamura and Jon Steinsson for pointing out this latter link to us.

Appendix B also shows that assuming the presence of real rigidities can alter one’s judgment regarding the empirical relevance of the Calvo versus the menu-cost models. As we shall see shortly, the empirical impulse response of import prices to an exchange rate movement is consistent with features of both models, in particular the initially rapid response of the index predicted by the menu-cost model, and the continued pass-through over medium-term horizons predicted by the Calvo model. The assumption of real rigidities slows predicted pass-through, which makes it more challenging for the Calvo model to match the initial import price index response to an exchange rate movement.

4 Empirical relevance of selective exits and entries

In order to assess the impact of selective exits and selective entries on standard estimates of exchange rate pass-through, one needs to form a view on several objects that are not directly observed, namely the type of pricing frictions giving rise to infrequent nominal adjustments, the extent of price change censoring through exits (e), and the prevalence of entries whose prices are relatively unresponsive to past exchange rate movements (n). In this section, we first argue that standard estimates of the import price response to exchange rate movements mix features of both the menu-cost and the Calvo models. We next simulate the models to derive bounds on the size of the biases over our forecast horizon. Finally, we present an alternative index construction method to purge standard pass-through estimates of much of the bias induced by selective entry.

4.1 Dynamic transmission of exchange rate shocks: data versus models

We focus our empirical analysis on finished goods categories, which account for about 60 percent of the total value of U.S. imports. They comprise automotive products, consumer goods, and capital goods. We leave aside fuel and material-intensive goods because the problems associated with selective exits and selective entries appear relatively benign for those categories given that (i) they are relatively homogeneous products, (ii) they tend to be traded between a large number of buyers and sellers, and (iii) their prices can often be readily observed in electronic trading platforms. In fact, the IPP obtains its crude oil import prices from a source outside of the sampling universe we observe for this paper, which altogether precludes an empirical discussion of exit and entry in that important category. Finally, for an economy as large as the United States, exchange rate movements and the price of fuel and material-intensive categories are arguably simultaneously determined to some degree,

which raises additional econometric issues.

Our estimation period begins in January 1994 and ends in March 2010. For each three-digit Enduse category (indexed by i), we construct a trade-weighted nominal exchange rate, $NEER_{i,t}$, and foreign producer price inflation, $\pi_{i,t}^*$. We then estimate by ordinary least squares the following equation,

$$\pi_{i,t} = \alpha + \sum_{l=0}^{24} b_{i,l} \Delta NEER_{i,t-l} + \sum_{l=0}^{24} c_{i,l} \pi_{i,t-l}^* + \varepsilon_{i,t}.$$

The number of lags is greater than is typically used in empirical pass-through literature. However, given the simulation results reported earlier, the additional lags seem to be an appropriate choice for robustness. The estimated impulse responses to a 1-percent depreciation of the U.S. dollar are presented in figure 5. The largest responses are found for machinery and equipment categories (Enduse 210, 211, 212, and 215), and, especially, for computers and semi-conductors (Enduse 213). Incidentally, this last category is also one for which the BLS makes special efforts to hedonically adjust prices. By contrast, some categories show little if any pass-through over our two-year horizon, notably automobiles and other vehicles (Enduse 300 and 301), apparel (Enduse 400), and home entertainment equipment (Enduse 412).

To compute a response for finished goods, we aggregate our three-digit category responses using 2006 trade weights. As shown in the lower-left panel, finished goods prices climb more than 0.1 percentage point in the first two months following a 1-percent exchange rate depreciation, another 0.1 percentage point over the remainder the first year, and a more modest 0.05 percentage point over the course of the second year. For reference, we also report results for when we regress the price index for finished goods on the exchange rate (the dashed line in the lower-left panel).²⁰ The responses are similar, although as indicated by the 95-percent confidence interval around the estimated responses for the finished goods price index (the gray area), the aggregated response is somewhat faster than the response estimated for the aggregate price index. The shape of the impulse response shares aspects of both the menu-cost and Calvo models. The initially rapid response is qualitatively similar to that in the menu-cost model, whereas the ensuing slow but steady increase is more akin to the protracted response in the Calvo model.²¹

²⁰Our simple specification does not allow for variation in the magnitude of the response over time, a restriction imposed in part due to the short period over which monthly import price data are available. Taking advantage of the longer time coverage of quarterly series, several authors have documented a decline in pass-through rates in recent decades (e.g., Marazzi et al. [2005]), including for finished goods (see Gust et al. [2010]). As mentioned earlier, we find little evidence that an increase in the occurrence of selective exits and entries could account for that pattern.

²¹Rather than using a mixture of Calvo and menu-cost models, an alternative approach would be to

We next compare the empirical response in each three-digit Enduse category to those generated by our baseline Calvo and menu-cost models. As section 3 showed, the pricing models have different implications for the size of the correction factors. As such, in deriving empirical bounds in the next section, we want to use our model-derived responses to make the best approximation of the richer observed empirical responses. Figure 6 directly compares the empirical responses in each three-digit Enduse category to those generated by the Calvo and menu-cost models. The models are calibrated to match category-level statistics as outlined in section 2.2, with the minor difference that we seek to match the observed cumulative rate of pass-through in the last quarter of the forecast horizon rather than some illustrative long-run value. Figure 6 also shows the linear combinations of model responses that minimize the Euclidian distance with the empirical response over the forecast horizon. Again, we find support for both models, with some Enduse categories clearly preferring one model over the other, and others being best represented by a mixture of the two models. On average, each model is attributed about half of the weight.²² To be clear, this exercise is not an attempt to make a definitive case of the relative merits of Calvo versus menu-cost models. Rather, the exercise is to estimate the best approximation of the richer observed empirical responses given our model-derived responses. In the next section, we use these approximations in calculating the bounds that are most appropriate given the observed responses.

4.2 Bounding standard pass-through estimates

To assess the quantitative importance of selective exits and entries, our next strategy is to derive three sets of bounds on the amount of exchange rate pass-through over the policy horizon. These bounds are related to the canonical cases discussed in sections 3.2 to 3.3, depending on whether we consider, respectively, the largest plausible number of selective exits and entries consistent with the data, the largest plausible number of selective exits in the presence of random entries, or the largest plausible number of selective entries in the presence of random exits.

Our worst case of selective exits assumes that all out-of-scope exits mask a price change. We treat exits for other reasons (as defined in table 1) as random because they typically are planned years in advance by the BLS and thus unlikely to be related to individual pricing decisions. Under these assumptions, we observe the rate of random exit, d , and the rate of selective exits, $ef(1 - d)$, as they correspond to the rate of out-of-scope and other exits

approximate the responses using heterogeneous Calvo models. However, the empirical evidence in Gagnon, López-Salido, and Vincent (2012) favors rules that are not purely time dependent.

²²Though the model responses displayed assume no selection effects, this finding is robust to assuming any degree of selective exits or selective entries in the calibration.

shown in table 1. Knowledge of these rates and of the observed frequency of price changes, $(1 - e) f / (1 - fe)$, is sufficient to identify d , e , and f in the model. Our worst case of selective entries occurs when all items added to the sample experience an unobserved price change upon entry (i.e., $n = 1$).

Given the sensitivity of biases to pricing assumptions, we derive our bounds under both Calvo and menu costs. Under Calvo, we compute the correction factors for the estimated cumulative response to an exchange rate movement directly from the analytical expressions for the estimated coefficients, shown in equations 5 to 7. Under menu costs, no such expressions are available, so we compute the corresponding correction factors through simulations. In particular, for each three-digit Enduse category, we select σ_ε , K , and β to match the observed frequency of price changes, the average absolute size of price changes, and the cumulative amount of pass-through by the last quarter of the forecast horizon following a 1-percent depreciation of the dollar. We then apply the correction factors to the estimated responses and aggregate them using 2006 trade shares to derive our bounds on the response of finished goods.

Our worst case of selective exits and entries is shown in the upper-left panel of figure 7. The estimated finished goods price response to a 1-percent depreciation of the dollar is 0.24 percent by the last quarter of the forecast horizon. After correcting for selective exits and selective entries, the figure for the menu-cost and Calvo models are in a similar range at 0.30 percent and 0.32 percent, respectively. If we instead assume that all entries are random (upper-right panel), then the corrected estimate in the last quarter falls to at most 0.26 percent under Calvo, and to at most 0.28 percent under menu costs. The bias is larger under menu costs in this case because the benefits from randomizing entries are largest when pass-through is slow, as in the Calvo model.²³ Under our worst case of selective entry and random exit (lower-left panel), the corrected response is very close to the actual response under menu costs, but remains somewhat higher (0.30 percent) than the uncorrected estimate (0.24 percent) by the last quarter of the forecast horizon.

Given our earlier evidence that features of both models are present in the data, we derive a fourth set of bounds under what we view as more plausible pricing assumptions. The three-digit Enduse responses of the Calvo and menu-cost models are first weighted according to the linear combination that provides the best fit of the empirical response, as we did in the previous section, and then aggregated using 2006 trade shares. We posit that all out-of-scope exits correspond to selective exits and selective entries, whereas all other exits are associated with random exits and random entries, so that about half of all exits and entries are selective.

²³We assume that individual (nonzero) price changes leading to exits have the same distribution as those for items remaining in the sample. Whether this is the case in practice remains an open question.

The corrected cumulative response under this last set of assumptions is presented in the lower-right panel of figure 7. Following a 1-percent depreciation of the dollar, the corrected cumulative response by the last quarter (0.28 percent) is above the estimated one (0.24 percent). Selective exits and selective entries contribute roughly equally to this difference, as hinted by the special case with only selective entries that is also displayed in the panel.

Summing up, our bound analysis suggests that, even under the stringent assumptions consistent with the microdata, the biases induced by selective exits and selective entries have a limited impact on standard pass-through estimates for finished goods over typical forecast horizons. Given that this conclusion contrasts, at least on the surface, with Nakamura and Steinsson’s (2012) assertion that the product replacement bias is quantitatively important, we next discuss how our respective findings can be reconciled.

4.3 Reconciling our empirical results with those of Nakamura and Steinsson (2012)

We highlight four differences between our empirical work and that of Nakamura and Steinsson (2012) informing our conclusion that downward biases due to item replacement are not as severe as asserted by these authors: the price index of interest, the time horizon of interest, the assumed distribution of heterogeneity in pricing decisions, and the extent of the selectivity in exits and entries. To facilitate comparisons, we begin by estimating an equation for empirical pass-through similar to equation 1 in their paper,

$$\Delta p_t^m - \Delta p_t^{cpi} = a + \sum_{l=0}^6 b_l \Delta s_{t-l} + \varepsilon_t. \quad (8)$$

We regress the difference between changes in the log of quarterly import prices excluding oil, Δp_t^m , and changes in the log of the CPI excluding food and energy, Δp_t^{cpi} , on a constant as well as the contemporaneous and first 6 lags of the change in the log of the Federal Reserve’s trade-weighted major currencies real exchange rate index, Δs_{t-l} . Our estimation period is 1995:Q1 to 2010:Q4 and the import price index is the deflator from the National Economic Accounts published by the U.S. Bureau of Economic Analysis.

As table 2 shows, our estimated import price response after 6 quarters is 0.41, a figure nearly identical to that reported by Nakamura and Steinsson (0.43) for the same horizon. Table 2 also presents standard pass-through estimates after 6 quarters for two special groups of products: finished goods (Enduse categories in the 200s, 300s, and 400s) and material-intensive goods excluding oil (Enduse categories below 200, excluding 100). The rate of pass-through after 6 quarters is much lower for finished goods prices (0.26) than for

material-intensive goods prices (0.93), highlighting the two groups' sharply different medium-term responses to exchange rate movements. The high pass-through estimate for material-intensive goods is perhaps not surprising given these items' frequent price adjustments, their high commodity content, and the strong comovement between commodity prices and the exchange rate. With pass-through nearly complete after 6 quarters and with their prices being frequently adjusted, it is unlikely that material-intensive goods are plagued by severe downward biases in measured pass-through due to item turnover, further justifying our focus on finished goods in earlier sections.

We momentarily assume that true long-run pass-through, β , is the same for all items in the IPP sample, leaving only the frequency of price changes free to vary across items. The assumption of identical long-run pass-through rates across product categories is admittedly unappealing given our evidence above for finished goods and material-intensive goods, but it makes it easier to illustrate the importance of other aspects of the data in comparing our findings with those of Nakamura and Steinsson (2012).²⁴ We set the monthly rate of item substitution for all Enduse categories to 2.5 percent, a figure in the range of rates reported in table 1.²⁵ For now, we also maintain our assumption that items within each three-digit Enduse categories share the same frequency of price changes. Finally, we posit that the data are generated by a Calvo pricing model.

We obtain corrected pass-through estimates for each of our canonical cases by applying a correction factor, Λ , to the sum of estimated pass-through coefficients from equation 8. Rather than reporting solely long-run correction factors, as Nakamura and Steinsson (2012) do, we also compute correction factors pertaining to a 6-quarter horizon (i.e., an 18-month horizon). The correction factor for measured pass-through after L months is given by

$$\Lambda(L)^{-1} = \sum_i w_i \left(\frac{\sum_{l=0}^L b_{l,i}^{biased}}{\sum_{l=0}^L b_{l,i}^{unbiased}} \right).$$

The right-hand side expression is the weighted share of true pass-through after L months measured by the econometrician, where w_i is the relative weight of Enduse i in the index. The term $b_{l,i}^{unbiased}$ is the theoretical unbiased pass-through coefficient on the l -th lag of exchange rate movements for Enduse i (given by equation 2). The term $b_{l,i}^{biased}$ is the corresponding biased coefficient in the presence of selective exits and selective entries (given by equation 4 for the general case). The only remaining step for computing $\Lambda(L)$ is to re-express $b_{l,i}^{unbiased}$

²⁴We note that Nakamura and Steinsson (2012) also consider the effect of cross-sectional heterogeneity in pass-through.

