Published in the *Quarterly Journal of Economics*, May 2010, 125(2): 675-727
http://dx.doi.org/10.1162/qjec.2010.125.2.675

Frequency of Price Adjustment and Pass-through*

Gita Gopinath
Department of Economics,
Harvard University and NBER

Oleg Itskhoki
Department of Economics,
Harvard University

May 17, 2009

Abstract

We empirically document using U.S. import prices that on average goods with a high frequency of price adjustment have a long-run pass-through that is at least twice as high as that of low-frequency adjusters. We show theoretically that this relationship should follow because variable mark-ups that reduce long-run pass-through also reduces the curvature of the profit function when expressed as a function of the cost shocks, making the firm less willing to adjust its price. Lastly, we quantitatively evaluate a dynamic menu-cost model and show that the variable mark-up channel can generate significant variation in frequency, equivalent to 37% of the observed variation in the data. On the other hand the standard workhorse model with constant elasticity of demand and Calvo or state dependent pricing has difficulty matching the facts.

*We wish to thank the international price program of the Bureau of Labor Statistics for access to unpublished micro data. We owe a huge debt of gratitude to our project coordinator Rozi Ulics for her invaluable help on this project. The views expressed here do not necessarily reflect the views of the BLS. We are grateful to Robert Barro, Elhanan Helpman and three anonymous referees for detailed comments. We thank participants at several venues for comments and Loukas Karabarbounis for excellent research assistance. This research is supported by NSF grant # SES 0617256.

1 Introduction

There is a current surge in research that investigates the behavior of prices using micro data with the goal of comprehending key aggregate phenomena such as the gradual adjustment of prices to shocks. A common finding across these studies is that there is large heterogeneity in the frequency of price adjustment even within detailed categories of goods. However, there is little evidence that this heterogeneity is meaningfully correlated with other measurable statistics in the data.[1] This makes it difficult to discern what the frequency measure implies for the transmission of shocks and which models of price setting best fit the data, all of which are important for understanding the effects of monetary and exchange rate policy.

In this paper we exploit the open economy environment to shed light on these questions. The advantage of the international data over the closed-economy data is that it provides a well-identified and sizeable cost shock, namely the exchange rate shock. We find that there is indeed a systematic relation between the frequency of price adjustment and long-run exchange rate pass-through. First, we document empirically that on average high-frequency adjusters have a long-run pass-through that is significantly higher than low-frequency adjusters. Next, we show theoretically that long-run pass-through is determined by primitives that shape the curvature of the profit function, primitives that also affect frequency and theory predicts a positive relation between the two in an environment with variable markups. Lastly, we calibrate a dynamic menu-cost model and show that the variable mark-up channel can generate significant variation in frequency, equivalent to 37% of the observed variation in the data. The standard workhorse model with constant elasticity of demand and Calvo or state dependent pricing generates long-run pass-through that is uncorrelated with frequency, contrary to the data.

We document the relation between frequency and long-run pass-through using micro-data on U.S. import prices at the dock.[2] *Long-run* pass-through is a measure of pass-through that does not compound the effects of nominal rigidity. We divide goods imported into the U.S. into frequency bins and use two non-structural approaches to estimate long-run exchange rate pass-through within each bin. One, we regress the cumulative change in the price of the good over its life in the sample, referred to as its life-long price change, on the exchange rate

[1] It is clearly the case that raw/homogenous goods display a higher frequency of adjustment than differentiated goods as documented in Bils and Klenow (2004) and Gopinath and Rigobon (2008). But outside of this finding, there is little that empirically correlates with frequency. Bils and Klenow (2004) and Kehoe and Midrigan (2007) are recent papers that make this point.

[2] The advantage of using prices at the dock is that they do not compound the effect of local distribution costs that play a crucial role in generating low pass-through into consumer prices.

movement over the same period. Two, we estimate an aggregate pass-through regression and compute the cumulative impulse response of the average monthly change in import prices within each bin to a change in the exchange rate over a 24 month period. Either procedure generates similar results: When goods are divided into two equal-sized frequency bins goods with frequency higher than the median frequency of price adjustment display, on average, long-run pass-through that is at least twice as high as goods with frequency less than the median frequency.

For the sample of firms in the manufacturing sector, high-frequency adjusters have a pass-through of 44% as compared to low-frequency adjusters with a pass-through of 21%. In the sub-sample of importers in the manufacturing sector from high income OECD countries, high-frequency adjusters have a pass-through of 59% compared to 25% for the low-frequency adjusters. This result similarly holds for the sub-sample of differentiated goods based on the Rauch (1999) classification. When we divide goods into frequency deciles so that frequency ranges between 3% and 100% per month, long-run pass-through increases from around 18% to 75% for the sub-sample of imports from high income OECD countries. Therefore, the data is characterized not only by a positive relationship between frequency and long-run pass-through, but also by a wide range of variation for both variables.

Both frequency and long-run pass-through depend on primitives that effect the curvature of the profit function. In Section 3 we show that it is indeed the case that higher long-run pass-through should be associated with a higher frequency of price adjustment. We analyze a static price setting model where long-run pass-through is incomplete and firms pay a menu cost to adjust preset prices in response to cost shocks.[3] We allow for two standard channels of incomplete long-run exchange rate pass-through: (i) variable mark-ups, and (ii) imported intermediate inputs.

A higher mark-up elasticity raises the curvature of the profit function with respect to prices, that is it reduces the region of non-adjustment. However, it also reduces the firms desired price adjustment, so that the firms price is more likely to stay within the bounds of non-adjustment. We show that this second effect dominates, implying that a higher mark-up elasticity both lowers pass-through and frequency. Alternatively, it reduces the curvature

[3]Our price setting model is closest in spirit to Ball and Mankiw (1994), while the analysis on the determinants of frequency relates closely to the exercise in Romer (1989) who constructs a model with complete pass-through (CES demand) and Calvo price setting with optimization over the Calvo probability of price adjustment. Other theoretical studies of frequency include Barro (1972); Sheshinski and Weiss (1977); Rotemberg and Saloner (1987) and Dotsey, King, and Wolman (1999). Finally, Devereux and Yetman (2008) study the relationship between frequency of price adjustment and short-run exchange rate pass-through in an environment with complete long-run pass-through.

of the profit function when expressed as a *function of the cost shocks*, generating lower frequency.

The positive relationship between frequency and long-run pass-through implies the existence of a selection effect, wherein firms that infrequently adjust prices are typically not as far from their desired price due to their lower desired pass-through of cost shocks. On the other hand, firms that have high desired pass-through drift farther away from their optimal price and, therefore, make more frequent adjustments. This potentially has important implications for the strength of nominal rigidities given the median duration of prices in the economy. It is important to stress that this selection effect is different from a classical selection effect of state-dependent models forcefully shown by Caplin and Spulber (1987), as it will be present in time-dependent models with optimally chosen periods of non-adjustment as in Ball, Mankiw, and Romer (1988).

In Section 4 we quantitatively solve for the industry equilibrium in a dynamic price-setting model. The standard model of sticky prices in the open economy assumes CES demand and Calvo price adjustment.[4] These models predict incomplete pass-through in the short-run when prices are rigid and set in the local currency, but perfect pass-through in the long-run. To fit the data we depart from this standard set-up. First, we allow for endogenous frequency choice via a menu cost model of state-dependent pricing.[5] Second, we allow for variable mark-ups, *à* la Dornbusch (1987) and Krugman (1987), which generates incomplete long-run pass-through. This source of incomplete pass-through has received considerable support in the open economy empirical literature as we discuss in Section 4. We examine how *variation* in the mark-up elasticity across firms effects the frequency of price adjustment.

We present four sets of results. One, variation in mark-up elasticity can indeed generate a strong positive relation between frequency and LRPT and can generate significant variation in frequency, equivalent to 37% of the observed variation in the data. The model generates a standard deviation in frequency across goods of 11%, as compared to 30% in the data. Two, a menu cost model that allows for joint cross-sectional variation in mark-up elasticity and menu costs can quantitatively account for both the positive slope between LRPT and frequency and the close to zero slope between size and frequency in the data. The model generates a slope of 0.55 between frequency and LRPT, while in the data it is 0.56. Similarly,

[4]See the seminal contribution of Obstfeld and Rogoff (1995) and the subsequent literature surveyed in Lane (2001). Recently, Midrigan (2007) analyzes an environment with state-dependent pricing, but assumes constant mark-ups and complete pass-through; Bergin and Feenstra (2001) allow for variable mark-ups in an environment with price stickiness, but they assume exogenous periods of non-adjustment.

[5]We could alternatively model this as a Calvo model where the Calvo parameter is chosen endogenously and this would deliver similar results

the slope coefficient for the relation between frequency and size is -0.05 in the model, close to the data estimate of -0.01. Further, it generates dispersion in frequency equivalent to 60% of the dispersion in the data. In both simulations the model matches the median absolute size of price adjustment of 7%. Third, we show that the non-structural pass-through regressions estimated in the empirical section recover the true underlying LRPT. Fourth, we verify that the observed correlation between frequency and LRPT cannot be explained by standard sticky price models with only exogenous differences in the frequency of price adjustment and no variation in LRPT.

Section 2 presents the empirical evidence. Section 3 presents the static model of frequency and LRPT and Section 4 describes the calibration of a dynamic model and its ability to match the facts. Section 5 concludes. All proofs are relegated to the Appendix.

2 Empirical Evidence

In this section we empirically evaluate the relation between the frequency of price adjustment of a good and the *long-run* response of the price of the good to an exchange rate shock. The latter, referred to as long-run exchange rate pass-through (LRPT) is defined to capture pass-through beyond the period when nominal rigidities in price setting are in effect. In the presence of strategic complementarities in price setting or other forms of real rigidities, this can require multiple rounds of price adjustment.[6] We use two non-structural approaches to estimate LRPT from the data. Our main finding is that goods whose prices adjust more frequently also have a higher exchange rate pass-through in the long-run as compared to low-frequency adjusters. In Section 4.3 we estimate the same regressions on data simulated from conventional sticky price models, and verify that both of these regressions indeed deliver estimates close to the true theoretical LRPT.

2.1 Data and Methodology

We use micro data on the prices of imported goods into the U.S. provided to us by the Bureau of Labor Statistics, for the period 1994-2005. The details regarding this dataset are provided in Gopinath and Rigobon (2008).

We focus on a subset of the data that satisfies the following criteria. First, we restrict

[6] Other sources of sluggish adjustment could include the presence of informational frictions or convex adjustment costs in price setting.

attention to market transactions and exclude intra-firm transactions, as we are interested in price-setting driven mainly by market forces.[7]

Second, we require that a good have at least one price adjustment during its life. This is because the goal of the analysis is to relate the frequency of price adjustment to the flexible price pass-through of the good and this requires observing at least one price change. In this database there are 30% of goods that have a fixed price during their life. For the purpose of our study these goods are not useful and are excluded from the analysis. We revisit this issue at the end of this section when we comment on item substitution.

Third, we restrict attention to dollar-priced imports in the manufacturing sector.[8] The restriction to manufactured goods allows us to focus on price setting behavior where firms have market power and goods are not homogenous. We restrict attention to dollar-priced goods, so as to focus on the question of frequency choice, setting aside the question of currency choice. This restriction does not substantially reduce the sample size since 90% of goods imported are priced in dollars. For the analysis of the relation between currency choice and pass-through see Gopinath, Itskhoki, and Rigobon (2007). The relation between the two papers is discussed in Section 2.4.

For each of the remaining goods we estimate the frequency of price adjustment following the procedure in Gopinath and Rigobon (2008). We then sort goods into high and low frequency bins, depending on whether the good's frequency is greater than or lower than the median frequency, and estimate LRPT within these bins.

The first approach estimates exchange rate pass-through over the life of the good in the BLS sample. Specifically, for each good we measure the cumulative change in the price of the good starting from its first observed new price to its last observed new price in the BLS data. We refer to this as the life-long change in price. We then relate it to the cumulative change in the exchange rate over this period. Specifically, *life-long pass-through*, β_L, is estimated from the following micro-level regression:

$$\Delta p_L^{i,c} = \alpha_c + \beta_L \Delta RER_L^{i,c} + \epsilon^{i,c}. \tag{1}$$

$\Delta p_L^{i,c}$ is equal to the life-long change in the good's log price relative to U.S. inflation, where i indexes the good and c the country. $\Delta RER_L^{i,c}$ refers to the cumulative change in the log

[7] A significant fraction of trade takes place intrafirm and these transactions constitute about 40% of the BLS sample. For empirical evidence on the difference between intra-firm and arms-length transactions, using this dataset, see Gopinath and Rigobon (2008) and Neiman (2007).

[8] That is, goods that have a one digit SIC code of 2 or 3. We exclude any petrol classification codes.

of the bilateral real exchange rate for country c over this same period.[9] The construction of these variables is illustrated in Figure 1. The real exchange rate is calculated using the nominal exchange rate and the consumer price indices in the two countries. An increase in the RER is a real depreciation of the dollar. Finally, α_c is a country fixed effect.

The second approach measures LRPT by estimating a standard aggregate pass-through regression. For each frequency bin, each country c and month t, we calculate the average price change relative to U.S. inflation, Δp_t^c, and the monthly bilateral real exchange rate movement vis-à-vis the dollar for that country, ΔRER_t^c . We then estimate a stacked regression where we regress the average monthly change in prices on monthly lags of the real exchange rate change:

$$\Delta p_t^c = \alpha_c + \sum_{j=0}^{n} \beta_j \Delta RER_{t-j}^c + \epsilon_t^c, \tag{2}$$

where α_c is a country fixed effect and n varies from 1 to 24 months. The *aggregate long-run pass-through* is then defined to be the cumulative sum of the coefficients, $\sum_{j=0}^{n} \beta_j$, at $n = 24$ months.

Before we proceed to describe the results we briefly comment on the two approaches. First, we use the real specification in both regressions to be consistent with the regressions run later on the model generated data in Section 4.3. However, the main empirical results from both the micro and aggregate regressions are insensitive to using a nominal specification, not surprisingly, given that the real and the nominal exchange rate move closely together at the horizons we consider. In our main tables we will report the estimates from both the real and the nominal specifications. In the nominal specification $\Delta p_L^{i,c}$ is the the life-long change in the log of the nominal price and β_L is the coefficient on the cumulative change in the log of the nominal exchange rate.[10] Similarly, we estimate the nominal equivalent of aggregate regression (2) and find very similar results to the real specification. These latter results are not reported for brevity.

Second, a standard assumption in the empirical pass-through literature is that movements in the real or nominal exchange rate are orthogonal to other shocks that effect the firm's pricing decision and are not affected by firm pricing. This assumption is motivated by the empirical finding that exchange rate movements are disconnected from most macro-variables at the frequencies studied in this paper. While this assumption might be more problematic

[9] The index i on the RER is to highlight that the particular real exchange rate change depends on the period when the good i is in the sample.

[10] In this specification we also include a control for the log of the change in the consumer price index for country c over the same duration for which the price change was estimated.

for commodities such as oil or metals and for some commodity-exporting countries such as Canada, it is far less restrictive for most differentiated goods and most developed countries. Moreover, our main analysis is to rank pass-through across frequency bins as opposed to estimating the true pass-through number. For this reason our analysis is less sensitive to concerns about the endogeneity of the real exchange rate.

Third, the life-long approach has an advantage in measuring LRPT in that it ensures that all goods have indeed changed their price. In the case of the second approach it is possible that even after 24 months some goods have yet to change price and consequently pass-through estimates are low. A concern however with the first approach is that since it conditions on a price change, estimates can be biased because while the exchange rate may be orthogonal to other shocks, when the decision to adjust is endogenous conditioning on a price change induces a correlation across shocks. The life-long regression addresses this selection issue by increasing the window of the pass-through regression to include a number of price adjustments that reduces the size of the selection bias. In Section 4.3 we confirm this claim via simulations.

2.2 Life-long Pass-through

In Table 1 and 2 we report the results from estimating the life-long equation (1). Panel A reports the evidence from the real specification and Panel B reports it for the nominal specification. In Table 1 the first price refers to the first observed price for the good and in Table 2 the first price refers to the first new price for the good. In both cases, the last price is the last new price.[11] The main difference between the results in the two tables relates to the number of observations, since there are goods with only one price adjustment during their life. Otherwise, the results are the same.

The first column reports the sub-sample of the analysis. The next six columns report the median frequency (Freq) within the low and high-frequency bins, the point estimate for LRPT (β_L^{Lo} and β_L^{Hi}) and the robust standard error ($s.e.(\beta_L^{Lo})$ and $s.e.(\beta_L^{Hi})$) for the estimate clustered at the level of country interacted with the BLS-defined Primary Strata Lower (PSL) of the good (mostly 2 to 4-digit harmonized codes). The next two columns report the difference in LRPT between high and low-frequency adjusters and the t-statistic associated with this difference. The number of observations, N_{obs}, and R^2 are reported in the last two columns.

[11] For the hypothetical item in Figure 1, Table 1 would use observations in $[0, t_2]$, while Table 2 would use observations only in $[t_1, t_2]$.

