Don't Panic: The Hitchhiker's Guide to Missing Import Price Changes

Etienne Gagnon, Benjamin R. Mandel, and Robert J. Vigfusson*†
Board of Governors of the Federal Reserve System

May 2011

Abstract

The response of an import price index to exchange rate movements may be mismeasured because some price changes are missed when constructing the index. Using two popular price-setting models, we investigate selection biases that arise when items experiencing a price change are especially likely to exit or to enter the index. Selective exits result in a downward bias in estimated pass-through for both the Calvo and menu-cost models. Selective entries also result in a downward bias, but the bias is initially small and its overall magnitude depends on how quickly prices respond to exchange rate movements. We calibrate these models using BLS micro data to derive empirical bounds on the magnitude of these biases. Our analysis suggests that the biases induced by selective exits and entries do not materially alter the literature's view that pass-through to U.S. import prices is low over typical forecast horizons.

*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, any other person associated with the Federal Reserve System, or the BLS. Kelsey Ayres, Benjamin Hopkins, and Gregory Howard provided superb research assistance.

†Comments and suggestions can be directed to etienne.gagnon@frb.gov, benjamin.r.mandel@frb.gov, robert.j.vigfusson@frb.gov.

1 Introduction

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. The most common way to measure exchange rate pass-through is based on regressions of changes in published import price indexes on changes in trade-weighted exchange rate indexes (along with other potentially important explanatory variables). Using these regressions, researchers have estimated low rates of exchange rate pass-through for the United States. Recent estimates (e.g., 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 percent in the contemporaneous quarter and an additional 2 percentage points over the next year, with little if any subsequent increases.

In this paper, we explore the possibility that such standard estimates are biased due to selection effects in the exit and entry of items in the index of import prices. In particular, we consider the possibility that items whose price is about to change are more likely than others to leave the index ("selective exits") and that items that recently experienced a price change are more likely than others to join the sample ("selective entries", a concept closely related to the "product replacement bias" of Nakamura and Steinsson (2009)). In both cases, an important fraction of micro price adjustments is taking place just outside the period over which item prices were recorded, thus lowering the measured response of prices to shocks. To throw these biases in stark relief to other index number issues associated with 'new' goods, we model exit and entry into the statistical agency's import price sample distinctly from the changing composition of the broader universe of traded items. As figure 1 shows, the composition of the universe of items constantly evolves as new items are brought to market and others are discontinued. For the BLS, this evolution poses the dual challenge of keeping its sample of import prices representative of the universe as well as confronting the thorny issues related to measuring the price inflation of items new to or discontinued from the universe. Abstracting from the latter issues, we focus on entry and exit into the sample.[1]

We show that, in principle, the biases resulting from entry and exit can be large, but that their magnitude is quite sensitive to one's assumptions about: (i) the choice of price-setting mechanism, (ii) the horizon over which pass-through is estimated, and (iii) the extent of selectivity in entry and exit.

First, we study the implications of selective exits and selective entries under two popular price-setting mechanisms, Calvo and menu-costs. We illustrate that the bias associated with

[1] The conference summaries of the Ottawa Group provide good overviews of on-going research conducted at statistical agencies worldwide on issues pertaining to items that are new to or discontinued from the universe. See also the BLS' "Measurement Issues in the Consumer Price Index" (1997), as well as the survey papers by Moulton (1996), Nordhaus (1998), and Gordon (2006).

selective entry is much smaller under menu-cost than Calvo pricing since, ultimately, it is the speed of pass-through rather than the frequency of individual price adjustments which determines the magnitude of this bias. In a Calvo model, the frequency of price updating and the speed of pass-through are tightly linked. Items that have a low frequency of updating respond very slowly to exchange rate movements. In a menu-cost model calibrated to the same low frequency of price changes, the response to an exchange rate shock is much more rapid, leaving little opportunity for the bias to accumulate. The bias associated with selective exit, in contrast, tends to be larger in the menu cost model when items are re-sampled at random from the universe.

To illustrate these points, we simulate our models under three extreme cases. In our first case, all exits and all entries into the import price sample are selective, which assumes that the largest possible fraction of price changes is censored. This case is related to the well-known *quality-change bias* by which statistical agencies have difficulties accounting for changes in quality from one vintage to the next, so that part or all of an item's effective price adjustment is censored. In this case, we show that true pass-though is underestimated by a similar factor under Calvo and menu-costs pricing. Moreover, the share of true pass-through left out depends little on the time horizon considered and the frequency of micro price adjustments.

Our second case, in which all exits are selective and all entering items are selected randomly from the universe, is related to the concept of *endogenous exits* wherein items with price adjustments are especially likely to exit the sample. As with our first case, the magnitude of the price index response is underestimated, even at short horizons, although by a smaller amount than when both exits and entries are selective. However, the overall bias relative to the first case is smaller because some of the randomly-selected entering items have not yet responded to current and past exchange rate movements. We find that randomizing entries in the presence of selective exits leads to a greater reduction in bias under Calvo pricing than under menu-cost pricing. The reason is that pass-through is relatively rapid under menu costs, so that only items added in recent periods contribute to the reduction in bias.

Our third case, in which all exits occur at random and all entries are selective, corresponds to the *product replacement bias* discussed by Nakamura and Steinsson (2009). A downward bias arises because price collectors systematically add observations to the sample that already have responded to current and past exchange rate innovations, making their next price change relatively insensitive to the history of exchange rate movements. Contrary to the above two cases, which involved selective exits, the estimated initial response of the price index suffers from little if any bias. Instead, the importance of the product replacement bias grows with the horizon considered. Over long horizons, Nakamura and Steinsson argue that this bias can be substantial. Using a Calvo model, they estimate that accounting for the product replacement bias would roughly double estimates of long-run exchange rate pass-through to non-oil U.S.

imports (from an elasticity of $0.2 - 0.4$ to $0.6 - 0.7$).

With these simulations in hand, we turn to assessing the empirical relevance of selective exits and entries. To begin, we note that the relevance of the associated biases for the purposes of policy makers varies by the amount of time it takes for the bias to become large. For instance, if it were to take 10 years for an exchange rate movement to transmit fully into import prices, then the far out lags (and any inaccuracies in their estimation) would have only miniscule effects on estimates of price inflation. Therefore, throughout our analysis, we focus on the response of import prices over the first two years following an exchange rate movement, which corresponds to the typical policy horizon of central bankers. The Federal Reserve staff forecast has a two-year horizon, and most central banks seek to achieve their inflation objective over a similar period.[2] Furthermore, this two-year horizon is likely the most relevant when using impulse responses to differentiate between models. For example, the implications of producer and local currency are most stark in the first two years after a shock. Likewise, the effects of adjustment costs in macro models are most apparent over relatively short horizons.

Measuring the empirical relevance of selective exits and selective entries over a typical forecast horizon is a difficult task because price collectors generally do not observe the reasons leading to an item's unplanned exit from the sample or its price history upon entry. We pursue several strategies to overcome these difficulties. We first review the methodology used by the BLS to deal with exits and entries. We argue that its sampling practices reduce the risk of selective exits and selective entries. We next use BLS micro data to derive empirical bounds on the price level response under various worst-case assumptions. To do so, we calibrate our Calvo and menu-cost models to match key features of individual import price adjustments and exchange rate movements. Our bounds suggest 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.

Finally, as an addition robustness exercise, we construct an alternative price index using BLS micro data that should in theory substantially mitigate the selective entry bias over typical forecast horizons. The constructed index delays the addition of sampled items to the index. When entries are selective, added items are too insensitive to past exchange rate movements. Simply delaying their entry in the index should therefore reduce this bias. However, when we estimate pass-through rates using these alternative price indexes, we do not find much evidence of bias reduction, casting further doubt on the importance of selective entry.

The remainder of the paper is structured as follows. Section 2 describes the sample of items used by the BLS to compute import price inflation and provides an overview of item entries and exits. Section 3 introduces the baseline Calvo and menu-cost models that we use to illustrate

[2] Other approaches in the literature consider the total effect on import prices of an exchange rate movement regardless of how long it take to materialize ("long-run pass-through") as well as the response between consecutive individual price adjustments ("medium-run pass-through," as defined by Gopinath, Itskhoki, and Rigobon, 2010).

the nature of the various biases and to gauge their quantitative importance. In Section 4, we discuss the possible biases associated with selection effects in sample exit and entry. Section 5 explores the empirical relevance of these biases by computing bounds on standard pass-through estimates and by constructing an alternative price index that should mitigates biases to some extent. Section 6 concludes.

2 Nature and occurrence of Item Exits and Entries

As stressed in the introduction, our study focuses on changes to the BLS import price sample, not on changes to the universe of items, although some of our results will have implications for them as well. For clarity, we reserve the terms "exit" and "entry" for changes in the composition of the sample. Throughout the presentation of the data below, as well as in the 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 enter 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. Our treatment of sample exits and entries is summarized in figure 2.

With the above terminology in mind, the remainder of this section provides some 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.

2.1 The International Price Program

Given identical data and similar methodology to Gopinath and Rigobon (2008) and Nakamura and Steinsson (2009), 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, though, import prices are collected through a monthly survey of U.S. establishments. We observe the price of approximately 20,000 imported items per month over our sample period, which runs from September 1993 to July 2007. The sample consists of rolling groups of items, each item having a sampling duration of about three years, on average.[3] The IPP chooses its firms and items based on a proportional-to-size sampling frame with some degree of over-

[3] For 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 ajdusting the last reported price by the average price change for the same firm and product category.

sampling of smaller firms and items.[4] Respondents must provide prices for actual transactions taking place as closely as possible to the first day of the month. 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. We also restrict our sample to U.S. dollar transactions, which account for about 90 percent of all observations.

2.2 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: (i) regular phaseout, (ii) accelerated phaseout, (iii) sample dropped, (iv) refusal, (v) firm out of business, (vi) out of scope, not replaced, and (vii) out of scope, replaced. Codes (i) through (iii) indicate that item exit is driven primarily by the phaseout schedule of the IPP sampling protocol. Codes (iv) and (v) 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 (vi) and (vii) are those instances in which price quotes are unavailable because the item ceases to be traded by importers.[5]

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 year and typically plans to retire items five years after their entry into the sample. The BLS may hasten the retirement if necessary, for example if an item is insufficiently traded. Given the planned nature of phaseouts, we see such exits as presenting minimal risks of being selective exits. Contrary to phaseouts, refusals and importers going out of business are not foreseen events. Nevertheless, we also see 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 making 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 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 no longer carry the item because domestic consumers stop placing orders.

It is worthwhile to note that exits are not generally accompanied by the simultaneous entry

[4] For 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.

[5] In some instances, the firm can provide an alternative item which meets BLS sampling needs (called 're-placed'), though that new item would still be recorded as a separate 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 sampling criteria of the exiting item. When it is possible, the BLS may link the price of the entering and exiting items through a one-time quality adjustment, in which case the change in the effective price is properly recorded. Such adjustments are relatively infrequent in practice, however. 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. For 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.[6]

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 entries 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 immediately after an unplanned exit when no quality adjustment is made. As mentioned above, such cases of rapid replacement are relatively infrequent. This assessment of the risk of selective entries stands in contrast with the working assumption in Nakamura and Steinsson (2009) that all entries are selective. For this reason, we will illustrate the magnitude of the bias in our quantitative analysis below under the full range - 0 to 100 percent - of possible selection effects.