²⁵Allowing the substitution rate to vary across 3-digit Enduse categories based on the entry rates displayed in table 1 has only a minor impact on our findings.

and $b_{i,i}^{biased}$ in terms of observables—namely \hat{f}_i and s_i —by replacing d , e , and f in equations 2 and 4 as needed. See each of our canonical cases for the mapping.

Table 2 reports corrected pass-through rates for our four canonical cases after six-quarters and in the long-run. When all exits and entries are random, both $\Lambda(18)$ and $\Lambda(\infty)$ are equal to one, so that no correction is required. When all exits and entries are selective, $\Lambda(18)$ and $\Lambda(\infty)$ are equal and corrected pass-through for all imported goods is largest at 0.55. For the other two canonical cases, $\Lambda(18)$ and $\Lambda(\infty)$ are not equal. When exits are selective and entries are random, $\Lambda(18)$ is larger than $\Lambda(\infty)$, and as such corrected pass-through declines from 0.48 to 0.45 as the horizon increases from 6 quarters to infinity. In reverse, when exits are random and entries are selective, $\Lambda(18)$ is less than $\Lambda(\infty)$, and corrected pass-through increases from 0.48 to 0.55 as the horizon lengthens to infinity. Hence, one important difference between our results and those of Nakamura and Steinsson (2012) is that we report correction factors consistent with different time horizons and with different explicit assumptions about the selectivity of exit and entry.

As noted above, because prices for material-intensive goods prices are frequently adjusted, there is little downward bias attributable to item replacement. Indeed, for all biased cases, the theoretical correction factors for material-intensive goods are quite small over the horizons considered. We also reiterate that statistical agencies can significantly mitigate downward biases by following sampling procedures that randomize entries. In particular, the long-run correction factor for all goods shrinks from 1.2 under the worst case of selective exits and selective entries to only 1.08 when entries are randomized.

Another important difference between our analysis and that of Nakamura and Steinsson is the treatment of heterogeneity in the frequency of price changes across items. Throughout our paper, we have assumed that the frequency of price changes is constant within three-digit Enduse categories, otherwise leaving the frequency free to vary across sectors of activity. This approach may leave some heterogeneity unaccounted for within three-digit Enduse categories, which could lead us to underestimate the magnitude of downward biases. Nakamura and Steinsson (2012) instead control for heterogeneity across firms by assuming that the observed frequency of price changes in the index is distributed according to a Beta distribution, which they estimate on observed individual frequencies. To contrast the two approaches, we adopt their parametrization of the Beta distribution (using $a = 0.44$ and $b = 3.50$ as parameters) and compute the long-run correction factors under the assumption of selective entries. (Recall from our discussion in section 3.5 that long-run pass-through under Calvo pricing does not depend on the prevalence of selective exits when all entries are selective.) As we change our distributional assumptions, the pass-through estimate for imported goods leaps from 0.55 to 0.70. This large correction is driven by a mass of observations at very low frequencies

of price changes. One salient feature of the Beta distribution estimated by Nakamura and Steinsson (2012) is the implication that a third of all observations have an observed monthly frequency below 2 percent, with the density of observations becoming arbitrarily large as the observed frequency approaches zero.²⁶ Or, the long-run correction factor under selective entries, $(\hat{f} + s - \hat{f}s) / \hat{f}$, is quite sensitive to the presence of very low frequencies because it becomes arbitrarily large as \hat{f} approaches zero.²⁷

To illustrate the sensitivity of the estimated long-run response to the mass of firms updating prices infrequently, we move all observations with a frequency below 2 percent to a mass point at 2 percent. The correction factor for the long-run response drops from 1.70 to 1.47. This decline should caution one against driving strong conclusions that depend on the very low frequencies of price changes, whose density is arguably difficult to estimate. In particular, the Beta distribution has only two parameters to capture the whole range of variation in the density of observed frequencies in the sample. It is thus conceivable that some of the very low frequencies implied by the calibration are off the empirical support.

Although the correction factors presented in table 2 are useful for exploring the sensitivity of the biases to alternative assumptions, we shall stress in concluding this comparison that we see most of them as implausibly large empirically because they assume rates of selective exits and selective entries that are overly severe. As we argued in section 1, not all exits appear to be selective; the substitution rate of 2.5 percent used our table would halved if we were using only out-of-scope exits as our measure of selective exits. Similarly, we found limited evidence in section 1.2 that the BLS methodology is conducive to systematically adding items recently repriced to the IPP sample. Regardless of the nature of exit, the correction factors shrink greatly as we lower the value of n , the fraction of entries that are selective. For example, under the assumption of random exit and assuming that 50 percent of entries are selective, then the correction factor is only 1.26 under Nakamura and Steinsson’s Beta distribution. As noted above, the randomization of entries can reduce downward biases much even in the presence of strong selectivity in exits. In short, we see the relatively tight bounds derived in section 4.2 as more indicative of the range of empirically plausible corrected pass-through estimates.

²⁶Whether this parametrization of the distribution of *observed* individual frequencies also implies that the distribution of *actual* frequencies of price changes has a large mass near zero depends on one’s views about the degree of selectivity in exits. In particular, if all exits are random, then the two distributions are identical, so that many price-setters rarely adjust prices. Nakamura and Steinsson (2012) are not explicit about their assumptions on this point.

²⁷In an appendix, Nakamura and Steinsson (2012) also estimate Beta distributions separately for 15 sectors. While this helps capture sectoral heterogeneity, it is still the case in each sector that the slope of the Beta distribution approaches infinity as the frequency of price changes falls to zero, which tends to boost the imputed aggregate correction factor.

4.4 Reducing biases through delayed entries

If the estimates were subject solely to a bias from selective entry, then one could use a simple trick to remove much of that bias over the policy-relevant horizon. Recall that selective entry arises because entering items are systematically less responsive to past exchange rate movements than items in the universe. Therefore, simply delaying the entry of substitutes into the index should reduce this bias. We show in the appendix that, when all exits are random and all entries are selective, the estimated (plim) coefficients in a Calvo model with an arbitrary M -period entry delay are given by

$$b_l = \begin{cases} f(1-f)^l \beta & \text{if } l \leq M \\ (1-s)^{l-M} f(1-f)^l \beta & \text{if } l > M \end{cases} .$$

Delaying entries thus eliminates the bias due to selective entry for the coefficients associated with the first M lags of exchange rate movements. The bias on subsequent lags is also reduced, with b_l representing a fraction $(1-s)^{l-M}$ of the true response when entries are delayed by M periods, compared to only $(1-s)^l$ when there is no entry delay.

The left panels of figure 8 show that delaying entries by 6 months can go a long way in correcting measured pass-through over policy-relevant horizons in the Calvo model when the only selection bias is in sample entry. The bias at the end of the horizon is negligible when prices are adjusted 20 percent of the time or more. Even at frequencies as low as 5 percent, the prediction over the first year of the forecast suffers little bias, while the accuracy of the response in the second year is greatly improved. The bias reduction is even larger in the menu-cost model (right panels). Delaying entries by 6 months virtually eliminates the downward bias at all frequencies considered. The consistency gains are especially large in the menu-cost model because delaying entries corrects most effectively biases associated with short lags of the exchange rate, which account for the bulk of the price level response.

The delay can also improve accuracy under our most biased canonical case with both selective exit and selective entry. In such a case, the delay again acts as a mechanism to randomize entry, which, as we illustrated in section 4.3, can go a long way in reducing the bias. Our simulations (not shown) indicate that the bias reduction achieved is sensitive to the true frequency of price changes and the length of the delay. For example, the bias is reduced by about 20 percent when the true frequency of price changes is 5 percent and the delay is 9 months, but by only half as much when the delay is 6 months. If the frequency were instead 15 percent, then a 9-month delay would eliminate over half of the bias.

It turns out that our trick of delaying entries can also mitigate biases when the only source of bias is selective exit. We show in the appendix for the general case with arbitrary

degrees of selective exit and selective entry that the (plim) regression coefficients under an M -period entry delay in the Calvo model are

$$b_l(M) = \begin{cases} \frac{(1-e)}{(1-fe)^{l+1}} f (1-f)^l \beta & \text{if } l \leq M \\ \frac{(1-e)}{(1-fe)^{M+1}} \left((1-d)^{l-M} + \frac{s(1-n)}{d} \left(1 - (1-d)^{l-M} \right) \right) f (1-f)^l \beta & \text{if } l > M \end{cases} .$$

We also prove that the bias diminishes as one increases the entry delay given any forecast horizon. As one delays entries by an arbitrary large number of periods, we have

$$\lim_{M \rightarrow \infty} \sum_{l=0}^{\infty} b_l(M) = \sum_{l=0}^{\infty} \frac{(1-e)}{(1-fe)^{l+1}} f (1-f)^l \beta = \beta.$$

In short, the estimated long-run pass-through in the Calvo model is unbiased in the limit, a result that holds whether exits are selective, entries are selective, or both. The short-run response remains downward biased in the presence of selective exit, however.

The intuition for why delaying entries can improve pass-through estimates when only exits are selective is somewhat subtle. Remember that, for a movement in the exchange rate l periods ago to have an impact on the index today, there must have been no price change over the past l periods. Delaying entries by M periods eliminates observations incorporated into the index in recent periods, leaving only those present in the index for at least M periods. Or, these surviving observations are less likely than observations in the universe to have experienced a price change over the past l period (since observations with a price change are more likely to have exited), meaning that they are relatively more likely to contribute to measured inflation today. Under Calvo pricing (left panels of figure 9), it turns out that this selection effect perfectly offsets the downward bias stemming from the censoring of price changes as we consider an arbitrarily long entry delay and forecast horizon. Under menu-cost pricing (right panels), the gains are negligible due to the greater mixing of observations.

Summing up, our analysis suggests that delaying entries is most effective at reducing biases associated with selective entry. If estimated pass-through over the forecast horizon increases much after delaying entries, then selective entry may be economically important, giving credence to the canonical cases implying the most severe biases over long horizons. By contrast, if estimated pass-through is insensitive to delaying entries, then selective entries may be unimportant. Using the BLS microdata, we have computed price indexes for the Enduse categories belonging to capital goods, automotive products, and consumer goods. We have constructed one index using the methodology introduced in section 1 (i.e., there is no entry delay and missing prices are carried forward). We have also constructed two alternative indexes that implement a 6-month and a 9-month entry delay, respectively. Figure 10 displays

the results. As the left-hand column shows, these alternative price indexes are somewhat more volatile than the corresponding published BLS index (the thick black line), especially when entry is delayed. However, estimated pass-through rates are very similar whether we use the published BLS index or our constructed indexes. If selective entry were quantitatively important, then the estimated pass-through rates for the constructed indexes with delayed entry should be noticeably greater than the pass-through rates for the published index or for the constructed index with no delay. Instead, the estimated pass-through rates are very similar. Thus, the available evidence suggests that selective entry is unlikely to be a large driver of biases in pass-through regressions.

5 Concluding remarks

We have investigated selection biases in standard exchange rate pass-through regressions that arise from missing price changes either due to item exit from or entry into the index. For both Calvo and menu-cost pricing models, we have shown that these selection effects lower the measured response of an import price index to exchange rate movements over typical policy horizons and that the magnitude of the biases can be sensitive to pricing assumptions.

In particular, in the presence of both selective exits and selective entries, the import price response is biased downward in both models. Assuming that entering items are sampled randomly from the universe alleviates some of the bias, especially under Calvo pricing. When entries are selective and exits occur at random, the downward bias tends to be small in the menu-cost model over any horizon, whereas the bias slowly grows from being negligible at short horizons to quite large over extended horizons in the Calvo model.

Assessing the quantitative importance of the biases is inherently challenging because selective exits and selective entries are, by their very nature, not observed. Our review of the BLS methodology suggests a moderate risk of such selection effects taking place in practice. We also argue that, under plausible assumptions about nominal price stickiness and the incidence of selective exits and selective entries, the presence of downward biases in standard pass-through regressions, although a concern, does not materially alter the literature's view that pass-through to U.S. import prices is low over typical forecast horizons. Even under our worst-case scenario, our estimated empirical bounds imply that at most about a third of an exchange rate shock is passed through to the price of imported finished goods after two years. We note that finished goods categories are precisely those for which one would expect the bias due to selective exit and selective entry to be most pronounced; material-intensive goods have fast and nearly complete pass-through, and hence little scope for large correction factors. Our judgment that biases are likely small empirically is further informed

by the insensitivity of measured pass-through to delaying the entry of items in the sample. Had the estimates been plagued by the most severe cases of downward bias, then measured pass-through would have risen as we delayed sample entry.

Although we have focused on import prices, our findings are relevant to the study of any price index subject to selection effects in sample exit and sample entry. The implications of selective exits and selective entries also extend to the measurement of the response of price indexes to aggregate and idiosyncratic shocks other than exchange rate movements.

Future research should aim at better identifying the causes of item exits as well as the characteristics of added items. Currently, the information contained in the IPP database provides useful and suggestive, but ultimately limited, guidance on these aspects.

References

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

Berger, David Berger, Jon Faust, John H. Rogers, and Kai Steverson (2009). "Border Prices and Retail Prices," International Finance Discussion Papers 972, Board of Governors of the Federal Reserve System.

Bergin, Paul, and Robert C. Feenstra (2009), "Pass-Through of Exchange Rates and Competition between Floaters and Fixers," *Journal of Money, Credit and Banking*, vol. 41, pages 35-70.

Bils, Mark (2009). "Do Higher Prices for New Goods Reflect Quality Growth or Inflation?," *Quarterly Journal of Economics*, Vol. 124(2), pages 637-675.

Bresnahan, Timothy F., and Robert J. Gordon (1997). "The Economics of New Goods," *Studies in Income and Wealth*, Volume 58, National Bureau of Economic Research.

Broda, Christian, and David E. Weinstein (2010). "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, vol. 100(3), pages 691-723.

Bureau of Labor Statistics. "Handbook of Methods," Washington DC, April 1997.