The main finding is that high-frequency adjusters have a life-long pass-through that is at least twice as high as low-frequency adjusters. In the low-frequency sub-sample, goods adjust prices on average every 14 months and pass-through only 21% in the long run. At the same time, in the high-frequency sub-sample, goods adjust prices every 3 months and pass-through 44% in the long-run. This is more strongly evident when we restrict attention to the high-income OECD sample: LRPT increases from 25% to 59% as we move from the low to the high-frequency sub-sample.

We also examine the sub-sample of manufactured goods that can be classified to be in the differentiated goods sector, following Rauch's classification.[12] For differentiated goods, moving from the low to high-frequency bin raises LRPT from 20% to 46% for goods from all source countries and from 25% to 59% in the high-income OECD sample. In all cases, the difference in pass-through across frequency bins is strongly statistically significant.

All the results hold for the nominal specification in Panel B. Similarly, the higher pass-through of high-frequency adjusters is evident in Table 2 where the first price is a new price. Since the results are similar for the case where we start with the first price as opposed to the first new price, for the remainder of the analysis we report the results for the former case, as it preserves a larger number of goods in the sample.[13]

As a sensitivity check we also restrict the sample to goods that have at least 3 or more price adjustments during their life. Results for this specification are reported in Table 3. As expected the median frequency of price adjustment is now higher, but the result that long-run pass-through is at least twice as high for the high-frequency bin as compared to the low-frequency bin still holds strongly and significantly.

We also estimate median quantile regressions to limit the effect of outliers and find that the results hold just as strongly. In the case of the all country sample, $\beta_L^{Lo} = 0.19$ and $\beta_L^{Hi} = 0.41$ with the difference having a t-statistic of 14.14. For the high-income OECD sub-

[12] Rauch (1999) classified goods on the basis of whether they were traded on an exchange (organized), had prices listed in trade publications (reference) or were brand name products (differentiated). Each good in our database is mapped to a 10 digit harmonized code. We use the concordance between the 10 digit harmonized code and the SITC2 (Rev 2) codes to classify the goods into the three categories. We were able to classify around 65% of the goods using this classification. Consequently, it must not be interpreted that the difference in the number of observations between all manufactured and the sub-group of differentiated represent non-differentiated goods. In fact, using Rauch's classification only a 100 odd goods are classified as non-differentiated.

[13] Since there can be months during the life of the good when there is no price information, as a sensitivity test, we exclude goods for whom the last new price had a missing price observation in the previous month to allow for the case that the price could have changed in an earlier month but was not reported. This is in addition to keeping only prices that are new prices (as in Table 2). We find that the results hold just as strongly in this case. The median frequency for the high (low) frequency goods is 0.35 (0.08) and the long-run pass-through is 0.61 (0.11) respectively. The t-stat of the difference in LRPT is 5.28.

sample the difference is 0.30 with a t-statistic of 12.97. We also verify that the results are not driven by variable pass-through rates across countries unrelated to frequency, by controlling for differential levels of pass-through across countries. We estimate the difference in the coefficient between high and low-frequency adjusters, within country, to be 27 percentage points with a t-statistic of 5.58. In Table 4 we allow for variation across countries in the difference ($\beta_L^{Hi} - \beta_L^{Lo}$) and again find that the relation between LRPT and frequency holds for goods from the same country/region as reported in Table 4.

Alternative Specifications: We now verify that the documented positive relationship between frequency and pass-through is not an artifact of splitting the items into two bins by frequency. First, we address this non-structurally by increasing the number of frequency bins. Specifically, we estimate the same regression across 10 frequency bins (deciles). The point estimates and 10% robust standard error bands are reported in Figures 2 for all manufactured goods and all manufactured goods from high-income OECD countries respectively. The positive relationship is evident in these graphs. For the high-income OECD sub-sample long-run pass-through increases from around 18% to 75%, as frequency increases from 0.03 to 1.[14] This wide range of pass-through estimates covers almost all of the relevant range of theoretical pass-through which for most specifications lies between 0 and 1. Furthermore, the positive relation between long-run pass-through and frequency is most evident for the higher frequency range, specifically among the goods that adjust every 8 months or more frequently and constitute a half of our sample. This fact assuages concerns that the relation between frequency and pass-through is driven by insufficient number of price adjustments for the very low-frequency goods.

As opposed to increasing the number of frequency bins, our second approach estimates the effect of frequency on long-run pass-through using a more structured specification. We estimate the following regression:[15]

$$\Delta p_L^{i,c} = \alpha_c + \beta_L \Delta RER_L^{i,c} + \delta_L \tilde{f}_{i,c} + \gamma_L \left(\tilde{f}_{i,c} \cdot \Delta RER_L^{i,c} \right) + \epsilon^{i,c}, \tag{3}$$

where $\tilde{f}_{i,c} \equiv f_{i,c} - \bar{f}_{i,c}$ is the demeaned frequency of the good relative to other goods in the sample. Therefore, coefficient β_L captures the average pass-through in the sample, while

[14] For the all country sample the long-run pass-through range is between 14% and 45%.

[15] This specification results from the following two-stage econometric model:

$$\Delta p_L^{i,c} = \alpha_c + \beta_L^{i,c} \Delta RER_L^{i,c} + \delta_L \tilde{f}_{i,c} + v^{i,c}$$
$$\beta_L^{i,c} = \beta_L + \gamma_L \tilde{f}_{i,c} + u^{i,c}.$$

Regression (3) consistently estimates γ_L provided that $u^{i,c}$ and $v^{i,c}$ are independent from $\Delta RER_L^{i,c}$ and $\tilde{f}_{i,c}$.

γ_L estimates the effect of frequency on long-run pass-through. The results from estimating this regression using both OLS and median quantile regressions are reported in Table 5. In the case of the OLS estimates robust standard errors clustered by country and PSL pair are reported. The lower panel presents the results for goods with at least 3 or more price changes. As is evident from the table, $\gamma_L > 0$ in all specifications. That is, goods that adjust prices more frequently also have higher LRPT. The reason the slope estimates vary across samples is partly driven by the fact that the relationship is non-linear as is evident in Figure 2. These results are also robust to including controls for differential pass-through rates across countries.

Between and Within-Sector Evidence: Does the relation between frequency and LRPT arise across aggregate sectors or is this a within-sector phenomenon?

To answer this we first perform a standard *variance decomposition* (see Theorem 3.3 in Greene, 2000, p. 81) for frequency:

$$S_f^T = S_f^B + S_f^W.$$

S_f^T is the total variance of frequency across all goods in the sample. S_f^B is the *between-sector* component of the variance, measured as the variance of frequency across the average goods in every sector. Finally, S_f^W is the *within-sector* component of variance, measured as the average variance of frequency across goods within sectors.

We perform the analysis both at the 2-digit and 4-digit sector level. At the 2-digit sector level (88 sectors), the fraction of total variance in frequency (equal to 0.073) explained by variation across 2-digit sectors is 15%, while the remaining 85% is explained by variation across goods within 2-digit sectors. At the 4-digit level (693 sectors) the between-sector component accounts for 30% of variation in frequency, while within-sector variation accounts for the remaining 70%. This evidence suggests that variation in frequency is driven largely by variation at highly disaggregated levels.

The second exercise we perform is to estimate the counterpart to equation (3) allowing for separate within and between-sector effects of frequency on pass-through.[16] The results are reported in Table 6. The within-sector estimates (γ_L^W) are positive and statistically

[16] Specifically, instead of $\gamma_L\big(\tilde{f}_{i,c} \cdot \Delta RER_L^{i,c}\big)$ we include two terms $\gamma_L^B\big(\bar{f}_{j(i),c} \cdot \Delta RER_L^{i,c}\big)$ and $\gamma_L^W\big(f_{i,c} - \bar{f}_{j(i),c}\big) \cdot \Delta RER_L^{i,c}$, where j indicates the sector which contains good i and $\bar{f}_{j(i),c}$ is the average frequency in sector j. Note that our earlier specification (3) is the restricted version of this regression under the assumption that $\gamma_L^W = \gamma_L^B$. Furthermore, the unconstrained specification allows for a formal decomposition of the effect of frequency on pass-through into within and between-sector contribution as discussed below.

significant in all specifications. The between-sector estimates (γ_L^B) are positive but the level of significance varies across specifications.

We can then quantify the contribution of the within-sector component to the relation between LRPT and frequency using the formula

$$\frac{\left(\gamma_L^W\right)^2 S_f^W}{\left(\gamma_L^W\right)^2 S_f^W + \left(\gamma_L^B\right)^2 S_f^B},$$

where the denominator is the total variance in LRPT explained by variation in frequency. Using the OLS (quantile regression) estimates, the within-sector contribution is 98% (78%) at the 2-digit level and 86% (62%) at the 4-digit level.

Therefore the relation between frequency and LRPT is largely a within-sector phenomenon, consistent with the evidence that most variation in frequency arises within sectors and not across aggregated sectors.[17]

2.3 Aggregate Regressions

The next set of results relates to the estimates from the aggregate pass-through regressions defined in (2). We again divide goods into two bins based on the frequency of price adjustment and estimate the aggregate pass-through regressions separately for each of the bins. We report the results only for the real specification, since the nominal specification delivers very similar results. The results are plotted in Figure 3. The solid line plots the cumulative pass-through coefficient, $\sum_j^n \beta_j$, as the number of monthly lags increases from 1 to 24. The dashed lines represent the 10% robust standard-error bands. The left column figures are for the all country sample and the right column figures are for the high-income OECD sub-sample; the top figures correspond to all manufactured goods, while the bottom figures correspond to the differentiated good sub-sample.

While pass-through at 24 months is lower than life-long estimates, it is still the case that high-frequency adjusters have a pass-through that is at least twice as high as low-frequency adjusters and this difference is typically significant. The results from this approach are therefore very much in line with the results from the life-long specification. In Figures 4 and

[17] This is not to say that there is no variation in frequency and pass-through across sectors. More homogenous sectors, such as 'Animal and Vegetable Products', 'Wood and Articles of Wood' and 'Base Metals and Articles of Base Metals', on average have higher frequency and higher long-run pass-through. More differentiated sectors have lower average frequency (with little variation across sectors) and lower long-run pass-through. However, the amount of variation across sectors is insufficient to establish a strong empirical relationship.

5 we report the results by country/region and for goods with 3 or more price adjustments. Here again we find similar results. The estimates in these sub-samples, however, become very noisy.

2.4 Additional Facts

In closing the empirical section, we discuss a number of additional relevant findings in the data:

Product Replacement: For the previous analysis we estimate LRPT for a good using price changes during the life of the good. Since goods get replaced frequently one concern could be the fact that goods that adjust infrequently have shorter lives and get replaced often and because we do not observe price adjustments associated with substitutions we might underestimate the true pass-through for these goods.[18] To address this concern we report in Table 7 the median life of goods within each frequency bin for the high-income OECD sample; very similar results are obtained for other sub-samples.

For each of the 10 frequency bins we estimate 2 measures of the life of the good. For the first measure we calculate for each good the difference between the discontinuation date and initiation date to capture the life of the good in the sample. 'Life 1' then reports the median of this measure for each bin. Goods get discontinued for several reasons. Most goods get replaced during routine sampling and some get discontinued due to lack of reporting. As a second measure we examine only those goods that were replaced either because the firm reported that the particular good was not being traded anymore and had/had not been replaced with another good in the same category or because the firm reports that it is going out of business.[19] This captures most closely the kind of churning one might be interested in and does not suffer from right censoring in measuring the life of the good. 'Life 2' is then the median of this measure within each bin. As can be seen, if anything, there is a negative relation between frequency and life: that is, goods that adjust infrequently have longer lives in the sample.

In the last two columns we report $[\text{Freq} + (1 - \text{Freq})/\text{Life}]$ for the two measures of 'Life' respectively. This corrects the frequency of price adjustment to include the probability of

[18]Note that substitutions pose a bigger concern only if there is reason to believe that pass-through associated with substitutions is different from that associated with price changes. Otherwise, our measures that condition on multiple rounds of price adjustment capture LRPT.

[19]Specifically this refers to the following discontinuation reasons reported in the BLS data: "Out of Business", "Out of Scope, Not replaced" and "Out of Scope, Replaced".

discontinuation. As is evident, the frequency ranking does not change when we include the probability of being discontinued using either measure. As mentioned earlier there are several goods that do not change price during their life and get discontinued. We cannot estimate pass-through for these goods. The median life of these goods is 20 months (using the second measure), which implies a frequency of 0.05. What this section highlights is that even allowing for the probability of substitution the benchmark frequency ranking is preserved.

Size of Price Adjustment: Figure 6 plots the median size of price adjustment by 10 frequency bins for the high-income OECD sub-sample. Median size is effectively the same across frequency bins, ranging between 6% and 7%.[20] This feature is not surprising given that size, unlike pass-through, is not scale independent and, for example, depends on the average size of the shocks. This illustrates the difficulty of using measures such as size in the analysis of frequency. We discuss this issue later in the paper.

Long-run versus Medium-run Pass-through: In this paper we estimate the long-run pass-through for a good. A separate measure of pass-through is pass-through conditional on only the first price adjustment to an exchange rate shock. In Gopinath, Itskhoki, and Rigobon (2007) we refer to this as medium-run pass-through (MRPT). As is well known, estimating pass-through conditional on only the first price adjustment may not be sufficient to capture LRPT due to staggered price adjustment by competitors among other reasons. These effects can be especially pronounced for goods that adjust prices more frequently relative to their average competitor.

In Gopinath and Rigobon (2008) we sort goods into different frequency bins and estimate MRPT within each bin, which is distinct from estimating LRPT. Secondly, we used both dollar (90% of the sample) and non-dollar (10% of the sample) priced goods. We document that goods that adjust less frequently have higher MRPT than goods that adjust more frequently. This result, relating to MRPT, was driven by the fact that goods that adjust less frequently were goods that were priced in a non-dollar currency. If the non-dollar goods are excluded from the sample there is no well-defined pattern in the relation between MRPT and frequency. This is further demonstrated in Figure 7 where we plot both LRPT and MRPT against frequency. Unlike LRPT, there is no relation between MRPT and frequency

[20] We also plot in this figure the 25% and 75% quantiles of the size of price adjustment distribution. Just as for median size, we find no pattern for the 25-th quantile, which is roughly stable at 4% across the 10 frequency bins. On opposite, 75-th quantile decreases from 15% to 10% as we move from low-frequency to high frequency bins.

for dollar priced goods. We also estimate equation (3) for the case where the left hand side variable conditions on first price adjustment instead of the life-long price change. The coefficient that estimates the effect of frequency on MRPT is -0.04 with a t-statistic of -0.9, confirming the result in Figure 7 that MRPT is unrelated to frequency in the dollar sample.

In Gopinath, Itskhoki, and Rigobon (2007) we present further systematic evidence on the relation between the currency in which goods are priced and MRPT. We argue theoretically that one should expect to find that goods priced in non-dollars indeed have a higher MRPT. In addition they will have longer price durations, conditioning on the same LRPT.

To clarify again, the measure of pass-through we estimate in this paper is a different concept from the main pass-through measures reported in Gopinath and Rigobon (2008) and Gopinath, Itskhoki, and Rigobon (2007). The evidence we find about the relation between frequency and pass-through relates to the long-run pass-through for dollar priced goods. As we argue below theoretically, the relevant concept for relating frequency to the structural features of the profit function is indeed long-run pass-through and that is why it is the focus of the current paper.

3 A Static Model of Frequency and Pass-through

In this section we investigate theoretically the relation between LRPT and frequency. Before constructing in the next section a fully-fledged dynamic model of staggered price adjustment we use a simple static model to illustrate the theoretical relationship between frequency of price adjustment and flexible price pass-through of cost shocks. The latter is the equivalent of LRPT in a dynamic environment and we will refer to it simply as pass-through. We show that, all else equal, higher pass-through is associated with a higher frequency of price adjustment. This follows because the primitives that reduce pass-through also reduce the curvature of the profit function in the space of the *cost shock*, making the firm less willing to adjust its price.

We consider the problem of a single monopolistic firm that sets its price before observing the cost shock.[21] Upon observing the cost shock the firm has an option to pay a menu cost to reset its price. The frequency of adjustment is then the probability with which the firm decides to reset its price upon observing the cost shock. We introduce two standard sources of incomplete pass-through into the model: variable mark-ups and imported inputs.

[21] Our modeling approach in this section is closest to Ball and Mankiw (1994), while the motivation of the exercise is closest to Romer (1989). References to other related papers can be found in the introduction.