2.3 Accounting for exit and entry

Every period, the number of observations in the sample used to construct import price inflation increases and decreases as items are added and dropped. More formally, let $exit(t)$ and $entry(t)$

[6] In those cases, where outgoing and incoming product groups are dissimilar, the benefit to overlapping their items is unclear.

be the number of price items exiting and entering the sample, respectively, in month t. 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 leaves in the next month. Analogously, the entry rate is measured as

$$entry\_rate\,(t) = \frac{entry(t)}{entry(t-1) + change(t-1) + no\_change(t-1)}.$$

The fourth through seventh columns of Table 1 show summary statistics about exits and entries over the period October 1995 to April 2005 for finished goods categories.[7] For industry groupings, we use the Bureau of Economic Analysis 3-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.[8] 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.[9]

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 course of the sample. However, the steadiness of the overall sample size hides a degree of heterogeneity in exit and entry rates

[7] Incomplete reporting for item discontintuation reasons in the IPP database 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 contain incomplete information about exits.

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

[9] An 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 product. The end-use weights for a given month would be the sum total of the individual item weights across items and HS10 products within that end-use. However, due to incomplete weight data for petroleum (end-use 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.

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 particularly high levels 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 3 presents the time series of the exit rate restricted to out-of-scope items along with an index of the broad nominal dollar.[10] 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 aggregate refusals and 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 real 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 their 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 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-

[10] The exit rates shown in the figure are 12-month moving averages.

through took place primarily in the 1990's, preceding the upticks in exit rates by quite a few years.

2.4 Micro price adjustments

As will be made clear in the next section, the quantitative implication of selective exits and selective entries can be sensitive to the price-setting frictions giving rise to infrequent and lumpy nominal price adjustments. It will be 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 is estimated to be 15.3 percent.[11] These levels are consistent with the weighted average of 14.1 percent in Nakamura and Steinsson (2009) and the median of 15 percent in Gopinath and Rigobon (2008). Treating all forced exits as price changes raises the frequency to about 16.5 percent over the sample period. The average absolute (nonzero) price change is 8.0 percent, in line with the mean overall estimate of 8.2 percent in Gopinath and Rigobon (2008). Here, again, there is significant dispersion across Enduse categories with items belonging to Enduse 101 (Fuels, n.e.s.-coal and gas) having an average price change of 13.4 percent, compared to 2.0 percent for Enduse 300 (Passenger cars, new and used).

2.5 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 simply reporting issues. Second, nearly half of all observations in the BLS sample refer to items that are traded between affiliates or entities of the same company. Although the BLS insists that these intra-company transfer prices be market-based or market-influenced, some have expressed concern 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

[11] Micro price studies differ in their usage of observations whose price is carried forward in the computation of the frequency. We choose to use such imputed price observations, consistent with the BLS practice for measuring inflation. One benefit of this approach 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 imputed prices, then our statistics would need to account for the fact that some quotes are inactive. Our decision makes price changes, exits, and entries slightly less frequent than if imputations were excluded. The broad findings of the paper do not hinge on this methodological choice, however.

changes. See Neiman (2010) and Gopinath and Rigobon (2008) for a comparison of intra-firm and market transactions.

3 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 various selection biases and how they interact with the frequency of prices change. As we will show, judgement on the quantitative importance of the biases is sensitive to the price-setting mechanism one sees as best representing the data-generating process. Although the Calvo and menu-cost models are only two of the many price setting mechanisms proposed in the literature, they are illustrative of the point that the severity of the biases often relates to the frequency of price changes and, more generally, to the speed at which exchange rate movements are passed-through to import prices.

3.1 Economic environment

We consider the following data-generating process for the change in the price (in logs) of an imported item i at period t,

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

Given the 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) the change in the exchange rate, Δx_t, and (c) the contribution of a (mean-zero) idiosyncratic factor, ε_{it}. The occurrence of a price change is marked by the indicator variable $\mathcal{I}_{it}^{f}$. The price deviation carried to the beginning of the next period is given by

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

If the firm does not change its price, then the aggregate and idiosyncratic shocks occurring in period t are simply added to the amount of price pressure that had already cumulated. If the firm adjusts its price, then the price is set to the optimum and no price pressure is carried into the next period.[12] 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,

[12] Our price-setting rule abstacts from forward-looking concerns, which greatly simplifies the exposition. Monthly exchange rate innovations are only weakly correlated, so our purely backward-looking rule should nevertheless capture central features of micro price adjustments in response to exchange rate movements.

to a firm's reset price. In what follows, we will 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 in general. To this end, we define aggregate price inflation as the average change in item prices,

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

Suppose that the econometrician estimates a linear model containing L lags of the aggregate variable,

$$\Delta p_t = a + \sum_{l=0}^{L} b_l \Delta x_{t-l} + r_t. \tag{1}$$

where r_t is an error term. In what follows, we explore how various assumptions about the timing of nominal adjustments impacts the econometrician's estimates of the regression coefficients.

3.1.1 Calvo model

In the Calvo model, the decision to change the price is exogenous to the firm. The indicator variable $\mathcal{I}_{it}^{f}$ is a random variable taking the value 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 1 shows (in a more general environment), the (plim) linear estimate of b_l is

$$b_l = f(1-f)^l \beta. \tag{2}$$

Intuitively, for a movement in the exchange l period 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 variable decays exponentially with the number of lags. Summing up the (plim) coefficients in the regression, we get

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

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

3.1.2 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. One can show that, to a first-order approximation, this range of inaction is symmetric around the price that sets the price pressure to zero (see Gopinath and Itskhoki, 2010, for a formal derivation). We thus approximate the decision to change the price as

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

Unfortunately, analytical results are challenging to derive 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. Moreover, note that the decision to change the price now depends on the value of β: The larger the pass-through coefficient for a given K, 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 benefit of adjusting the price), the more likely is a firms 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.

3.2 Calibration of the models

We first set the mean, standard deviation, and autoregressive coefficient of exchange rate innovations to match the corresponding moment 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 over time ($\rho = 0.19$). We report results for $\beta = 0.3$, which is in-line with recent estimates in the literature (e.g., Marazzi et al (2005), Gopinath, Itskhoki, 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 (2009), the median size of individual price changes rather insensitive to the frequency of price change changes, hovering between 6 and 7 percent. In the case of the Calvo model, we set the probability of a price change equal to a given frequency and calibrate the variance of individual innovations (which is assumed to be Gaussian) to match a median size of price changes of 6.5 percent. In the case of the menu-cost model, we choose the menu cost K and the standard deviation of ε_{it} to match both the median size and the average frequency of price changes. We make the additional assumption that ε_{it}. is normally distributed with mean zero. 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.

3.3 Impulse response to an exchange rate movement

Our exercise illustrates the point that the choice of a particular model can have important consequences for the dynamic response of the import price index. Although the Calvo and menu-cost models are calibrated to the same (in steady-state) frequency of price change and the long-run pass-through coefficient, the dynamic transmission of the exchange rate shock is markedly different between the models, with faster rates of pass-through at short horizons in the menu-cost model than in the Calvo model.[13]

The response of import price inflation to an exchange rate movement in the Calvo and menu-cost models are shown in the upper, middle, and bottom panels of figure 4 for (steady-state) frequencies of price changes of 5 percent, 20 percent, and 35 percent, respectively. In the case of 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 of price changes (upper panel), the exchange rate movement has not yet fully diffused by the end of the forecast horizon, although the impact on import price inflation is rather small. For a frequency of price changes of 20 percent (middle panel), the shock is almost entirely passed-through by the end of the forecast period, with negligible amount of trade price inflation left. Higher frequencies of price changes lead to even faster pass-through. The cumulative response of the import price index in the Calvo model can be seen in the left panels of figure 5 as the sum of the dark, medium, and light bars. For example, when the frequency of price changes is 5 percent, just over 70 percent of the long-run response of the import price index has taken place after two years, leaving almost 30 percent of the price response beyond the forecast horizon. By contrast, the transmission of the exchange rate shock is virtually complete after two years at frequencies of 20 percent or higher.

The speed of pass-through is markedly higher in the menu-cost model at all frequencies

[13] In practice, the frequency of price changes and the degree of exchange rate pass-through interrelated. 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.

(sum of dark, medium, and light bars in right panels of figure 5). Under our low-frequency calibration, there is negligible amount of import price inflation as a result of the shock after a year, even for frequencies as low as 5 percent, well over 90 percent of the long-term response of the price level has already taken place after a year. The speed of transmission is even higher for higher-frequency calibrations, with the bulk of the price level response taking place over just a handful of months.

4 Selection Effects in Item Exits and Entries

We now expand the baseline model to allow for the exit and entry of items in the index. As was the case earlier, we assume that the universe of items available for purchase is constant over time. Prices are collected at the end of the period after nominal adjustments, exits, and entries have taken place. Items entering or exiting the index thus cannot be used to compute inflation because either their past or current prices are unknown to the statistical agency.

Our model is summarized in figure 2. Exits occur through two channels. First, items face a exogenous probability d of dropping out of the sample every period (the "random exit" channel). These exits do not depend on the behavior of firms and are thus akin to the sample rotation performed by the BLS. 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 an exogenous probability e of exit the sample. Such a situation could occur if, for example, price collectors failed to hedonically adjust an item's price after a change in its characteristics, treating instead the old and new prices as unrelated exits and entries. In total, a fraction $s_t = d + (1 - d)\, e f_t$ of items exits the sample every period. Our model is not properly one in which some exits from the index are "endogenous" since the decision to exit is always exogenous to firms. Nevertheless, it has the feature that some exits partly censor the adjustment of the price index.

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 the same as for the entire universe, with some fraction f_t of deviations having their price reset during the period. Another fraction n of entering items *systematically* are sampled 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 (i.e., $u_{it} = 1$). Note the symmetry between the selective exit and selective entry channels: They both occur because items experiencing a price change in the current period are more likely to either exit or enter the index.

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

$$b_l = f\,(1-f)^l \left(\frac{1-e}{1-fe}\right)\left((1-d)^l + \frac{s\,(1-n)}{d}\left(1-(1-d)^l\right)\right)\beta. \qquad (3)$$

Relative to equation 2, the above expression has two new terms , $\frac{1-e}{1-fe}$ and $(1-d)^l + \frac{s(1-n)}{d}\left(1-(1-d)^l\right)$, which 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.

4.1 All exits and entries are random

When all entries and exits are exogenous (i.e., $s = d$ and $n = 0$), the (plim) coefficients in the Calvo model with *iid* exchange rate innovations are

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

In short, standard pass-through regressions are unbiased even though, every period, an arbitrary fraction $s = d$ (with $d < 1$) of items in the basket is replaced. Intuitively, items in the index have the same distribution of deviations from the optimum price as items in the universe; the only impact of exits and entries is to alter the number of observations usable to compute inflation at any point in time. For the same reason, biases are absent when exchange rate innovations are correlated and in the menu-cost model.

4.2 All exits and entries are selective

Consider now the case when all exits and entries are selective (i.e., $s = fe$ and $n = 1$). This case is related to the well-known "quality-change bias" by which statistical agencies have difficulties accounting for changes in quality from one vintage to the next, so that part or all of an item's effective price adjustment is censored. In our example, the price change is fully censored, the disappearance of the old vintage and the arrival of the new one being recorded as unrelated exits and entries.[14]

[14] In principle, mismeasured changes in quality can result in either upward or downward biases, 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 cheaper components). In practice, the quality change bias is associated with a systematic overestimation of inflation, which, in the case of the U.S. CPI, is estimated at around half a percentage point per year (see BLS (1997b)). By contrast, our canonical case with all exits and entries being selective, 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.

In the Calvo model with *iid* exchange rate innovations, we have

$$b_l = f\,(1-f)^l \left(\frac{1-e}{1-fe}\right)\beta. \tag{4}$$

All coefficients are downwardly biased by the same factor $(1-e)\,/\,(1-fe)$ relative to the true response. Note that $\hat{f} = \frac{(1-e)}{(1-ef)}f$ 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 (i.e., $s = ef$) is low because what crucially matters is the prevalence of exits among price changes rather than among observations in the index.