Bureau of Labor Statistics. "Measurement Issues in the Consumer Price Index," Washington D.C., June 1997, <http://www.bls.gov/cpi/cpigm697.htm> (last access: March 2010).

Boskin, Michael J., E. Dulberger, R. Gordon, Z. Griliches, and D. Jorgenson (1996). "Toward a More Accurate Measure of the Cost of Living," Final Report to the Senate Finance Committee, December 4.

Campa, José Manuel and Linda S. Goldberg (2005). "Exchange Rate Pass-Through into Import Prices," *Review of Economics and Statistics*, vol. 87(4), pages 679-690.

- Danziger, Leif (1999). “A Dynamic Economy with Costly Price Adjustments,” *American Economic Review*, vol. 89(4), pages 878-901.
- Gagnon, Etienne, David López-Salido, and Nicolas Vincent (2012). “Individual Price Adjustment along the Extensive Margin,” *NBER Macroeconomics Annual 2012*, vol. 27.
- Gertler, Mark and John Leahy (2008). “A Phillips Curve with an *Ss* Foundation,” *Journal of Political Economy*, vol. 116(3), pages 533-572.
- 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*, vol. 101(3), pages 200-205.
- Goldberg, Pinelopi and Michael M. Knetter (1997). “Goods Prices and Exchange Rates: What Have We Learned?,” *Journal of Economic Literature*, vol. 35(3), pages 1243-1272.
- Gopinath, Gita and Oleg Itskhoki (2010). “Frequency of Price Adjustment and Pass-Through,” *Quarterly Journal of Economics*, Vol. 125(2), pages 675–727.
- Gopinath, Gita and Oleg Itskhoki (2011). “In Search of Real Rigidities,” *NBER Macroeconomics Annual*, 2010, vol. 25, , pages 261-309.
- Gopinath, Gita and Roberto Rigobon (2008), “Sticky Borders,” *Quarterly Journal of Economics*, vol. 123(2), pages 531-75.
- Gopinath, Gita, Oleg Itskhoki, and Roberto Rigobon (2010). “Currency Choice and Exchange Rate Pass-Through,” *American Economic Review*, vol. 100(1), pages 304-36.
- Gordon, Robert J. (2006). “The Boskin Commission Report: A Retrospective One Decade Later,” *International Productivity Monitor*, vol. 12, pages 7-22.
- Gust, Christopher, Sylvain Leduc, and Robert J. Vigfusson (2010). “Trade Integration, Competition, and the Decline in Exchange-Rate Pass-Through,” *Journal of Monetary Economics*, vol. 57(3), pages 309-324.
- Marazzi, Mario, Nathan Sheets, Robert J. Vigfusson, Jon Faust, Joseph Gagnon, Jaime Marquez, Robert F. Martin, Trevor Reeve, John Rogers (2005). “Exchange Rate Pass-Through to U.S. Import Prices: Some New Evidence,” *International Finance Discussion Papers 833*, Board of Governors of the Federal Reserve System.
- 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*, vol. 26, pages 924-47.
- Midrigan, Virgiliu (2011). “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, vol. 79, pages 1139–1180.
- Nakamura, Emi and Jón Steinsson (2012). “Lost in Transit: Product Replacement Bias and Pricing to Market,” *American Economic Review*, vol. 102(7), pages 3277–3316.

Neiman, Brent (2010). “Stickiness, Synchronization, and Passthrough in Intrafirm Trade Prices,” *Journal of Monetary Economics*, April 2010.

Sheshinski, Eytan and Yoram Weiss (1977). “Inflation and Costs of Price Adjustment,” *Review of Economic Studies*, vol. 44(2), pages 287-303.

Weber, Henning (2012). “Product Replacement Bias in Inflation and Its Consequences for Monetary Policy,” *Journal of Money, Credit and Banking*, vol. 44(2-3), pages 255-299.

A Regression coefficients in the Calvo model

This appendix derives analytical expressions for measured pass-through coefficients when the data are generated by a Calvo model with selection biases in sample exit and entry. The environment is as described in section 3 with the extra simplifying assumption that exchange rate innovations are uncorrelated over time. We begin by describing the general case. We then investigate how delaying the entry of items in the index affect the regression coefficients. We finally prove that the bias on the coefficients declines as one delays the entry of items in the index.

A.1 General case

Let \mathcal{I}_{it}^f , \mathcal{I}_{it}^d , and \mathcal{I}_{it}^e be indicator variables that an item i present in the sample at the beginning of period t has experienced, respectively, a price change, a random exit, and a selective exit (the latter being conditional on a price change and no random exit). For any exiting item, we also define an indicator variable \mathcal{I}_{it}^n that the corresponding entry is selective. For convenience, let also $\mathcal{I}_{it}^s = \mathcal{I}_{it}^d + (1 - \mathcal{I}_{it}^d) \mathcal{I}_{it}^f \mathcal{I}_{it}^e$ be an indicator that an item has exited during the period, either through a random exit (probability d) or a selective exit (probability $(1 - d)fe$).

We first derive an expression for the contemporaneous impact of an exchange rate movement on the price index. Using the covariance approach, we have

$$\begin{aligned} b_0 &= \frac{\text{cov} \left(\int \Delta p_{it} di, \Delta x_t | \mathcal{I}_{it}^s = 0 \right)}{\text{var} (\Delta x_t)} = \frac{\int \text{cov} (\Delta p_{it}, \Delta x_t | \mathcal{I}_{it}^s = 0) di}{\text{var} (\Delta x_t)} \\ &= \frac{\text{cov}(u_{it} + \beta \Delta x_t + \varepsilon_{it}, \Delta x_t | \mathcal{I}_{it}^s = 0, \mathcal{I}_{it}^f = 1)}{\text{var} (\Delta x_t)} \Pr \left[\mathcal{I}_{it}^f = 1 | \mathcal{I}_{it}^s = 0 \right] \\ &= \frac{(1 - e) f}{1 - fe} \beta. \end{aligned}$$

The covariance term is conditioned on $\mathcal{I}_{it}^s = 0$ because, among observations present in the sample at the beginning of the period, only those that do not exit are used to compute inflation. These usable observations either had no price change and no exit (probability $(1 - d)(1 - f)$) or a price change and no exit (probability $(1 - d) f (1 - e)$). Only the latter observations, which account for a share $(1 - e) f (1 - fe)$ of usable observations, have a nonzero contribution to inflation.

Proceeding similarly with b_1 ,

$$\begin{aligned} b_1 &= \frac{\int \text{cov}(\Delta p_{it}, \Delta x_{t-1} | \mathcal{I}_{it}^s = 0) di}{\text{cov}(\Delta x_{t-1})} \\ &= \frac{(1-e)f}{1-fe} \frac{\text{cov}(u_{it} + \beta \Delta x_t + \varepsilon_{it}, \Delta x_{t-1} | \mathcal{I}_{it}^s = 0, \mathcal{I}_{it}^f = 1)}{\text{cov}(\Delta x_t)}. \end{aligned}$$

Since Δx_t and ε_{it} are assumed to be independent of Δx_{t-1} , the covariance term is impacted solely through the interaction between Δx_{t-1} and the cumulated price pressure u_{it} . Conditioning on past realizations of the indicator variables, there are five distinct cases:

$$u_{it} = \begin{cases} u_{it-1} + \beta \Delta x_{t-1} + \varepsilon_{it-1} & \text{if } \left\{ \mathcal{I}_{it-1}^s = 0, \mathcal{I}_{it-1}^f = 0 \right\} \\ 0 & \text{if } \left\{ \mathcal{I}_{it-1}^s = 0, \mathcal{I}_{it-1}^f = 1 \right\} \\ u_{it-1} + \beta \Delta x_{t-1} + \varepsilon_{it-1} & \text{if } \left\{ \mathcal{I}_{it-1}^s = 1, \mathcal{I}_{it-1}^n = 0, \mathcal{I}_{it-1}^f = 0 \right\} \\ 0 & \text{if } \left\{ \mathcal{I}_{it-1}^s = 1, \mathcal{I}_{it-1}^n = 0, \mathcal{I}_{it-1}^f = 1 \right\} \\ 0 & \text{if } \left\{ \mathcal{I}_{it-1}^s = 1, \mathcal{I}_{it-1}^n = 1 \right\} \end{cases}. \quad (9)$$

Consequently,

$$\begin{aligned} b_1 &= \frac{(1-e)f}{1-fe} \frac{\text{cov}(\beta \Delta x_{t-1}, \Delta x_{t-1})}{\text{cov}(\Delta x_t)} \left(\Pr \left[\mathcal{I}_{it-1}^s = 0, \mathcal{I}_{it-1}^f = 0 \right] + \Pr \left[\mathcal{I}_{it-1}^s = 1, \mathcal{I}_{it-1}^n = 0, \mathcal{I}_{it-1}^f = 0 \right] \right) \\ &= \frac{(1-e)f}{1-fe} (1-f) \beta ((1-d) + (d + (1-d)fe)(1-n)) \end{aligned}$$

Intuitively, among all observations usable to compute inflation, only those with a price change in the current period (marginal probability $\frac{(1-e)f}{1-fe}$) and either no price change and no exit in the previous period (marginal probability $(1-d)(1-f)$) or an exit accompanied by an entry with no price change (marginal probability $(d + (1-d)fe)(1-n)(1-f)$) feature a nonzero contribution of Δx_{t-1} to inflation.

The general case with b_l is illustrated in the upper panel of figure 11. The figure shows the various states that are usable in the computation of inflation at each period, along with their associated marginal probability period by period. The arrows indicate the paths through which an exchange rate movement in the period $t-l$ is reflected as a nonzero price change in period t . Observations that have not yet responded to an exchange rate movement at period $t-l$ can find their way in the index either by having been present in the sample prior to period $t-l$ or by entering the sample through a substitution. The marginal probability from period $t-l$ to period $t-1$ associated with the former event (no price change and no exit) is $(1-d)(1-f)$ for each period. The marginal probability associated with the addition of

an item in period $t - k$ whose last price change was prior to period $t - l$ is the product of the probability of having a substitution (probability $d + fr(1 - d)$), a random selection from the universe (probability $1 - n$), and no price change for $l - k + 1$ periods (probability $(1 - f)^{l-k+1}$). Summing up across all usable paths, we have a general expression for the (plim) regression coefficient

$$b_l = \left(\frac{1 - e}{1 - fe} \right) \left((1 - d)^l + \frac{s(1 - n)}{d} \left(1 - (1 - d)^l \right) \right) f (1 - f)^l \beta.$$

A.2 Delayed entries

We next assume that the econometrician only uses observations that have been in the sample for more than M periods in the computation of inflation. This assumption is made in section 4.4 to argue that delaying the entry of items in the basket can mitigate some of the biases associated with selective exits and entries. We distinguish between two cases: $l \leq M$ and $l > M$.

The first case is illustrated in the middle panel of figure 11. Because entries are delayed by more periods than the number of lags in the exchange rate movement considered, all observations contributing to inflation and reflecting Δx_{t-l} must have been in the index continuously since before period $t - l$. For this situation to occur, we must have had no price change from period $t - l$ to $t - 1$ (marginal probability $(1 - f)/(1 - fe)$ during l periods), and a price change at t (marginal probability $(1 - e)f/(1 - fe)$). The resulting (plim) regression coefficient is

$$b_l(l \leq M) = \frac{(1 - e)}{(1 - fe)^{l+1}} f (1 - f)^l \beta. \quad (10)$$

The case of $l > M$ is illustrated at the bottom of figure 11. It mixes elements of the general case with no delay (upper panel) and the case with $l \leq M$ (middle panel). Prior to period $t - M$, observations that have not yet responded to the exchange rate movement at period $t - l$ could have found their way in the index either through a substitution or by having been present in the sample before period $t - l$. From period $t - M$ onward, only observations that are continuously present in the index from the end of period $t - M - 1$ onward can be used to compute inflation. Summing up the probabilities over all possible paths and simplifying, we get

$$b_l(l > M) = \frac{(1 - e)}{(1 - fe)^{M+1}} \left((1 - d)^{l-M} + \frac{s(1 - n)}{d} \left(1 - (1 - d)^{l-M} \right) \right) f (1 - f)^l \beta. \quad (11)$$

A.3 Proof that biases are declining in the entry delay

We conclude this appendix by proving that delaying the entry of items in the index always improves pass-through estimates. We assume that the number of lags in the regression is at least as large as the forecast horizon, T , a condition typically satisfied in standard pass-through regressions. Let $b_l(M)$ be the (plim) coefficient associated with the l -th lag of the exchange rate and an entry delay of M periods. The proof proceeds in two steps. We first prove that $b_l(M+1) \geq b_l(M)$, so that delaying entries by an extra period always (weakly) increases the size of the (plim) regression coefficients. We then show that the cumulative response over any forecast horizon remains bounded above by the true response.