3.1 Demand and Costs

Consider a single price setting firm that faces a residual demand schedule $Q = \varphi(P|\sigma, \varepsilon)$, where P is its price and $\sigma > 1$ and $\varepsilon \geq 0$ are two demand parameters.[22] We denote the price elasticity of demand by

$$\tilde{\sigma} \equiv \tilde{\sigma}(P|\sigma, \varepsilon) = -\frac{\partial \ln \varphi(P|\sigma, \varepsilon)}{\partial \ln P}$$

and the *super-elasticity* of demand (in the terminology of Klenow and Willis, 2006), or the elasticity of elasticity, as

$$\tilde{\varepsilon} \equiv \tilde{\varepsilon}(P|\sigma, \varepsilon) = \frac{\partial \ln \tilde{\sigma}(P|\sigma, \varepsilon)}{\partial \ln P}.$$

Here $\tilde{\sigma}$ is the *effective* elasticity of demand for the firm that takes into account both direct and indirect effects from price adjustment.[23] Note that we introduce variable mark-ups into the model by means of variable elasticity of demand. This should be viewed as a reduced form specification for variable mark-ups that would arise in a richer model due to strategic interactions between firms.[24]

We impose the following normalization on the demand parameters: When the price of the firm is unity ($P = 1$), elasticity and super-elasticity of demand are given by σ and ε respectively (that is, $\tilde{\sigma}(1|\sigma, \varepsilon) = \sigma$ and $\tilde{\varepsilon}(1|\sigma, \varepsilon) = \varepsilon$). Moreover, $\tilde{\sigma}(\cdot)$ is increasing in σ and $\tilde{\varepsilon}(\cdot)$ is increasing in ε for any P. Additionally, we normalize the level of demand $\varphi(1|\sigma, \varepsilon)$ to equal 1 independently of the demand parameters σ and ε (see Section 4 for an example of such a demand schedule). These normalizations prove to be useful later when we approximate the solution around $P = 1$.

The firm operates a production technology characterized by a constant marginal cost:

$$MC \equiv MC(a, e; \phi) = (1 - a)(1 + \phi e)c,$$

where a is an idiosyncratic productivity shock and e is a real exchange rate shock. We will refer to the pair (a, e) as the cost shock to the firm. We further assume that a and e are independently distributed with $\mathbb{E}a = \mathbb{E}e = 0$ and standard deviations denoted by σ_a and σ_e respectively. Parameter $\phi \in [0, 1]$ determines the sensitivity of the marginal cost to the

[22] Since this is a partial equilibrium model of the firm, we do not explicitly list the prices of competitors or the sectoral price index in the demand functions. An alternative interpretation is that P stands for the relative price of the firm.

[23] For example, in a model with large firms, price adjustment by the firm will also affect the sectoral price index which may in turn indirectly affect the elasticity of demand.

[24] The Atkeson and Burstein (2007) model is an example: in this model the effective elasticity of residual demand for each monopolistic competitor depends on the primitive constant elasticity of demand, the market share of the firm and the details of competition between firms.

exchange rate shock and can be less than 1 due to the presence of imported intermediate inputs in the cost function of firms.

We normalize the marginal cost so that $MC = c = (\sigma - 1)/\sigma$ when there is no cost shock ($a = e = 0$). Under this normalization, the optimal flexible price of the firm when $a = e = 0$ is equal to 1, since the marginal cost is equal to the inverse of the mark-up. This normalization is therefore consistent with a symmetric general equilibrium in which all firms relative prices are set to 1 (for a discussion see Rotemberg and Woodford, 1999).

Finally, the profit function of the firm is given by:

$$\Pi(P|a,e) = \varphi(P)\big(P - MC(a,e)\big), \tag{4}$$

where we suppress the explicit dependence on parameters σ, ε and ϕ. We denote the *desired price* of the firm by $P(a,e) \equiv \arg\max_P \Pi(P|a,e)$ and the maximal profit by $\Pi(a,e) \equiv \Pi\big(P(a,e)|a,e\big)$.

3.2 Price Setting

For a given cost shock (a,e), the desired flexible price maximizes profits (4) so that[25]

$$P_1 \equiv P(a,e) = \frac{\tilde{\sigma}(P_1)}{\tilde{\sigma}(P_1) - 1}(1-a)(1+\phi e)c, \tag{5}$$

and the corresponding maximized profit is $\Pi(a,e)$. Denote by $\bar{P}_0$ the price that the firm sets prior to observing the cost shocks (a,e). If the firm chooses not to adjust its price, it will earn $\Pi(\bar{P}_0|a,e)$. The firm will decide to reset the price if the profit loss from non-adjusting exceeds the menu cost, κ:

$$\mathrm{L}(a,e) \equiv \Pi(a,e) - \Pi(\bar{P}_0|a,e) > \kappa.$$

Define a set of shocks upon observing which the firm decides not to adjust its price by

$$\Delta \equiv \{(a,e):\ \mathrm{L}(a,e) \le \kappa\}.$$

Note that the profit-loss function $\mathrm{L}(a,e)$ and, hence, Δ depend on the preset price $\bar{P}_0$.

The firm sets its initial price, $\bar{P}_0$, to maximize expected profits where the expectation is taken conditional on the realization of the cost shocks (a,e) upon observing which the firm

[25] The sufficient condition for maximization is $\tilde{\sigma}(P_1) > 1$ provided that $\tilde{\varepsilon}(P_1) \ge 0$. We assume that these inequalities are satisfied for all P.

does not reset its price:[26]

$$\bar{P}_0 = \arg\max_P \int_{(a,e)\in\Delta} \Pi(P|a,e)\mathrm{d}F(a,e),$$

where $F(\cdot)$ denotes the joint cumulative distribution function of the cost shock (a,e). Using the linearity of the profit function in costs, we can re-write the *ex ante* problem of the firm as

$$\bar{P}_0 = \arg\max_P \Big\{\varphi(P)\big(P - \mathbb{E}_\Delta\{(1-a)(1+\phi e)\}\cdot c\big)\Big\}, \tag{6}$$

where $\mathbb{E}_\Delta\{\cdot\}$ denotes the expectation condition on $(a,e)\in\Delta$. We prove the following:

Lemma 1 $\bar{P}_0 \approx P(0,0) = 1$, *up to second order terms.*

Proof: See the working paper version, Gopinath and Itskhoki (2008). ■

Intuitively, a firm sets its ex ante price as if it anticipates the cost shock to be zero ($a = e = 0$), i.e. equal to its unconditional expected value. This will be an approximately correct expectation of the shocks (a,e) over the region Δ, if this region is nearly symmetric around zero and the cost shocks have a symmetric distribution, as we assume. The optimality condition (5) implies that, given our normalization of the marginal cost and elasticity of demand, $P(0,0) = 1$.

3.3 Pass-through

Using Lemma 1 we can prove:

Proposition 1 *(i) The following first order approximation holds,*

$$\frac{P(a,e)-\bar{P}_0}{\bar{P}_0} \approx \Psi\cdot\big(-a+\phi e\big), \qquad \text{where} \qquad \Psi \equiv \frac{1}{1+\frac{\varepsilon}{\sigma-1}}. \tag{7}$$

(ii) Exchange rate pass-through equals

$$\Psi_e = \phi\Psi = \frac{\phi}{1+\frac{\varepsilon}{\sigma-1}}. \tag{8}$$

[26]We implicitly assume, as is standard in a partial equilibrium approach, that the stochastic discount factor is constant for the firm.

Lemma 1 allows us to substitute $\bar{P}_0$ with $P(0,0) = 1$. Then, a and ϕe constitute proportional shocks to the marginal cost and the desired price of the firm responds to them with elasticity Ψ. This pass-through elasticity can be smaller than one because mark-ups adjust to limit the response of the price to the shock. The *mark-up elasticity* is given by

$$\frac{\partial \tilde{\mu}(P)}{\partial \ln P}\bigg|_{P=1} = -\frac{\tilde{\varepsilon}(P)}{\tilde{\sigma}(P)-1}\bigg|_{P=1} = -\frac{\varepsilon}{\sigma-1},$$

where $\tilde{\mu}(P) \equiv \ln\left[\tilde{\sigma}(P)/(\tilde{\sigma}(P)-1)\right]$ is the log mark-up. A higher price increases the elasticity of demand, which in turn, leads to a lower optimal mark-up.

Mark-up elasticity depends on both the super-elasticity and elasticity of demand: it is increasing in the super-elasticity of demand ε and decreasing in the elasticity of demand σ provided that $\varepsilon > 0$. Exchange rate pass-through, Ψ_e, is the elasticity of the desired price of the firm with respect to the exchange rate shock. It is increasing in cost sensitivity to the exchange rate, ϕ, and decreasing in the mark-up elasticity, $\varepsilon/(\sigma-1)$.[27]

3.4 Frequency

In this static framework, we interpret the probability of resetting price in response to a cost shock (a,e) as the frequency of price adjustment. Formally, frequency is defined as

$$\Phi \equiv 1 - \Pr\{(a,e) \in \Delta\} = \Pr\{\mathrm{L}(a,e) > \kappa\}, \tag{9}$$

where the probability is taken over the distribution of the cost shock (a,e).

To characterize the region of non-adjustment, Δ, we use:

Lemma 2 *The following second order approximation holds:*

$$\mathrm{L}(a,e) \equiv \Pi(a,e) - \Pi(\bar{P}_0|a,e) \approx \frac{1}{2}\frac{\sigma-1}{\Psi}\left(\frac{P(a,e)-\bar{P}_0}{\bar{P}_0}\right)^2,$$

where Ψ is again as defined in (7)*.*

Note that Lemma 2 implies that the curvature of the profit function with respect to prices is proportional to

$$\frac{\sigma-1}{\Psi} = (\sigma-1)\left[1+\frac{\varepsilon}{\sigma-1}\right],$$

[27] In the working paper version (Gopinath and Itskhoki, 2008) we also allowed for variable marginal costs as an additional channel of incomplete pass-through. In this case, the effect of σ on Ψ can be non-monotonic. While greater elasticity of demand limits the variable mark-up channel, it amplifies the variable marginal cost channel.

and increases in both σ and ε. That is, higher elasticity of demand and higher mark-up elasticity increases the curvature of the profit function. Holding pass-through (i.e., the response of desired price to shocks) constant this should lead to more frequent price adjustment. However, greater mark-up elasticity also limits desired pass-through which, as we show below, more than offsets the first effect.

This is seen when we combine the results of Proposition 1 and Lemma 2 and arrive at the final approximation to the profit loss function:

$$\mathrm{L}(a,e) \approx \frac{1}{2}(\sigma-1)\Psi\big(-a+\phi e\big)^2, \tag{10}$$

which again holds up to third order terms. This expression makes it clear that forces that reduce pass-through (i.e., decrease Ψ and ϕ) also reduce the profit loss from not adjusting prices and, as a result, lead to lower frequency of price adjustment. Note that in the space of the cost shock, the curvature of the profit loss function decreases as pass-through elasticity Ψ decreases.

Alternatively, primitives that lower Ψ reduce the region of non-adjustment in the price space (Lemma 2). However, a lower Ψ implies that the desired price adjusts by less and therefore is more likely to remain within the bounds of non-adjustment thus reducing the frequency of price adjustment. This second effect always dominates (equation 10).

Combining (9) and (10) we have

$$\Phi \approx \Pr\left\{|X| > \sqrt{\frac{2\kappa}{(\sigma-1)\Psi\Sigma}}\right\}, \tag{11}$$

where $X \equiv \Sigma^{-1/2}\cdot(-a+\phi e)$ is a standardized random variable with zero mean and unit variance and $\Sigma \equiv \sigma_a^2 + \phi^2\sigma_e^2$ is the variance of the cost shock $(-a+\phi e)$. This leads us to the following:

Proposition 2 *The frequency of price adjustment decreases with mark-up elasticity and increases with the sensitivity of costs to exchange rate shocks. It also decreases with the menu cost and increases with the elasticity of demand and the size of shocks.*

Taken together the results on pass-through and frequency in Proposition 1 and 2 imply that:

Proposition 3 *(i) Higher mark-up elasticity as well as lower sensitivity of cost to exchange rate shocks reduce both frequency of price adjustment and pass-through; (ii) Higher menu costs and smaller cost shocks decrease frequency, but have no effect on pass-through.*

Proposition 3 is the central result of this section. It implies that as long as mark-up elasticity varies across goods, we should observe a positive cross-sectional correlation between frequency and pass-through. Similarly, variation across goods in cost sensitivity to exchange rate, ϕ, can also account for the positive relationship between frequency and pass-through.[28] Furthermore, other sources of variation in frequency do not affect pass-through and hence cannot account for the observed empirical relationship between the two variables.

As for the average absolute size of price adjustment, conditional on adjusting price, the effect of mark-up elasticity and the volatility of cost shocks works through two channels. The direct effect of lower mark-up elasticity or more volatile shocks is to increase the change in the desired price, while their indirect effect is to increase the frequency of price adjustment. The first effect increases the average size of price adjustment while the second reduces it. In Gopinath and Itskhoki (2008) we discuss conditions under which the direct effect dominates, such that lower mark-up elasticity ($\varepsilon/(\sigma-1)$) or larger cost shocks (Σ) are associated with a larger absolute size of price adjustment. On the other hand, the average absolute size of price adjustment increases with the size of the menu cost (κ) as frequency of price adjustment decreases. Consequently, as long as there is variation across goods in both κ and ε or Σ, one should not expect to see a robust correlation between frequency and size.

4 Dynamic Model

We now consider a fully dynamic specification with state dependent pricing and variable mark-ups and quantitatively solve for the industry equilibrium in the U.S. market. First, we show that cross-sectional *variation in mark-up elasticity* can generate a positive relation between frequency and LRPT in a dynamic setting and can generate significant variation in frequency, equivalent to 37% of the observed variation in the data. Further, a menu cost model that allows for joint cross-sectional variation in mark-up elasticity and menu costs can quantitatively account for both the positive slope between LRPT and frequency and the close to zero slope between size and frequency in the data. Second, we show that the pass-through regressions estimated in Section 2 recover the true underlying LRPT. Third, we verify that the observed correlation between frequency and LRPT cannot be explained by standard sticky price models with only exogenous differences in frequency of price adjustment and no variation in LRPT.

[28] Quantitatively, however, the effect of ϕ on frequency is limited by the ratio of the variances of the exchange rate and idiosyncratic shocks, σ_e^2/σ_a^2, since ϕ affects frequency through $\Sigma = \sigma_a^2(1+\phi^2\sigma_e^2/\sigma_a^2)$. We calibrate this ratio in Section 4 and show that the effect of ϕ on frequency is negligible.

The importance of the variable mark-up channel of incomplete pass-through, as argued for theoretically by Dornbusch (1987) and Krugman (1987), has been documented in the empirical evidence of Knetter (1989), Goldberg and Knetter (1997), Fitzgerald and Haller (2008), Burstein-Jaimovich among others.[29]

We model the variable mark-up channel of incomplete pass-through using the Kimball (1995) kinked demand. Our setup is most comparable to Klenow and Willis (2006) with two distinctions. First, we have exchange rate shocks that are more idiosyncratic than the aggregate monetary shocks typically considered in the closed economy literature. Second, we examine how *variation* in mark-up elasticity across goods affects the frequency of price adjustment.[30]

4.1 Setup of the Model

In this subsection we lay out the ingredients of the dynamic model. Specifically, we describe demand, the problem of the firm and the sectoral equilibrium.

4.1.1 Industry Demand Aggregator

The industry is characterized by a continuum of varieties indexed by j. There is a unit-measure of U.S. varieties and a measure $\omega < 1$ of foreign varieties available for domestic consumption. The smaller fraction of foreign varieties captures the fact that not all varieties of the differentiated good are internationally traded in equilibrium.

The technology of transforming the intermediate varieties into the final good is characterized by the Kimball (1995) aggregator:

$$\frac{1}{|\Omega|}\int_{\Omega}\Upsilon\left(\frac{|\Omega|C_j}{C}\right)\mathrm{d}j = 1 \tag{12}$$

[29] Unlike the evidence for import prices, Bils, Klenow, and Malin (2009) find that variable mark-ups play a limited role in consumer/retail price data. These two sets of findings can be potentially reconciled by the evidence in Goldberg and Hellerstein (2007), Burstein and Jaimovich (2009) and Gopinath, Gourinchas, Hsieh, and Li (2009) who find that variable mark-ups play an important role at the whole-sale cost level and a limited role at the level of retail prices. Atkeson and Burstein (2007) also argue for the importance of variable mark-ups in matching empirical features of whole-sale traded goods prices.

[30] Klenow and Willis (2006) argue that the levels of mark-up elasticity required to generate sufficient aggregate monetary non-neutrality generates price and quantity behavior that is inconsistent with the micro data on retail prices. We differ from this analysis in the following respects. One, we match facts on import prices, which has different characteristics such as the median frequency of price adjustment. Two, we calibrate the mark-up elasticity to match the evidence on the micro-level relationship between frequency and pass-through for whole-sale prices and as discussed in footnote 36 this implies a lower mark-up elasticity for the median good and the model-generated data is consistent with the micro facts.

with $\Upsilon(1) = 1$, $\Upsilon'(\cdot) > 0$ and $\Upsilon''(\cdot) < 0$. C_j is the quantity demanded of the differentiated variety $j \in \Omega$, where Ω is the set of available varieties in the home country with measure $|\Omega| = 1 + \omega$. Individual varieties are aggregated into sectoral final good demand, C, which is implicitly defined by (12).