The left and right panels of figure 5 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. We tentatively assumed that a quarter of all price changes are accompanied by an exit, a proportion roughly equal to the median across 3-digit Enduse categories of the worse-case probability of exit (0.28) that we estimate later in section 5.2. We leave the other model parameters unchanged relative to the base case described in section 3.2. In addition to $n = 1$, the figure shows the special case $n = 0$ (no selective exit), which we will considered shortly. As noted earlier, the censoring of price changes reduces the frequency of price changes observed by the econometrician. For underlying frequencies of 5, 20 and 35 percent in the population of items, the econometrician would report frequencies of about 4, 16, and 29 percent, respectively.

In our calibrated Calvo and menu-cost models, the size of the bias created by selective exit is somewhat large over the forecast horizon at all frequencies considered when the price of entering items has been optimized. For low frequencies of price changes, the bias is roughly equaled to e, the probability of an item exit conditional on a price change, which we set to a quarter in the simulations. The bias declines somewhat as we consider higher frequencies, reaching about 20 percent of the long-run response when the underlying frequency is 35 percent.

4.3 All exits are selective and all entries are random

It can be challenging for price collectors to know if exits are selective or random as they have to press respondents for information about the circumstances in which they take place. Price collectors 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 $n = 0$).

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

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

The size of the bias depends on the relative strength of two opposite 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 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 arbitrary large number, so that individual coefficients are always upwardly biased at sufficiently long lags. Nevertheless, the cumulative response remains downwardly biased because the coefficients converge rapidly to zero.[15]

As shown in the left panels of figure 5, 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 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 5, 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.

4.4 All exits are random and all entries are selective (product replacement bias)

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 special case corresponds to the environment considered by Nakamura and Steinsson (2009). Under these assumptions, the (plim) coefficient on the $l-th$ lag of exchange innovations in our baseline Calvo model is

$$b_l = f(1-f)^l(1-s)^l\beta. \tag{6}$$

[15] When all exits are selective and all entries are random, the long-run response of the index is given by $\frac{1-fe-(1-f)e^2}{1-fe}\beta < \beta$. Randomizing entries thus reduces the downward bias without eliminating it entirely. The same conclusion applies for the general case, for which the long-run response is $\left(\frac{1-e}{1-fe}\right)\left(\frac{f+(d+(1-d)ef)(1-n)(1-f)}{f+d-fd}\right)\beta$. One can show that this response is always smaller than β but greater than the case with selective entries.

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 $(1-s)^l$, which Nakamura and Steinsson (2009) designate as "product replacement bias."

A few comments are worth making. First, the nature of the product replacement bias is that items entering the basket *systematically* are less sensitive to past movements in the exchange rate than items in general. Their inclusion of these entering items in pass-through regressions thus lowers the effective estimates. 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, the expression "product replacement bias" is slightly misleading. The product replacement bias derived by Nakamura and Steinsson (2009) is associated with a specific type of replacement, one that includes a random exit and a selective entry. As section 4.3 shows, assuming that exits are selective rather than random leads to different expressions for b_l. Also, even with no items exiting the basket, and hence no replacement, any enlargement of the index results in a product "replacement" bias when prices of added items systematically are less responsive to past movements in the exchange rate than other prices.

Fourth, as stressed by Nakamura and Steinsson, the bias is most important for product categories with very low frequency of price changes. The left panels of figure 6 illustrate the bias over the policy-relevant horizon under Calvo pricing 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-third of the actual cumulative pass-through is correctly estimated at the 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 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 since s is small, meaning that the product replacement 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 product replacement 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 6 show the estimated cumulative contribution of the regression coefficients on the various lags of exchange rate movements (the dark-shaded bars), along with the product replacement bias left out by the econometrician (the light-shaded bars), under menu-cost pricing. The product replacement 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 product replacement bias does not have time to cumulate to something large.

Finally, our figure depicts a the worst-case assumption that all entries are selective ($n = 1$), the assumption maintained by Nakamura and Steinsson. As shown in section 4.1, there would be no bias if price collectors were replacing exiting items by observations randomly selected from the population ($n = 0$). In the generic case in which all exits are random and a fraction n of entries are selective, we have

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

Departing from the extreme case of $n = 1$ can substantially reduce the size of the product replacement bias. As a rule of thumb, the reduction in the bias by the end of the forecast horizon is roughly proportional to $1 - n$, so that, for example, setting $n = 0.5$ would roughly halve the area represented by the light bars.

5 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 price-setting 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 mixes elements 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 a method that purges standard pass-through estimates of much of the product replacement bias. Overall, our findings suggest that the biases induced by selective exits and selective entries, while a concern, do not materially alter the literature's view that pass-through to U.S. import prices is low over typical forecast horizons.

5.1 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 entry and exit 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 an appropriate choice for robustness. The estimated impulse responses to a 1-percent depreciation of the U.S. dollar are presented in figure 7. 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. We obtain a similar response when we regress the index for finished goods on the exchange rate (the dashed line in the lower-left panel).[16] The shape of the impulse response shares features 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.

Figure 8 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 3.2, with the minor difference that we seek to match the observed cumulative rate of pass-though in the last quarter of the forecast horizon rather than some illustrative long-run value. Figure 8 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. Though 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.

5.2 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 4.2 to 4.3, 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 forced exits mask a price change. We dismiss voluntary exits as *not* being selective because they typically are planned years in advance by the BLS and thus are unrelated 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 voluntary and forced exits 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., $n=1$), as posited by Nakamura and Steinsson (2009).

[16] Our 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 occurence of selective exits and entries could account for that pattern.

Given the sensitivity of biases to price-setting 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 4 to 6. 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 9. 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, this figure could be as large as 0.30 percent in the menu-cost model, and 0.32 percent in the Calvo model. The slightly wider bound under Calvo pricing is due to the larger correction for selective entries in that model. 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. Under the worst case of product replacement bias (lower-left panel), the corrected response is very close to the actual response under menu costs, but remains somwhat higher (0.30 percent) that the uncorrected estimate (0.24 percent) by the last quarter of the forecast horizon. Even in this case, the near-term estimate remains relatively precise, however.

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 the most plausible price-setting 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 forced exits correspond to selective exits and a selective entry, whereas all voluntary exits are associated with random exits and random entries, so that about half of all exits and entries are selective. If one believed that the BLS' general sampling procedure avoided selective entries, then the resulting impulse response would provide a tighter upper bound on the actual response. The corrected cumulative response under this last set of assumptions is presented in the lower-right panel of figure 9. 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, while the potential for large sampling biases certainly exists under our worst case assumptions, the bounds analysis most in line with the industry-level micro data suggests a limited impact of these biases on standard estimates of pass-through over typical forecast horizons.

5.3 Reducing biases through delayed entries

If the estimates were subject solely to a product replacement bias, as postulated by Nakamura and Steinsson (2009), then one could use a simple trick to remove much of that bias over the policy-relevant horizon. Remember that the product replacement bias arises because entering items systematically are less responsive to past exchange rate movements than items in the universe. Therefore, simply delaying the entry of substitutes in the index reduces 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

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

Delaying entries entirely thus eliminates the product replacement bias for the coefficients associated with the current and 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 10 shows that delaying entries by 6 months can go a long way in reducing the product replacement bias over the policy-relevant horizon in the Calvo model. 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 product replacement 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 product replacement 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.

It turns out that our trick of delaying entries can also mitigate biases in the presence of selective exits. We show in the appendix that, with both selective exits and selective entries that the (plim) regression coefficients under an M-period entry delay in the Calvo model are

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

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

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

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 why delaying entries can improve pass-through estimates in the presence of selective exits 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 inflation today. Under Calvo pricing (left panels of figure 11), it turns out that this selection effect perfectly offsets the downward bias stemming from the censoring of price changes as we consider an arbitrary long entry delay and forecast horizon. Under menu-cost pricing (right panels), the gains are negligible due to the greater mixing of observations.

Our simulations suggest that delaying entries is most effective at reducing biases associated with selective entries over typical forecast horizons, a finding that is robust across pricing mechanisms and degrees of price flexibility. Thus, delaying entries can help us shed light on the empirical importance of selective entries: If estimated pass-through over the forecast horizon increases much after delaying entries, then selective entries may be economically important. By contrast, if estimated pass-through is insensitive to delaying entries, then selective entries may be a marginal phenomenon.

Using the BLS microdata, we have computed price indexes for the enduse categories of capital goods, automotive products and consumer goods. We have constructed one index that should replicate the BLS published index (ie no delay). We have also constructed two indexes that implement the 6-month and 9-month entry delay, respectively. Figure 14 shows the results. As shown in the left-hand column, these alternative price indexes are more volatile than the corresponding published BLS index (the thick black line). However, estimated pass-through rates are very similar, whether we use the published BLS index or these constructed indexes. If selective entry were important, then the estimated pass-through rates for the constructed indexes with delayed entry should be 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. As such, the available evidence suggests a limited role for selective entry.

6 Concluding Remarks

In this paper, we have investigated selection biases in standard exchange rate pass-through regressions arising when micro price adjustments tend to occur just after an item has exited the price index (selective exit), or just before it has entered the index (selective entry). For both Calvo and menu-cost price-setting 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 price-setting assumptions.

In particular, in the presence of both selective exits and selective entries, the import price response is biased downward in both the Calvo and menu-cost 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 case considered in Nakamura and Steinsson (2009), the downward bias tends to be small in the menu-cost model over any horizon, whereas it slowly grows from being negligible at short horizons to possibly being important 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 some 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, while 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 model simulations 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. Furthermore, our worst-case scenario is likely too severe as our constructed alternative price indexes suggest a limited role for selective 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. Similarly, the implications of selective exits and selective entries extend to measurement of the response of price indexes to other shocks than exchange rate movements, be they aggregate or idiosyncratic.

Finally, we believe that 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 but limited guidance on these aspects.

References

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.

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.

Gertler, Mark and John Leahy (2008). "A Phillips Curve with an *Ss* Foundation," *Journal of Political Economy*, vol. 116(3), pages 533-572.

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 and Rigobon (2008): "Sticky Borders," Quarterly Journal of Economics, 123(2): 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. "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, April 2005.

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, pp. 924-47.

Moulton, Brent R. (1996). "Bias in the Consumer Price Index: What Is the Evidence?," *Journal of Economic Perspectives*, vol. 10(4), pages 159-77.

Nakamura, Emi, and Jón Steinsson. "Lost in Transit: Product Replacement Bias and Pricing to Market," mimeo, Columbia University, June 2009.

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

Nordhaus, William D. (1998). "Quality Change in Price Indexes," *Journal of Economic Perspectives*, vol. 12(1), pages 59-68, Winter.

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

Appendix 1: Regression Coefficients in the Calvo Model

In this appendix, we derive analytical expressions for pass-through coefficients when the data are generated by a Calvo model with selection biases in the exit and entry of items. The environment is as described in section 4 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 provide a proof that the bias on the coefficients declines as one delays the entry of items in the index.

6.1 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 (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 $\mathcal{I}_{it}^{s}=\mathcal{I}_{it}^{d}+\left(1-\mathcal{I}_{it}^{d}\right)\mathcal{I}_{it}^{f}\mathcal{I}_{it}^{e}$ be an indicator variable 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

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

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 can be 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, can have a nonzero contribution to inflation.

Proceeding similarly with b_1,

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

Since Δx_t and ε_{it} are assumed to be independent of Δx_{t-1}, the covariance term is impacted solely through the possible interactions between Δx_{t-1} and the cumulated price pressure u_{it}.