A.3.1 Step 1: $b_l(M+1) \geq b_l(M)$

We distinguish between three cases: $l < M+1$, $l = M+1$, and $l > M+1$. When $l < M+1$, the plim coefficients are given by equation 10 whether the delay is M periods or $M+1$ periods, so that $b_l(M) = b_l(M+1)$. When $l = M+1$, $b_{M+1}(M)$ is given by equation 11 and $b_{M+1}(M+1)$ is given by equation 10. For $b_{M+1}(M+1) \geq b_{M+1}(M)$ to be true in this case, we must have

$$\frac{1}{1-fe} \geq 1-d+s(1-n).$$

Note that if the above equation holds for $n = 0$, then it holds for all $n \in [0, 1]$. Imposing $n = 0$ and using $s = d + (1-d)fe$, we have

$$\frac{1}{1-fe} \geq 1-d+ef,$$

which is always satisfied. Finally, we want to show that $b_l(M+1) \geq b_l(M)$ when $l > M+1$. The plim coefficients are given by equation 11. Note that

$$\frac{\partial}{\partial n} (b_l(M+1) - b_l(M)) = -\Omega \left(\frac{1 - (1-d)^{l-M-1}}{1-fe} - \left(1 - (1-d)^{l-M}\right) \right),$$

where Ω is some positive constant. The difference between $b_l(M+1)$ and $b_l(M)$ is thus linear in n and either always increasing or always decreasing in n . By showing that $b_l(M+1) \geq b_l(M)$ for $n = 0$ and $n = 1$, we will have proven that the result holds for the worse scenario under either case. Consider first

$$b_l(M|n=1) = f(1-f)^l \frac{(1-e)}{(1-fe)^{M+1}} (1-d)^{l-M} \beta.$$

We have $b_l(M+1|n=1) \geq b_l(M|n=1)$ if and only if $(1-d)^{l-M-1} \geq (1-fe)(1-d)^{l-M}$, which is always true. Consider next

$$b_l(M|n=0) = \frac{f(1-f)^l(1-e)}{(1-fe)^{M+1}} \left((1-d)^{l-M} + \frac{s}{d} \left(1 - (1-d)^{l-M} \right) \right) \beta.$$

We have $b_l(M+1|n=0) \geq b_l(M|n=0)$ if and only if

$$d(1-d)^{l-M-1} + s \left(1 - (1-d)^{l-M-1} \right) \geq (1-fe) \left(d(1-d)^{l-M} + s \left(1 - (1-d)^{l-M} \right) \right),$$

which can be shown to hold if and only if

$$fes \left(1 - (1-d)^{l-M} \right) \geq 0,$$

a condition that is always satisfied. Summing up, the individual coefficients are increasing in the entry delay, so that cumulative pass-through over any forecast horizon also is increasing in the entry delay.

A.3.2 Step 2: the cumulative estimated response is bounded above by true response

To complete the proof, we show that the estimated pass-through under delayed entries never exceeds true pass-through over any forecast horizon. The true pass-through after L periods is

$$\sum_{l=0}^L b_l = \sum_{l=0}^L f(1-f)^l \beta = \left(1 - (1-f)^{L+1} \right) \beta.$$

Because $b_l(M+1) \geq b_l(M)$, the estimated pass-through is largest when $M \geq L$, which is associated with

$$\sum_{l=0}^L b_l(M \geq L) = \sum_{l=0}^L \frac{(1-e)}{(1-fe)^{l+1}} f(1-f)^l \beta = \left(1 - \left(\frac{1-f}{1-fe} \right)^{L+1} \right) \beta.$$

Comparing the above expression to equation 3, it is immediate that cumulative pass-through under delayed sample entry is bounded above by the unbiased case.

B Real rigidities

To explore the incidence of real rigidities on our theoretical correction factors, we follow Gopinath and Itskhoki (2011) in adopting a flexible specification for the evolution of the firm's reset price,

$$\Delta p_{it}^* = (1 - \alpha) (\beta \Delta x_t + \varepsilon_{it}) + \alpha (\Delta m_t).$$

The parameter $\alpha \in [0, 1]$ controls the extent of real rigidities by capturing the emphasis placed by the firm on matching movements in the average price of competing importers, Δm_t , when resetting its own price. As was the case earlier, firms fully release the pressure accumulated since their last adjustment when updating their price and a fraction β of exchange rate movements is passed-through in the long run.

Under Calvo pricing, it is possible to derive expressions for actual pass-through to the universe of prices and for measured pass-through in the import price index. Abstracting from idiosyncratic disturbances, which have no influence on the index, the initial effect on reset prices of a one-time jump in the exchange rate is the sum of its direct effect on $\Delta p_{i,t}^*$ and of its indirect effect through the price of competing importers,

$$\Delta p_{i,t}^* = (1 - \alpha) \beta \Delta x_t + \alpha \Delta m_t.$$

The impact in subsequent periods depends only on the evolution of the price of competing importers,

$$\Delta p_{i,t+l}^* = \alpha \Delta m_{t+l}.$$

The price of competing importers evolves according to

$$\Delta m_{t+l} = \sum_{k=0}^l f (1 - f)^{l-k} \Delta p_{i,t+k}^*. \quad (12)$$

Intuitively, the fraction f of firms updating their price in period $t + l$ takes into account $\Delta p_{i,t+l}^*$ in setting $\Delta p_{i,t+l}$. Of these firms, a fraction $1 - f$ did not update in period $t + l - 1$ and will also take into account $\Delta p_{i,t+l-1}^*$. More generally, a fraction $(1 - f)^s$ of firms updating in period $t + l$ will take into account $\Delta p_{i,t+l-s}^*$ when setting $\Delta p_{i,t+l}$.

The above three equations allow us to solve recursively for $\Delta p_{i,t+l}^*$ and Δm_{t+l} . In the initial period,

$$\frac{\Delta p_{i,t}^*}{\Delta x_t} = \left(\frac{1 - \alpha}{1 - \alpha f} \right) \beta \quad (13)$$

and

$$\frac{\Delta m_t}{\Delta x_t} = \left(\frac{1 - \alpha}{1 - f\alpha} \right) f\beta. \quad (14)$$

Subsequently, one can show that

$$\frac{\Delta p_{i,t+l}^*}{\Delta x_t} = \frac{\alpha(1 - \alpha)}{(1 - f\alpha)^{l+1}} f(1 - f)^l \beta \quad (15)$$

and

$$\frac{\Delta m_{t+l}}{\Delta x_t} = \frac{(1 - \alpha)}{(1 - f\alpha)^{l+1}} f(1 - f)^l \beta$$

for $l > 0$. Absent real rigidities ($\alpha = 0$), the reset price would initially move by $\beta\Delta x_t$ and remain unchanged thereafter. With real rigidities ($\alpha > 0$), the reset price moves gradually toward its long-run level so that the response to an exchange rate movement is spread over several price adjustments. The cumulative response of the price of competing importers after L period is given by

$$\sum_{l=0}^L \frac{\Delta m_{t+l}}{\Delta x_t} = \left(1 - \left(\frac{1 - f}{1 - \alpha f} \right)^{L+1} \right) \beta, \quad (16)$$

which implies a slower rate of pass-through than the case with $\alpha = 0$. The cumulative change in the target price is a weighted sum of the long-run change and the cumulative change in the price of competing importers,

$$\sum_{l=0}^L \frac{\Delta p_{i,t+l}^*}{\Delta x_t} = \left(1 - \left(\frac{1 - f}{1 - \alpha f} \right)^{L+1} \right) \alpha\beta + (1 - \alpha)\beta.$$

Turning to measured pass-through in the sample, using the arguments made in section 3 to derive equation (4), we can express the change in the import price index as

$$b_l = \sum_{k=0}^l \frac{(1 - e)}{1 - fe} \left((1 - d)^k + \frac{s(1 - n)}{d} \left(1 - (1 - d)^k \right) \right) f(1 - f)^k \frac{\Delta p_{i,t+l-k}^*}{\Delta x_t}. \quad (17)$$

The variable b_l corresponds to the (plim) regression coefficient on the l -th lag of the exchange rate. The above equation, along with equations (13) and (15), allow us to compute the b_l at all lags. To gain intuition about the role of real rigidities, we shall highlight three results.

First, the correction factor on the initial response of the index is independent of the extent of real rigidities. This can be seen by using equations (13), (14), and (17) to write

$$b_0 = \left(\frac{1 - e}{1 - fe} \right) \frac{\Delta m_t}{\Delta x_t}.$$

The inverse of $(1 - e)/(1 - fe)$ is the correction factor applicable to the initial index response. It does not depend on α . Second, the correction factor for long-run pass-through is also independent of the extent of real rigidities. This can be proven by considering the cumulative sum of b_l in equation (17) and taking its limit as the horizon is extended. We obtain

$$\lim_{L \rightarrow \infty} \sum_{l=0}^L b_l = \frac{(1 - e)}{1 - fe} \left(\frac{f + s(1 - n)(1 - f)}{f + d - fd} \right) \beta.$$

The inverse of the expression in front of β is the correction factor on long-run pass-through, which is again independent of α . Third, when both exits and entries are selective, the correction factors are independent of real rigidities at all horizons. In such a case, a constant fraction $(1 - e)/(1 - fe)$ of $\Delta m_{t+l}/\Delta x_t$ is missed every period.

In the general case, the correction factors on the cumulative index response may depend on α at horizons beyond the initial period but finite. However, we find that these correction factors are rather insensitive to the extent of real rigidities. This aspect is illustrated in figure 12 for a Calvo model with a frequency of price changes of 15 percent and a rate of substitution of 5 percent. The parameter α is set to either 0 (no real rigidities) or 0.6 (with real rigidities). The latter value corresponds to the empirical estimate of Gopinath and Itskhoki (2011). The correction factors are identical for the canonical cases in which all exits and entries are either selective (upper-left panel) or random (upper-right panel), as these cases correspond to the absence of a bias and to the censoring of a constant fraction of the index response, respectively. When all exits are selective and all entries are random (lower-left panel) or all exits are random and all entries are selective (lower-right panel), the presence of real rigidities only has a minor impact on the size of the correction factor. This finding is robust to considering smaller or larger frequency of price changes.

Finally, we note a trade-off between the extent of real rigidities and the degree of state dependence in the calibrated model. For many of the three-digit Enduse finished good categories considered in figure 6, we find an initially rapid response of the import price index to an exchange rate movement that is inconsistent with the slow response predicted by the Calvo model. For these categories, the addition of real rigidities worsens the fit of the Calvo model by further slowing its predicted response. As an experiment, we assumed $\alpha = 0.6$ and recomputed the linear combinations of Calvo and menu-cost impulse responses that minimize the Euclidian distance with the empirical response over our forecast horizon. For 9 out of the 16 categories, the weight placed on the menu-cost model rose as we added real rigidities. A strong preference for one model over the other was unchanged for the other categories. This finding cautions against concluding that evidence of real rigidities in the data, in the form of pass-through over more than one price adjustment, necessarily implies slow pass-through.

Table 1: Exit rate, entry rate, and mean frequency and absolute size of individual price changes in the IPP import price sample

	Enduse Category	Relative Weight	Exit Rate			Entry Rate	Mean Freq. of Price Changes	Mean Absolute Size of Price Change
			all reasons	out-of-scope	others reasons			
Foods, Feeds & Beverages	000 Green coffee, cocoa beans, cane sugar	0.3	2.8	1.1	1.7	3.0	47.0	8.7
	001 Other agricultural foods	2.8	2.4	0.8	1.5	2.6	21.1	9.4
	010 Nonagricultural products	1.1	2.2	0.8	1.4	2.4	20.5	7.1
Industrial Supplies and Materials	100 Petroleum & products, excluding gas	16.7	3.6	1.9	1.6	2.5	38.0	11.7
	101 Fuels, n.e.s.-coal & gas	1.8	3.4	2.1	1.2	4.0	55.9	13.4
	110 Paper base stocks	0.2	3.6	1.6	2.0	3.2	36.5	6.1
	111 Newsprint & other paper products	0.7	3.3	1.2	2.0	2.5	19.6	5.0
	120 Agricultural products	0.5	2.2	0.6	1.5	2.1	28.4	7.4
	121 Textile supplies & related materials	0.7	2.6	0.8	1.7	2.5	8.0	6.8
	125 Chemicals, excl. meds., food additives	3.4	2.4	0.7	1.7	2.5	11.2	7.3
	130 Lumber & unfinished building materials	1.1	2.4	0.9	1.5	2.9	33.2	7.7
	131 Building materials, finished	1.0	2.6	0.8	1.8	2.9	10.4	5.7
	140 Steelmaking materials-unmanufactured	0.4	2.4	0.9	1.5	2.0	21.2	7.1
	141 Iron & steel mill products-semifinished	1.3	3.5	1.3	2.1	3.9	15.3	21.4
	142 Major non-Fe metals-crude & semifin.	2.7	3.0	1.4	1.5	2.5	43.4	5.8
	150 Iron & steel products, ex. advanced mfg.	0.5	2.6	0.8	1.9	2.4	9.5	7.1
	151 Iron & steel mfg.-advanced	0.4	2.9	0.7	2.2	2.4	13.2	7.2
	152 Fin. metal shapes & adv. mfg., ex. steel	0.9	2.7	0.6	2.0	2.4	13.3	5.5
	161 Finished	1.5	2.7	1.1	1.6	2.8	7.1	7.9
Capital Goods	210 Oil drilling, mining & const. machinery	1.1	2.5	0.9	1.6	2.9	6.9	6.6
	211 Industrial & service machinery, n.e.c.	6.2	2.5	0.8	1.7	2.5	6.3	6.7
	212 Agricultural machinery & equip.	0.4	2.7	1.2	1.5	3.1	8.9	5.3
	213 Computers, periph. & semiconductors	7.5	3.7	2.2	1.5	5.0	9.7	9.6
	214 Telecommunications equip.	2.3	3.4	1.8	1.6	3.6	5.8	8.9
	215 Business mach. & equip., ex. Computers	0.5	3.3	1.5	1.8	2.5	5.2	6.3
	216 Scientific, hospital & medical machinery	1.5	3.1	1.2	1.9	3.2	4.9	6.9
Auto-motive	300 Passenger cars, new & used	8.0	2.8	1.6	1.1	3.5	5.3	2.0
	301 Trucks, buses, & special-purp. vehicles	1.4	2.8	1.9	0.9	3.9	5.8	2.9
	302 Parts, engines, bodies, & chassis	5.4	2.8	1.2	1.6	3.0	8.0	7.1
Consumer Goods (Ex. Food & Auto.)	400 Apparel, footwear, & household goods	6.6	3.5	1.7	1.8	3.6	3.9	7.6
	401 Other consumer nondurables	5.0	2.4	0.8	1.5	2.7	6.0	7.7
	410 Household goods	6.1	2.9	1.2	1.7	3.0	4.6	6.2
	411 Recreational equip. & materials	2.3	3.2	1.8	1.5	3.1	4.8	5.7
	412 Home entertainment equip.	3.1	3.7	2.2	1.5	4.1	5.6	5.8
	413 Coins, gems, jewelry, & collectibles	1.3	3.1	1.1	1.9	3.1	6.9	5.9
	500 Imports, N.E.S.	3.5	2.7	0.6	2.1	1.8	5.2	12.5
Total	100.0	3.0	1.4	1.6	3.1	15.3	8.0	

Notes: These statistics are computed by first applying uniform weights to items within each Enduse-month combination and then averaging the resulting monthly statistics. The sample period is from October 1995 to April 2005. Missing item prices are imputed by their last observed price and used in the computation of the above statistics. The table also shows the relative 2006 import value shares used to aggregate three-digit Enduse statistics.

