

Frequency of Price Adjustment and Pass-through*

Gita Gopinath

Department of Economics,
Harvard University and NBER

Oleg Itskhoki

Department of Economics,
Harvard University

June 30, 2008

Abstract

A common finding across empirical studies of price adjustment is that there is large heterogeneity in the frequency of price adjustment. However, there is little evidence of how distant prices are from the desired flexible price. Without this evidence, it is difficult to discern what the frequency measure implies for the transmission of shocks or to understand why some firms adjust more frequently than others. We exploit the open economy environment, which provides a well-identified and sizeable cost shock namely the exchange rate shock to shed light on these questions. First, we empirically document that high frequency adjusters have a long-run pass-through that is at least twice as high as low frequency adjusters in the data. Next, we show theoretically that long-run pass-through is determined by the same primitives that shape the curvature of the profit function and, hence, also affect frequency. In an environment with variable mark-ups or variable marginal costs, theory predicts a positive relation between frequency and pass-through, as documented in the data. Consequently, estimates of long-run pass-through shed light on the determinants of the duration of prices. 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. Lastly, we calibrate a dynamic menu-cost model and show that variable mark-ups chosen to match the variation in pass-through in the data can generate substantial variation in price duration, equivalent to one third of the observed variation in the data.

*We wish to thank the international price program of the Bureau of Labor Statistics for access to unpublished micro data.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 also thank Loukas Karabarbounis for excellent research assistance. We thank participants at several venues for comments. This research is supported by NSF grant # SES 0617256.

1 Introduction

A common finding across all empirical studies of price adjustment is that there is large heterogeneity in the frequency of price adjustment across 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 or to understand why some firms adjust more frequently than others, 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. When we move to this environment we find that there is indeed a systematic relation between the frequency of price adjustment and exchange rate pass-through. First, we empirically document that high frequency adjusters have a long-run pass-through that is at least twice as high as low frequency adjusters in the data. 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 mark-ups or variable marginal costs, as documented in the data. Consequently, estimates of long-run pass-through shed light on the determinants of the duration of prices. 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. Lastly, we calibrate a dynamic menu-cost model and show that variable mark-ups chosen to match the pass-through in the data can generate substantial variation in price duration, equivalent to one third of the observed variation in 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 estimate the long-run exchange rate pass-through within each bin. We do this in two ways. One, we regress the life-long change in the price of

¹It 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 (2007). 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.

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

the good (relative to U.S. inflation) on the real exchange rate movement over the same period. Two, we estimate an aggregate pass-through regression and estimate the cumulative impulse response of the average monthly change in import prices (relative to U.S. inflation) within each bin to a change in the real exchange rate over a 24 month period. Either procedure generates similar results: When goods are divided into two equal-sized frequency bins, high-frequency adjusters display long-run pass-through that is at least twice as high as low-frequency adjusters.

For the sample of firms in the manufacturing sector, high-frequency adjusters have a pass-through of 40% as compared to low-frequency adjusters with a pass-through of 20%. In the sub-sample of importers in the manufacturing sector from high income OECD countries, high frequency adjusters have a pass-through of 58% compared to 27% for the low frequency adjusters. This result similarly holds for the sub-sample of differentiated goods according to the Rauch (1999) classification. When we split 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.³ Therefore, the data is characterized not only by a positive relationship between frequency and pass-through, but also by a wide range of variation for both variables. Empirically, it is as hard to identify the factors behind the variation in pass-through as it is to explain frequency. Our findings suggest that the variation in exchange rate pass-through can be largely driven by the same unobservable primitives that determine the frequency of price adjustment.

The positive relationship between frequency and pass-through implies the existence of a selection effect. 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. This potentially has important implications for the strength of nominal rigidities given the median durations 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). For instance, the effect we highlight will be present in time-dependent models with optimally chosen periods of non-adjustment as in Ball, Mankiw, and Romer (1988).

Next we analyze the theoretical relation between frequency and long-run pass-through in a static price setting model where long-run pass-through is incomplete and firms pay a menu

³For the all countries sub-sample, pass-through increases from 14% to 45%.

cost to adjust prices in response to cost shocks.⁴ We allow for three standard channels of incomplete long-run exchange rate pass-through: (i) variable mark-ups, (ii) variable marginal costs and (iii) imported intermediate inputs. The first two channels increase the curvature of the profit function. Holding pass-through (i.e., the response of the desired price to shocks) constant this leads to more frequent price adjustments. However, these two channels also limit pass-through which more than offsets the effect of increased curvature of the profit function. Consequently, all else equal, higher long-run pass-through is associated with a higher frequency of price adjustment. The imported intermediate inputs channel reduces the sensitivity of firms to exchange rate shocks and reduces estimated exchange rate pass-through and frequency, all else equal.

The simple analytical model of Section 3 is a useful tool to study the qualitative relationship between variables. However, to assess the quantitative importance of these mechanisms and to evaluate the ability of different models to match the empirical facts we construct and calibrate, in Section 4, a dynamic price-setting model.

The standard model of sticky prices in the open economy assumes CES demand and Calvo price adjustment.⁵ 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 need to depart from this standard set-up. Firstly, we need to allow for endogenous frequency choice: specifically, we construct a menu cost model of state-dependent pricing.⁶ Secondly, we need a source of heterogeneous long-run pass-through that does not arise in the standard CES set-up. The departure we focus on is in the tradition of Dornbusch (1987) and Krugman (1987), which generates incomplete long-run pass-through via the channel of variable mark-ups.⁷ We then quantitatively analyze the performance of a model with these two features in matching the facts in the data. Our setup is most comparable with Klenow and Willis (2006) who in a closed economy model introduce state-

⁴Our 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); Rotemberg and Saloner (1987) and Dotsey, King, and Wolman (1999).

⁵See 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; Atkeson and Burstein (2005) and Gust, Leduc, and Vigfusson (2006) consider an environment with variable mark-ups to examine exchange rate pass-through, but they assume flexible pricing.

⁶We could alternatively model this as a Calvo model where the Calvo parameter is chosen endogenously and this would deliver similar results

⁷This source of incomplete pass-through has received considerable support in the empirical literature, such as Knetter (1989) and other evidence summarized in the paper by Goldberg and Knetter (1997).

dependent pricing and Kimball preferences to generate variable mark-ups. We view this model as an approximation to a setting in which strategic interactions between firms lead to mark-up variability and incomplete pass-through of shocks.