The associated demand schedules with aggregator (12) are given by:

$$C_j = \psi\left(D\frac{P_j}{P}\right) \cdot \frac{C}{|\Omega|}, \qquad \text{where} \qquad \psi(\cdot) \equiv \Upsilon'^{-1}(\cdot), \tag{13}$$

P_j is the price of variety j, P is the sectoral price index and $D \equiv \int_\Omega \Upsilon'\left(\frac{|\Omega|C_j}{C}\right)\frac{C_j}{C}dj$. The sectoral price index satisfies

$$PC = \int_\Omega P_j C_j \mathrm{d}j \tag{14}$$

since the aggregator in (12) is homothetic.

4.1.2 Firm's Problem

Consider a home firm producing variety j. Everything holds symmetrically for foreign firms and we superscript foreign variables with an asterisk. The firm faces a constant marginal cost:

$$MC_{jt} = \frac{W_t^{1-\phi}(W_t^*)^\phi}{A_{jt}}. \tag{15}$$

A_j denotes the idiosyncratic productivity shock that follows an autoregressive process in logs:[31]

$$a_{jt} = \rho_a a_{j,t-1} + \sigma_a u_{jt}, \qquad u_{jt} \sim iid\, \mathcal{N}(0,1).$$

W_t and W_t^* denote the prices of domestic and foreign inputs respectively and we will interpret them as wage rates. Parameter ϕ measures the share of foreign inputs in the cost of production.[32]

The profit function of a firm from sales of variety j in the domestic market in period t is:

$$\Pi_{jt}(P_{jt}) = \left[P_{jt} - \frac{W_t^{1-\phi}(W_t^*)^\phi}{A_{jt}}\right]C_{jt},$$

where demand C_{jt} satisfies (13). Firms are price setters and must satisfy demand at the posted price. To change the price, both domestic and foreign firms must pay a menu cost κ_j.

[31] In what follows corresponding small letters denote the logs of the respective variables.

[32] The marginal cost in (15) can be derived from a constant returns to scale production function which combines domestic and foreign inputs.

Define the state vector of firm j by $\mathbb{S}_{jt} = (P_{j,t-1}, A_{jt}; P_t, W_t, W_t^*)$. It contains the past price of the firm, the current idiosyncratic productivity shock and the aggregate state variables, namely, sectoral price level and domestic and foreign wages. The system of Bellman equations for the firm is given by:[33]

$$\begin{cases} V^N(\mathbb{S}_{jt}) &= \Pi_{jt}(P_{j,t-1}) + \mathbb{E}\big\{Q(\mathbb{S}_{j,t+1})V(\mathbb{S}_{j,t+1})\big|\mathbb{S}_{jt}\big\}, \\ V^A(\mathbb{S}_{jt}) &= \max_{P_{jt}} \Big\{\Pi_{jt}(P_{jt}) + \mathbb{E}\big\{Q(\mathbb{S}_{j,t+1})V(\mathbb{S}_{j,t+1})\big|\mathbb{S}_{jt}\big\}\Big\}, \\ V(\mathbb{S}_{jt}) &= \max\big\{V^N(\mathbb{S}_{jt}), V^A(\mathbb{S}_{jt}) - \kappa_j\big\}, \end{cases} \tag{16}$$

where $V^N(\cdot)$ is the value function if the firm does not adjust its price in the current period, $V^A(\cdot)$ is the value of the firm after it adjusts its price and $V(\cdot)$ is the value of the firm making the optimal price adjustment decision in the current period; $Q(\cdot)$ represents the stochastic discount factor.

Conditional on price adjustment, the optimal resetting price is given by

$$\bar{P}(\mathbb{S}_{jt}) = \arg\max_{P_{jt}} \Big\{\Pi_{jt}(P_{jt}) + \mathbb{E}\big\{Q(\mathbb{S}_{j,t+1})V(\mathbb{S}_{j,t+1})\big|\mathbb{S}_{jt}\big\}\Big\}.$$

Therefore, the policy function of the firm is:

$$P(\mathbb{S}_{jt}) = \begin{cases} P_{j,t-1}, & \text{if } V^N(\mathbb{S}_{jt}) > V^A(\mathbb{S}_{jt}) - \kappa_{jt}, \\ \bar{P}(\mathbb{S}_{jt}), & \text{otherwise.} \end{cases} \tag{17}$$

In the first case the firm leaves its price unchanged and pays no menu cost, while in the second case it optimally adjusts its price and pays the menu cost.

4.1.3 Sectoral Equilibrium

Sectoral equilibrium is characterized by a path of the sectoral price level, $\{P_t\}$, consistent with the optimal pricing policies of firms given the exogenous paths of their idiosyncratic productivity shocks and wage rates in the two countries $\{W_t, W_t^*\}$. We define $E_t \equiv W_t^*/W_t$ to be the wage-based (real) exchange rate.

We assume that all prices are set in the domestic unit of account, consistent with the evidence of dollar (local currency) pricing documented in the data. We further assume that the value of the domestic unit of account is stable relative to movements in the exchange rate and the domestic real wage is also stable. These assumptions that domestic price

[33] In general, one should condition expectations in the Bellman equation on the whole history $(\mathbb{S}_{jt}, \mathbb{S}_{j,t-1}, \ldots)$, but in our simulation procedure we assume that $\mathbb{S}_{jt}$ is a sufficient statistic following Krusell and Smith (1998).

and real wage inflation are negligible relative to real exchange rate fluctuations are a good approximation for the U.S. ¿From the modeling perspective this amounts to setting $W_t \equiv 1$ and assuming that all shocks to the (real) exchange rate arise from movements in W_t^*.

The simulation procedure of the model is discussed in detail in Appendix B, while below we discuss the calibration of the model and simulation results.

4.2 Calibration

We adopt the Klenow and Willis (2006) specification of the Kimball aggregator (12) that results in

$$\psi(x_j) = \left[1 - \varepsilon \ln\left(\frac{\sigma x_j}{\sigma - 1}\right)\right]^{\sigma/\varepsilon}, \qquad \text{where} \qquad x_j \equiv D\frac{P_j}{P}. \tag{18}$$

This demand specification is conveniently governed by two parameters, $\sigma > 1$ and $\varepsilon > 0$, and the elasticity and super-elasticity are given by:

$$\tilde{\sigma}(x_j) = \frac{\sigma}{1 - \varepsilon \ln\left(\frac{\sigma x_j}{\sigma-1}\right)} \qquad \text{and} \qquad \tilde{\varepsilon}(x_j) = \frac{\varepsilon}{1 - \varepsilon \ln\left(\frac{\sigma x_j}{\sigma-1}\right)}.$$

Note that this demand function satisfies all normalizations assumed in Section 3. When ε goes to 0 it results in CES demand with elasticity σ .

The calibrated values for the parameters $\{\beta, \kappa, \sigma_a, \rho_a, \sigma_e, \rho_e, \omega, \phi, \varepsilon, \sigma\}$ are reported in Table 8. The period of the model corresponds to one month and as is standard in the literature we set the discount rate to equal 4% on an annualized basis ($\beta = 0.96^{\frac{1}{12}}$).

To calibrate the share of imports, $\omega/(1+\omega)$, we use measures of U.S. domestic output in manufacturing from the Bureau of Economic Analysis input output tables and subtract U.S. exports in manufacturing. We calculate U.S. imports in manufacturing using U.S. trade data available from the 'Center for International Data at UC Davis' website. The 4 year average import share for the period 1998-2001 is close to 17% implying an $\omega = 0.2$.

To calibrate the cost sensitivity to the exchange rate shock, ϕ and ϕ^*, requires detailed information on the fraction of imports used as inputs in production by destination of output, the currency in which these costs are denominated, among others. Such information is typically unavailable. For our purposes we use the OECD input-output tables to calculate the ratio of imports to industrial output and find that the share for manufacturing varies between 8% and 27% across high income OECD countries. For our benchmark calibration we use a value of $1 - \phi^* = 0.25$ for foreign firms and $\phi = 0$ for domestic firms since almost all imports into the U.S. are priced in dollars. This implies that for an average firm in the industry the sensitivity of the marginal cost to the exchange rate is $\bar{\phi} = \phi^* \cdot \omega/(1+\omega) = 12.5\%$.

The (log of) the real exchange rate, $e \equiv \ln\left(W_t^*/W_t\right)$ is set to follow a very persistent process with an autocorrelation of around 0.985 and the monthly innovation to the real exchange rate is calibrated to equal 2.5%. These parameter values are consistent with the empirical evidence for developed countries. For instance for countries in the Euro zone the standard deviation of the change in the bilateral monthly real exchange rate ranges between 2.6% and 2.8%. The persistence parameter is in the mid-range of estimates reported in Rogoff (1996).

We set the steady state elasticity of demand, $\sigma = 5$ that implies a markup of 25%. This value is in the middle of the range for mean elasticity estimated by Broda and Weinstein (2006) using U.S. import data for the period 1990-2001. They estimate the mean elasticity for SITC-5 to be 6.6 and for SITC-3 to be 4.0.

We simulate the model for different values of the super-elasticity of demand, ε. The range of values is chosen to match the range of LRPT for the high-income OECD sample of 10% to 75%. Note that $\phi = 0.75$ bounds the long-run pass-through from above. The specific values used are $\varepsilon \in \{0, 2, 4, 6, 10, 20, 40\}$. Our baseline good that matches the middle of the LRPT range has $\varepsilon = 4$.[34] Given the value of $\sigma = 5$ this implies that the range of mark-up elasticity is between 0 and 10.

The parameters κ, σ_a, ρ_a are jointly calibrated to match moments of the data on the median duration of price adjustment, the median size of price adjustment and the autocorrelation of new prices in the data. This is done holding other parameters constant and with $\varepsilon = 4$, as for the baseline good that has a price duration of 5 months in the data. The implied menu cost, κ, in the baseline calibration is equal to 2.5% of the revenues conditional on adjustment. We set the standard deviation of the innovation in productivity to 8.5% and the persistence of the idiosyncratic shock to 0.95. These parameter values allow us to match the median absolute size of price change of 7% conditional on adjustment and the autocorrelation of new prices in the data. More precisely, we estimate a pooled regression in the data where we regress the new price of each good on its lagged new price, allowing for a fixed effect for each good. The autocorrelation coefficient is estimated to be 0.77.[35] For the model to match this we need a high persistence rate for the idiosyncratic shock of at least 0.95.[36] Note that the data moments imply that the ratio of variance of the innovation to the

[34] See Dossche, Heylen, and Van Den Poel (2006) for a discussion of calibrations of ε in the literature and for empirical evidence on the importance of non-zero ε.

[35] We also estimate the regression with a fixed effect for every (2-digit sector$\times$date) pair to control for sectoral trends, in addition to the fixed effect for each good. The autocorrelation coefficient in this case is 0.71.

[36] Note that in our calibration we need to assume neither very large menu costs, nor very volatile idiosyn-

exchange rate shock relative to the idiosyncratic shock is low, $\sigma_e^2/\sigma_a^2 = (2.5\%/8.5\%)^2 < 0.1$.

4.3 Quantitative Results

The first set of simulation results is about the relation between frequency and LRPT in the model-simulated data. We simulate the dynamic stationary equilibrium of the model for each value of the super-elasticity of demand and compute the frequency of price adjustment and LRPT for all firms in the industry and then separately for domestic and foreign firms. Figure 8 plots the resulting relationship between frequency and LRPT for these three groups of firms. The LRPT estimates are computed using the life-long regression (1). An exchange rate shock has two effects: a direct effect that follows from the firms costs changing and an indirect effect that follows from the change in the sectoral price index, as other firms respond to the cost shock. The magnitude of the second effect depends on the fraction of firms that are affected by the exchange rate shock and consequently, the extent of pass-through depends also on how widespread the shock is. In our calibration this is determined by $\bar{\phi}$.

In Figure 8 the line marked 'Foreign' captures the strong positive relation between LRPT and frequency in the model simulated import prices into the U.S. The variation in ε that matched the empirical range in LRPT generates a range in frequency of approximately 0.10 to 0.35, equivalent to durations of 3 to 10 months. On the other hand, the relation between frequency and pass-through is effectively absent for domestic firms, as well as for the sample of all firms in the industry: as frequency varies between 0.1 and 0.35, the LRPT estimate for domestic firms fluctuates between 0% and 5%, while for the full sample of firms it lies between 5% and 12%. For domestic firms the non zero LRPT is because of the indirect effect working through the sectoral price index. The overall low LRPT estimates for the industry are consistent with empirical estimates of low pass-through into consumer prices. Goldberg and Campa (2008) estimate exchange rate pass-through into import prices and consumer prices for a large sample of developed countries and document that pass-through into consumer prices is far lower than pass-through into import prices across all countries. For the case of the U.S., for the period 1975-2003, they estimate LRPT into import price to be 42% and into consumer prices to be 1%. This supports our emphasis on at-the-dock prices and why international data provides a meaningful environment for our study.

cratic shocks, as opposed to Klenow and Willis (2006). There are a few differences between our calibration and that of Klenow and Willis (2006). First, we assume a smaller mark-up elasticity: Our baseline good has a super-elasticity of 4, as opposed to 10 in their calibration. Second, they assume a much less persistent idiosyncratic shock process and match the standard deviation of relative prices rather than the average size of price adjustment.

A second set of results relates to the performance of the two LRPT estimators employed in the empirical section— the life-long regression (1) and the aggregate regression (2) — when applied to the simulated panel of firm prices. Figure 9 plots the relationship between frequency and three different measures of long-run pass-through for the exercise described above. The first measure of long-run pass-through ('Aggregate') is the 24-month cumulative pass-through coefficient from the aggregate pass-through regression. The second measure ('Life-Long 1') corresponds to the life-long micro-level regression in which we control for firm idiosyncratic productivity. This ensures that the long-run pass-through estimates do not compound the selection effect present in menu cost models. This type of regression is however infeasible to run empirically since firm-level marginal costs are not observed. The third measure ('Life-Long 2') corresponds to the same life-long micro-level regression, but without controlling for firm idiosyncratic productivity. This estimate is the counterpart to the empirical life-long estimates we presented in Section 2. We observe from the figure that all three measures of LRPT produce the same qualitative patterns and very similar quantitative results. In addition, all the estimates are close to the theoretical long-run (flexible-price) pass-through.[37] We conclude that within our calibration both estimators produce accurate measures of long-run pass-through. Specifically, 24 months is enough for the aggregate regression to capture the long-run response of prices and life-long regressions do not suffer from significant selection bias.

We next plot for illustration in Figure 10 the aggregate pass-through coefficients for different horizons up to 24 months. These are the counterparts to our empirical results in Figures 3-5. We do this for three values of the super-elasticity of demand, $\varepsilon = 6$, 10 and 20. For these three parameter values aggregate pass-through converges respectively to 31%, 21% and 12% at the 24 month horizon.[38] The frequency of price adjustment in these three cases is 0.20, 0.17 and 0.13 respectively. Note, importantly, that this figure illustrates a clear ranking in pass-through even prior to convergence to the long-run pass-through.

In a third set of simulation results we argue that variation in ϕ or κ alone cannot quantitatively explain the findings in the data. This far we only considered variations in ε. For this exercise we instead set a fixed $\varepsilon = 4$ and first vary ϕ between 0 and 1 for the baseline

[37] One can show that the theoretical flexible price pass-through is approximated by $\bar{\phi} + \Psi(\phi - \bar{\phi})$, where $\bar{\phi}$ is the average sectoral sensitivity of the firms marginal cost to exchange rate and Ψ is the pass-through elasticity as defined in Section 3. Note the close relation between this expression and the one in equation (8).

[38] Note that in our menu cost model the convergence is very fast and is almost over by the end of 6 months. This contrasts with much slower dynamics in the data. A menu cost model has a hard time generating the observed short-run behavior of prices as it predicts very fast adjustment to shocks. This is less of a concern for our purposes, since we study the long-run relationship between frequency and pass-through.

value of $\kappa = 2.5\%$ and then vary κ between 0.5% and 7.5% for the baseline value of $\phi = 0.75$. Figure 11 plots the results. First, observe that variation in ϕ indeed generates a positive relationship between frequency and LRPT, however, the range of variation in frequency is negligible, as predicted in Section 3 (see footnote 28). As ϕ increases from 0 to 1, LRPT increases from 0 to 55%, while frequency increases from 0.20 to 0.23 only. An important implication of this is that frequency of price adjustment should not be very different across domestically produced and imported goods, despite their ϕ being very different. This is indeed the case in the data. Gopinath and Rigobon (2008) document that for categories of goods in the import price index that could be matched with the producer price index and using the duration measures from Nakamura and Steinsson (2008) for producer prices the mean duration for the import price index is 10.3 months and it is 10.6 months for the producer price index.