Table 2: Reconciliation of our results with those of Nakamura and Steinsson (2011)

	All goods excluding oil	Finished goods	Material-intensive goods excluding oil
Standard pass-through estimates after 6 quarters			
Gagnon, Mandel, and Vigfusson	0.41	0.26	0.93
Nakamura and Steinsson	0.43	n.a.	n.a.
Item frequencies are constant within 3-digit Enduse categories			
<i>Case 1: Exits and entries are random</i>			
Corrected pass-through (6 quarters)	0.41	0.26	0.93
<i>Correction factor</i>	1.00	1.00	1.00
Corrected pass-through (long-run)	0.41	0.26	0.93
<i>Correction factor</i>	1.00	1.00	1.00
<i>Case 2: Exits and entries are selective</i>			
Corrected pass-through (6 quarters)	0.55	0.37	1.04
<i>Correction factor</i>	1.32	1.40	1.12
Corrected pass-through (long-run)	0.55	0.37	1.04
<i>Correction factor</i>	1.32	1.40	1.12
<i>Case 3: Exits are selective and entries are random</i>			
Corrected pass-through (6 quarters)	0.48	0.32	0.96
<i>Correction factor</i>	1.16	1.21	1.04
Corrected pass-through (long-run)	0.45	0.29	0.95
<i>Correction factor</i>	1.08	1.10	1.02
<i>Case 4: Exits are random and entries are selective</i>			
Corrected pass-through (6 quarters)	0.48	0.31	1.02
<i>Correction factor</i>	1.16	1.19	1.10
Corrected pass-through (long-run)	0.55	0.37	1.04
<i>Correction factor</i>	1.32	1.40	1.12
Item frequencies follow Beta distribution			
<i>All entries are selective</i>			
Corrected pass-through (long-run)	0.70	n.a.	n.a.
<i>Correction factor</i>	1.70	n.a.	n.a.
Corrected pass-through (long-run, no $f < 0.02$)	0.61	n.a.	n.a.
<i>Correction Factor</i>	1.47	n.a.	n.a.
Addendum: NS' corrected estimates			
Corrected pass-through (long-run)	0.67	n.a.	n.a.
<i>Correction factor</i>	1.71	n.a.	n.a.

Notes: Standard pass-through estimates after 6 quarters are obtained from an OLS estimation of equation 8. All correction factors assume Calvo pricing, a monthly frequency of item substitution of 2.5 percent, and uniform long-run pass-through rates across items. Corrected 6-quarter and long-run pass-through estimates are obtained by multiplying standard pass-through estimates after 6 quarters by a correction factor consistent with an 18-month horizon and an infinite horizon, respectively. The first set of correction factors assumes a constant frequency of price changes across items within each three-digit Enduse category. The second set uses a Beta distribution with parameters 0.44 and 3.50, as estimated by Nakamura and Steinsson (2012). For reference, the addendum cites the correction factor for the case with “local currency pricing imports” and “no heterogeneity in pass-through” from table IV in Nakamura and Steinsson (2012).

Figure 1: Exit rate, entry rate, and the dollar





Figure 2: Coefficients on lags of the exchange rate in pass-through regressions in baseline Calvo and menu-cost models

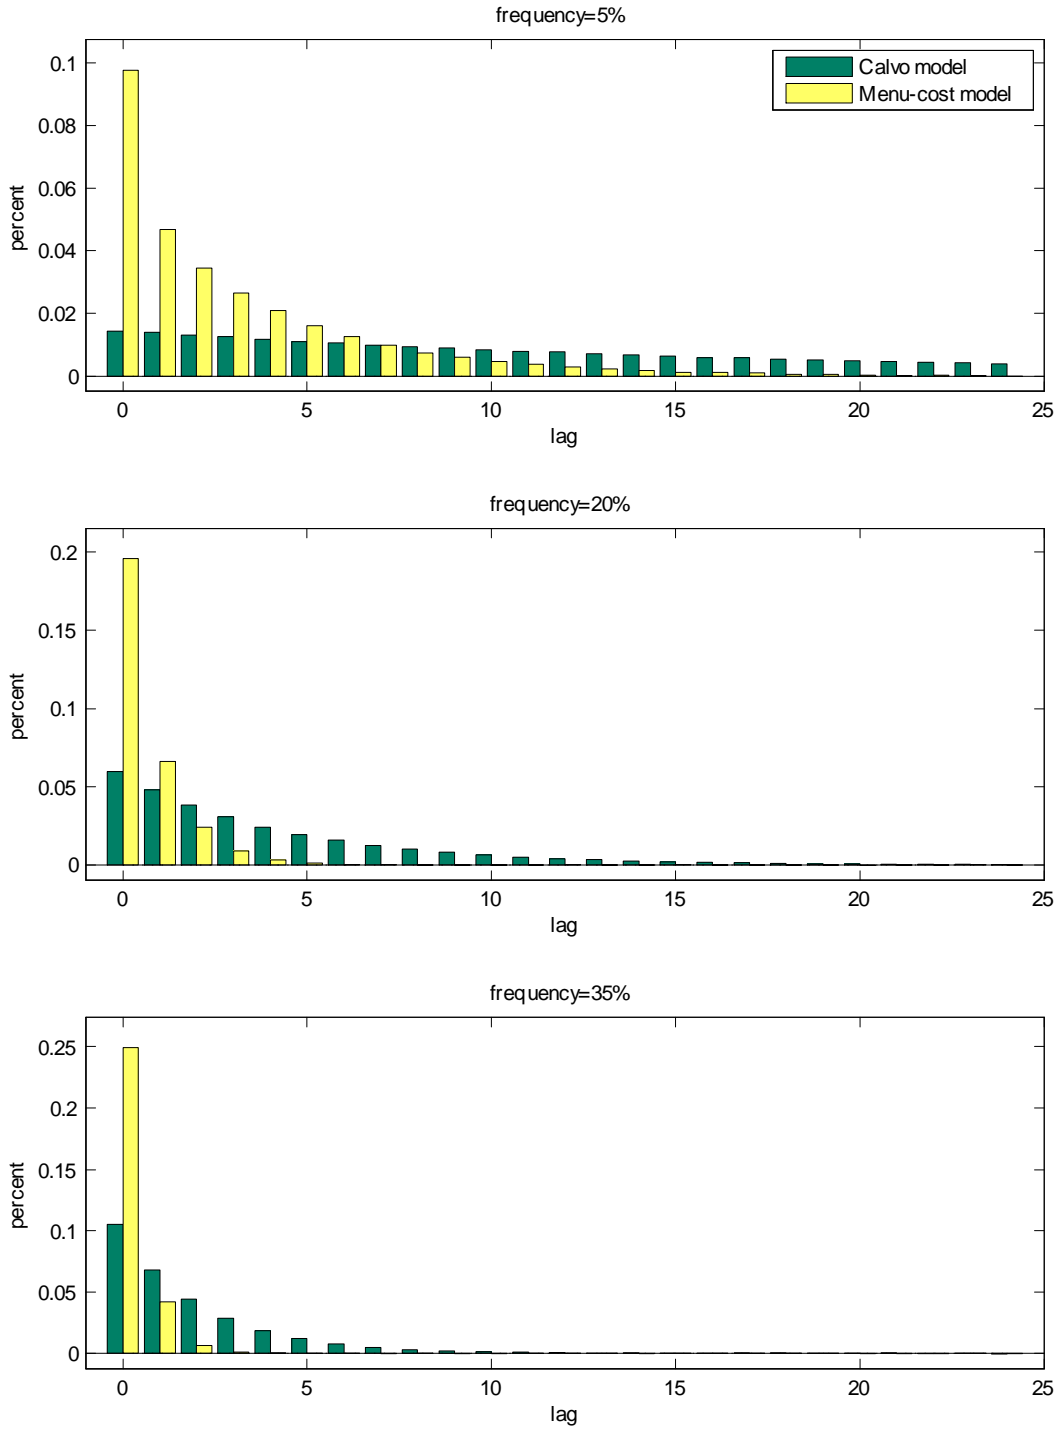
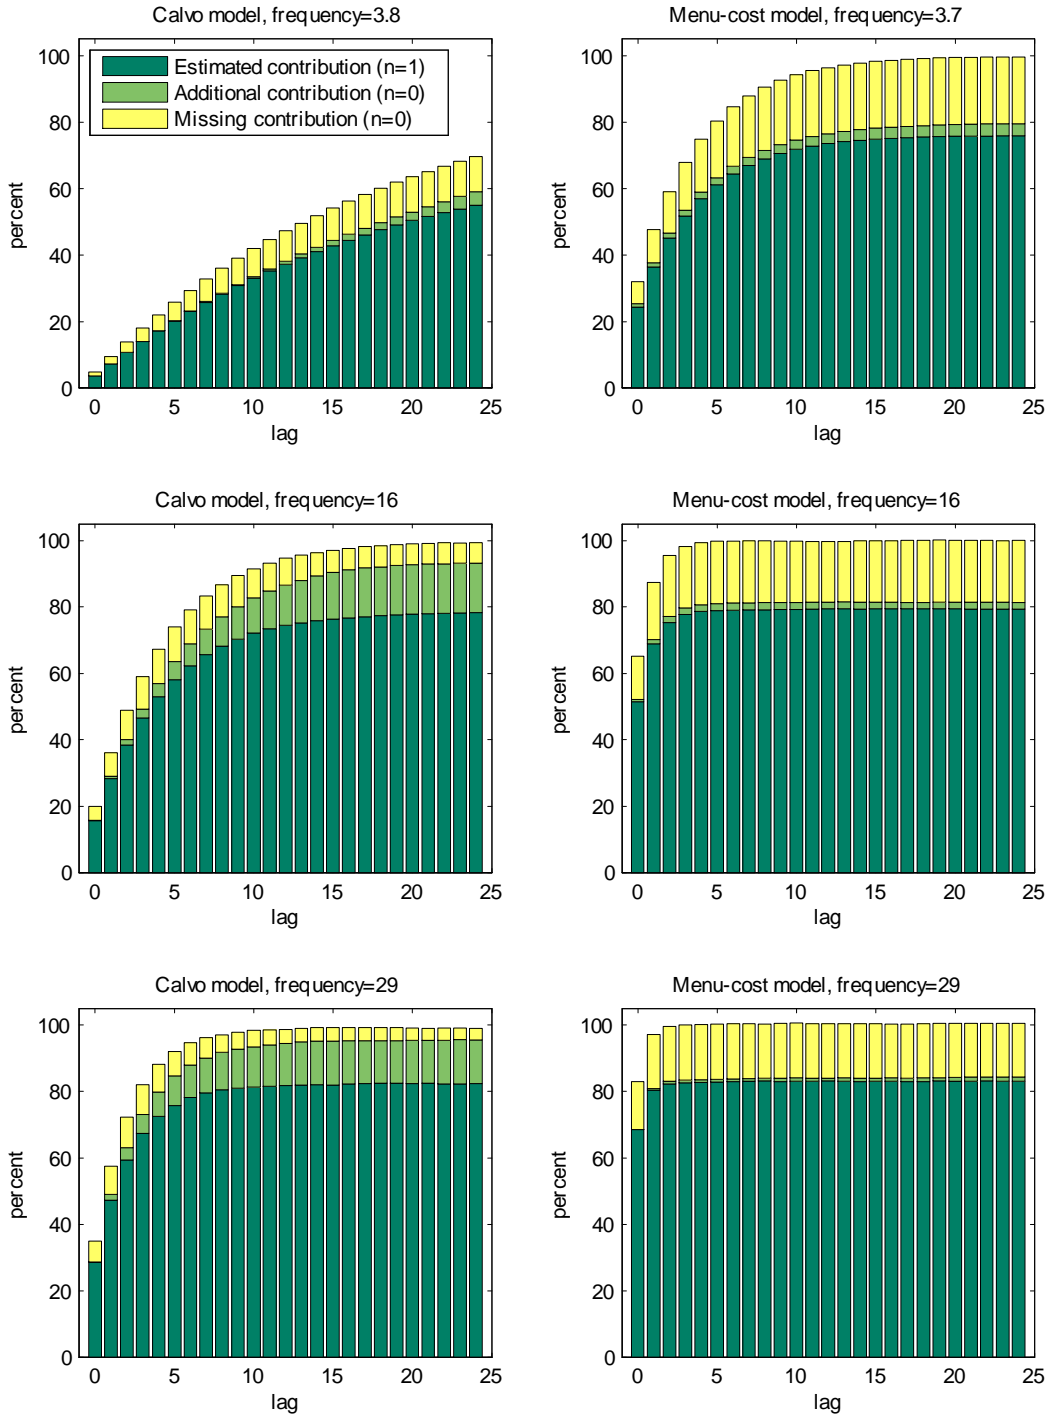


Figure 3: Cumulative contribution of coefficients on lagged exchange rate variables under selective exit ($e=0.25$)



Notes: The figure illustrates the case with selective entries ($n = 1$) and the case with fully-randomized entries ($n = 0$). The observed frequencies of price changes listed in the top, middle, and bottom rows correspond to true frequencies of 5 percent, 20 percent, and 35 percent, respectively.