Our calibration exercise confirms that the theoretical link between frequency and pass-through illustrated by the simple two period model of Section 3 holds in a fully dynamic menu cost model. Moreover, we find that variable mark-ups can indeed generate quantitatively large effects and explain a significant share of variation in the frequency of price adjustment. Specifically, when we vary the amount of mark-up variability (by changing the curvature of demand) to match the range of observed long-run pass-through coefficients (10% to 70%) the model predicts a wide range of frequencies which corresponds to variation in price durations between 10 and 3 months. In other words, variation in mark-up variability alone can explain about one third of the observed variation in frequency of price adjustment.⁸

Finally, by estimating the same empirical regressions on the model-generated data, we show that a mechanical relationship between frequency and pass-through while present is extremely weak and, hence, the pure variation in frequency of price adjustment when long-run pass-through is complete cannot account for the observed empirical relationship in most standard models of price setting.

Section 5 concludes. All proofs and a detailed description of the simulation procedure are relegated to the Appendix.

2 Empirical Evidence

In this section we document that firms that adjust prices more frequently have a higher exchange rate pass-through in the long-run. Long-run pass-through is defined to capture pass-through beyond the period when nominal rigidities in price setting are in effect.

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 1994-2005.⁹ Since we are interested in prices that serve an allocative role, we will restrict attention to market transactions and exclude intra-firm

⁸Introducing additional variation in the size of menu costs and the size of cost shocks allows us to fully match the joint behavior of exchange rate pass-through and frequency and size of price adjustment.

⁹For details regarding this data see Gopinath and Rigobon (2007)

transactions. The goal of this analysis is to relate the frequency of price adjustment to the flexible price pass-through of the good, which is the long-run pass-through. For this purpose we need to observe at least one price change. In this database there are 30% goods that have a fixed price during their life. For the purpose of our study these goods are not very useful.¹⁰ For each of the remaining goods we estimate the frequency of price adjustment following the procedure in Gopinath and Rigobon (2007). We then sort goods into high and low frequency bins and estimate long-run pass-through within these bins.

We restrict attention to dollar priced imports in the manufacturing sector.¹¹ Since we are interested in price-setting behavior we will restrict attention to the manufacturing sector 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. 90% of goods imported are priced in dollars. For an 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 5.

To estimate long-run pass-through we use two approaches. The first approach estimates, for each good, 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. $\Delta p_{LR}^{i,c}$ is then defined as the log of this price change relative to U.S. inflation over the same period, where i indexes the good and c the country. $\Delta RER_{LR}^{i,c}$ refers to the cumulative change in the log of the bilateral real exchange rate for country c over this same period. 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.¹² An increase in the RER is a real depreciation of the dollar. Life-long pass-through, β_{LR} , is estimated by the following regression

$$\Delta p_{LR}^{i,c} = \alpha_c + \beta_{LR} \Delta RER_{LR}^{i,c} + \epsilon^{i,c} \quad (1)$$

where α_c is a country fixed effect. A similar regression was used to estimate long-run pass-through in Gopinath, Itskhoki, and Rigobon (2007). Estimating pass-through conditional on only the first price adjustment may not be sufficient to capture long-run pass-through, especially for the higher frequency adjusters, due to the interaction between nominal and real rigidities, which is why we use the concept of life-long pass-through that conditions on multiple rounds of adjustment.

¹⁰We will revisit this at the end of this section when we comment on item substitution.

¹¹Goods that have a one digit SIC code of 2 or 3. We exclude any petrol classification codes.

¹²The 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.

The second approach measures long-run pass-through 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, $\Delta p_{c,t}$, and the monthly bilateral real exchange rate movement vis-à-vis the dollar for that country, $\Delta RER_{c,t}$. We then estimate a stacked regression where we regress the average monthly change in prices on monthly lags of the real exchange rate change:

$$\Delta p_{c,t} = \alpha_c + \sum_{j=0}^n \beta_j \Delta RER_{c,t-j} + \epsilon_{c,t}, \quad (2)$$

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

Before we proceed to describing the results we briefly comment on the two approaches. First, we use the real specification in both regressions to be consistent with the regressions we run on the model generated data in Section 4.3. However, the empirical results are insensitive to alternative specifications such as regressing the nominal price change on the change in the nominal exchange rate with controls for foreign and U.S. inflation.

Second, a standard assumption in the empirical pass-through literature is that movements in the real 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 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 long-run pass-through 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 change prices is chosen endogenously, 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 which reduces the size of the

selection bias. In Section 4.3 we estimate the same regressions on the data generated from conventional models of sticky prices, both menu cost and Calvo, and verify that both of these regressions indeed deliver estimates close to the true theoretical long-run pass-through.

2.2 Evidence

In Table 1 we report the results from estimating the life-long equation (1), when the goods are sorted into high and low frequency bins. In Panel A, the first price refers to the first observed price for the good and in Panel B, the first price refers to the first new price for the good. In both cases, the last price is the last new price.¹³ The main difference in the results between Panel A and B relates to the number of observations, since there are goods with only one price adjustment during their life. Otherwise, the main results are unchanged. For each frequency bin, the second column of Table 1 reports the median frequency, the third column reports the long-run pass-through estimate (β_{LR}) and the fourth column reports the robust standard error for this estimate ($\sigma(\beta_{LR})$) clustered at the level of country interacted with the BLS defined primary strata of the good.¹⁴ Finally, N in the fifth column is the number of goods in each sub-sample.

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-samples, goods adjust prices on average every 14 months and pass-through only 20% in the long run; at the same time, in the high-frequency sub-sample, goods adjust prices every 3 months and pass-through 40% in the long-run. This is more strongly the case when we restrict attention to the high-income OECD sample: long-run pass-through increases from 27% to 58% as we move from the low to the high frequency sub-sample. We also look at the manufacturing goods sub-sample that can be classified to be in the differentiated goods sector, following Rauch’s classification.¹⁵ For differentiated goods, moving from low to high frequency bin raises long-run pass-through from 19% to 40% for goods from all source countries and from 26% to 58% in the high-income OECD sample.¹⁶ In all cases, pass-through estimates across

¹³For the hypothetical item in Figure 1, Panel A would use observations in $[0, t_2]$, while Panel B would use observations only in $[t_1, t_2]$.

¹⁴This refers to mostly 3 and 4 digit harmonized codes

¹⁵Rauch (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.

¹⁶Only a subset of manufactured goods can be classified using the Rauch classification. Consequently, it must not be interpreted that the difference in the number of observations between manufactured and the sub-group of manufactured and differentiated represent non-differentiated goods. In fact, using Rauch’s

the frequency bins are statistically different at conventional levels of significance. All the results hold similarly for the Panel B regressions, with a somewhat larger difference in pass-through estimates across the frequency bins. Since the results are similar for the case where we start with the first price as opposed to the first new price, for the rest of the specifications we report the results for the Panel A case only.

Since, it can be argued that a single price adjustment may be insufficient to capture long-run pass-through, especially in a world with real rigidities, as a sensitivity check we 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 2. As expected the median frequency of 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.