Conditioning on past realizations of the indicator variables, there are five distinct cases:

$$u_{it} = \begin{cases} u_{it-1} + \beta \Delta x_{t-1} + \varepsilon_{it-1} & \text{if } \left\{ \mathcal{I}^s_{it-1} = 0, \mathcal{I}^f_{it-1} = 0 \right\} \\ 0 & \text{if } \left\{ \mathcal{I}^s_{it-1} = 0, \mathcal{I}^f_{it-1} = 1 \right\} \\ u_{it-1} + \beta \Delta x_{t-1} + \varepsilon_{it-1} & \text{if } \left\{ \mathcal{I}^s_{it-1} = 1, \mathcal{I}^n_{it-1} = 0, \mathcal{I}^f_{it-1} = 0 \right\} \\ 0 & \text{if } \left\{ \mathcal{I}^s_{it-1} = 1, \mathcal{I}^n_{it-1} = 0, \mathcal{I}^f_{it-1} = 1 \right\} \\ 0 & \text{if } \left\{ \mathcal{I}^s_{it-1} = 1, \mathcal{I}^n_{it-1} = 1 \right\} \end{cases}. \tag{7}$$

Consequently,

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

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 no price change (marginal probability $(d + (1-d) fe)(1-n)(1-f)$) have a nonzero contribution of Δx_{t-1} to inflation.

The general case with b_l is illustrated in figure 12. It shows the various states that are usable in the computation of inflation at each period along with their marginal probabilities. The arrows indicate the paths through which an exchange rate movement in the period $t-l$ is reflected in a nonzero price change in period t. The figure also shows, for each period, the marginal probabilities associated with the paths that can be used to compute inflation. Observations that have not 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 item 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 $(1-d)(1-f)$. The marginal probability associated with the entry of an item that simple

Summing up the probabilities across all usable paths, we have

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

6.2 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 5.3 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 12. Because entries are delayed by more periods than the lag of the exchange rate movement considered, all observations used to compute inflation have been continuously in the index since before period $t-l$. For an exchange rate movement at period $t-l$ to be reflected in a nonzero price change today, we must have had had no price change from period $t-l$ to $t-1$ (marginal probability $(1-f)/(1-fe)$), and a price change at t (marginal probability $(1-e) f/(1-fe)$). The resulting coefficient is

$$b_l (l \leq M) = f (1-f)^l \frac{(1-e)}{(1-fe)^{l+1}} \beta. \tag{8}$$

The case of $l > M$ is illustrated at the bottom of figure X. 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 already present in the index at the end of period $t-M-1$ can be used to compute inflation. Summing up the probabilities over all possible paths and simplifying, we get

$$b_l (l > M) = f (1-f)^l \frac{(1-e)}{(1-fe)^{M+1}} \left((1-d)^{l-M} + \frac{s(1-n)}{d} \left(1 - (1-d)^{l-M} \right) \right) \beta. \tag{9}$$

6.3 Proof that biases are declining in the entry delay

We conclude this appendix by showing that delaying the entry of items in the index always improves pass-though 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 coefficients 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 nevertheless bounded above by the true response.

6.3.1 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 8, so that $b_l(M) = b_l(M+1)$. When $l = M+1$, $b_{M+1}(M)$ is given by equation 9 and $b_{M+1}(M+1)$ is given by equation 8. For $b_{M+1}(M) \geq b_{M+1}(M)$ 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 9. Note that

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

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 prove that the result hold 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 the cumulative pass-through over any forecast horizon also is increasing in the entry delay.

6.3.2 Step 2: estimated cumulative response is bounded above by true response

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

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

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

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

It is immediate that the above expression is bounded above by the unbiased case.

Table 1: The frequency and size of import price changes and IPP item entry and exit rates

	Enduse	Enduse Weight	Freq. of Change Std.	Freq. of Change Alt.	Entry	Exit All	Exit Selective	Exit Random	Absolute Size
Capital Goods	210 Oil drilling, mining & const. machinery	1.9	6.9	7.7	2.9	2.5	0.9	1.6	6.6
	211 Industrial & service machinery, n.e.c.	10.6	6.3	7.0	2.5	2.5	0.8	1.7	6.7
	212 Agricultural machinery & equip.	0.7	8.9	10.0	3.1	2.7	1.2	1.5	5.3
	213 Computers, periph. & semiconductors	12.7	9.7	11.7	5.0	3.7	2.2	1.5	9.6
	214 Telecommunications equip.	4.0	5.8	7.4	3.6	3.4	1.8	1.6	8.9
	215 Business mach. & equip., ex. Computers	0.9	5.2	6.6	2.5	3.3	1.5	1.8	6.3
	216 Scientific, hospital & medical machinery	2.6	4.9	6.0	3.2	3.1	1.2	1.9	6.9
Auto-motive	300 Passenger cars, new & used	13.6	5.3	6.9	3.5	2.8	1.6	1.1	2.0
	301 Trucks, buses, & special-purp. vehicles	2.4	5.8	7.6	3.9	2.8	1.9	0.9	2.9
	302 Parts, engines, bodies, & chassis	9.3	8.0	9.2	3.0	2.8	1.2	1.6	7.1
Consumer Goods (Ex. Food & Auto.)	400 Apparel, footwear, & household goods	11.2	3.9	5.6	3.6	3.5	1.7	1.8	7.6
	401 Other consumer nondurables	8.6	6.0	6.8	2.7	2.4	0.8	1.5	7.7
	410 Household goods	10.4	4.6	5.7	3.0	2.9	1.2	1.7	6.2
	411 Recreational equip. & materials	3.9	4.8	6.5	3.1	3.2	1.8	1.5	5.7
	412 Home entertainment equip.	5.2	5.6	7.7	4.1	3.7	2.2	1.5	5.8
	413 Coins, gems, jewelry, & collectibles	2.2	6.9	8.0	3.1	3.1	1.1	1.9	5.9
	Total	100.0	6.2	7.6	3.4	3.0	1.5	1.5	6.7

Notes: Shown are average frequencies of price changes, exits, and entries for the IPP sample of U.S. imports. The frequency is the number of items changing price divided by the number of items in the sample in both the current and previous month. Entry and exit represent the first and last period in which an item enters the IPP sample. Selective exits are instances where firms report an item to be out of scope as the reason for discontinuation. The absolute size of a price change is the average percentage change for a given item's price conditional on a price change. Each entry in the table averages across items within a 3-digit enduse category for the months October 1995 through April 2005. Individual items are unweighted; enduse-month pairs are weighted by their 2006 import value.