Figure 4: Cumulative contribution of coefficients on lagged exchange rate variables model under random exits and selective entries ($n=1, s=0.05$)

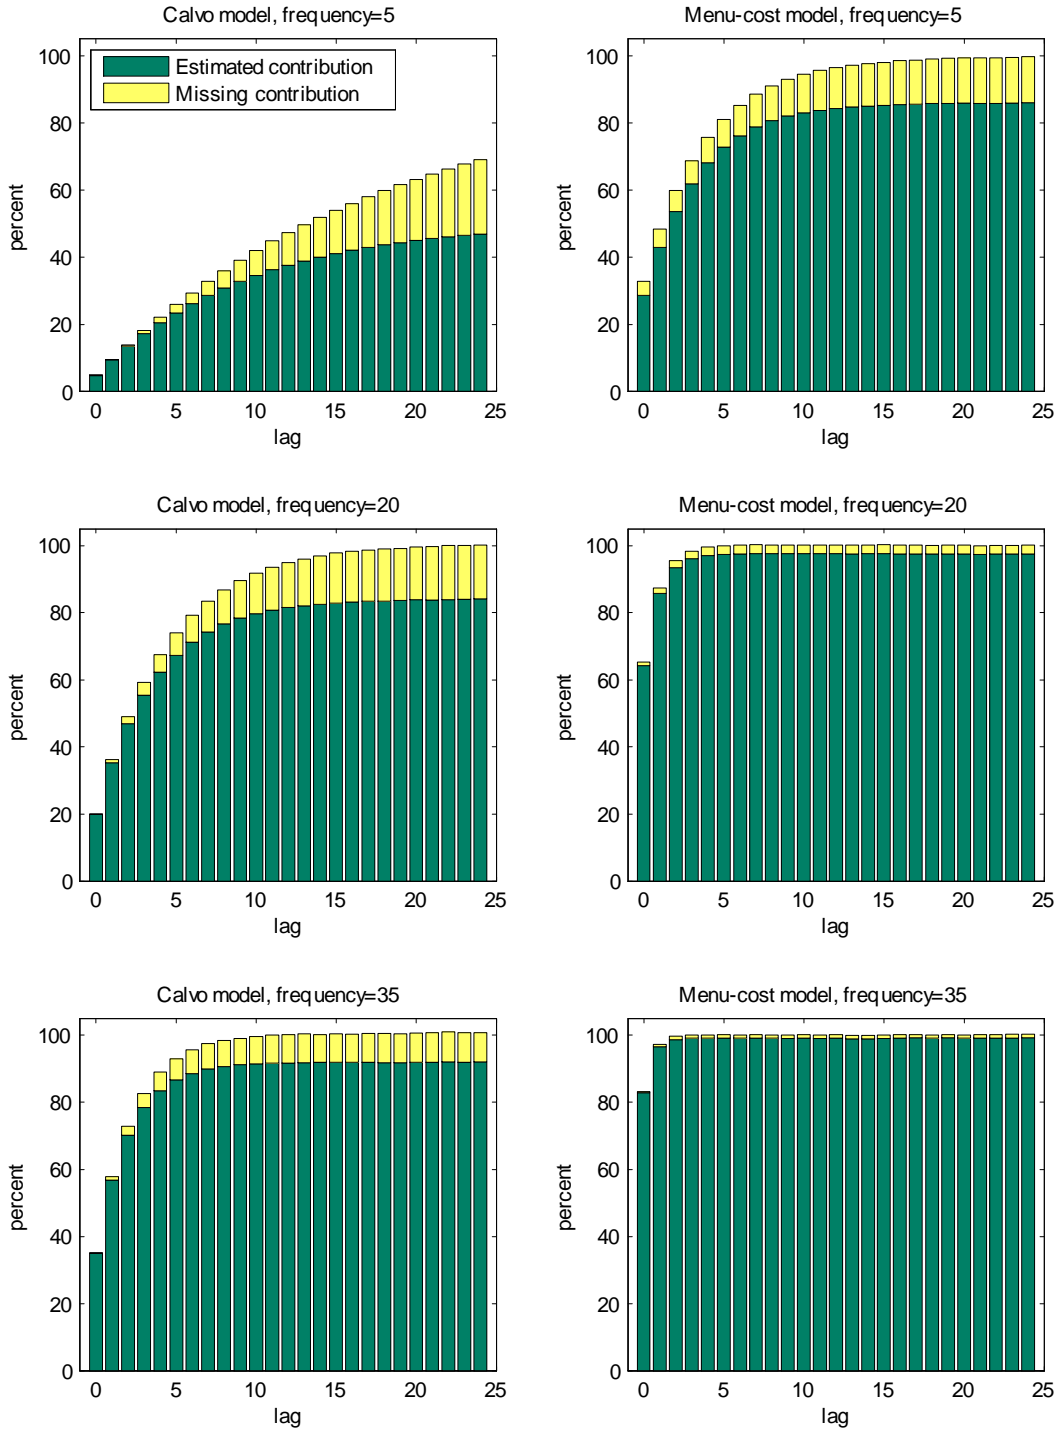


Figure 5: Pass-through to imported finished goods prices following a 1-percent depreciation of the dollar (by 3-digit Enduse categories)

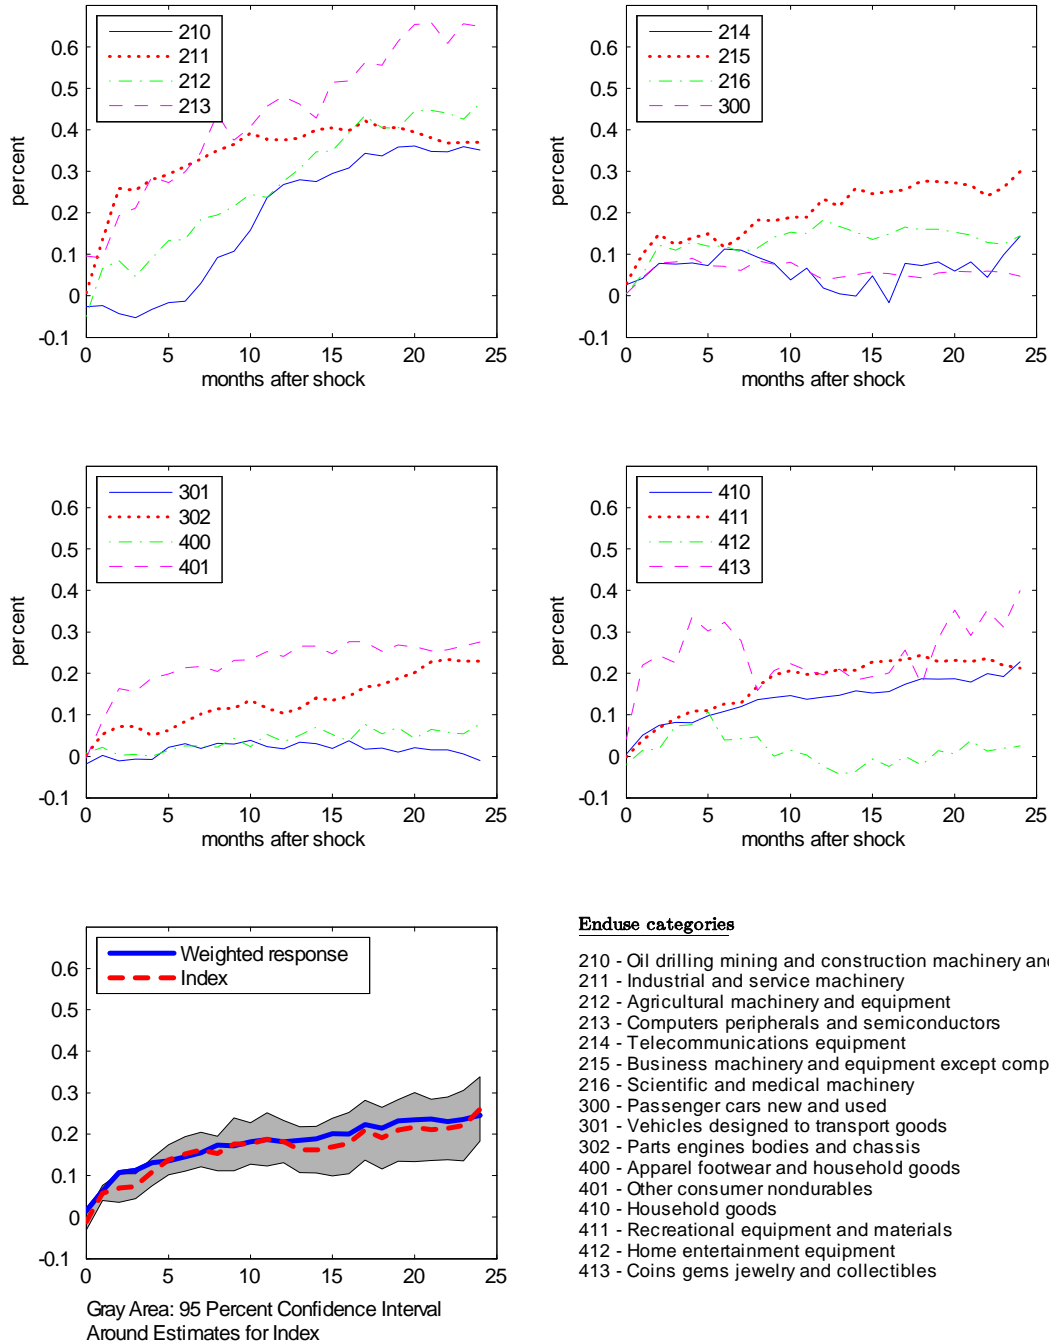
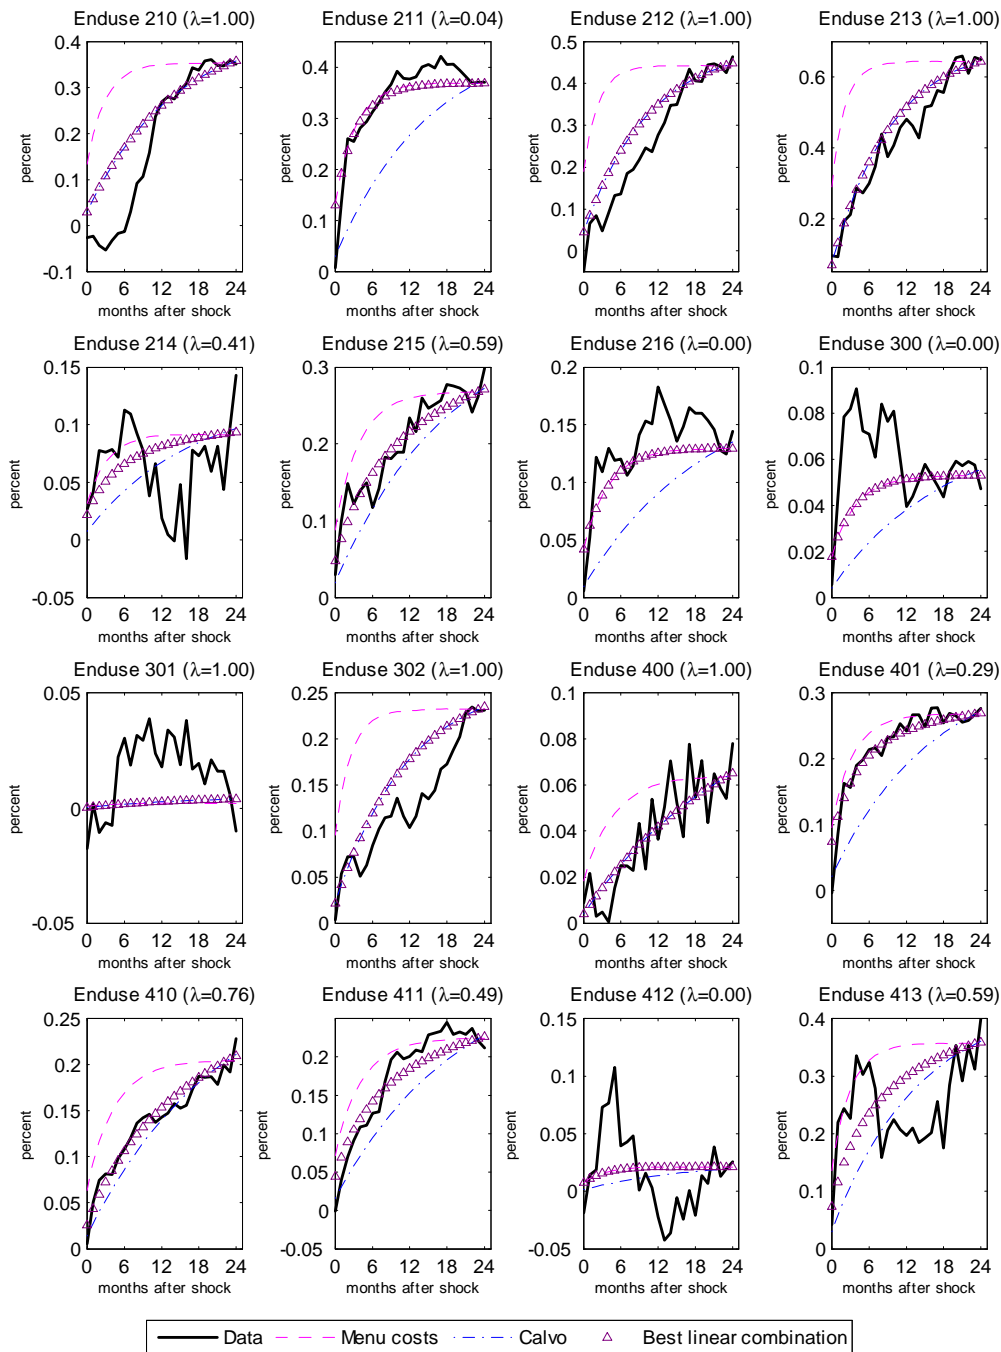


Figure 6: Pass-through to imported finished goods prices following a 1-percent depreciation of the dollar: models versus data



Notes: The figure shows the empirical impulse response for three-digit Enduse categories along with calibrated impulse responses in the Calvo and the menu-cost models. The figure also shows the combination of models minimizing $\|\lambda IRF_{Calvo} + (1 - \lambda) IRF_{MC} - IRF_{data}\|$ over $\lambda \in [0, 1]$, where $\|\cdot\|$ is the Euclidian distance and IRF_{Calvo} , IRF_{MC} , and IRF_{data} are the impulse responses in the Calvo model, the menu-cost model, and the data.