In Table 3 we perform the analysis at the country/region level. Here again we note the twice higher long-run pass-through for the high frequency adjusters as compared to the low frequency adjusters. Since the samples get much smaller the significance levels drop.¹⁷

To ensure that this positive relationship between frequency and pass-through exists even when the number of bins is increased we estimate the same regression with 10 bins. The point estimates and 10% 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 and most strongly for the high-income OECD sample. For the high-income OECD sub-sample long-run pass-through ranges from around 18% to 75%, as frequency ranges from 0.03 to 1.¹⁸ 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. Secondly, 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 the concerns that the relation between frequency and pass-through is driven by insufficient number of price adjustments for the low-frequency goods.

The next set of results relates to the estimates from the aggregate pass-through regressions defined in (2). We again split the goods into two bins based on frequency of price adjustment

classification only a 100 odd goods are classified as non-differentiated.

¹⁷The non-high income OECD sample has a sizeable number of observations, nevertheless the difference in pass-through is not significant. This highlights the fact that the main result is most strongly evident for the high income OECD sub-sample.

¹⁸For the all country sample the long-run pass-through range is between 14% and 45%.

and estimate the aggregate pass-through regressions separately for each of the bins. 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% 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 5 we report the results by country/region and for goods with 3 or more price adjustments. Here again we find the same result. The estimates in these samples, however, become very noisy.

Product Replacement: For the above analysis we estimated long-run pass-through 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. To address this concern we report in Table 4 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 look only at those goods that got 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 measure within each bin. As can be seen, if anything, there is a negative relation

¹⁹Specifically this refers to the following discontinuation reasons: “Out of Business”, “Out of Scope, Not replaced” and “Out of Scope, Replaced”.

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 adjusts the frequency of price adjustment to include the probability of 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. 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. This illustrates the difficulty of using measures such as size in the analysis of frequency. We discuss this issue later in the paper.

3 A Static Model of Frequency and Pass-through

In this section we investigate theoretically the relation between pass-through 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 pass-through of cost shocks. We consider the problem of a single firm that fixes 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 three standard sources of incomplete pass-through into the model: variable mark-ups, variable marginal costs and imported inputs. We show that all else equal, higher pass-through is associated with a higher frequency of price adjustment.

²⁰We 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.

²¹Our 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 we introduce the *super-elasticity*²³ of demand, or the elasticity of elasticity, as

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

$\tilde{\sigma}(\cdot)$ is the effective elasticity of demand for the firm that takes into account both direct and indirect effects from price adjustment.²⁴ Note that we introduce mark-up variability 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.²⁵

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 .²⁶ These normalizations will be useful later when we approximate the solution around $p = 1$.

The firm operates a production technology characterized by the cost function:

$$C(q|a, e; \eta, \phi) = (1 - a)(1 + \phi e)c(q|\eta),$$

where a is an idiosyncratic productivity shock and e is a real exchange rate shock. We will refer to the pair (a, e) as a marginal cost shock of 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 exchange rate shock and η is a parameter governing the degree of returns to scale of the

²²Since 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.

²³We use the terminology of Klenow and Willis (2006).

²⁴For 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.

²⁵Atkeson and Burstein (2005) 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 the firms.

²⁶Klenow and Willis (2006) design an example of such a demand function: $\varphi(p|\sigma, \varepsilon) = A[1 - \varepsilon \ln p]^{\sigma/\varepsilon}$.

production technology. The larger is η , the stronger are the diminishing returns to scale in production and, hence, the more convex is the cost function. The marginal cost of the firm is then given by

$$MC(q|a, e; \eta, \phi) = (1 - a)(1 + \phi e)mc(q|\eta),$$

where $mc(q|\eta) \equiv \partial c(q|\eta)/\partial q$. We denote the elasticity of marginal cost with respect to quantity by:

$$\tilde{\eta} \equiv \tilde{\eta}(q|\eta) = \frac{\partial \ln mc(q|\eta)}{\partial \ln q}.$$

We introduce the following normalization for the marginal cost schedule: When the quantity produced by the firm is equal to one ($q = 1$), $\tilde{\eta}(\cdot)$ is equal to η (that is, $\tilde{\eta}(1|\eta) = \eta$) and $\tilde{\eta}(\cdot)$ is increasing in η for all q . Additionally, we normalize the level of the marginal cost so that $mc(1|\eta) = (\sigma - 1)/\sigma$. Under this normalization, the optimal flexible price of the firm when $a = e = 0$ is equal to 1, as we show below. Intuitively, the marginal cost is set 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 more on this see Rotemberg and Woodford, 1999).

Finally, the profit function of the firm is given by:²⁷

$$\Pi(p|a, e) = p\varphi(p) - C(\varphi(p)|a, e). \quad (3)$$

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 $\Pi(a, e) \equiv \Pi(p(a, e)|a, e)$.

3.2 Price Setting

For a given cost shock (a, e) , the desired flexible price maximizes profits (3) so that²⁸

$$p_1 \equiv p(a, e) = \frac{\tilde{\sigma}(p_1)}{\tilde{\sigma}(p_1) - 1} (1 - a)(1 + \phi e)mc(\varphi(p_1)), \quad (4)$$

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 :

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

²⁷From now on we suppress the explicit dependence on parameters σ , ε , η and ϕ .

²⁸Note that this condition constitutes a fixed point problem for p_1 . The sufficient condition for maximization is $\tilde{\sigma}(p_1) > 1$ provided that $\tilde{\varepsilon}(p_1) \geq 0$ and $\tilde{\eta}(\varphi(p_1)) \geq 0$. We assume that these inequalities are satisfied for all p .

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

$$\Delta \equiv \Delta_\kappa = \{(a, e) : L(a, e) \leq \kappa\}.$$

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 does not reset its price:²⁹

$$\bar{p}_0 = \arg \max_p \int_{(a,e) \in \Delta} \Pi(p|a, e) dF(a, e),$$

where $F(\cdot)$ denotes the joint cumulative distribution function of the cost shock (a, e) .³⁰ Using the linearity of the profit function in costs, we can re-write the *ex ante* problem of the firm as

$$\bar{p}_0 = \arg \max_p \left\{ p\varphi(p) - \mathbb{E}_\Delta \{(1-a)(1+\phi e)\} \cdot C(\varphi(p)) \right\}, \quad (5)$$

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.³¹

Proof: See Appendix. ■

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. The optimality condition (4) implies that, given our normalization of the marginal cost and elasticity of demand, $p(0, 0) = 1$. This condition also results in:

Lemma 2 *The following first order approximation holds,*

$$\frac{p(a, e) - \bar{p}_0}{\bar{p}_0} \approx \Psi \cdot (-a + \phi e), \quad (6)$$

where

$$\Psi \equiv \frac{1}{1 + \frac{\epsilon}{\sigma-1} + \sigma\eta}. \quad (7)$$

²⁹We implicitly assume, as is standard in a partial equilibrium approach, that the stochastic discount factor is constant for the firm.