Figure 1: Sampling from the universe of items available for purchase

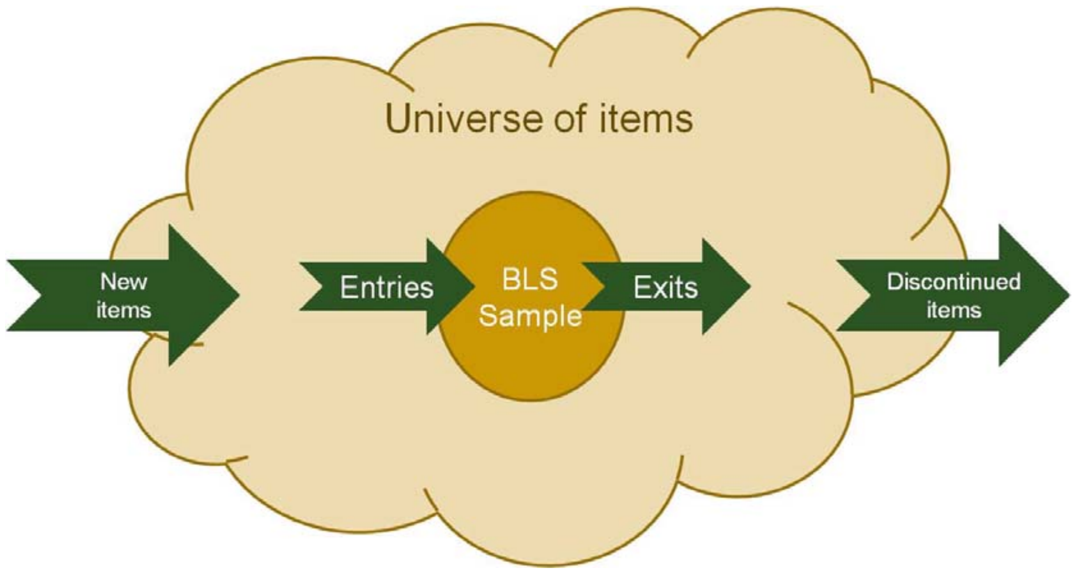

Figure 2: Selection effects in item exit and entries in the model

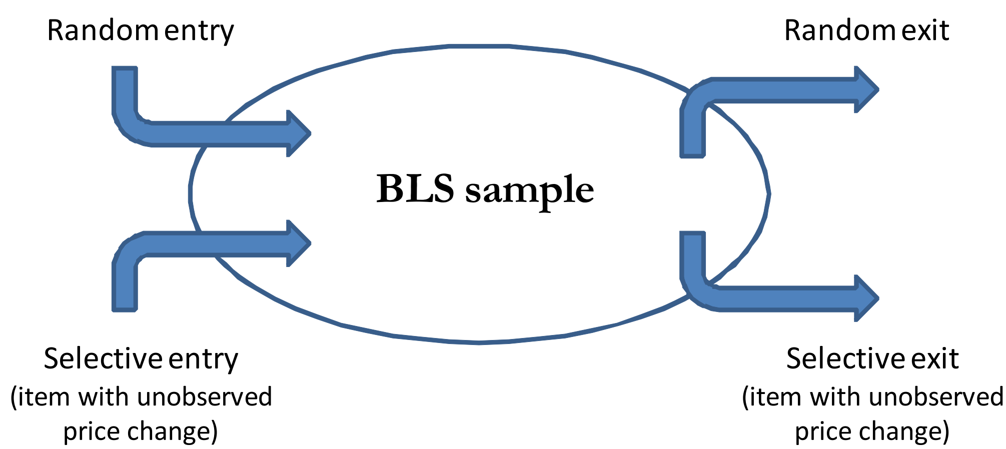

Figure 3: Exit rate, entry rate, and the dollar

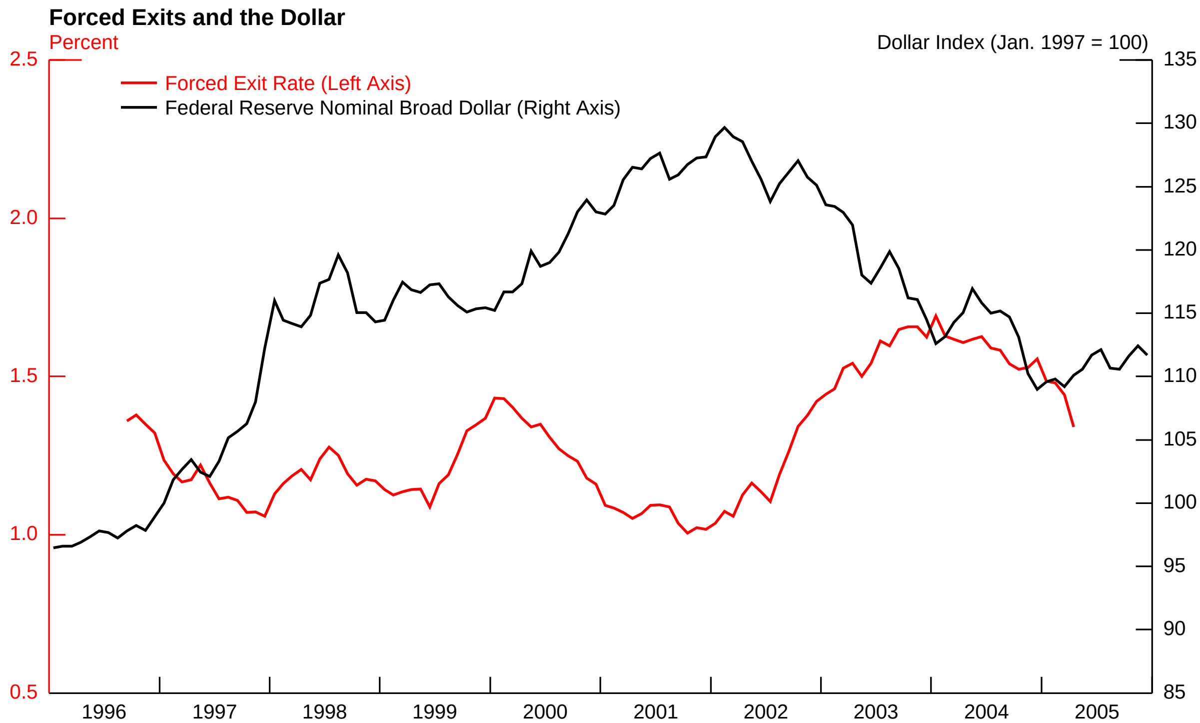

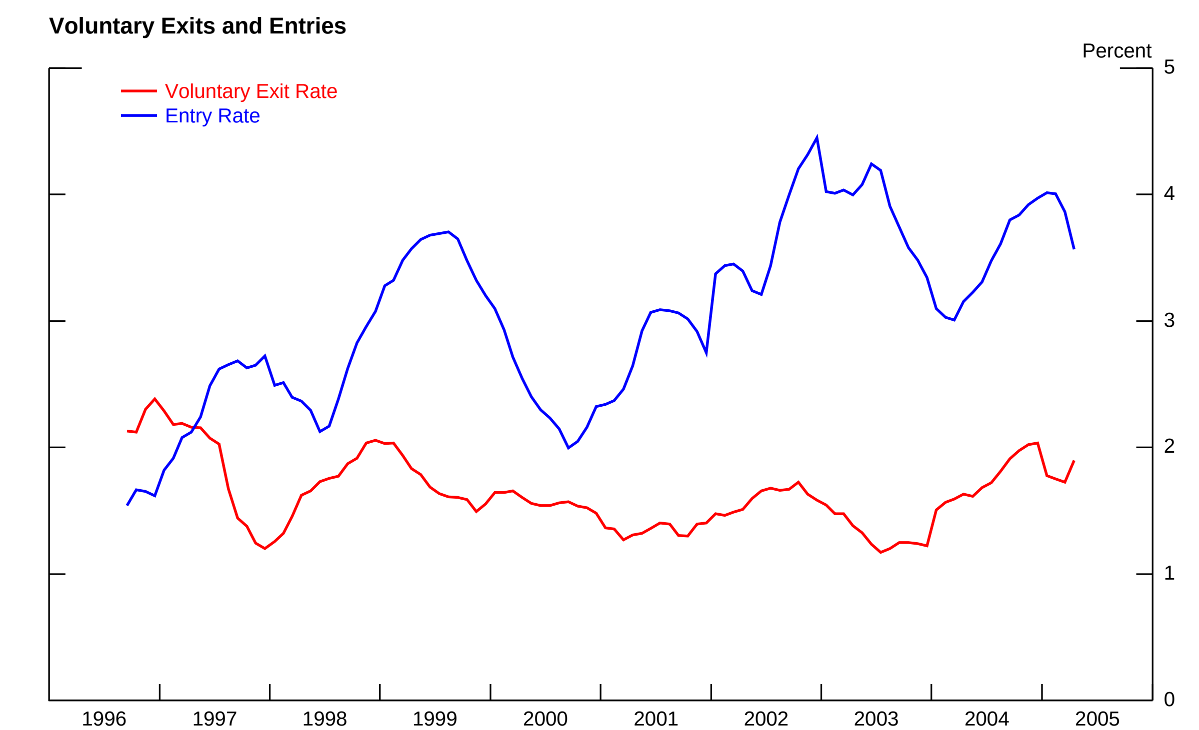

Figure 4: Coefficients on lags of the exchange rate in pass-through regressions in baseline Calvo and menu-cost models

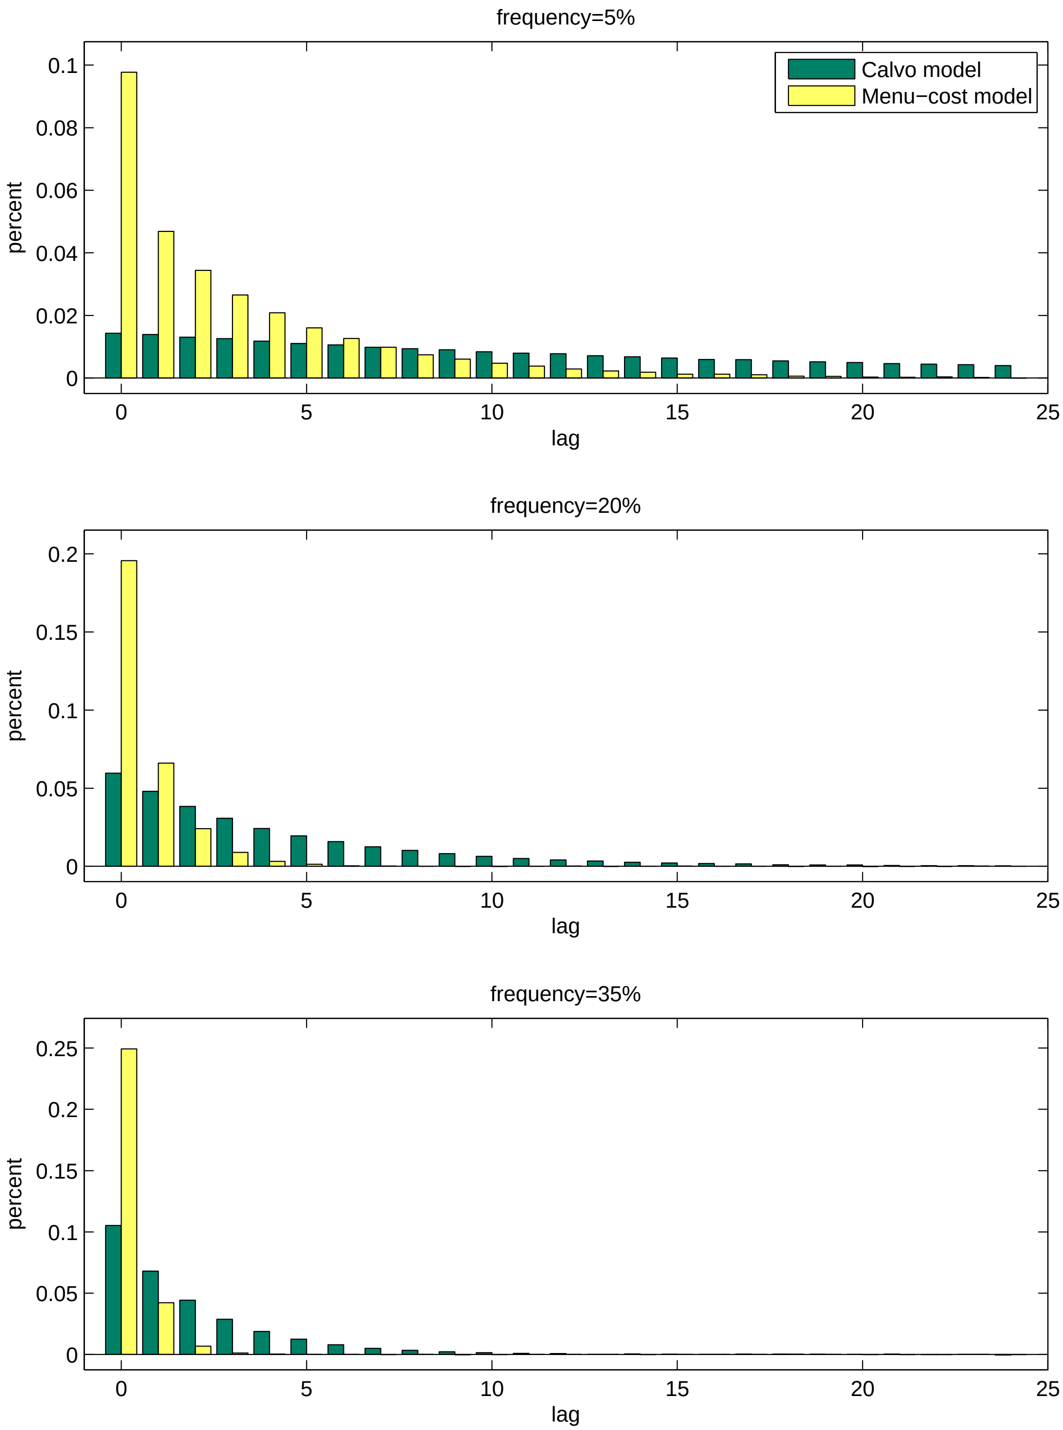

Figure 5: Cumulative contribution of coefficients on lagged exchange rate variables under selective exit ($e = 0.25$)