Figure 7: Upper bounds on exchange rate pass-through to imported finished goods

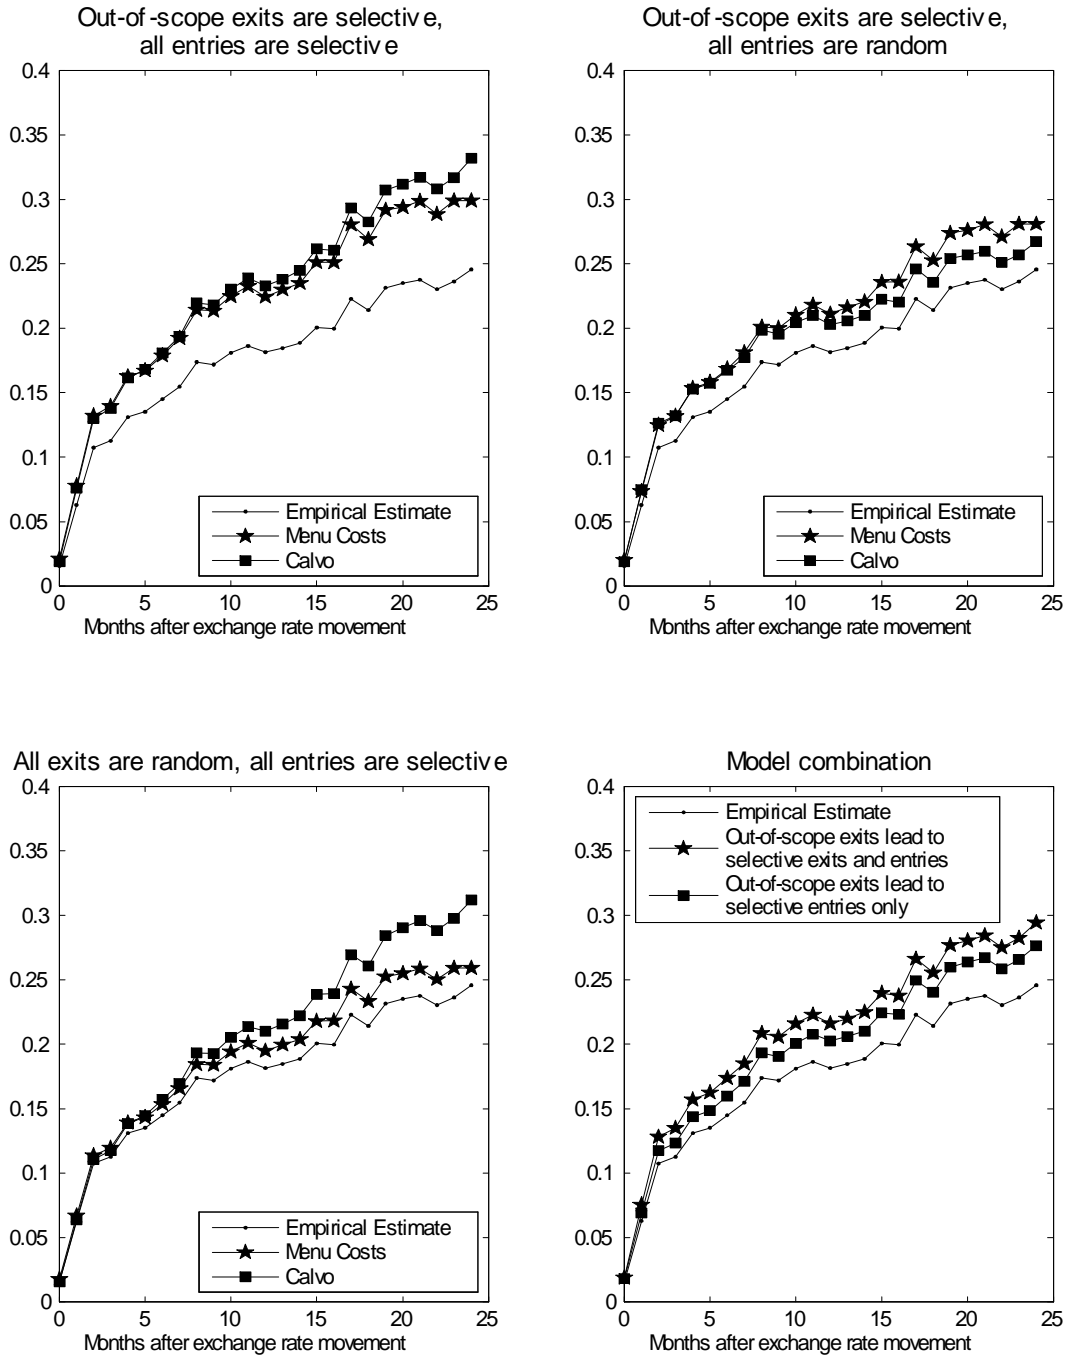


Figure 8: Impact of delaying entries on cumulative contribution of coefficients on lagged exchange rate variables under random exits and selective entries ($n=1, s=0.05$)

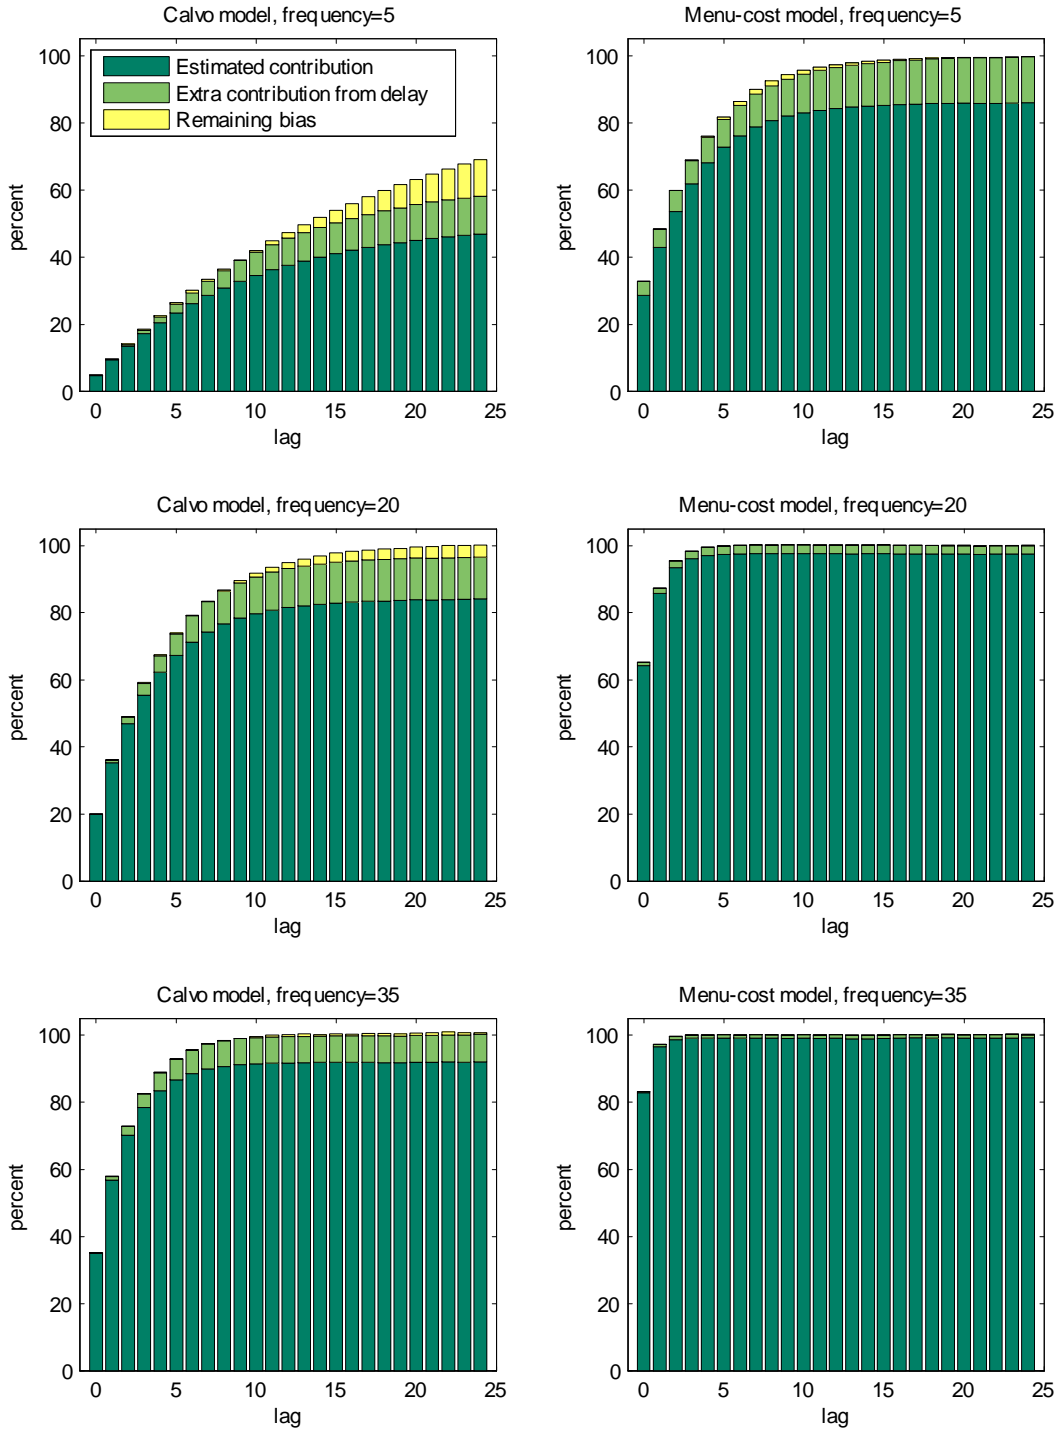
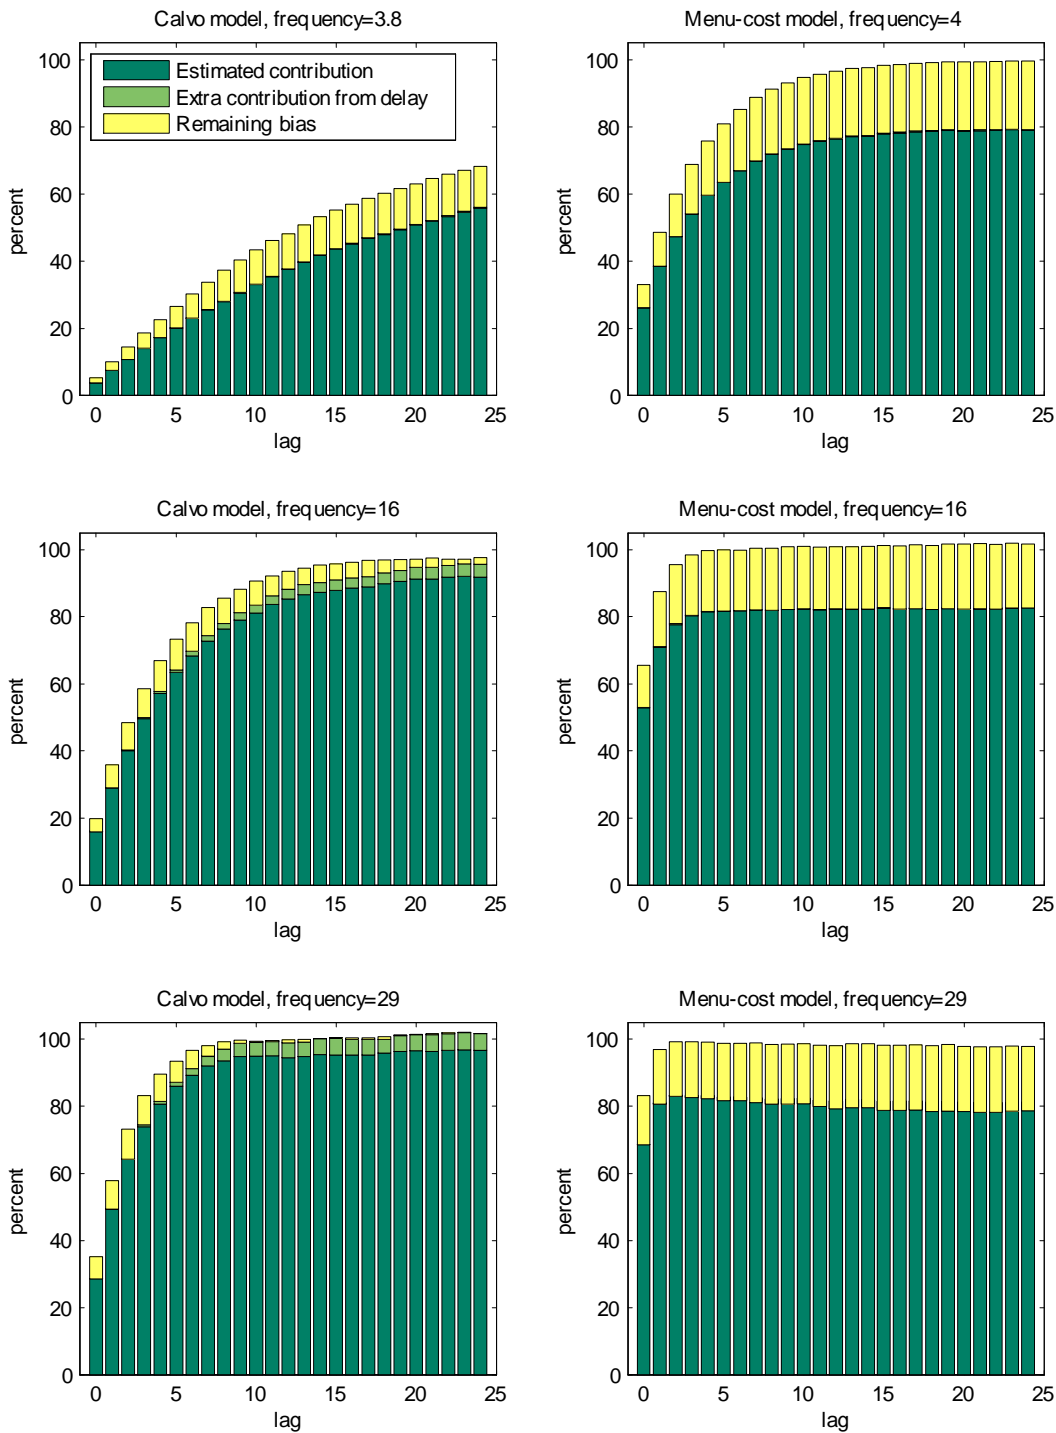


Figure 9: Impact of delaying entries on cumulative contribution of coefficients on lagged exchange rate variables under selective exits and random entries ($n=0$, $e=0.25$)



Notes: The observed frequencies of price changes listed in the top, middle, and bottom rows correspond to true frequencies of 5 percent, 20 percent, and 35 percent, respectively.