³⁰Formally, $L(a, e)$ and, hence, Δ depend on the preset price \bar{p}_0 . Therefore, this expression for \bar{p}_0 is implicit and constitutes a fixed point problem.

³¹The appendix makes precise what these second order terms are. The standard size of the shocks, as well as the size of the menu cost are natural benchmarks as they determine how far a price can be from its desired level. All approximations in this section become exact as typical cost shocks and menu costs tend to zero.

Proof: See Appendix. ■

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 elasticity can be smaller than one because either mark-ups adjust to limit the response of the price to the shock, or the marginal cost adjusts to limit the movement of the cost. The elasticity of mark-up with respect to price is given by

$$\left. \frac{\partial \tilde{\mu}(p)}{\partial \ln p} \right|_{p=1} = - \left. \frac{\tilde{\varepsilon}(p)}{\tilde{\sigma}(p) - 1} \right|_{p=1} = - \frac{\varepsilon}{\sigma - 1},$$

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

$$\left. \frac{\partial \ln mc(\varphi(p))}{\partial \ln p} \right|_{p=1} = \left. \frac{\partial \ln mc(q)}{\partial \ln q} \right|_{q=1} \cdot \left. \frac{\partial \ln \varphi(p)}{\partial \ln p} \right|_{q=\varphi(p)} = \tilde{\eta}(\varphi(p)) \cdot \left. (-\tilde{\sigma}(p)) \right|_{p=1} = -\sigma\eta.$$

Higher price reduces demand and, therefore, reduces marginal cost if there are decreasing returns to scale. Overall, a one percent increase in price leads to a $\varepsilon/(\sigma - 1)$ percent reduction in desired mark-up and a $\sigma\eta$ percent reduction in marginal cost. As a result, the desired price increases only by Ψ percent in response to a 1 percent cost shock.

3.3 Pass-through and Frequency of Adjustment

We now introduce exchange rate pass-through, which is the elasticity of the firm's desired price with respect to the exchange rate shock. Formally, it is defined as³²

$$\Psi_e \equiv \left. \frac{\partial \ln p(a, e)}{\partial \ln(1 + e)} \right|_{a=e=0}.$$

Then Lemma 2 has the following

³²One can use an alternative — empirically-motivated — definition of pass-through. If one observes desired prices for all values of cost shocks, then pass-through can be defined as

$$\hat{\Psi}_e \equiv \text{cov}(\ln p(a, e) - \ln \bar{p}_0, e) / \text{var}(e).$$

Lemma 2 and the symmetry of exchange rate shocks distribution imply that these two definitions of pass-through are first-order equivalent, i.e. $\hat{\Psi}_e \approx \Psi_e$ holds up to second-order terms. If, however, desired prices are observable only conditional on adjustment, this induces a negative correlation between a and e (see (12) below) — a selection effect — which biases upwards the regression based pass-through (conditional on adjustment). The way we deal with the selection issues in the data is by increasing the window of the pass-through regression to include a number of price adjustments; mean reversion of productivity shocks assures then that the selection bias is small. We verify that this is the case when we estimate the empirical regressions on the model-generated data in Section 4.

Corollary 1 *Exchange rate pass-through equals*

$$\Psi_e = \phi\Psi = \frac{\phi}{1 + \frac{\varepsilon}{\sigma-1} + \sigma\eta}. \quad (8)$$

Proof: See Appendix (proof of Lemma 2'). ■

Intuitively, this corollary is a direct implication of (6). Observe from (8) that exchange rate pass-through is increasing in cost sensitivity to the exchange rate, ϕ , and decreasing in the super-elasticity of demand, ε , and elasticity of the marginal cost, η . These summarize the three channels of incomplete pass-through in the model. Ψ_e is in general non-monotonic in the elasticity of demand, σ . Specifically, higher elasticity of demand leads to higher pass-through if and only if

$$\frac{\varepsilon}{\sigma-1} > (\sigma-1)\eta, \quad (9)$$

Higher σ amplifies the marginal cost channel and attenuates the mark-up channel which results in the non-monotonic effect. We summarize these findings in:

Proposition 1 *Exchange rate pass-through, Ψ_e , depends uniquely on $\{\sigma, \varepsilon, \eta, \phi\}$. It is increasing in ϕ and decreasing in ε and η . It increases in σ if and only if condition (9) is satisfied.*

We now examine how variation in these parameters affects 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

$$\Phi \equiv 1 - \Pr\{(a, e) \in \Delta\} = \Pr\{L(a, e) > \kappa\}, \quad (10)$$

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

To make further progress in characterizing the region of non-adjustment, Δ , we use the second order approximation to $L(a, e)$ provided in

Lemma 3 *The following second order approximation holds:*

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

where Ψ is again as defined in (7).

Proof: See Appendix. ■

Note that Lemma 3 implies that the curvature of the profit function,

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

increases in σ , ε and η . That is, higher elasticity of demand, higher variability of mark-ups and of marginal costs increase 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 variability of mark-ups and marginal costs also limits pass-through which may more than offset the effect of increased curvature of the profit function. Indeed, combining the results of Lemmas 2 and 3, we arrive at our final approximation to the profit loss function:

$$L(a, e) \approx \frac{1}{2}(\sigma - 1)\Psi(-a + \phi e)^2, \quad (11)$$

which again holds up to third order terms (see Appendix). This expression makes it clear that similar forces that reduce pass-through (i.e., decrease Ψ and ϕ) also reduce the curvature of the profit function (in the space of the primitive shocks) and, thus, limit the profit loss from not adjusting prices and, as a result, lead to lower frequency of price adjustment. We illustrate these effects in Figure 7 for particular demand and cost functions, but without recurring to approximations.

Combining (11) and (10) we have

$$\Phi \approx \Pr \left\{ |-a + \phi e| > \sqrt{2\kappa / [(\sigma - 1)\Psi]} \right\},$$

or equivalently,

$$\Phi \approx \Pr \left\{ |X| > \sqrt{\frac{2\kappa}{(\sigma - 1)\Psi\Sigma}} \right\}, \quad (12)$$

where $X = (-a + \phi e)/\sqrt{\Sigma}$ is the 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, Φ , decreases with the variability of mark-ups ε and the degree of decreasing returns to scale η . It increases with the sensitivity of costs to exchange rate shocks ϕ . It also decreases with the menu cost κ and increases with the elasticity of demand σ and the size of the shocks σ_a and σ_e .*

Combining the results of Propositions 1 and 2, we conclude that:

Proposition 3 (i) *Variable mark-ups and marginal costs, as well as lower sensitivity of cost to exchange rate shocks reduce both frequency of price adjustment and pass-through;* (ii) *Higher elasticity of demand increases frequency of price adjustment and may increase or decrease pass-through;* (iii) *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 variation in mark-up and marginal cost variability across the goods is important, we should observe a positive cross-sectional correlation between frequency and pass-through.^{33,34}

4 Dynamic Model

We now consider a fully dynamic specification with state dependent pricing and variable mark-ups. We adopt a partial equilibrium approach by focusing on the industry equilibrium in the U.S. market. We show that the positive relation between frequency and long-run pass-through is obtained in the dynamic setting and when we choose parameters to match the variation in pass-through observed in the data we obtain price durations that range from 3 months to 10 months — about one third of the empirical variation in frequency documented in Figure 2 and Table 4. We then verify that the observed correlation between frequency and pass-through cannot be explained by models with exogenous differences in frequency of price adjustment and no variation in long-run pass-through.