Next observe that the assumed range of variation in the menu cost, κ, easily delivers large range of variation in frequency, as expected. However, it produces almost no variation in LRPT, which is stable around 39%. Not surprisingly, in a menu cost model, exogenous variation in the frequency of price adjustment cannot generate a robust positive relationship between frequency and measured long-run pass-through.[39]

The fourth set of results reported in Figure 12 presents a robustness check by examining if a Calvo model with large exogenous differences in the flexibility of prices can induce a positive correlation between frequency and measured LRPT even though the true LRPT is the same. We simulate two panels of firm prices — one for a sector with low probability of price adjustment (0.07) and another for a sector with high probability of price adjustment (0.28), the same as in Table 1. We set $\varepsilon = 0$ (CES demand) and keep all parameters of the model as in the baseline calibration. The figure plots both pass-through estimates from aggregate regressions at different horizons, as well as life-long pass-through estimates (depicted as the circle and the square over the 36 months horizon mark). Aggregate pass-through at 24 months is 0.52 for low-frequency adjusters, while it is 0.70 for high-frequency adjusters. At the same time, the life-long pass-through estimates are 0.61 and 0.70 respectively. As is well known, the Calvo model generates much slower dynamics of price adjustment as compared to the menu cost model. This generates a significant difference in aggregate pass-through even at the 24 months horizon, however, this difference is far smaller than the one documented in Section 2. The difference in pass-through is yet much smaller for the Calvo model if we

[39] A similar pattern emerges for variation in the elasticity of demand σ when the super-elasticity ε is set to zero (i.e., CES demand): variation in σ leads to variation in frequency with long-run pass-through stable at ϕ.

consider the life-long estimates of long-run pass-through. Further, as Figure 2 suggests, the steep relation between frequency and pass-through arises once frequency exceeds 0.13. A Calvo model calibrated to match a frequency of at least 0.13 converges sufficiently rapidly and there is no bias in the estimates. Therefore, we conclude that standard sticky price models with exogenous differences in the frequency of adjustment would have difficulty matching the empirical relationship between frequency and pass-through.

In the last set of simulation results we evaluate the quantitative performance of the model in matching the following three moments of the data: the standard deviation of frequency, the slope coefficients in a regression of LRPT on frequency and of size on frequency. The last two represent the slopes in figures 2 (HIOECD) and 6 respectively. We estimate these moments in the model simulated series, just as we do in the data, by sorting goods into 10 bins based on their frequency of adjustment and then estimate LRPT within each bin. Table 9 describes the results. The second column provides the moments in the data.

The third column presents the results for the case when there is only variation in ε. Variation in ε alone can generate 37% of the empirical dispersion in frequency. The slope coefficient while large and positive (1.86) is much steeper than in the data (0.56). The range in pass-through is however comparable by design. The model with only variation in ε generates a positive relation between size and frequency, as discussed in section 3, with a wide range of sizes that varies from 3.8% to 11.8% as ε decreases. In the data, however, this slope is close to 0 and the size range is small, between 5.4%-7.4%. This is not surprising given that in the data we sort goods based on frequency and consequently a high-frequency bin combines goods that adjust frequently either because they have a low super-elasticity of demand, ε, or because they have a low menu cost, κ. A model that allows for such additional sources of variation across goods should improve the fit of the model.

We show that this is indeed the case by introducing variation in κ alongside variation in ε. Specifically, for each ε we have firms facing a range of menu cost parameters from 0.1% to 20% of steady state revenues conditional on adjustment.[40] The cross-sectional distribution of menu costs is independent of the distribution of ε. This variation in κ can match the two slope statistics almost perfectly, as reported in column 5 of Table 9. The model generates a slope of 0.55 between frequency and LRPT, while in the data it is 0.56 and also matches the close to zero slope between size and frequency in the data, -0.05 versus -0.01. The fraction of the standard deviation in frequency explained by the model increases to 60%. Lastly, for

[40] Note that a firm with a 20% menu cost adjusts at most once a year implying an annual cost of adjustment of less that 2% of revenues.

the fourth column we shut down variation in ε, so all of the variation is in κ. The model with only κ performs poorly for reasons described earlier as it predicts no slope between frequency and pass-through (as emphasized in Figure 11) and a sizeable negative slope for the size-frequency relationship (as seen in Figure 14).

In Figure 13 and Figure 14 we plot the relationship between frequency and LRPT and between frequency and size for the case with variation in ε alone and for the case with joint variation in ε and κ against the data series from Figure 2 and Figure 6. As follows from the results in Table 9, a menu cost model that allows for joint cross-sectional variation in mark-up elasticity and menu costs can quantitatively account for both the LRPT-frequency and size-frequency relationships and generate large dispersion in frequency across goods. Further, the variable mark-up channel is essential to matching the relation between LRPT and frequency and can generate a significant fraction of the variation in frequency observed in the data.

5 Conclusion

We exploit the open economy environment with an observable and sizeable cost shock, namely the exchange rate shock, to shed light on the mechanism behind the sluggish response of aggregate prices to cost shocks. We find that firms that adjust prices infrequently also pass-through a lower amount even after several periods and multiple rounds of price adjustment, as compared to high-frequency adjusters. In other words, firms that infrequently adjust prices are typically not as far from their desired price due to their lower desired pass-through of cost shocks. On the other hand, firms that have high pass-through drift farther away from their optimal price and, therefore, make more frequent adjustments.

We also show evidence that there is interesting variation within sectors in the frequency of price adjustment that is linked to LRPT. This within-sector variation is not surprising in models with variable mark-ups. In a model where the strategic interactions between firms is explicitly modeled, as in Feenstra, Gagnon, and Knetter (1996) and Atkeson and Burstein (2007), mark-up elasticity will, among other things, depend on the market share of the firm and this relationship is non-monotonic. So firms in a sector with the same elasticity of substitution across goods will differ in their mark-up elasticity depending on their market share. Consequently, one should expect to see differences in mark-up elasticity, pass-through and frequency of price adjustment even within the same disaggregated sector.

We have evaluated the empirical evidence through the lens of standard dynamic pricing models and find that a menu cost model with variation in the mark-up elasticity can match the facts in the data, while models with exogenous frequency of price adjustment and no variation across goods in LRPT have a difficult time in matching the facts.

A Proofs for Section 3

Proof of Proposition 1 The desired flexible price of the firm (5) can be rewritten in logs as

$$\ln P(a,e) = \tilde{\mu}\big(P(a,e)\big) + \ln(1-a) + \ln(1+\phi e),$$

where $\tilde{\mu} \equiv \tilde{\mu}(P) = \ln\big[\tilde{\sigma}(P)/(\tilde{\sigma}(P)-1)\big]$ is the log mark-up. Taking a first order Taylor approximation around the point $a = e = 0$ gives

$$\left(1 - \frac{\partial\tilde{\mu}(P(0,0))}{\partial \ln P}\right)\cdot\big[\ln P(a,e) - \ln P(0,0)\big] + O\big(P(a,e)-P(0,0)\big)^2 = \big(-a+\phi e\big) + O(\|a,e\|)^2,$$

where $O(x)$ denotes the same order of magnitude as x and $\|\cdot\|$ is some norm in $\mathbb{R}^2$.

Using the definitions of $\tilde{\mu}$, $\tilde{\sigma}$ and $\tilde{\varepsilon}$, we obtain: $\partial\tilde{\mu}(P)/\partial\ln P = -\tilde{\varepsilon}(P)/\big(\tilde{\sigma}(P)-1\big)$. Given our demand and cost normalization, we have $P(0,0) = 1$ and hence $\tilde{\sigma}\big(P(0,0)\big) = \sigma$ and $\tilde{\varepsilon}\big(P(0,0)\big) = \varepsilon$. This allows us to rewrite the Taylor expansion as:

$$\ln P(a,e) - \ln P(0,0) + O\big(P(a,e)-P(0,0)\big)^2 = \Psi\big(-a+\phi e\big) + O(\|a,e\|)^2,$$

where $\Psi \equiv \Big[1+\frac{\varepsilon}{\sigma-1}\Big]^{-1}$. This immediately implies that $O\big(P(a,e)-P(0,0)\big) = O\big(\|a,e\|\big)$, so that we can rewrite:

$$\ln P(a,e) - \ln P(0,0) = \Psi\big(-a+\phi e\big) + O(\|a,e\|)^2.$$

The final step makes use of Lemma 1 which states that $\bar{P}_0 = P(0,0) + O\big(\|a,e\|\big)^2$ and allows us to substitute $\bar{P}_0$ for $P(0,0)$ in the Taylor expansion above without affecting the order of approximation:

$$\ln P(a,e) - \ln \bar{P}_0 = \Psi\big(-a+\phi e\big) + O(\|a,e\|)^2.$$

Lastly, for small shocks, the difference in logs is approximately equal to the percentage change, which results in expression (7) in the text of the proposition. Pass-through elasticity is defined as $\partial P(0,0)/\partial e = \phi\Psi$, which follows from the above approximation. ■

Proof of Lemma 2 Taking a second order Taylor approximation to the profit-loss function around the desired price $P(a,e)$ results in,

$$L(a,e) \equiv \Pi(a,e) - \Pi(\bar{P}_0|a,e) = -\frac{1}{2}\frac{\partial^2\Pi(a,e)}{\partial P^2}\big(P(a,e)-\bar{P}_0\big)^2 + O(\|a,e\|)^3$$

where the first order term is zero due to the FOC of profit maximization and $O\big(P(a,e)-\bar{P}_0\big) = O(\|a,e\|)$ from Lemma 1 and the proof of Proposition 1. The second derivative of the profit function with respect to price is

$$\frac{\partial^2\Pi(P|a,e)}{\partial P^2} = \varphi'(P)\frac{\partial\Pi(P|a,e)}{\partial P} - \frac{\tilde{\sigma}(P)\varphi(P)}{P}\left\{\tilde{\varepsilon}(P) + \frac{mc}{P}\left[1-\tilde{\varepsilon}(P)\right]\right\}.$$

Evaluating this expression at $P_1 = P(a,e)$, we have:

$$\frac{\partial^2 \Pi(a,e)}{\partial P^2} = -\frac{\big(\tilde{\sigma}(P_1)-1\big)\cdot\varphi(P_1)}{P_1}\left[1+\frac{\tilde{\varepsilon}(P_1)}{\tilde{\sigma}(P_1)-1}\right],$$

where we used the first order condition which implies $mc/P_1 = \big(\tilde{\sigma}(P_1)-1\big)/\tilde{\sigma}(P_1)$. Note that $\tilde{\sigma}(P_1) > 1$ and $\tilde{\varepsilon}(P_1) \geq 0$ are indeed sufficient conditions for profit maximization at P_1.

Assuming that $\tilde{\varepsilon}(\cdot)$ is a smooth function, we can use the following approximation:

$$\frac{\partial^2 \Pi(a,e)}{\partial P^2} = \frac{\partial^2 \Pi(0,0)}{\partial P^2} + O(\|a,e\|) = -\frac{\sigma-1}{\Psi} + O(\|a,e\|),$$

where the second equality evaluates the second derivative of the profit function at $a = e = 0$, taking into account that $P(0,0) = 1$ and $\varphi(1) = 1$, $\tilde{\sigma}(1) = \sigma$ and $\tilde{\varepsilon}(1) = \varepsilon$ due to our normalizations.

Combining the above results and the implication of Lemma 1 that $\bar{P}_0 = P(0,0) + O(\|a,e\|)^2 = 1 + O(\|a,e\|)^2$, we can rewrite the approximation to the profit loss function as

$$\mathrm{L}(a,e) \equiv \Pi(a,e) - \Pi(\bar{P}_0|a,e) = \frac{1}{2}\frac{\sigma-1}{\Psi}\left(\frac{P(a,e)-\bar{P}_0}{\bar{P}_0}\right)^2 + O(\|a,e\|)^3. \qquad \blacksquare$$

B Simulation Procedure

To simulate the dynamic model of Section 4, we first need to make a few approximations. As described in the text, the demand for a variety j is a function of the normalized relative price, $x_{jt} = D_t P_{jt}/P_t$, where the general expression for the normalization parameter D_t was provided in the text. In the case of the Klenow and Willis (2006) demand specification, this expression becomes:

$$D_t = \frac{\sigma-1}{\sigma}\int_\Omega \frac{C_{jt}}{C_t}\exp\left\{\frac{1}{\varepsilon}\left[1-\left(\frac{|\Omega|C_{jt}}{C_t}\right)^{\varepsilon/\sigma}\right]\right\}\mathrm{d}j, \tag{19}$$

where

$$\exp\left\{\frac{1}{\varepsilon}\left[1-\left(\frac{|\Omega|C_{jt}}{C_t}\right)^{\varepsilon/\sigma}\right]\right\} = \Psi'\left(\frac{|\Omega|C_{jt}}{C_t}\right) = \psi^{-1}\left(\frac{|\Omega|C_{jt}}{C_t}\right).$$

In a symmetric steady state $P_j = P = \bar{P}$, $C_j = \bar{C}/|\Omega|$ for all j, so that $x_j \equiv \bar{D} = (\sigma-1)/\sigma$, and the elasticity and super-elasticity of demand equal σ and ε respectively. Equations (14) and (13) in the text together with our demand specification imply the following (implicit) expression for the sectoral price level:

$$P_t = \frac{1}{|\Omega|}\int_\Omega P_{jt}\left[1-\varepsilon\ln\left(\frac{\sigma D_t}{\sigma-1}\frac{P_{jt}}{P_t}\right)\right]^{\sigma/\varepsilon}\mathrm{d}j. \tag{20}$$

We can now prove the following

Lemma 3 *(i) The first-order deviation of D_t from $\bar{D} = (\sigma-1)/\sigma$ is nil. (ii) The geometric average provides an accurate first-order approximation to the sectoral price level:*

$$\ln P_t \approx \frac{1}{|\Omega|}\int_\Omega \ln P_{jt}\mathrm{d}j.$$

In both cases the order of magnitude is $O\big(\|\{\hat{P}_{jt}\}_{j\in\Omega}\|\big)$, where $\|\cdot\|$ is some vector norm in L^∞ and the hat denotes log deviation from steady state value, $\hat{P}_{jt} \equiv \ln P_{jt} - \ln \bar{P}$.

Proof: Writing (19) and (20) in log deviations from the symmetric steady state we obtain two equivalent first-order accurate representations for $\hat{D}_t$:

$$\frac{\sigma}{\sigma-1}\hat{D}_t = \frac{1}{|\Omega|}\int_\Omega \big(\hat{C}_{jt} - \hat{C}_t\big)\mathrm{d}j = -\frac{1}{|\Omega|}\int_\Omega \big(\hat{P}_{jt} - \hat{P}_t\big)\mathrm{d}j,$$

Now using the definition of the Kimball aggregator (12) and the fact that $\Psi(\cdot)$ is a smooth function, we have

$$\frac{1}{|\Omega|}\int_\Omega \big(\hat{C}_{jt} - \hat{C}_t\big)\mathrm{d}j = 0$$

up to second-order terms, specifically, the approximation error has the order of $O\big(\|\{\hat{C}_{jt}\}_{j\in\Omega}\|^2\big)$. Combining the two results together immediately implies that $\hat{D}_t = 0$ and $\hat{P}_t = |\Omega|^{-1}\int_\Omega \hat{P}_{jt}\mathrm{d}j$ up to the second-order terms, $O\big(\|\{\hat{C}_{jt}\}_{j\in\Omega}\|^2\big)$. Taking the log-differential of the demand equation (18), we have

$$\hat{C}_{jt} - \hat{C}_t = -\sigma\big(\hat{D}_t + \hat{P}_{jt} - \hat{P}_t\big),$$

which allows us to conclude that $O\big(\|\{\hat{C}_{jt}\}_{j\in\Omega}\|\big) = O\big(\|\{\hat{P}_{jt}\}_{j\in\Omega}\|\big)$. Finally, note that the expression for $\hat{P}_t$ is equivalent to

$$\ln P_t = \frac{1}{|\Omega|}\int_\Omega \ln P_{jt}\mathrm{d}j$$

since in a symmetric steady state $P_j = P = \bar{P}$. ■

This result motivates us to make the following assumptions: In our simulation procedure we set $D_t \equiv \bar{D} = (\sigma-1)/\sigma$ and compute the sectoral price index as the geometric average of individual prices. Lemma 3 ensures that these are accurate first-order approximations to the true expressions and using them speeds up the simulation procedure considerably as we avoid solving for another layer of fixed point problems.[41] We verify, however, that computing the sectoral price index according to the exact expression (20) does not change the results.

Additionally, we introduce two more approximations. First, we set the stochastic discount factor to be constant and equal to the discount factor $\beta < 1$. Second, we set the sectoral consumption index to $C_t \equiv 1$. Both of these assumptions are in line with our partial equilibrium approach and they introduce only second-order distortions to the price-setting problem of the firm.

The assumptions in Section 4.1.3 allow us to reduce the state space for each firm to $\mathbb{S}_{jt} = (P_{j,t-1}, A_{jt}; P_t, e_t)$, where $e_t = \ln(W_t^*/W_t) = \ln W_t^*$. We iterate the Bellman operator (16) on a logarithmic grid for each dimension of $\mathbb{S}_{jt}$. Specifically, the grid for individual price P_{jt} is chosen so that an increment is no greater than a 0.5% change in price (typically, around 200 grid points).

[41] This is also an assumption adopted by Klenow and Willis (2006).