Figure 10: Alternative price indexes and measured pass-through

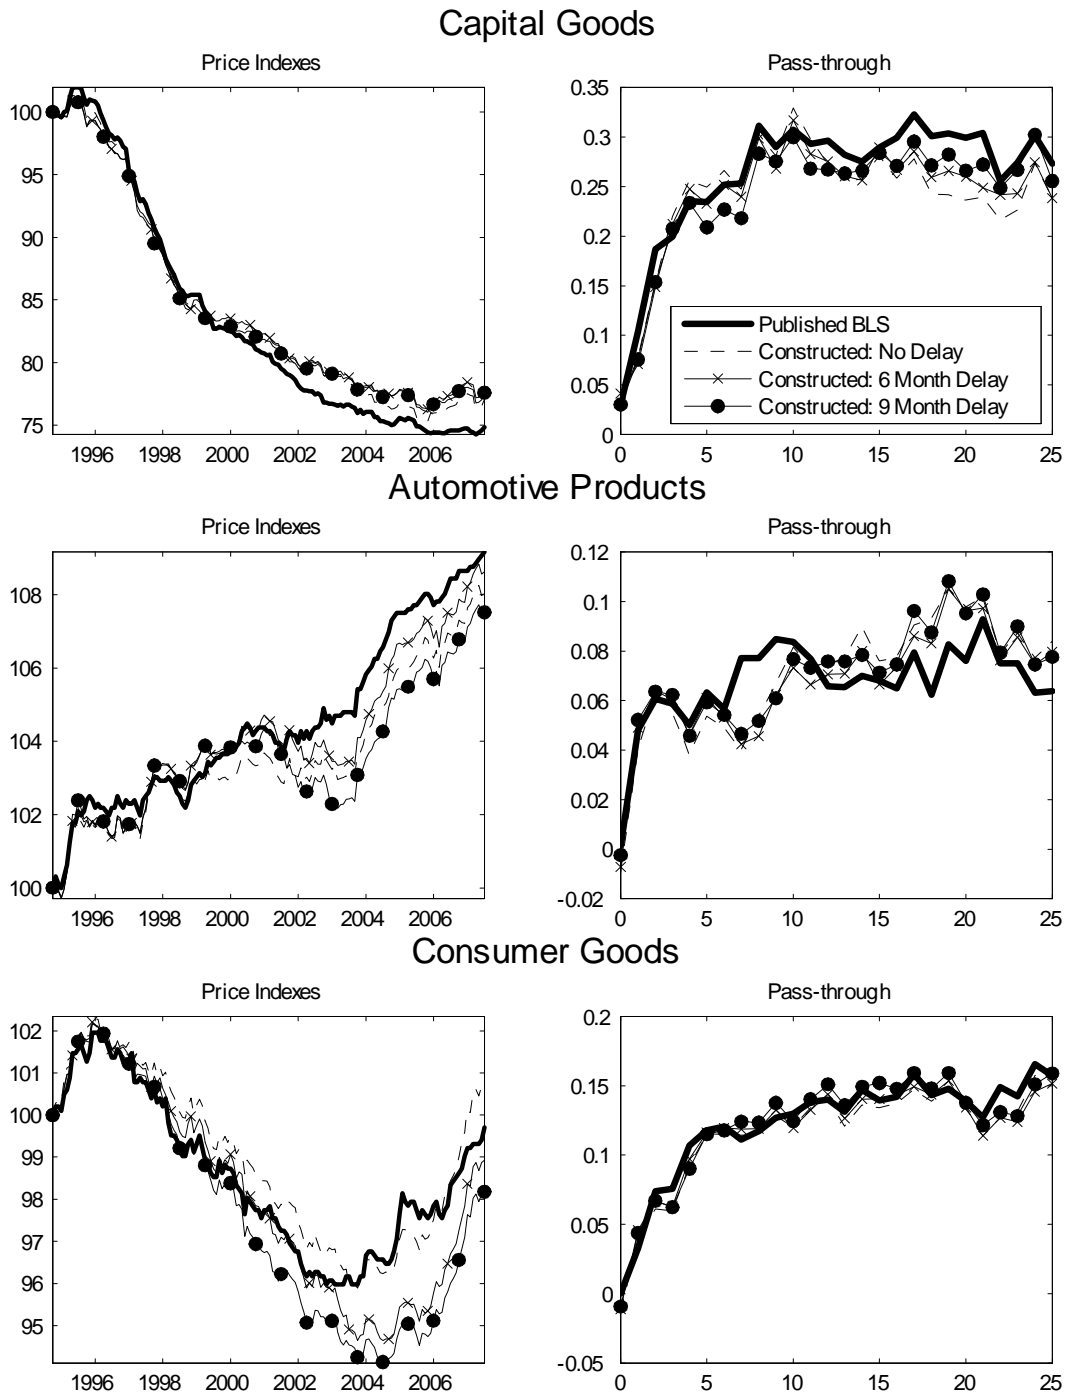
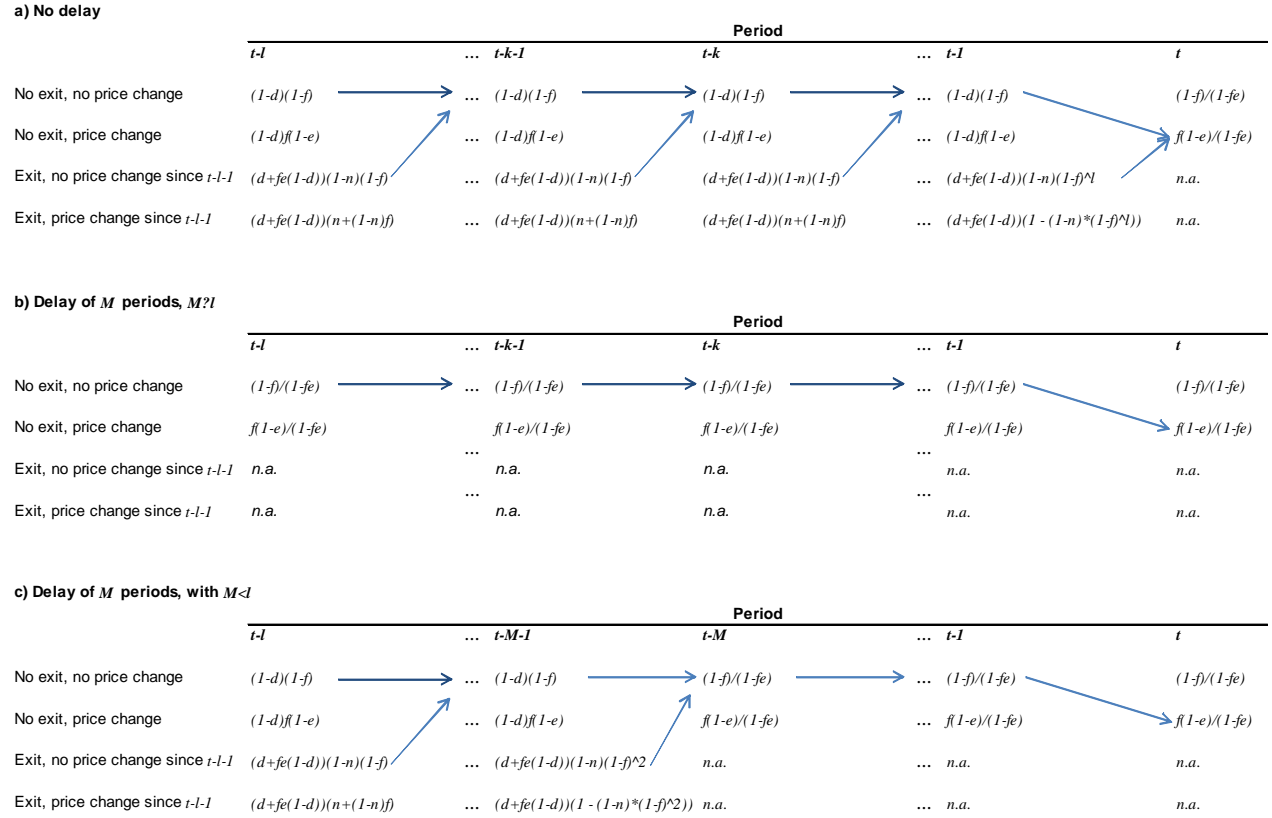
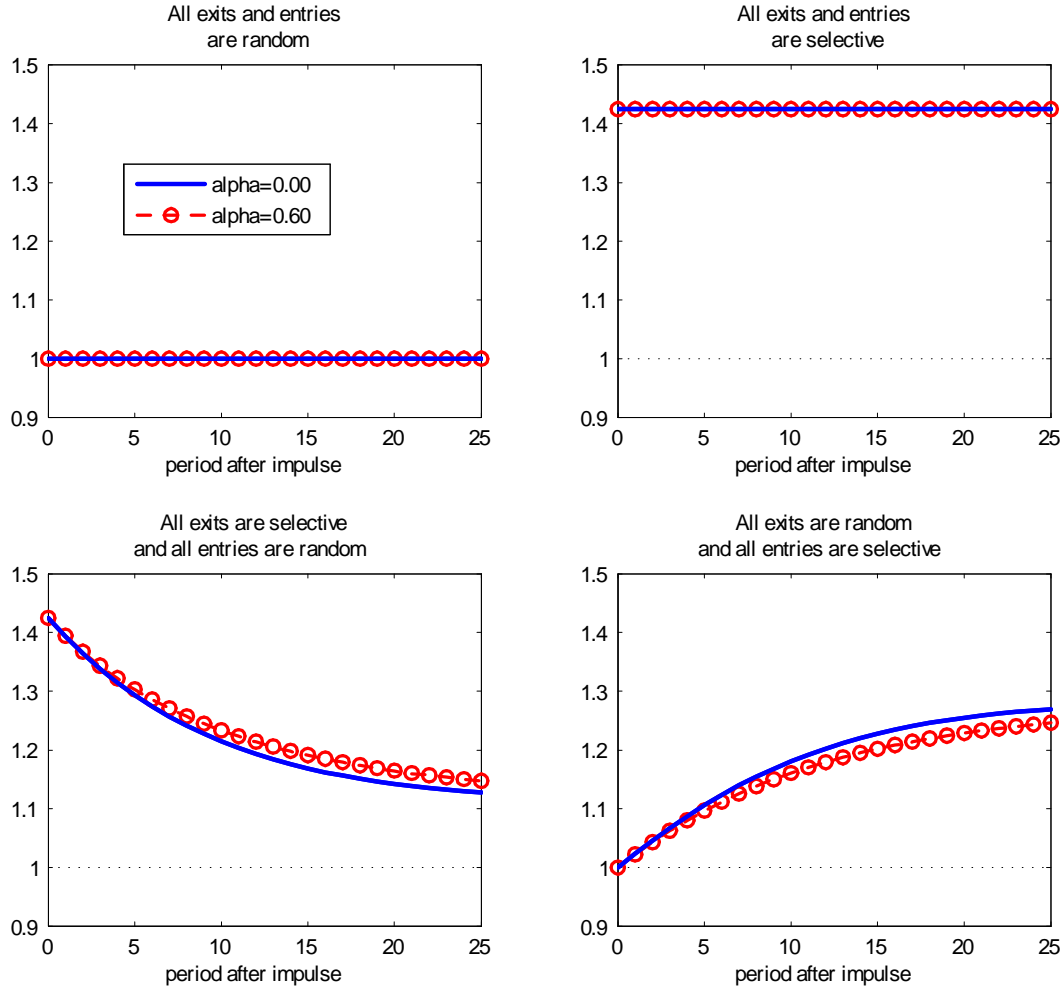


Figure 11: Marginal probabilities of observations usable to computed inflation in period t in the Calvo model with *iid* exchange rate innovations



Notes: This figure shows the marginal probabilities in periods $t - l$ to t of items whose price can be used to compute inflation in period t . The arrows illustrate the various paths through which a movement in the exchange rate in period $t - l$ could be reflected as a nonzero contribution to inflation in period t . The upper, middle, and lower panels show the case in which observations entering the sample are delayed by 0, $M \geq l$, and $M < l$ period(s), respectively.

Figure 12: Correction factors for the cumulative index response to an exchange rate shock in a Calvo model with real rigidities



Notes: This figure shows the theoretical correction factors one would apply to the cumulative import price index response to a one-time jump in the exchange rate. The model assumes Calvo pricing with an exogenous frequency of price changes of 15 percent and a substitution rate of 5 percent. The real rigidity parameter is specified as in appendix B.