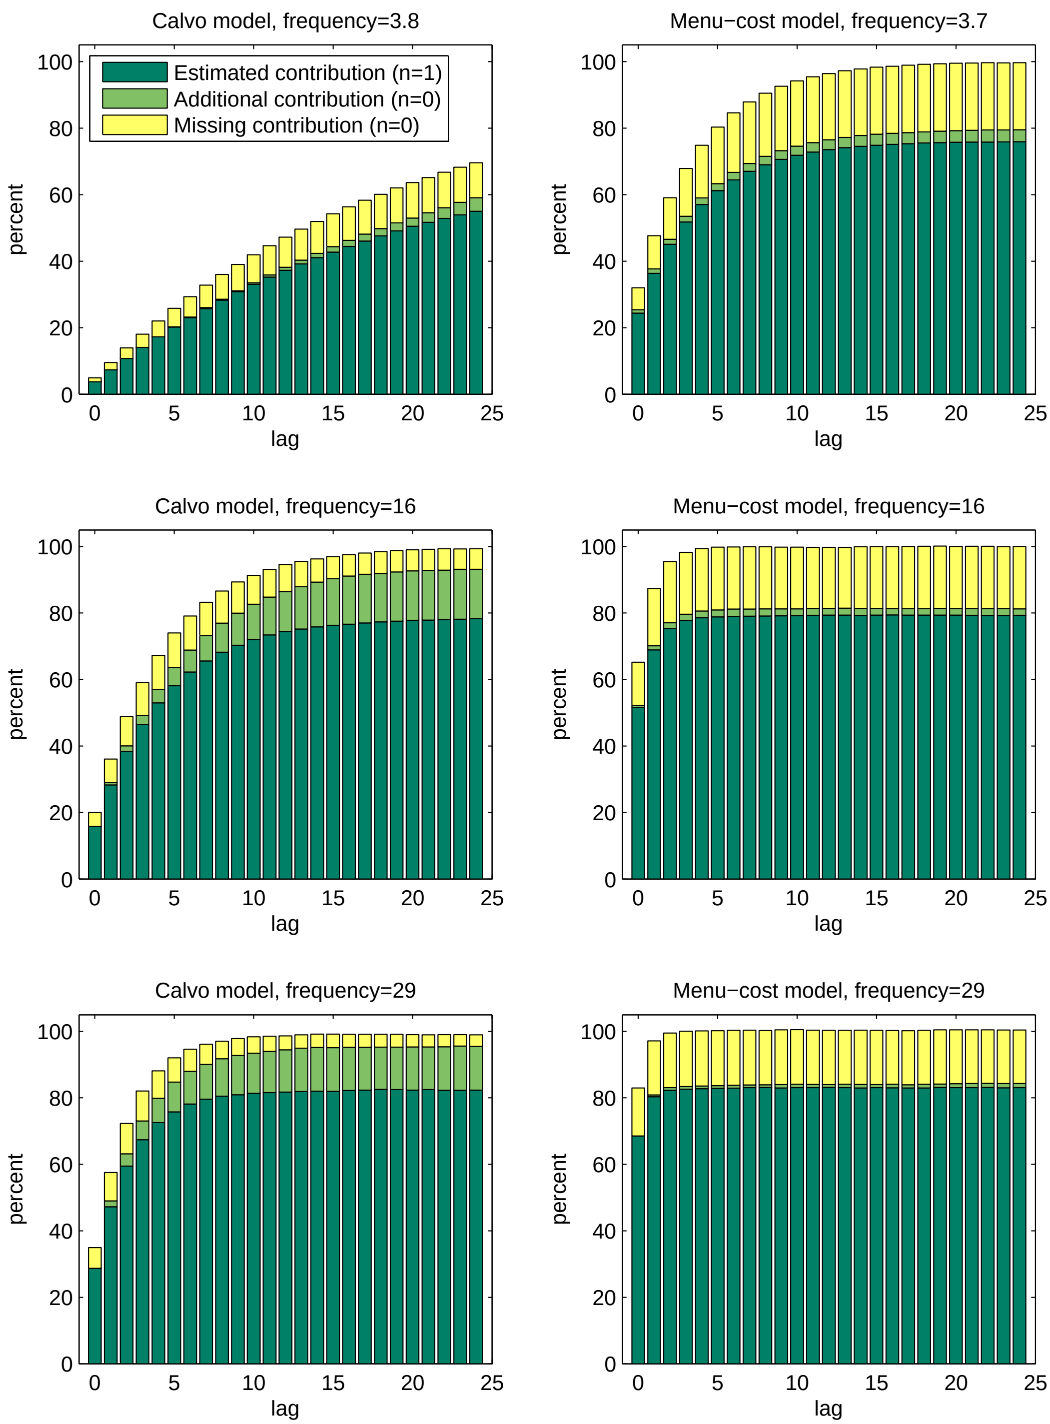

Figure 6: Cumulative contribution of coefficients on lagged exchange rate variables in Calvo model under severe product replacement bias ($n = 1$, $s = 0.05$)

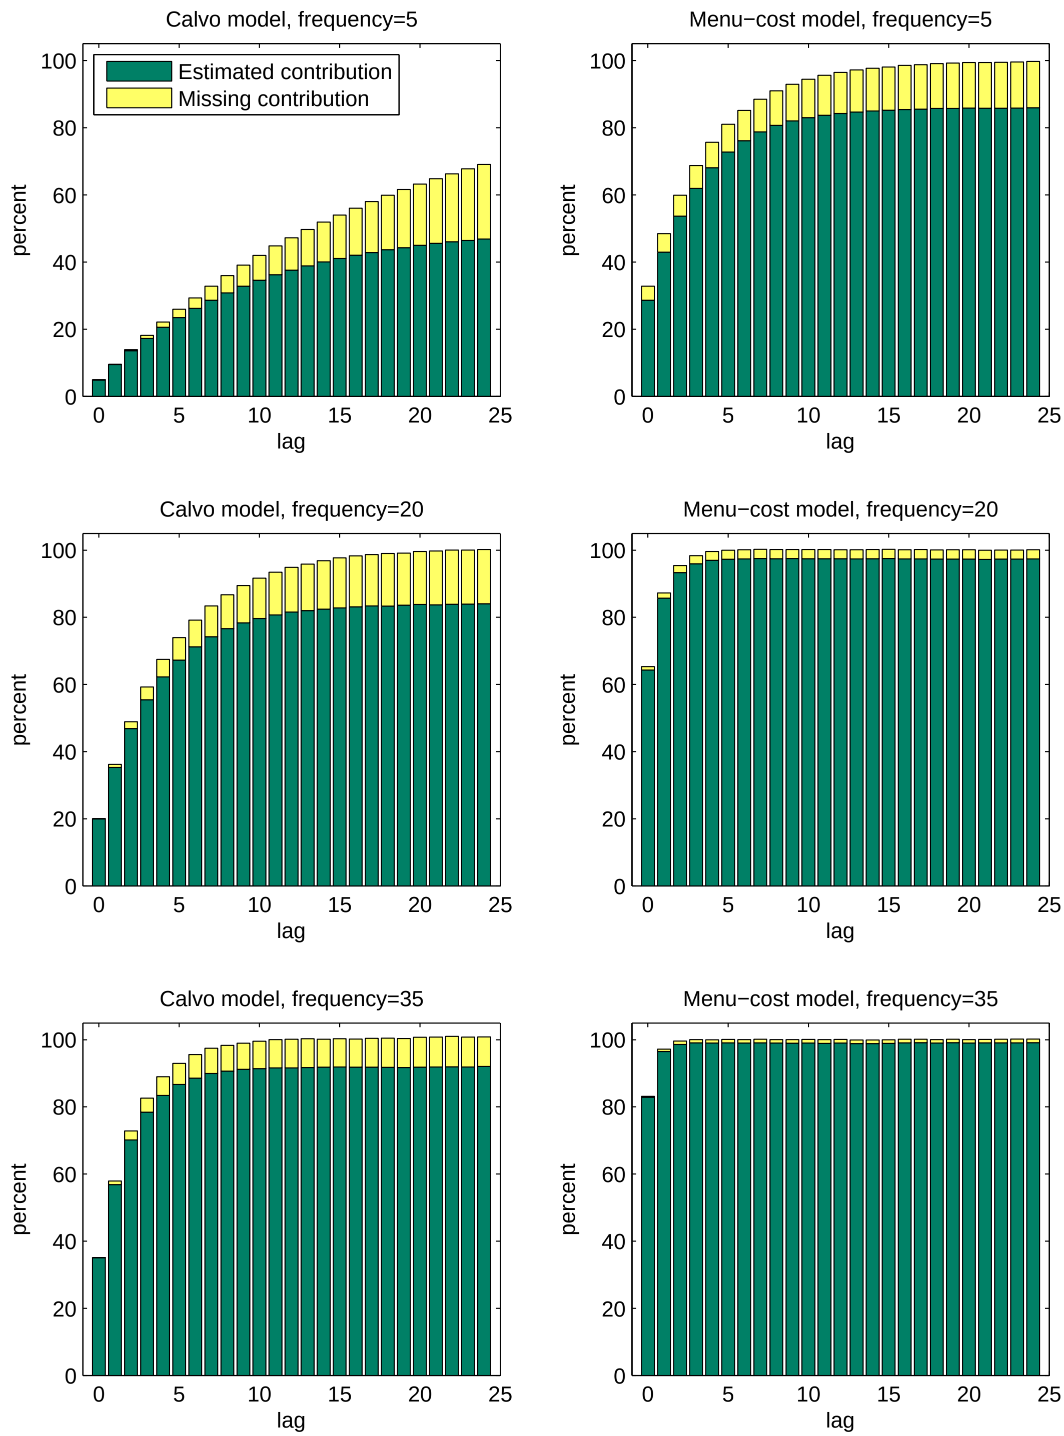

Figure 7: Pass-through to imported finished goods prices following a 1-percent depreciation of the dollar (by 3-digit Enduse categories)

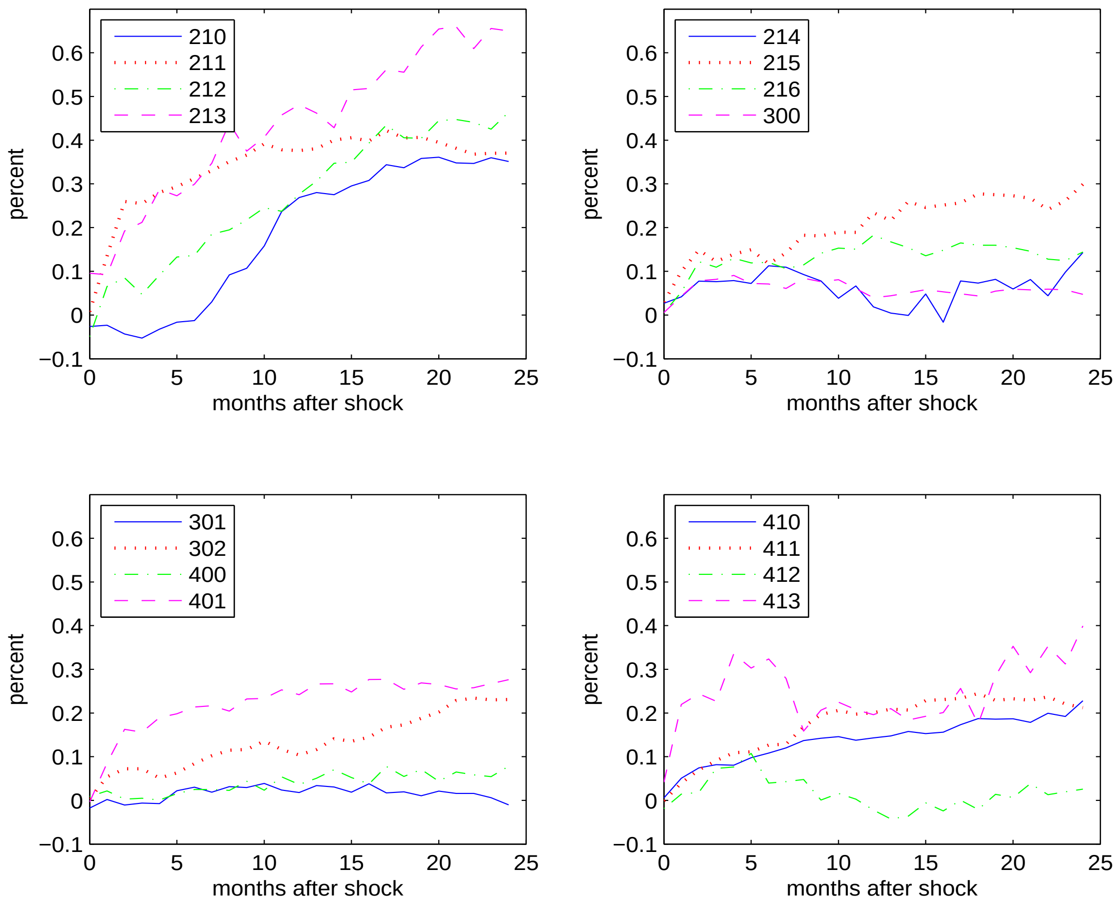

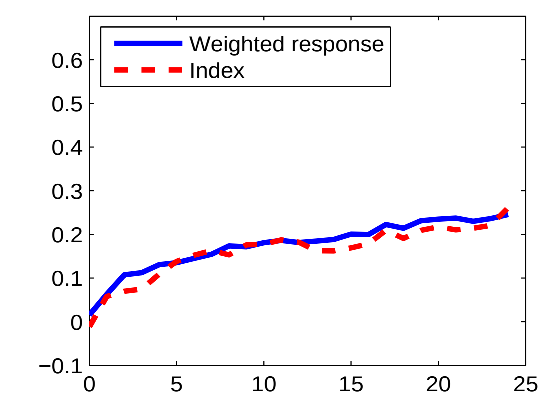

Enduse categories

210 – Oil drilling mining and construction machinery and equipment
211 – Industrial and service machinery
212 – Agricultural machinery and equipment
213 – Computers peripherals and semiconductors
214 – Telecommunications equipment
215 – Business machinery and equipment except computers
216 – Scientific and medical machinery
300 – Passenger cars new and used
301 – Vehicles designed to transport goods
302 – Parts engines bodies and chassis
400 – Apparel footwear and household goods
401 – Other consumer nondurables
410 – Household goods
411 – Recreational equipment and materials
412 – Home entertainment equipment
413 – Coins gems jewelry and collectibles

Figure 8: Pass-through to imported finished goods prices following a 1-percent depreciation of the dollar: models versus data