The grid for idiosyncratic shock A_{jt} contains 11 grid points and covers ± 2.5 unconditional standard deviations for the stochastic process. The grid for the sectoral price level P_t is such that an increment is no greater than a 0.2% change in the price level (typically, around 30 grid points). Finally, the grid for the real exchange rate e_t has 15 grid points with increments equal to σ_e and covering $\pm 7\sigma_e$.[42]

To iterate the Bellman operator (16), a firm needs to form expectations about the future path of the exogenous state variables, (A_{jt}, e_t, P_t). Since A_{jt} and e_t follow exogenous stochastic processes specified above, the conditional expectations for these variables are immediate.[43] The path of the sectoral price index, P_t, however, is an endogenous equilibrium outcome and to set prices the firm needs to form expectations about this path. This constitutes a fixed point problem: the optimal decision of a firm depends on the path of the price level and this optimal decision feeds into the determination of the equilibrium path of the price level.[44] Following Krusell and Smith (1998), we assume that firms base their forecast on the restricted set of state variables, specifically:

$$\mathbb{E}_t \ln P_{t+1} = \gamma_0 + \gamma_1 \ln P_t + \gamma_2 e_t.$$

In principle, the lags of $(\ln P_t, e_t)$ can also be useful for forecasting $\ln P_{t+1}$, however, in practice, $(\ln P_t, e_{t+1})$ alone explains over 95% of variation in $\ln P_{t+1}$ and e_t is a sufficient statistic to forecast e_{t+1}. The firms use the forecasting vector $(\gamma_0, \gamma_1, \gamma_2)$ consistent with the dynamics of the model. This reflects the fact that they form rational expectations given the restricted set of state variables which they condition on. To implement this, we:

(i) Start with an initial forecasting vector $\left(\gamma_0^{(0)}, \gamma_1^{(0)}, \gamma_2^{(0)}\right)$;

(ii) Given the forecasting vector, iterate the Bellman equation till convergence to obtain policy functions for price setting;

(iii) Using the policy functions, simulate M dynamic paths of the sectoral price level $\left\{P_t^{(m)}\right\}_{t=0}^{T}$, where in every period we make sure that P_t is consistent with the price setting of firms;

(iv) For each simulation estimate $\left(\hat{\gamma}_0^{(0,m)}, \hat{\gamma}_1^{(0,m)}, \hat{\gamma}_2^{(0,m)}\right)$ from regressing $\ln P_{t+1}$ on $\ln P_t$, e_t and a constant and obtain $\left(\gamma_0^{(1)}, \gamma_1^{(1)}, \gamma_2^{(1)}\right)$ by taking the median of $\left(\hat{\gamma}_0^{(0,m)}, \hat{\gamma}_1^{(0,m)}, \hat{\gamma}_2^{(0,m)}\right)$;

[42] We let the (log of) the real exchange rate, e_t, follow a binomial random walk process within wide boundaries. Specifically, its value each period either increases or decreases by σe with equal probabilities and reflects from the boundaries of a grid. We use this procedure to numerically generate a highly persistent process for the exchange rate, which is harder to obtain using the Tauchen routine.

[43] Recall that A_{jt} follows a first-order autoregressive process; we discretize it using the Tauchen routine.

[44] In fact, there are two distinct fixed point problems, one static and one dynamic. The price that the firm sets today, P_{jt}, depends both on the price level today, P_t, and the expectation of the price level in the future, $\mathbb{E}_t P_{t+1}$. The static problem is easy to solve: holding the expectations constant, we find P_t consistent with

$$\ln P_t = \frac{1}{|\Omega|} \int_{\Omega} \ln P_{jt}(P_t) \mathrm{d}j,$$

where $P_{jt}(P_t)$ underlines the dependence of the individual prices on the sectoral price level.

(v) Iterate this procedure till joint convergence of the forecasting equation coefficients.

This constitutes a reasonable convergence procedure in a stochastic environment.

Once the forecasting vector is established, we iterate the Bellman operator to find policy functions for domestic and foreign firms in every state. This then allows us to simulate the panel of individual prices similar to the one we use in the empirical section. Specifically, we simulate a stationary equilibrium with $12,000$ domestic and $2,400$ foreign firms operating in the local market. We simulate the economy for 240 periods and then take the last 120 periods. During this time interval each firm appears in the sample for on average 35 consecutive months (on average 4 price adjustments for each firm), which generates an unbalanced panel of firm price changes, as we observe in the data.[45] On this simulated data, we estimate the same regressions (1) and (2) as we do on the BLS dataset. We repeat the simulation B times and take the median of all statistics across these B simulations to exclude noise in the estimates.

In the final simulation exercise we allow firms to be heterogeneous in both their menu cost κ and super-elasticity of demand ε. Specifically, for every value of ε we have firms with different values of κ distributed uniformly on $\{0.1\%, 0.5\%, 1.25\%, 2.5\%, 5\%, 10\%, 20\%\}$. We simulate price data for a panel of firms using the procedure described above. We then estimate frequency for every firm and sort firms into frequency deciles. For each frequency decile we estimate life-long pass-through and median size of price adjustment.

[45]Specifically, for each firm we choose a random time interval during which its price is observed by the econometrician, although the good exists in all time periods. This captures the feature that in the data the BLS only observes price changes when the good is in the sample and only a few price changes are observed.

References

Atkeson, A., and A. Burstein (2007): "Pricing-to-Market, Trade Costs and International Relative Prices," *American Economic Review*, (Forthcoming).

Ball, L., and G. Mankiw (1994): "Asymmetric Price Adjustment and Economic Fluctuations," *Economic Journal*, 104, 247–261.

Ball, L., N. G. Mankiw, and D. Romer (1988): "The New Keynesian Economics and the Output-Inflation Trade-off," *Brookings Papers on Economic Activity*, 1, 1–65.

Barro, R. J. (1972): "A Theory of Monopolistic Price Adjustment," *Review of Economic Studies*, 39(1), 17–26.

Bergin, P., and R. Feenstra (2001): "Pricing to Market, Staggered Contracts and Real Exchange Rate Perisistence," *Journal of International Economics*, 54(2), 333–359.

Bils, M., and P. Klenow (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112(5), 947–85.

Bils, M., P. Klenow, and B. Malin (2009): "Reset Price Inflation and the Impact of Monetary Policy Shocks," NBER Working paper No. 14787.

Broda, C., and D. E. Weinstein (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121(2), 541–585.

Burstein, A., and N. Jaimovich (2009): "Understanding Movements in Aggregate and Product-level Real Exchange Rates," Working paper.

Caplin, A. S., and D. F. Spulber (1987): "Menu Costs and the Neutrality of Money," *Quarterly Journal of Econmics*, 102(4), 703–26.

Devereux, M. B., and J. Yetman (2008): "Price Adjustment and Exchange Rate Pass-through," Working paper.

Dornbusch, R. (1987): "Exchange Rate and Prices," *American Economic Review*, 77(1), 93–106.

Dossche, M., F. Heylen, and D. Van Den Poel (2006): "The Kinked Demand Curve and Price Rigidity : Evidence from Scanner Data," Working Papers of Faculty of Economics and Business Administration, Ghent University, Belgium.

Dotsey, M., R. King, and A. Wolman (1999): "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics*, 114(2), 655–690.

Feenstra, R. C., J. E. Gagnon, and M. M. Knetter (1996): "Market Share and Exchange Rate Pass-through in World Automobile Trade," *Journal of International Economics*, 40(1-2), 187–207.

Fitzgerald, D., and S. Haller (2008): "Exchange Rates and Producer Prices: Evidence from Micro Data," Stanford working paper.

Goldberg, L. S., and J. M. Campa (2008): "The Sensitivity of the CPI to Exchange Rates: Distribution Margins, Imported Inputs, and Trade Exposure," *Review of Economics and Statistics*, forthcoming.

Goldberg, P., and R. Hellerstein (2007): "A Framework for identifying the Sources of Local-Currency Price Stability with an Empirical Application," NBER Working paper No. 13183.

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

Gopinath, G., P.-O. Gourinchas, C.-T. Hsieh, and N. Li (2009): "Estimating the Border Effect: Some New Evidence," NBER Working paper No. 14938.

Gopinath, G., and O. Itskhoki (2008): "Frequency of Price Adjustment and Pass-through," NBER working paper no. 14200.

Gopinath, G., O. Itskhoki, and R. Rigobon (2007): "Currency Choice and Exchange Rate Pass-through," Harvard University.

Gopinath, G., and R. Rigobon (2008): "Sticky Borders," *Quarterly Journal of Economics*, 123(2), 531–575.

Greene, W. H. (2000): *Econometric Analysis*. Prentice-Hall, NJ, 4th edn.

Kehoe, P., and V. Midrigan (2007): "Sticky Prices and Real Exchange Rates in the Cross-section," NYU.

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

Klenow, P., and J. Willis (2006): "Real Rigidities and Nominal Price Changes," Stanford Mimeo.

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

Krugman, P. (1987): "Pricing to Market when the Exchange Rate Changes," in *Real Financial Linkages among Open Economies*, ed. by S. Arndt, and J. Richardson, pp. 49–70. MIT Press, Cambridge.

Krusell, P., and A. A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 867–896.

Lane, P. R. (2001): "The New Open Economy Macroeconomics: A Survey," *Journal of International Economics*, 54(2), 235–66.

Midrigan, V. (2007): "International Price Dispersion in State-dependent pricing models," NYU.

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

Neiman, B. (2007): "Multinationals, Intrafirm Trades, and International Macroeconomic Dynamics," Working Paper.

Obstfeld, M., and K. Rogoff (1995): "Exchange Rate Dynamics Redux," *Journal of Political Economy*, 103, 624–60.

Rauch, J. (1999): "Networks versus Markets in International Trade," *Journal of International Economics*, 48, 7–35.

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

Romer, D. (1989): "Staggered Price Setting with Endogenous Frequency of Adjustment," *NBER Working Paper 3134*.

Rotemberg, J. J., and G. Saloner (1987): "The Relative Rigidity of Prices," *American Economic Review*, 77(5), 917–26.

Rotemberg, J. J., and M. Woodford (1999): "The Cyclical Behavior of Prices and Costs," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford. North-Holland.

Sheshinski, E., and Y. Weiss (1977): "Inflation and Costs of Price Adjustment," *Review of Economic Studies*, 44(2), 287–303.

Table 1: Life-long Pass-through

	Low Frequency			High Frequency			Difference		N_{obs}	R^2
	Freq.	β_L^{Lo}	s.e.(β_L^{Lo})	Freq.	β_L^{Hi}	s.e.(β_L^{Hi})	$\beta_L^{Hi} - \beta_L^{Lo}$	t-stat		
					Panel A: Real Specification					
All Countries										
– Manufacturing	0.07	0.21	0.03	0.38	0.44	0.05	0.23	4.18	14227	0.07
– Differentiated	0.07	0.20	0.04	0.29	0.46	0.06	0.26	3.89	7870	0.08
High-Income OECD										
– Manufacturing	0.07	0.25	0.04	0.40	0.59	0.06	0.34	4.43	5988	0.08
– Differentiated	0.07	0.25	0.07	0.33	0.59	0.08	0.34	3.50	2982	0.10
					Panel B: Nominal Specification					
All Countries										
– Manufacturing	0.07	0.18	0.03	0.38	0.36	0.05	0.18	3.42	14227	0.06
– Differentiated	0.07	0.16	0.04	0.29	0.37	0.06	0.21	3.22	7870	0.07
High-Income OECD										
– Manufacturing	0.07	0.22	0.04	0.40	0.55	0.07	0.32	4.19	5888	0.07
– Differentiated	0.07	0.23	0.07	0.33	0.54	0.08	0.31	3.20	2982	0.08

Note: Robust standard errors clustered by coutnry*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes.

Table 2: Life-long Pass-through: Starting with a New Price

	Low Frequency			High Frequency			Difference		N_{obs}	R^2
	Freq.	β_L^{Lo}	s.e.(β_L^{Lo})	Freq.	β_L^{Hi}	s.e.(β_L^{Hi})	$\beta_L^{Hi} - \beta_L^{Lo}$	t-stat		
				Panel A: Real Specification						
All Countries										
– Manufacturing	0.10	0.20	0.03	0.47	0.46	0.06	0.26	3.71	9002	0.09
– Differentiated	0.09	0.16	0.05	0.33	0.48	0.08	0.32	3.68	4708	0.10
High-Income OECD										
– Manufacturing	0.10	0.19	0.05	0.50	0.67	0.06	0.49	6.00	4079	0.10
– Differentiated	0.09	0.17	0.08	0.40	0.60	0.09	0.43	3.73	1942	0.11
				Panel B: Nominal Specification						
All Countries										
– Manufacturing	0.10	0.16	0.03	0.47	0.39	0.06	0.23	3.42	9002	0.08
– Differentiated	0.09	0.13	0.05	0.33	0.41	0.08	0.28	3.45	4708	0.10
High-Income OECD										
– Manufacturing	0.10	0.14	0.05	0.50	0.65	0.06	0.50	5.94	4079	0.09
– Differentiated	0.09	0.13	0.08	0.40	0.57	0.09	0.44	3.82	1942	0.10

Note: Robust standard errors clustered by country*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes.

Table 3: Life-long Pass-through: 3 or More Price Changes

	Low Frequency			High Frequency			Difference		N_{obs}	R^2
	Freq.	β_L^{Lo}	s.e.(β_L^{Lo})	Freq.	β_L^{Hi}	s.e.(β_L^{Hi})	$\beta_L^{Hi} - \beta_L^{Lo}$	t-stat		
				Panel A: Real Specification						
All Countries										
– Manufacturing	0.13	0.25	0.04	0.57	0.48	0.07	0.23	3.06	6111	0.11
– Differentiated	0.11	0.19	0.06	0.42	0.56	0.09	0.38	3.99	3031	0.16
High-Income OECD										
– Manufacturing	0.12	0.32	0.07	0.60	0.70	0.08	0.38	4.00	2856	0.11
– Differentiated	0.11	0.28	0.11	0.50	0.71	0.09	0.43	3.20	1312	0.16
				Panel B: Nominal Specification						
All Countries										
– Manufacturing	0.13	0.22	0.04	0.57	0.42	0.07	0.20	2.70	6111	0.10
– Differentiated	0.11	0.15	0.06	0.42	0.49	0.09	0.34	3.59	3031	0.13
High-Income OECD										
– Manufacturing	0.12	0.28	0.07	0.60	0.67	0.08	0.39	3.85	2856	0.10
– Differentiated	0.11	0.25	0.12	0.50	0.68	0.10	0.43	3.08	1312	0.14

Note: Robust standard errors clustered by country*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes.

Table 4: Life-long Pass-through: Regions

	Low Frequency			High Frequency			Difference		N_{obs}	R^2
	Freq.	β_L^{Lo}	s.e.(β_L^{Lo})	Freq.	β_L^{Hi}	s.e.(β_L^{Hi})	$\beta_L^{Hi} - \beta_L^{Lo}$	t-stat		
				Panel A: Real Specification						
Japan	0.07	0.31	0.07	0.27	0.62	0.15	0.31	1.81	1418	0.07
Euro Area	0.07	0.24	0.09	0.31	0.49	0.10	0.25	2.03	1802	0.08
Canada	0.10	0.38	0.12	0.87	0.74	0.23	0.36	1.58	1150	0.07
Non-HIOECD	0.07	0.17	0.03	0.36	0.34	0.06	0.17	2.44	8239	0.06
				Panel B: Nominal Specification						
Japan	0.07	0.23	0.07	0.27	0.55	0.15	0.32	1.78	1418	0.05
Euro Area	0.07	0.23	0.09	0.31	0.47	0.09	0.24	1.97	1802	0.07
Canada	0.10	0.35	0.13	0.87	0.79	0.24	0.44	1.93	1150	0.07
Non-HIOECD	0.07	0.13	0.03	0.36	0.24	0.06	0.11	1.71	8239	0.05

Note: Robust standard errors clustered by coutnry*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes. 5% threshold for t-statistic for a one sided test of $\beta_L^{Hi} = \beta_L^{Lo}$ against $\beta_L^{Hi} > \beta_L^{Lo}$ is 1.65.

Table 5: Slope coefficient for frequency–pass-through relation

	OLS					Median Quantile Regression					
	β_L	s.e.(β_L)	γ_L	s.e.(γ_L)	t-stat	β_L	s.e.(β_L)	γ_L	s.e.(γ_L)	t-stat	N_{obs}
All countries											
– Manufacturing	0.33	0.03	0.40	0.14	2.79	0.31	0.01	0.43	0.03	14.09	14227
– Differentiated	0.33	0.04	0.70	0.18	3.96	0.32	0.01	0.72	0.05	12.57	7870
High-Income OECD											
– Manufacturing	0.42	0.04	0.63	0.15	4.35	0.35	0.02	0.57	0.04	13.27	5988
– Differentiated	0.43	0.05	0.95	0.15	6.17	0.39	0.02	0.84	0.08	10.83	2982
3 or more price changes											
All countries											
– Manufacturing	0.37	0.04	0.36	0.17	2.12	0.37	0.02	0.37	0.06	6.43	6111
– Differentiated	0.38	0.06	0.88	0.20	4.35	0.40	0.02	0.92	0.08	10.94	3031
High-Income OECD											
– Manufacturing	0.51	0.06	0.60	0.18	3.42	0.45	0.02	0.47	0.08	6.30	2856
– Differentiated	0.52	0.07	0.94	0.19	4.86	0.51	0.04	0.90	0.13	6.83	1312

Note: Robust standard errors clustered by coutnry*PSL pair are reported for the OLS regressions. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes. β_L estimates average LRPT for the sub-sample; γ_L estimates the relationship between frequency and LRPT and the t-statistic is reported for the hypothesis $\gamma_L = 0$.