The variable mark-ups channel of incomplete pass-through is motivated by the theoretical work of Dornbusch (1987) and Krugman (1987) and the empirical evidence supporting this channel in Knetter (1989) and Goldberg and Knetter (1997).^{35,36} We introduce the variable

³³Similarly, variation in sensitivity of marginal cost to exchange rate, ϕ , can also account for the positive relationship between frequency and pass-through. However, the effect of ϕ on frequency is limited by the ratio of the variances of the exchange rate shock and productivity shock, σ_e^2/σ_a^2 . To see this note from (12) that frequency increases in Σ and $\Sigma = \sigma_a^2(1 + \phi^2\sigma_e^2/\sigma_a^2)$. Empirically, σ_e^2/σ_a^2 is small; see calibration of a dynamic model in the next section, where we show that the effect of ϕ on frequency is negligible.

³⁴In the Appendix we show additionally that the average size of price adjustment is generally increasing in κ , Ψ and Σ . Recall that frequency is decreasing in κ , but increasing in Ψ and Σ . Therefore, as long as there is variation across goods in both κ and Ψ or Σ , one should not expect to see a robust correlation between frequency and size.

³⁵This channel has been further explored in recent quantitative work in the open economy literature by Atkeson and Burstein (2005) and Gust, Leduc, and Vigfusson (2006). However, these papers assume flexible price setting. Bergin and Feenstra (2001) allow for variable mark-ups in an environment with price stickiness, but they assume exogenous periods of non-adjustment.

³⁶We shut down the variable marginal cost channel of incomplete pass-through. The rationale for this is the following: variable marginal cost channel is observationally equivalent to the variable markups channel from the point of view of pass-through, however, it does not generate law of one price violations which is a

mark-up channel of incomplete pass-through using Kimball (1995) kinked demand which we view as an approximation to a setting in which strategic interactions between large firms lead to mark-up variability and incomplete pass-through of shocks. Our setup is most comparable with Klenow and Willis (2006) with the distinction that we have exchange rate shocks that are more idiosyncratic than the aggregate shocks typically considered in the literature.

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 sectoral equilibrium.

4.1.1 Industry Demand Aggregator

The industry is characterized by a continuum of varieties indexed by j . There is a measure 1 of U.S. varieties and a measure $\omega < 1$ of foreign varieties available for domestic consumption. The smaller share of foreign varieties captures the feature in the data of home-bias in consumption.

The Kimball (1995) consumption aggregator is given by

$$\frac{1}{|\Omega|} \int_{\Omega} \Psi \left(\frac{|\Omega| C_j}{C} \right) dj = 1 \quad (13)$$

with $\Psi(1) = 1$, $\Psi'(\cdot) > 0$ and $\Psi''(\cdot) < 0$. C_j is the consumption of the differentiated variety $j \in \Omega$, where Ω is the set of varieties available for consumption in the home country with measure $|\Omega| = 1 + \omega$. Individual varieties are aggregated into sectoral consumption level, C , which is implicitly defined by (13).

Consumers maximize C given the prices of varieties $\{P_j\}$ and income E allocated for industry consumption. The demand function for individual varieties is then given by:

$$\Psi' \left(\frac{|\Omega| C_j}{C} \right) = D \frac{P_j}{P},$$

where $D \equiv \int_{\Omega} \Psi' \left(\frac{|\Omega| C_j}{C} \right) \frac{C_j}{C} dj$ and P is the sectoral price index that satisfies the condition

$$E = PC = \int_{\Omega} P_j C_j dj \quad (14)$$

salient feature of international price data.

since the aggregator in (13) is homothetic. The demand for a particular variety can be expressed as

$$C_j = \psi \left(D \frac{P_j}{P} \right) \cdot \frac{C}{|\Omega|}, \quad \psi(\cdot) \equiv \Psi'^{-1}(\cdot). \quad (15)$$

4.1.2 Firm's Problem

Consider a representative home firm j . Everything holds symmetrically for foreign firms and we superscript foreign variables with an asterisk. In each period the firm produces a unique variety j of the differentiated good given a constant marginal cost

$$MC_t = \frac{W_t^{1-\phi} (W_t^*)^\phi}{A_t}. \quad (16)$$

A_j denotes the idiosyncratic productivity shock which follows an autoregressive process in logs:³⁷

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

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.³⁸

The profit function of the home firm producing variety j in period t is:

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

where demand C_{jt} is given by (15). Firms are price setters and must satisfy demand at the posted price. In what follows we will interpret the domestic wage, W_t , as the numéraire and assume that both domestic and foreign firms set prices in the units of the domestic wages. This is the model equivalent of local currency pricing in a world without money. To change the price, both domestic and foreign firms must pay a menu cost κ , also in terms of domestic wages.

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

³⁷In what follows corresponding small letters denote the logs of the variables.

³⁸The marginal cost in (16) can be derived from a constant returns to scale production function which combines domestic and foreign inputs.

equations for the firm is given by³⁹

$$\begin{cases} V^N(\mathbb{S}_{jt}) &= \Pi(P_{j,t-1}) + \mathbb{E}\{Q(\mathbb{S}_{j,t+1})V(\mathbb{S}_{j,t+1})|\mathbb{S}_{jt}\}, \\ V^A(\mathbb{S}_{jt}) &= \max_{P_{jt}} \left\{ \Pi(P_{jt}) + \mathbb{E}\{Q(\mathbb{S}_{j,t+1})V(\mathbb{S}_{j,t+1})|\mathbb{S}_{jt}\} \right\}, \\ V(\mathbb{S}_{jt}) &= \max \{V^N(\mathbb{S}_{jt}), V^A(\mathbb{S}_{jt}) - \kappa_{jt}\}. \end{cases} \quad (17)$$

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

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

Therefore, the policy function of the firm—the optimal price adjustment policy—is:

$$P(\mathbb{S}_{jt}) = \begin{cases} P_{j,t-1}, & \text{if } V^N(\mathbb{S}_{jt}) \geq V^A(\mathbb{S}_{jt}) - \kappa_{jt}, \\ \bar{P}(\mathbb{S}_{jt}), & \text{otherwise.} \end{cases} \quad (18)$$

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

We assume an exogenous process for the domestic and foreign wage rates, W_t and W_t^* and define W_t^*/W_t to be the (wage-based) real exchange rate.⁴⁰ The domestic wage is assumed to be the numéraire. The sectoral equilibrium is then determined by the equilibrium 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 $\{W_t, W_t^*\}$. The simulation procedure of the model is discussed in detail in Appendix B.