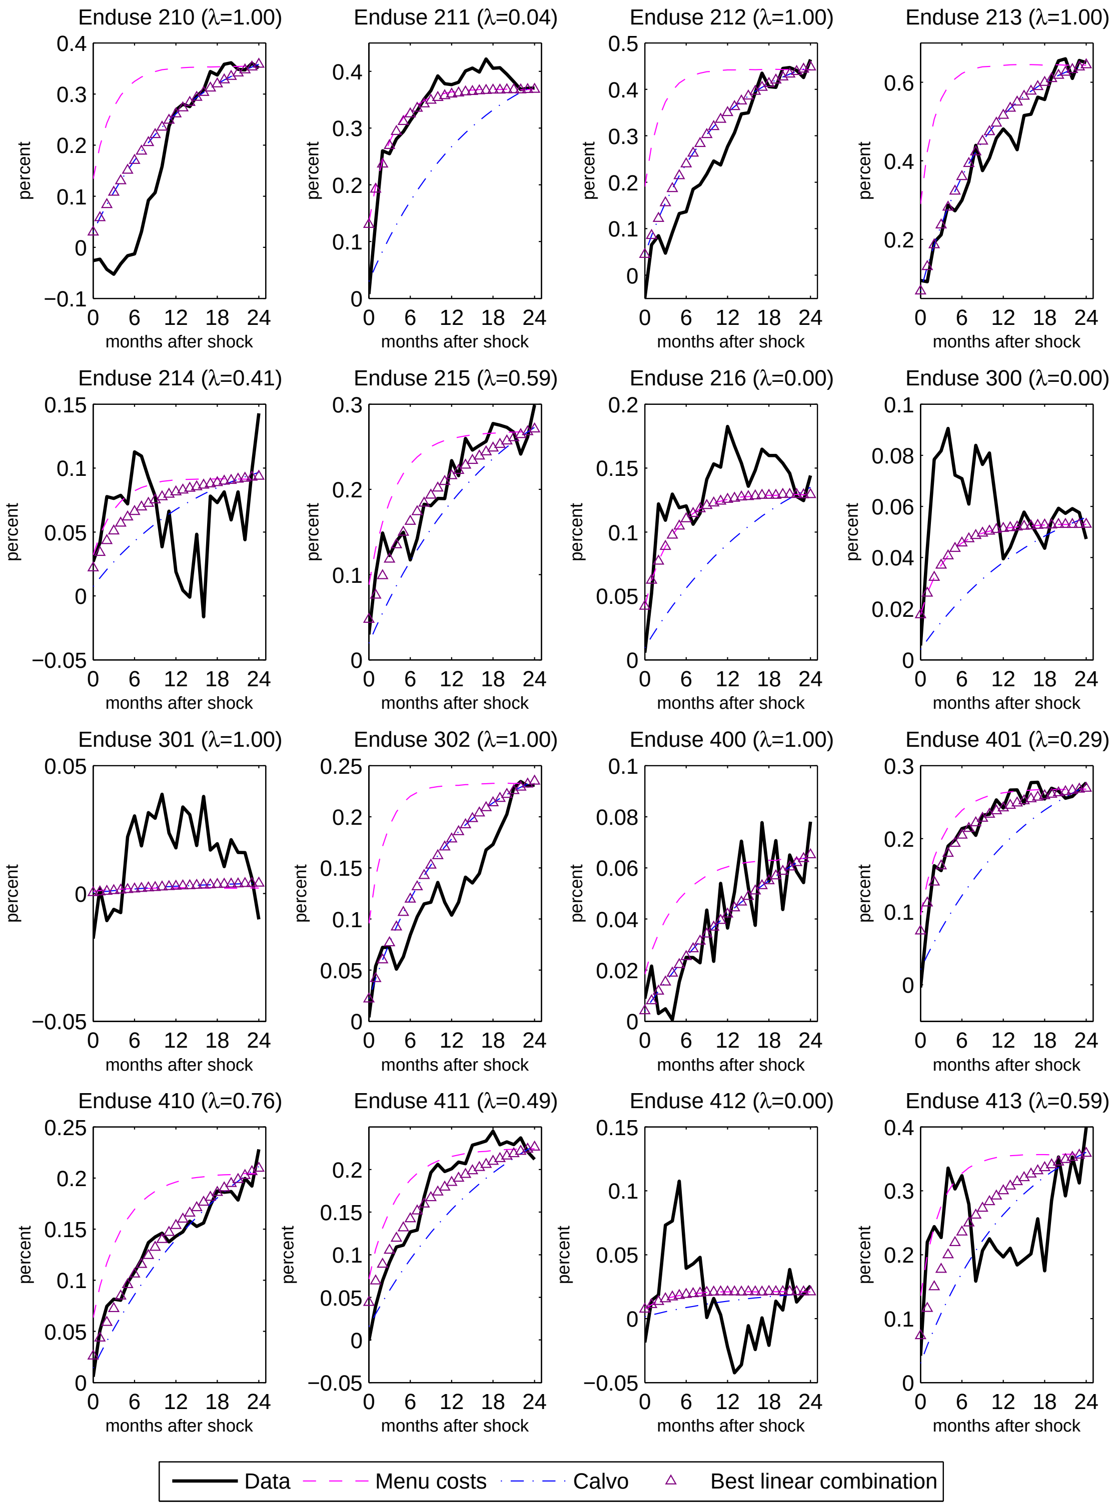

Figure 9: Upper bounds on exchange rate pass-through to finished goods

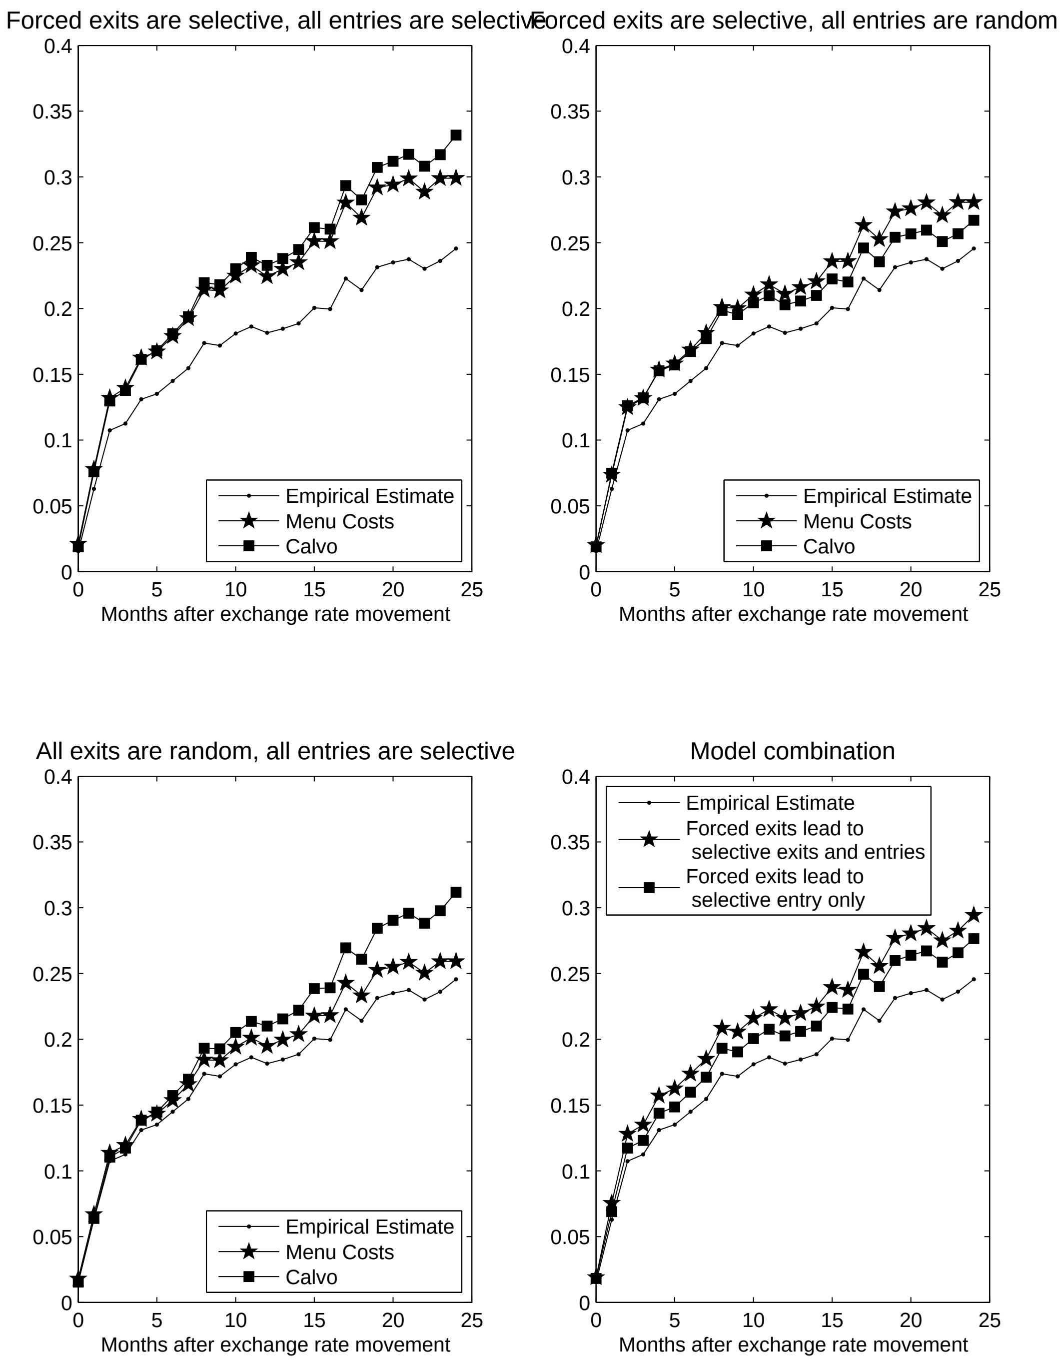

Figure 10: Impact of delaying entries on cumulative contribution of coefficients on lagged exchange rate variables under severe product replacement bias ($n = 1$, $s = 0.05$)