Table 6: Frequency-LRPT relation: Within versus between-sector evidence

	Freq	OLS			Median Quantile Regression		
	Within	γ_L^W	γ_L^B	Within	γ_L^W	γ_L^B	Within
2-digit	85%	0.49	0.17	98%	0.40	0.51	78%
		(0.09)	(0.12)		(0.03)	(0.05)	
4-digit	70%	0.48	0.30	86%	0.40	0.48	68%
		(0.07)	(0.08)		(0.04)	(0.04)	

Note: Robust standard errors in brackets. Second column reports contribution of the within-sector component to total variation in frequency of price adjustment. The estimates γ_L^W and γ_L^B are from the unrestricted version of equation (3) where we allow for separate coefficients on average sectoral frequency (γ_L^B) and on the deviation of the good's frequency from the sectoral average (γ_L^W), as described in footnote 16. Columns 5 and 8 report the contribution of the within-sector component to the relation between frequency and pass-through according to the formula in the text.

Table 7: Substitutions

Decile	Freq	Life 1	Life 2	Eff Freq 1	Eff Freq 2
1	0.03	59	42	0.05	0.05
2	0.05	50	34	0.07	0.08
3	0.07	52	32	0.09	0.10
4	0.10	55	36	0.12	0.13
5	0.13	52	33	0.15	0.16
6	0.18	49	32	0.20	0.21
7	0.29	50	26	0.30	0.31
8	0.44	51	34	0.46	0.46
9	0.67	52	33	0.67	0.68
10	1.00	43	30	1.00	1.00

Note: Effective Frequency (Eff Freq) corrects the measure of frequency during the life of the good (Freq) for the probability of the good's discontinuation/replacement according to [Freq+(1-Freq)/Life] for the two measures of the life of the good in the sample (Life).

Table 8: Parameter Values

Parameter	Symbol	Value	Source
Discount Factor	β	$0.96^{1/12}$	Annualized interest rate of 4%
Fraction of imports	$\omega/(1+\omega)$	16.7%	BEA input-output table and U.S.import data
Cost sensitivity to ER shock			OECD input-output tables
Foreign firms	ϕ^*	0.75	
U.S. firms	ϕ	0	
Menu cost	κ	2.5%	Price durations of 5 months when $\varepsilon = 4$
Demand elasticity	σ	5	Broda and Weinstein (2006)
Exchange rate process, e_t			
Std.dev. of ER shock	σ_e	2.5%	U.S.-Euro bilateral RER
Persistence of ER	ρ_e	0.985	Rogoff (1996)
Idiosyncratic productivity process, a_t			
St.dev. of shock to a_t	σ_a	8.5%	Absolute size of price adjustment of 7%
Persistence of a_t	ρ_a	0.95	Autocorrelation of new prices of 0.77

Table 9: Quantitative Results: Model Vs. Data

	Data	Variation in		
		Super-elasticity, ε	Menu Cost, κ	ε and κ
Slope(Freq., LRPT)	0.56	1.86	0.03	0.55
Min LRPT	0.06	0.13	0.44	0.22
Max LRPT	0.72	0.76	0.46	0.57
Slope(Freq., size)	-0.01	0.23	-0.15	-0.05
Min size	5.4%	3.8%	4.8%	5.8%
Max size	7.4%	11.8%	12.2%	8.2%
Std. dev. of Freq.	0.30	0.11	0.17	0.18
Min freq.	0.03	0.07	0.06	0.05
Max freq.	1.00	0.44	0.59	0.61

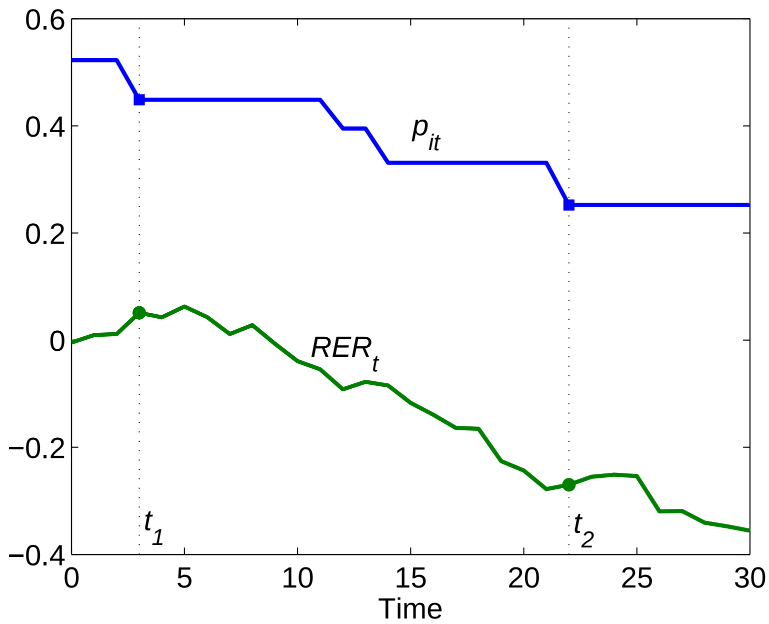

Figure 1: Illustration of Life-long pass-through

Note: In this hypothetical example we observe the price of good i from $t = 0$ to $t = 30$. The figure plots the observed price and the corresponding bilateral real exchange rate for the same period, both in logs. The first observed new prices is set on $t_1 = 3$ and the last observed new price is set on $t_2 = 22$. Therefore, for this good we have $\Delta p_L^{i,c} = p_{it_2} - p_{it_1}$ and $\Delta RER_L^{i,c} = RER_{t_2} - RER_{t_1}$. In the baseline specification we additionally adjust $\Delta p_L^{i,c}$ by U.S. CPI inflation over the same period.