4.2 Calibration

In this section we briefly discuss the calibration of the parameters of the model, while the details are reported in Appendix B. We adopt the Klenow and Willis (2006) specification of

³⁹We abuse the notation here somewhat since, in general, one should condition expectations in the Bellman equation on the whole history $(\mathbb{S}_{jt}, \mathbb{S}_{j,t-1}, \dots)$. In our simulation procedure we will assume that \mathbb{S}_{jt} is a sufficient statistic.

⁴⁰In a model with a large share of non-tradable goods in consumption, this measure of the real exchange rate will be close to a CPI-based real exchange rate.

the Kimball aggregator (13) which results in

$$\psi(x_j) = \left[1 - \varepsilon \ln \left(\frac{\sigma x_j}{\sigma - 1} \right) \right]^{\sigma/\varepsilon}, \quad x_j \equiv D \frac{P_j}{P}. \quad (19)$$

This demand specification is conveniently governed by two parameters, $\sigma > 1$ and $\varepsilon > 0$, and the elasticity and super-elasticity are given by:⁴¹

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

Note that this demand function satisfies all normalizations assumed in Section 3.

We now briefly discuss the calibration of the main parameters of the model. The calibrated parameters are also summarized in Table 5. The period in the model corresponds to one month. We set the measure of foreign firms, ω , to equal 0.2 which implies that around 17% (i.e., $\omega/(1 + \omega)$) of the goods in the domestic consumption bundle are imported. This number is consistent with the share of imports from U.S. input-output tables. We calibrate the fraction of foreign costs that are not sensitive to the real exchange rate movement to be $\phi = 0.75$, which is consistent with OECD input-output tables. The idea is to allow for the fact that a fraction of foreign firm inputs can be priced and stable in dollars and therefore not sensitive to exchange rate movements. For the U.S. firms we set $\phi = 0$ to capture the fact that 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 $\bar{\phi} = \phi \cdot \omega/(1 + \omega) = 12.5\%$ which also is the long-run exchange rate pass-through into the sectoral price level (see below).

The menu cost, κ , in the baseline calibration is set to 2.5% of the revenues conditional on adjustment. In the simulation, firms adjust prices on average once in 8 months, which means that overall menu costs constitute around 0.2% of revenues on an annual basis, well within the range used in the literature. The (log of) the real exchange rate, $e \equiv \ln(W_t^*/W_t)$, follows a very persistent process with the autocorrelation of around 0.97.⁴² The monthly innovation to the real exchange rate is calibrated to equal $\sigma_e = 2.5\%$. This values for both persistence and volatility of the real exchange rate are consistent with the empirical patterns for the developed countries. We next calibrate the variance and persistence of the idiosyncratic shock process. Since we model productivity shocks, we calibrate them to be fairly persistent with a monthly auto-regression coefficient of $\rho_a = 0.95$.⁴³ We set the standard deviation of the

⁴¹A useful feature of this demand specification is that it converges to CES demand with elasticity σ when $\varepsilon \rightarrow 0$.

⁴²Since we use the domestic wage as a numéraire and unit of account, we need to specify only the process for the relative wage, W_t^*/W_t , without specifying the processes for the individual wages, W_t and W_t^* .

⁴³This is as in Bilal and Klenow (2004).

innovation in productivity to 8.5%. This is about 3 times as large as shocks to productivity at the monthly frequency. The standard deviation is chosen to match the median size of price adjustment of 7% conditional on adjustment. Since the shocks that impact a firm include both cost and demand shocks it is reasonable to set the standard deviation at this higher number.⁴⁴ Note that our calibration implies a low relative variance of the exchange rate shock, $\sigma_a^2/\sigma_e^2 \ll 1$.

Finally, we discuss the demand parameters. In the baseline calibration, we set the steady state elasticity of demand, σ , to 5, which implies a steady state markup of 25%. This parameterization is consistent with the literature that estimates demand elasticity at the disaggregated sectoral level. Since there are no available measures of super-elasticity of demand, our approach is to choose a range of values for ε to match the observed pass-through range in the data. Specifically, we choose the following values for the super-elasticity of demand: $\{0, 2, 4, 6, 10, 20, 40\}$. This implies the variation in long-run pass-through between roughly 10% and 75%, taking into account that $\phi = 0.75$ bounds the long-run pass-through from above.

4.3 Simulation Results

In this section we report the results from simulating the model. Recall that we choose values of ε to match the observed range of long-run pass-through in the data ($[0.1, 0.75]$), $\varepsilon \in \{0, 2, 4, 6, 10, 20, 40\}$. We then simulate the dynamic stationary equilibrium of the model for each value and compute the frequency of price adjustment and long-run pass-through for all firms in the industry and then separately for domestic and foreign firms. Figure 8 plots the resulting relationship between frequency and pass-through for these three groups of firms.⁴⁵

From Figure 8 we observe that for both domestic and foreign firms the range of frequencies corresponding to the assumed variation in super-elasticity of demand, ε , is approximately 10 to 3 months. This amounts to roughly one third of the variation in frequency in the data. Next, note the strong positive relationship between frequency and long-run pass-through for the foreign firms: pass-through increases from below 10% to over 75% as frequency increases

⁴⁴Note that in our calibration we need to assume neither very large menu costs, nor very volatile idiosyncratic shocks, as opposed to Klenow and Willis (2006). There are a few differences between our calibration and that of Klenow and Willis (2006). They assume a much less persistent idiosyncratic shock process and match the standard deviation of relative prices rather than the average size of adjustment. The interaction of these two deviations appears to drive the differences in the results.

⁴⁵For the long-run pass-through estimates we plot the coefficients from the life-long regression (1). They are, however, very close to the estimates from the aggregate regression (2). See discussion below.

from 0.10 to 0.35.

We further observe that the relationship between frequency and pass-through is effectively absent for domestic firms, as well as within the full sample of all firms in the industry: as frequency varies between 0.1 and over 0.3, long-run pass-through estimate for domestic firms fluctuates between 0% and 5%, while for the full sample of firms it lies between 5% and 12%. This finding is consistent with empirical estimates of pass-through into producer and consumer prices. This supports our emphasis on at-the-dock prices and why international data provides a meaningful environment in which to study the determinants of frequency.