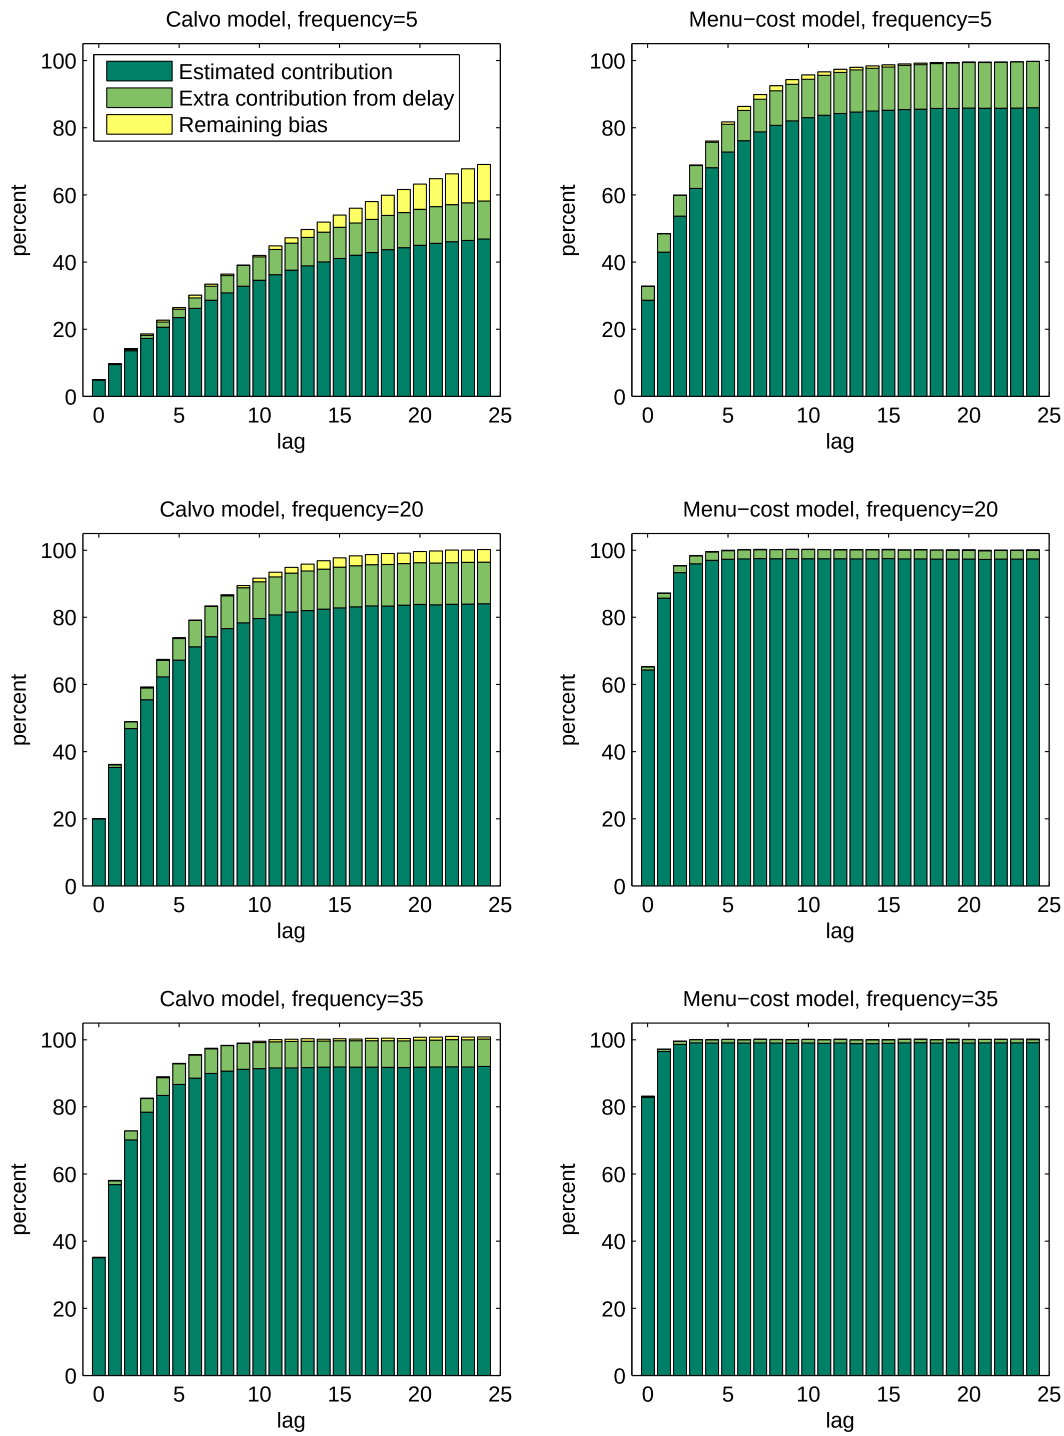

Figure 11: Impact of delaying entries on cumulative contribution of coefficients on lagged exchange rate variables under selective exits and random entries ($n = 0$, $e = 0.25$)

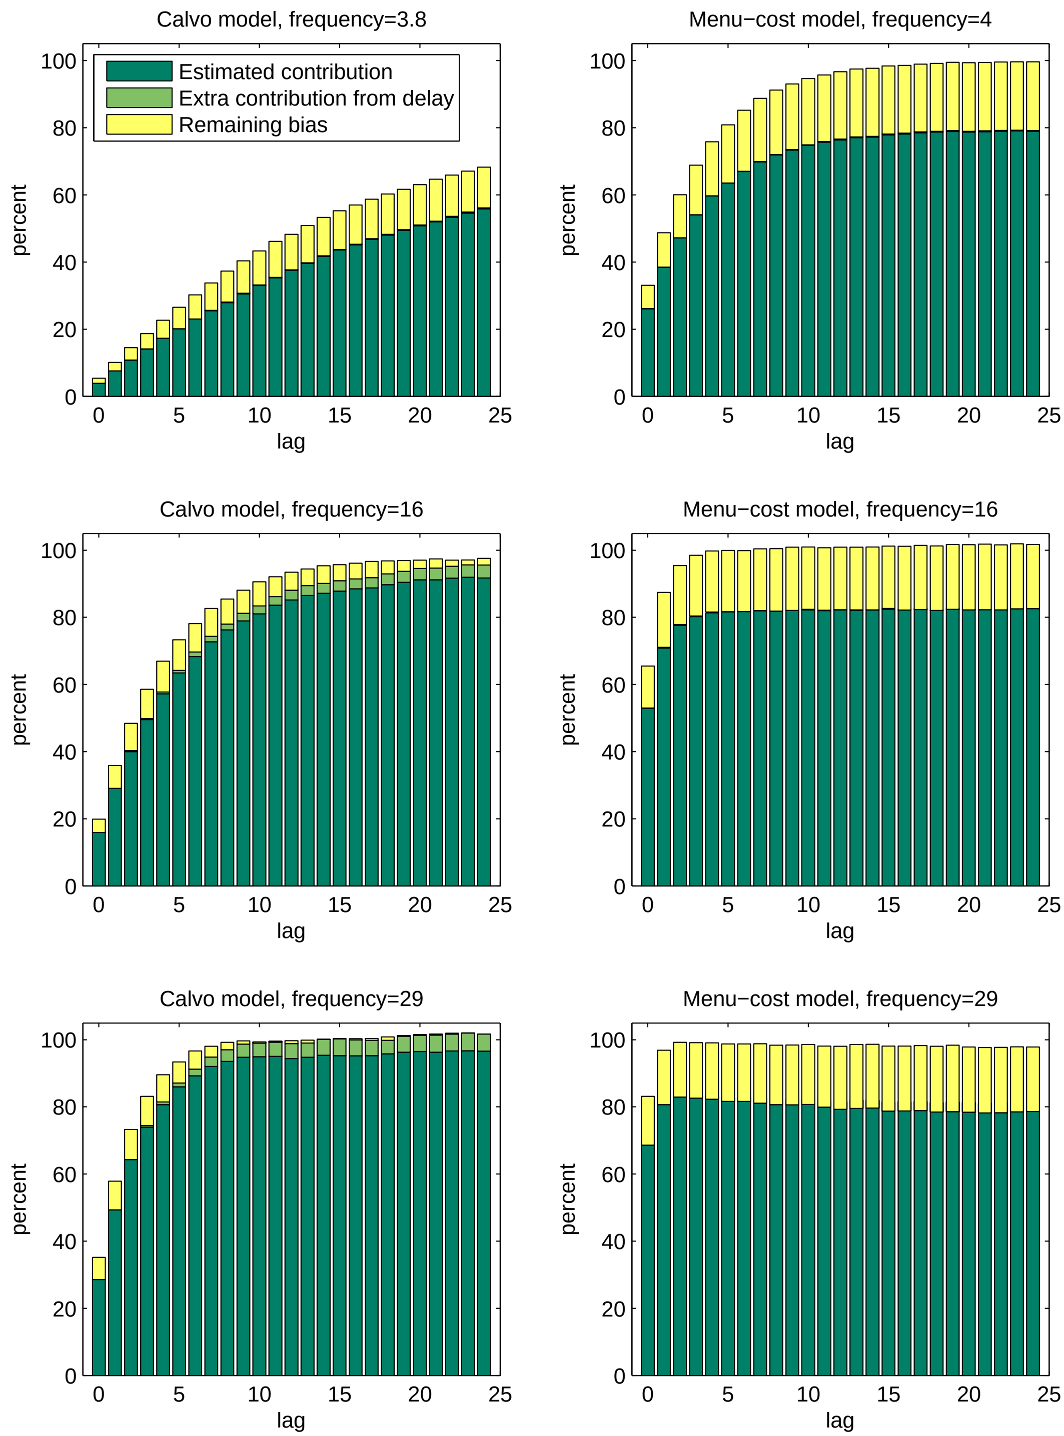

Figure 12: Marginal probabilities of observations usable to computed inflation in period t in the Calvo model with *iid* exchange rate innovations

1. General Case

	Period						
	$t-l$	...	$t-k-1$	$t-k$	...	$t-1$	t
No exit; No price change	$(1-d)(1-f)$	...	$(1-d)(1-f)$	$(1-d)(1-f)$	...	$(1-d)(1-f)$	$\frac{1-f}{1-fe}$
No exit; Price change	$f(1-d)(1-e)$	...	$f(1-d)(1-e)$	$f(1-d)(1-e)$	...	$f(1-d)(1-e)$	$f\frac{1-e}{1-fe}$
Exit; No price change since *t-l-1*	$(d+fe(1-d))*(1-n)(1-f)$	...	$(d+fe(1-d))*(1-n)(1-f)^{l-k}$	$(d+fe(1-d))*(1-n)(1-f)^{l-k+1}$	...	$(d+fe(1-d))*(1-n)(1-f)^{l}$	$n.a.$
No exit; Price change since *t-l-1*	$(d+fe(1-d))*(n+(1-n)f)$	...	$(d+fe(1-d))*(1-(1-n)(1-f)^{l-k})$	$(d+fe(1-d))*(1-(1-n)(1-f)^{l-k+1})$	...	$(d+fe(1-d))*(1-(1-n)(1-f)^{l})$	$n.a.$

2. M-period entry delay, $M \geq l$

	Period						
	$t-l$	...	$t-k-1$	$t-k$	...	$t-1$	t
No exit; No price change	$\frac{1-f}{1-fe}$	...	$\frac{1-f}{1-fe}$	$\frac{1-f}{1-fe}$	...	$\frac{1-f}{1-fe}$	$\frac{1-f}{1-fe}$
No exit; Price change	$f(1-d)(1-e)$	...	$f(1-d)(1-e)$	$f\frac{1-e}{1-fe}$	...	$f\frac{1-e}{1-fe}$	$f\frac{1-e}{1-fe}$
Exit; No price change since *t-l-1*	$n.a.$	...	$n.a.$	$n.a.$	...	$n.a.$	$n.a.$
No exit; Price change since *t-l-1*	$n.a.$	...	$n.a.$	$n.a.$	...	$n.a.$	$n.a.$

3. M-period entry delay, $M < l$

	Period						
	$t-l$	...	$t-M-1$	$t-M$	...	$t-1$	t
No exit; No price change	$(1-d)(1-f)$	...	$(1-d)(1-f)$	$\frac{1-f}{1-fe}$	...	$\frac{1-f}{1-fe}$	$\frac{1-f}{1-fe}$
No exit; Price change	$f(1-d)(1-e)$	...	$f(1-d)(1-e)$	$f\frac{1-e}{1-fe}$	...	$f\frac{1-e}{1-fe}$	$f\frac{1-e}{1-fe}$
Exit; No price change since *t-l-1*	$(d+fe(1-d))*(1-n)(1-f)$	...	$(d+fe(1-d))*(1-n)(1-f)^{l-M}$	$n.a.$	...	$n.a.$	$n.a.$
No exit; Price change since *t-l-1*	$(d+fe(1-d))*(n+(1-n)f)$	...	$(d+fe(1-d))*(1-(1-n)(1-f)^{l-M})$	$n.a.$	...	$n.a.$	$n.a.$

Notes: The 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 panel show the case in which observations entering the sample are delayed by 0, $M \geq l$, and $M < l$ period(s), respectively.

Figure 13: Results for Constructed Alternative Price Indexes