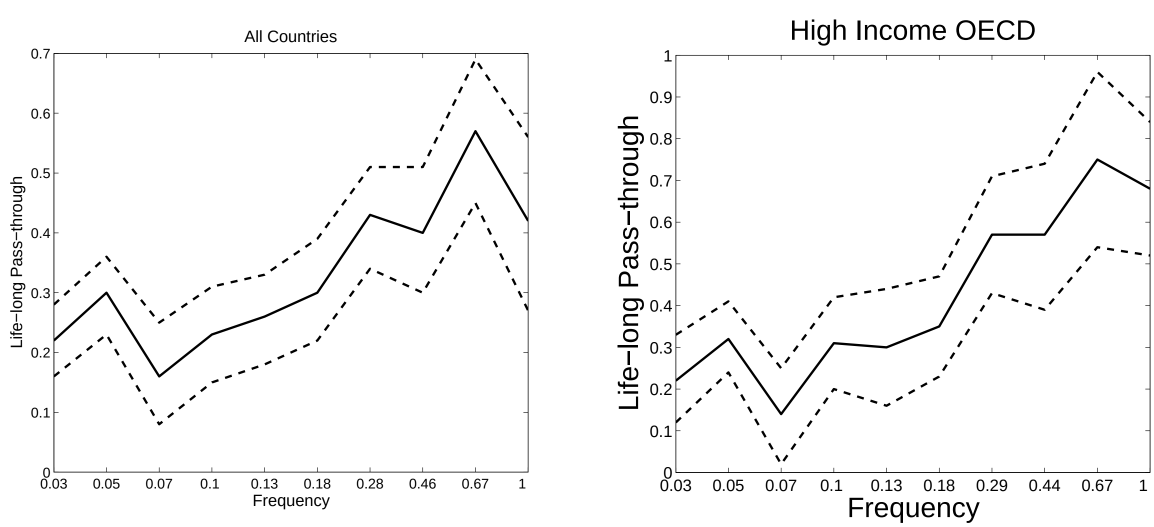

Figure 2: Life-long Pass-through across Frequency Deciles

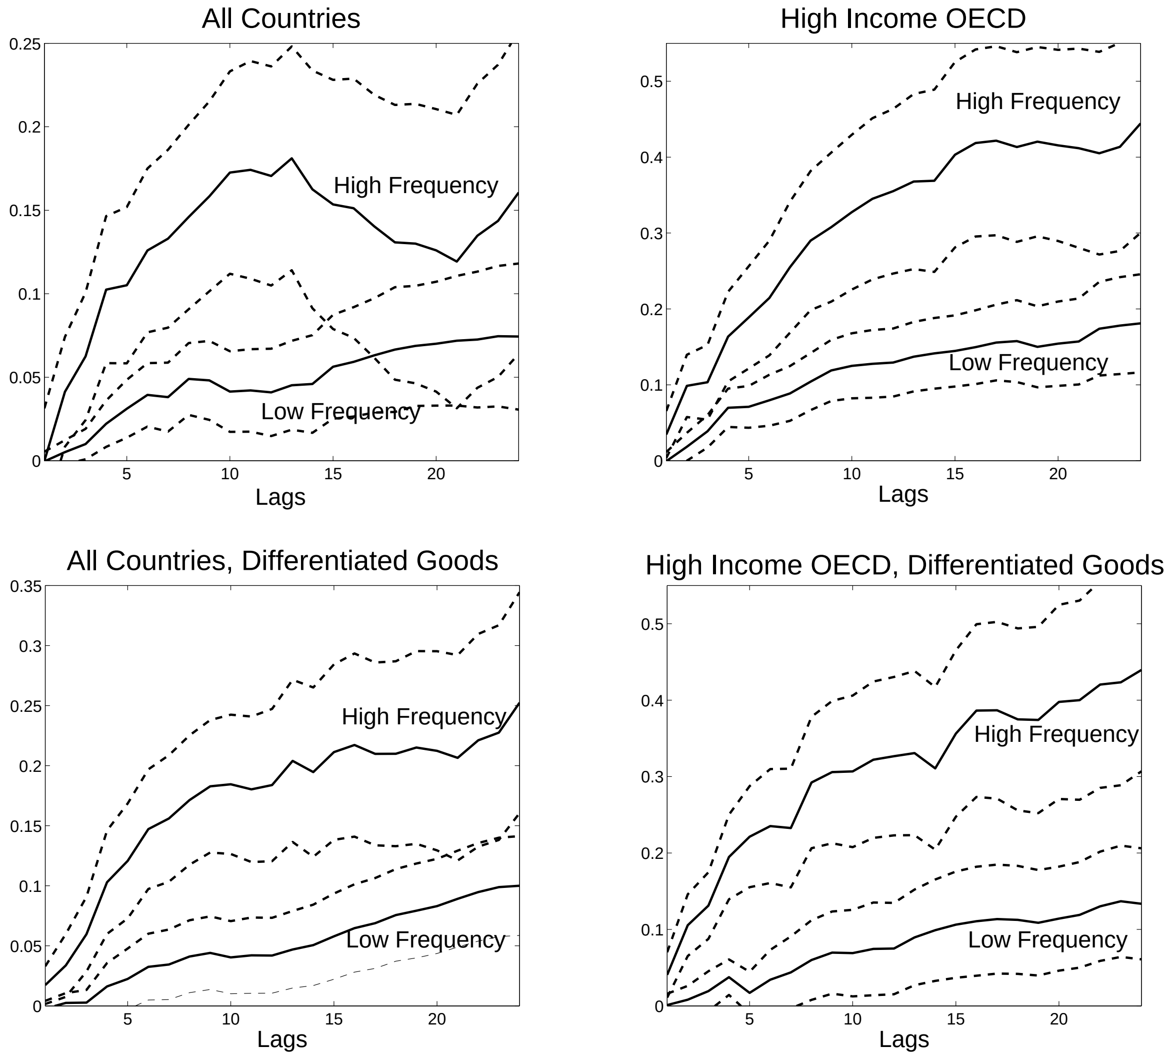

Figure 3: Aggregate Pass-through Regressions

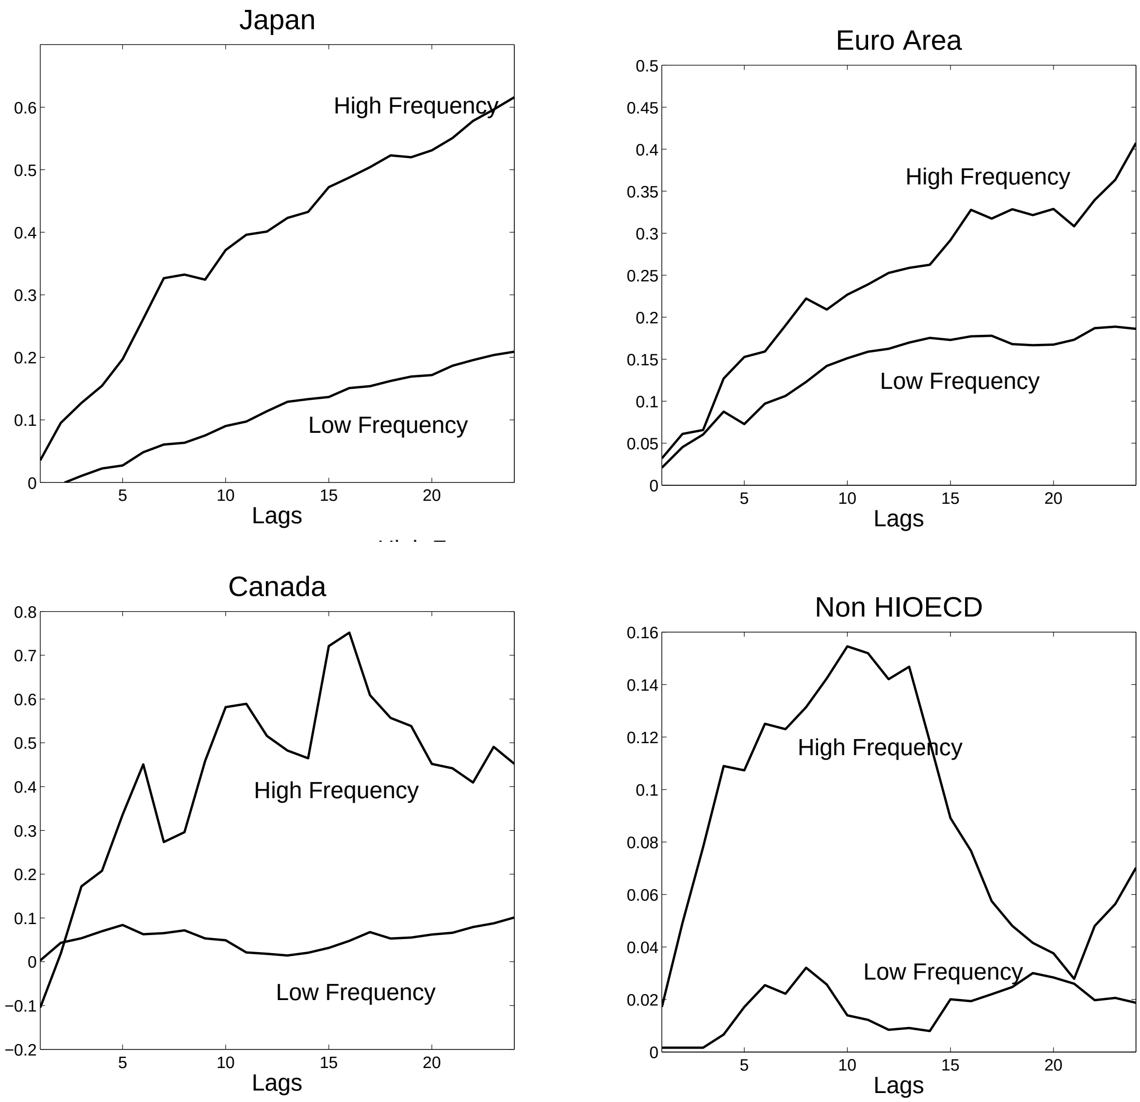

Figure 4: Aggregate Pass-through Regressions: Regions

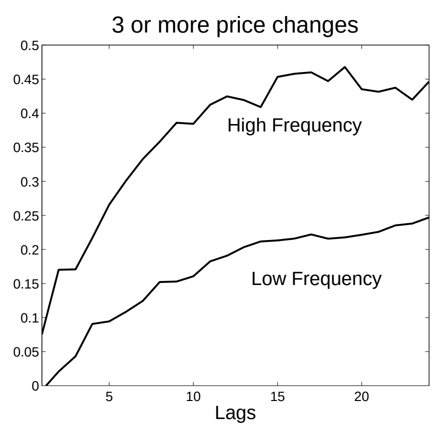

Figure 5: Aggregate Pass-through Regressions: 3 or more price changes

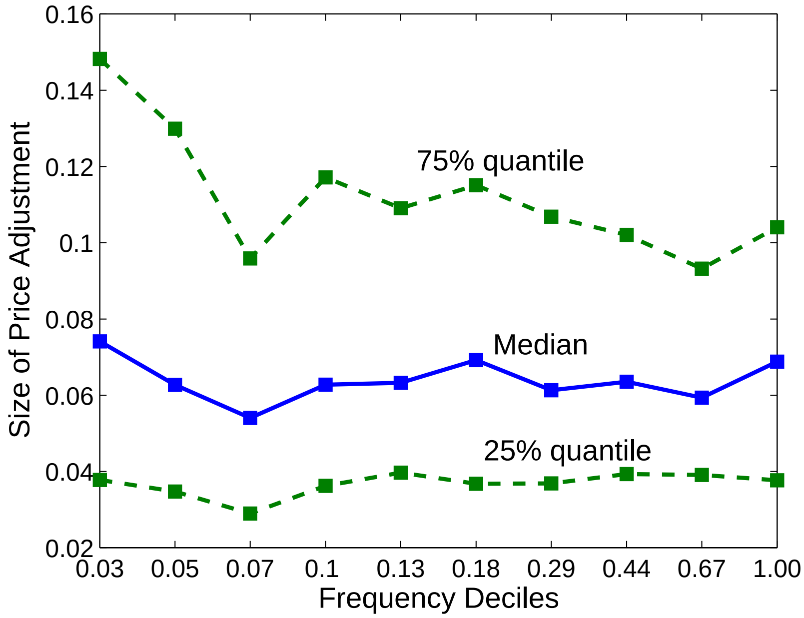

Figure 6: Size of Price Adjustment

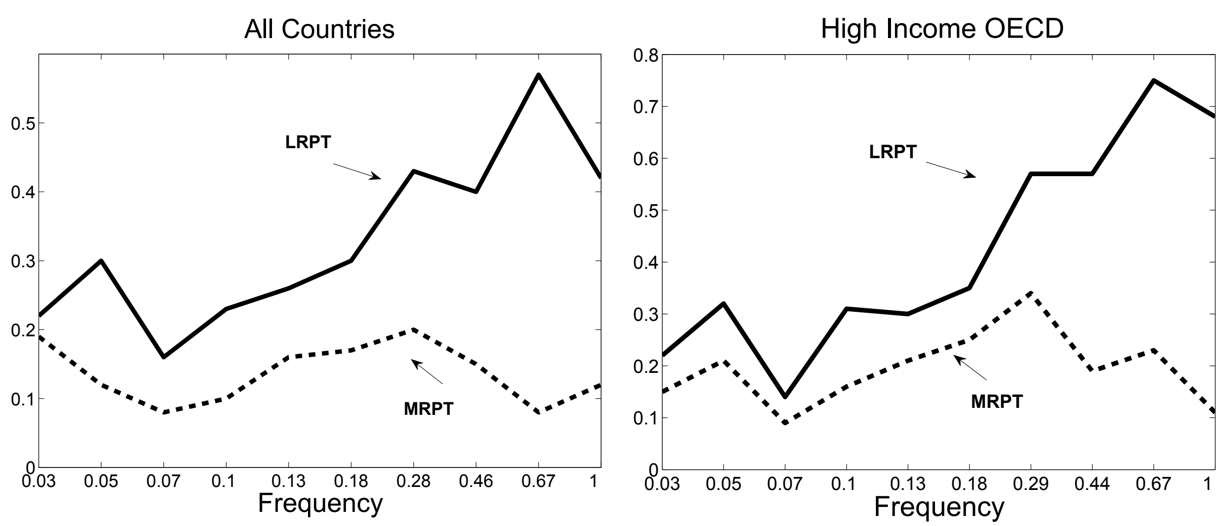

Figure 7: Long-Run versus Medium-Run Pass-through

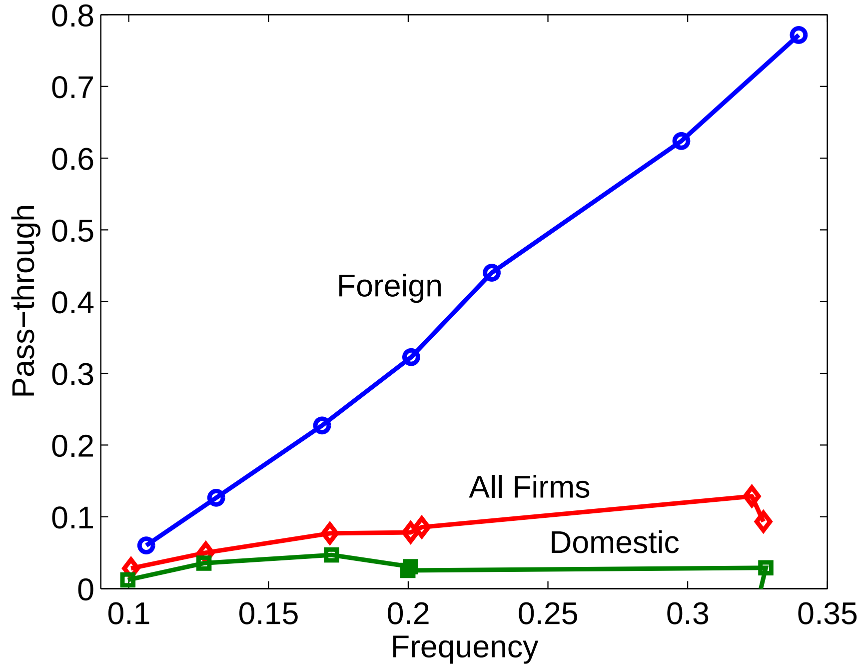

Figure 8: Frequency and Pass-through in the model (variation in ε)

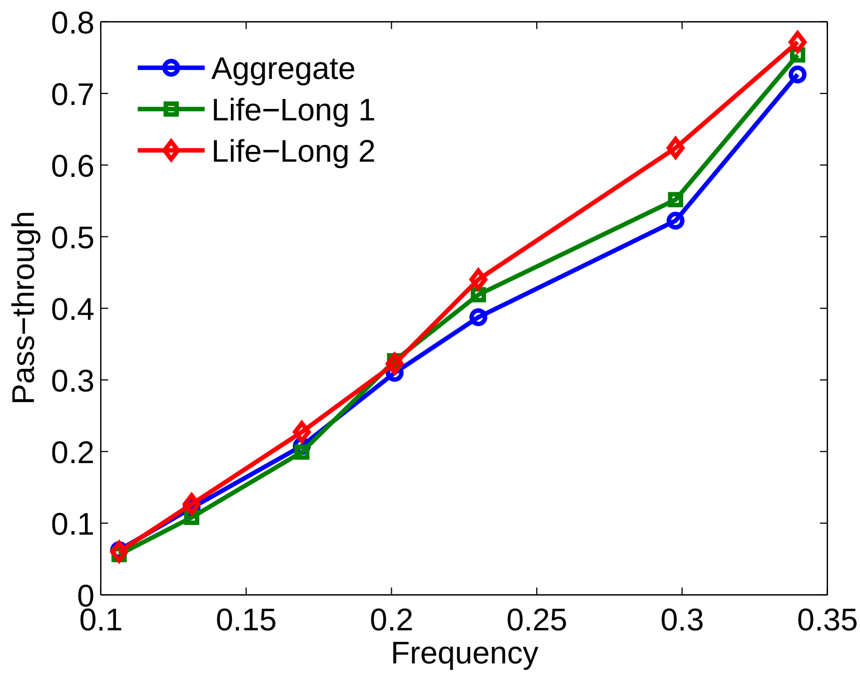

Figure 9: Measures of Long-run Pass-through

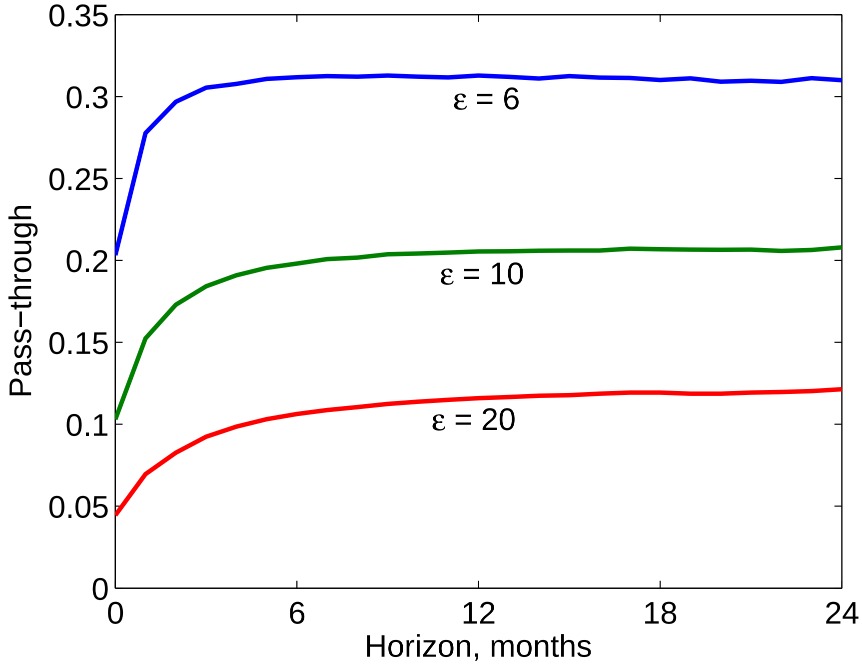

Figure 10: Aggregate Pass-through Regressions (variation in ε)

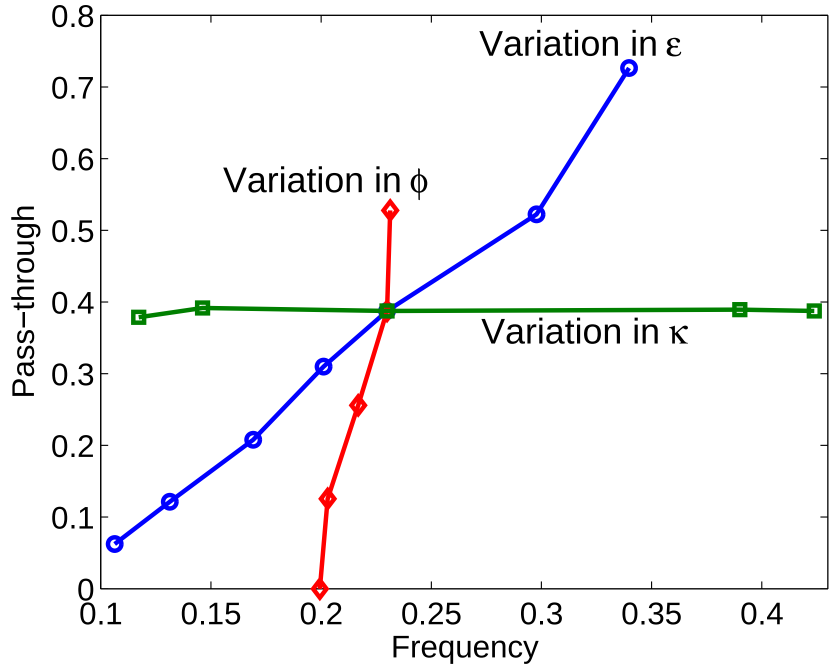

Figure 11: Frequency and Long-run Pass-through: variation in ϕ and κ

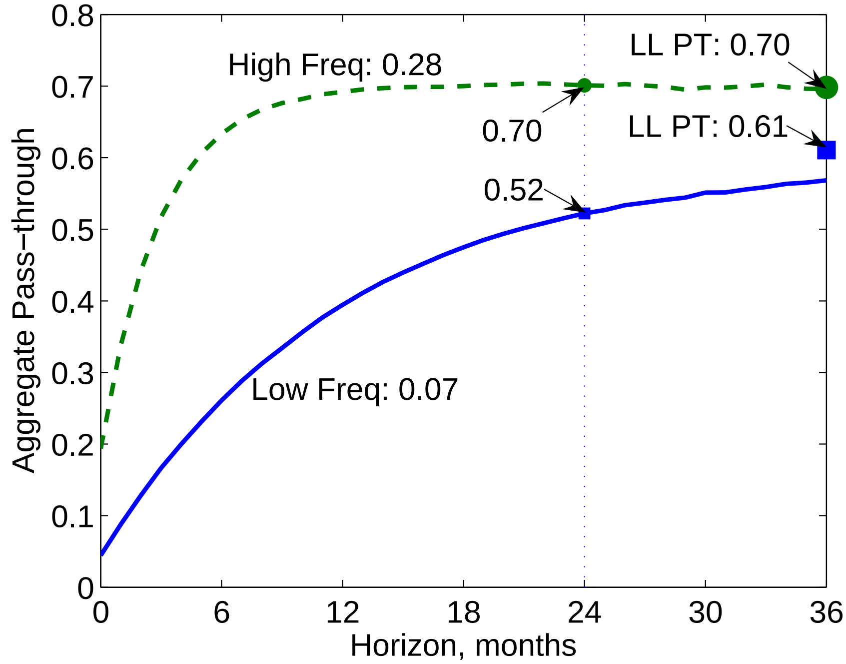

Figure 12: Frequency and Long-run Pass-through in a Calvo model

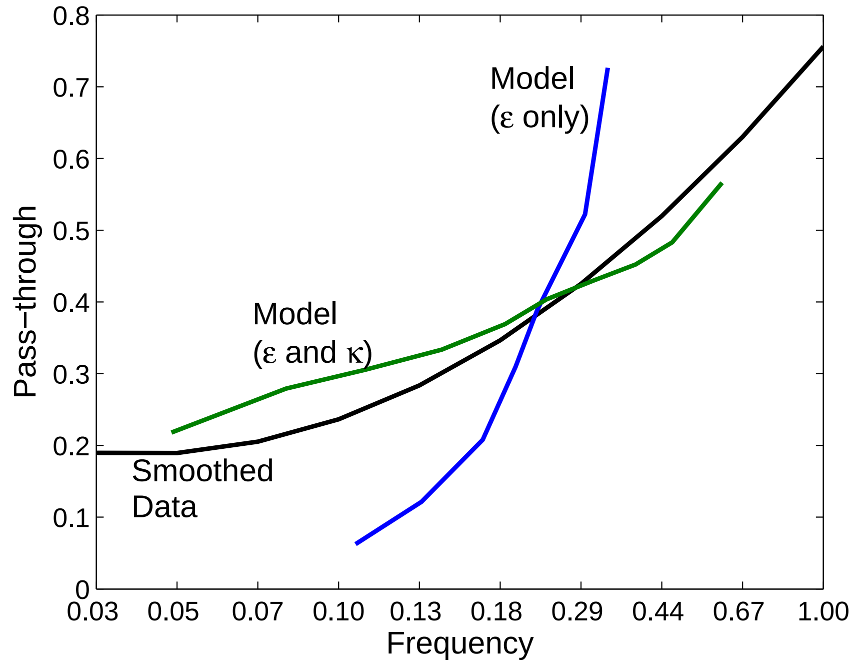

Figure 13: LRPT and frequency: Model against the Data

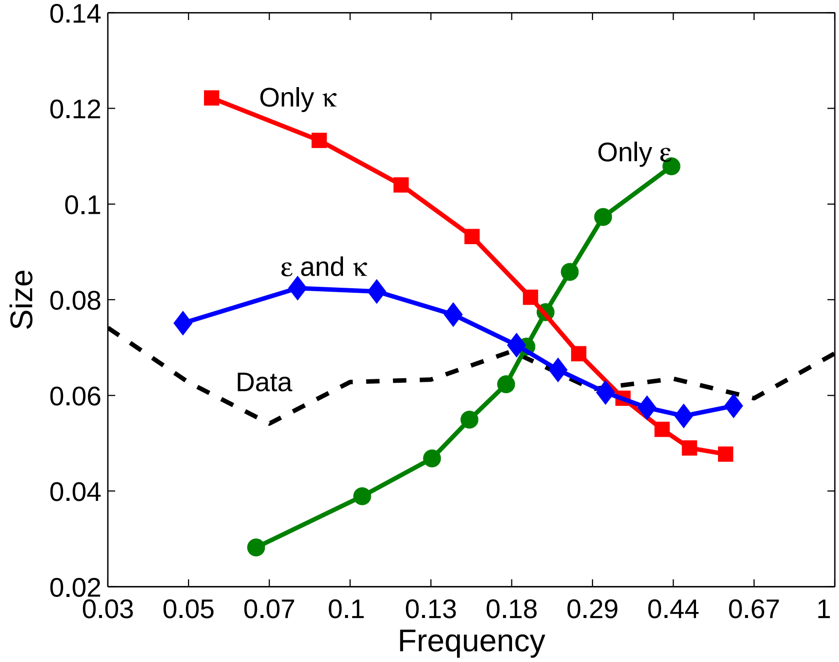

Figure 14: Size and frequency: Model against the Data