We now briefly discuss the performance of the two long-run pass-through estimators — the one based on life-long regression (1) and the other based on 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 life-long pass-through 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.⁴⁶ 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.⁴⁷ The frequency of price adjustment in these three

⁴⁶The theoretical flexible price pass-through can be approximated by $\bar{\phi} + \Psi(\phi - \bar{\phi})$, where $\bar{\phi} \equiv \omega\phi/(1+\omega)$ is the average sectoral sensitivity of the firms marginal cost to exchange rate and Ψ is as defined in Section 3.

⁴⁷Note that in our menu cost model the convergence is very fast and is almost over by the end

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 other words, variation in mark-up elasticity holding other things equal leads to a positive correlation between frequency and pass-through even in the short-run.

This far we only considered variations in ε . We next argue that variations in ϕ and in κ alone cannot quantitatively explain the findings in the data. For this exercise we set $\varepsilon = 4$ and first vary ϕ between 0 and 1 for the baseline 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 pass-through, however, the range of variation in frequency is negligible, as predicted in Section 3 (see footnote 33).⁴⁸ As ϕ increases from 0 to 1, long-run pass-through increases from 0 to 55%, while frequency increases from 0.20 to 0.23 only. 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 long-run pass-through, 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.

We now contrast the predictions of our model with the data. Specifically, Figure 12 plots the relationship between frequency and long-run pass-through from our simulation in which we vary the super-elasticity of demand, ε , and the smoothed data series from Figure 2 in which we plotted the empirical estimated long-run pass-through by 10 frequency bins. Despite the fact that the model generates substantial variation in frequency, the relationship between frequency and pass-through in the model is much steeper than in the data. This indicates that additional exogenous sources of variation in frequency, such as differences in menu costs or the sizes of idiosyncratic shocks, are required to match the data. The additional sources will flatten the model relationship between frequency and pass-through and bring the model closer to the data.^{49,50}

of 6 months. This contrasts with much slower dynamics in the data. As we already emphasized in Gopinath, Itskhoki, and Rigobon (2007), 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.

⁴⁸Recall that the impact of ϕ on frequency is bounded by the ratio of the variances of exchange rate and idiosyncratic productivity shocks which has to be small to match the moments in the data (i.e., real exchange rate volatility and average size of price adjustment).

⁴⁹The reason is that in the data we sort the goods based on frequency. Therefore, 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 . This flattens out the relationship between frequency and pass-through.

⁵⁰Also note that a model with joint variation in menu cost and super-elasticity of demand can easily match the flat average size of price adjustment across frequency bins. When the only variation in frequency

Lastly, we present 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 long-run pass-through even though the true long-run pass-through 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 (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 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 exogenous differences in the frequency of adjustment would have difficulty in matching the facts in standard model environments.

5 Discussion and Conclusion

To conclude, we exploit the open economy environment with an observable and sizeable cost shock, namely the exchange rate shock to shed light on the question of what drives the frequency of price adjustment and how distant prices are from their desired levels. 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

comes from the menu cost, the size of price adjustment *decreases* from 12% to 6% as frequency increases. On opposite, when the variation comes only from super-elasticity of demand, the size of price adjustment *increases* from 5% to 14% as frequency increases. Joint variation in menu cost and super-elasticity of demand leads to this two opposing effect on size canceling out and allows to match a nearly flat size of price adjustment of around 7%. These results are available from authors upon request.

have high pass-through drift farther away from their optimal price and, therefore, make more frequent adjustments. The implication of this finding for the quantitative importance of nominal rigidities given the median durations of prices in the economy is left to future research.

In this paper we only considered dollar priced goods and endogenous frequency choice. The currency choice decision was the subject of the analysis in Gopinath, Itskhoki, and Rigobon (2007). We briefly summarize here the link between the two papers. The main finding of Gopinath, Itskhoki, and Rigobon (2007) is that non-dollar priced goods display higher pass-through conditional on the first adjustment to the exchange rate shock, as compared to dollar priced goods. In addition, Gopinath and Rigobon (2007) document that non-dollar priced goods have on average longer duration of prices (14 vs. 11 months) as well as less variation in duration as compared to dollar priced goods. These two features explain the negative correlation between frequency and pass-through documented in Table 11 of Gopinath and Rigobon (2007) as it is mainly driven by the pass-through difference between dollar and non-dollar priced goods.⁵¹ If the model in the current paper were extended to include endogenous currency choice, then the data would imply that non-dollar priced goods have higher menu costs or smaller sizes of idiosyncratic shocks. This way, non-dollar priced goods would have longer price durations despite their higher desired pass-through which determines their currency choice. Finally, the findings in this paper further corroborate the result in Gopinath, Itskhoki, and Rigobon (2007) that what matters for currency choice is medium-run pass-through, while in the long-run pass-through can be high even for the dollar-priced goods.

⁵¹Another difference is that in Gopinath and Rigobon (2007) pass-through is determined conditional on only the first adjustment as opposed to the long-run pass-through we estimate in this paper.

A Results for Section 3

In this appendix we first formally prove Lemmas 1–3 introduced in Section 3 and then provide some extra results concerning the average size of price adjustment.

A.1 Proofs of Lemmas 1–3

Recall that Lemma 1 states that the optimally preset price \bar{p}_0 is approximately equal to $p(0, 0) = 1$, the desired price when the cost shock is nil ($a = e = 0$). This is a very intuitive result since \bar{p}_0 is the optimal price for the expected value of the cost shock for which the firm chooses not to adjust its price (see (5)), given the symmetry of the cost shock distribution and the approximate symmetry of the non-adjustment set (Lemma 3). Nevertheless, the proof of this lemma is quite involved due to the fixed point nature of the problem and relies on some results from functional analysis. Further, if one takes $\bar{p}_0 \approx p(0, 0) = 1$ as given, the proofs of Lemmas 2 and 3 are straightforward and follow from standard first and second order Taylor approximations. Therefore, a reader who is not interested in the technical details of the proof of Lemma 1 and is willing to take it for granted, can simply look at the proofs of Lemmas 2' and 3' that follow below and are the counterparts to Lemmas 2 and 3.

Before proceeding with the proofs, we introduce some additional notation. We will denote by o the variables that have the same or smaller order of magnitude as the cost shock (a, e) . Formally, $o \equiv O(\|a, e\|_2)$, where $\|a, e\|_2 = \sqrt{a^2 + e^2}$ is the Euclidian vector norm in \mathbb{R}^2 and $O(\cdot)$ satisfies $\lim_{x \rightarrow 0} |O(x)/x| < \infty$ and denotes same (or lower) order of magnitude. Note that o is the order of magnitude of a particular realization of (a, e) . We also need a probabilistic order of magnitude $o_p \equiv O_p(\|a, e\|_2)$ which satisfies $p \lim_{x \rightarrow 0} |O_p(x)/x| < \infty$. We have then $o_p = \sqrt{\sigma_a^2 + \sigma_e^2}$ and $o = o_p$ with probability 1.

Next, since the cost shock (a, e) enters the marginal cost in a separable way, we can introduce a univariate sufficient statistic

$$1 + \delta = (1 - a)(1 + \phi e),$$

which reflects the overall size of the cost shock. Note that directly from the definition, we have

$$\delta = -a + \phi e + o^2,$$

which also implies $\delta = o$ and $O_p(\delta) = o_p$. We will also write p_δ and $p(\delta)$ interchangeably instead of $p(a, e)$; $\Pi(p|\delta)$ instead of $\Pi(p|a, e)$ and $\Pi(\delta)$ instead of $\Pi(a, e)$. Similarly, the profit loss from non-adjusting prices is

$$L(\delta|\bar{p}_0) = \Pi(\delta) - \Pi(\bar{p}_0|\delta),$$

where we now make the dependence on the preset price, \bar{p}_0 , explicit. Finally, the region of non-adjustment (with some abuse of notation) is still

$$\Delta \equiv \Delta(\bar{p}_0) = \{\delta : L(\delta|\bar{p}_0) \leq \kappa\}.$$

As we will show below, $L(\delta|\bar{p}_0) = o_p^2$. Therefore, a natural requirement is that $\kappa = o_p^2$, or equivalently, $\kappa = O(\sigma_a^2 + \sigma_e^2)$. This requirement ensures that region Δ is a non-trivial subset of the range of shocks δ .

As explained in the text, the optimal preset price maximizes expected profits conditional on non-adjustment upon the realization of the shock:

$$\bar{p}_0 \equiv \arg \max_p \mathbb{E}_\Delta \Pi(p|\delta),$$

where

$$\mathbb{E}_\Delta f(\delta) = \int_\Delta f(\delta) dG_\Delta(\delta), \quad G_\Delta(\delta) \equiv \frac{G(\delta)}{\Pr\{\delta \in \Delta\}}$$

and $G(\cdot)$ is the cumulative distribution function of δ implied by respective *cdf* $F(\cdot)$ for (a, e) . We assume that the original distribution of (a, e) is symmetric, so that the distribution of δ , $G(\cdot)$, is also symmetric. Note that, since $\Delta = \Delta(\bar{p}_0)$, the definition of \bar{p}_0 is in fact a fixed point problem. Finally, since profit function is linear in the cost function, we have

$$\bar{p}_0 = \arg \max_p \{p\varphi(p) - (1 + \bar{\delta}_\Delta)c(\varphi(p))\}, \quad \bar{\delta}_\Delta \equiv \mathbb{E}_\Delta \delta.$$

This immediately implies that $\bar{p}_0 = p(\bar{\delta}_\Delta) = \arg \max_p \Pi(p|\bar{\delta}_\Delta)$. Lemma 1 states that \bar{p}_0 is close to $p(0)$, which is the case whenever $\bar{\delta}_\Delta$ is close to 0.⁵² We introduce

$$\alpha \equiv p_0 - \bar{p}_0,$$

and to prove Lemmas 1–3 we will need to evaluate (the upper bound on) the order of magnitude of α . To anticipate the result, we will show that $\alpha = O(\kappa) = o_p^2$.

We are now ready to prove the lemmas.

Proof of Lemma 1: Taking the first order condition for maximization of profit in (3), we have:

$$\frac{\partial \Pi(p|\delta)}{\partial p} = \varphi(p) \left[(1 + \delta) \frac{\tilde{\sigma}(p)(1 + \delta)mc(\varphi(p))}{p} - (\tilde{\sigma}(p) - 1) \right] = 0, \quad (20)$$

which implies

$$p_\delta \equiv \arg \max_p \Pi(p|\delta) = \frac{\tilde{\sigma}(p_\delta)}{\tilde{\sigma}(p_\delta) - 1} (1 + \delta)mc(\varphi(p_\delta)).$$

⁵²Note that the unconditional expectation of δ , as well as expectation of δ conditional on any symmetric region, equals zero due to the symmetry of the distribution: $\mathbb{E}\delta = 0$ and $\mathbb{E}_{[-a, a]}\delta = 0$.

Substituting in $\delta = 0$, we can directly verify that $p_0 = 1$ is a solution since according to our normalization $\tilde{\sigma}(1) = \sigma$ and $mc(\varphi(1)) = (\sigma - 1)/\sigma$.⁵³ Noting that $a = e = 0$ is a special case that delivers $\delta = 0$, implies $p(0, 0) = 1$ as stated in Lemma 1.

We next introduce the following intermediate result:

Lemma 2': *First order Taylor approximation to the desired price around $\delta = 0$ implies:*

$$\begin{aligned}\ln p_\delta - \ln p_0 &= \Psi\delta + O(\delta^2), \\ \frac{p_\delta - p_0}{p_0} &= \Psi\delta + O(\delta^2).\end{aligned}$$

Proof: Taking the logs in the expression for the desired price, p_δ , and rearranging terms, we have

$$\ln p_\delta - \ln p_0 - \tilde{\mu}(p_\delta) - \ln mc(\varphi(p_\delta)) = \ln(1 + \delta), \quad (21)$$

where

$$\tilde{\mu}(p) \equiv \ln \left(\frac{\tilde{\sigma}(p)}{\tilde{\sigma}(p) - 1} \right)$$

is the log desired mark-up.

We now take the first order Taylor approximations to (21) on both sides: around $\delta = 0$ on the right-hand side and respectively around $p_0 = 1$ on the left-hand side. The Taylor approximation to the RHS is immediate:

$$\ln(1 + \delta) = \delta + O(\delta^2) = \delta + o^2.$$

Now consider the LHS. Note that

$$\frac{\partial \tilde{\mu}(p)}{\partial \ln p} = -\frac{1}{\tilde{\sigma}(p)(\tilde{\sigma}(p) - 1)} \frac{\partial \tilde{\sigma}(p)}{\partial \ln p} = -\frac{\tilde{\varepsilon}(p)}{\tilde{\sigma}(p) - 1}$$

and

$$\frac{\partial \ln mc(\varphi(p))}{\partial \ln p} = \frac{\partial \ln mc(\varphi(p))}{\partial \ln q} \cdot \frac{\partial \ln \varphi(p)}{\partial \ln p} = \tilde{\eta}(\varphi(p)) \cdot (-\tilde{\sigma}(p)).$$

Evaluating these expressions at $p = p_0 = 1$, we can write the first order Taylor approximation to the left-hand side of (21) as

$$\ln p_\delta - \ln p_0 - \tilde{\mu}(p_\delta) - \ln mc(\varphi(p_\delta)) = \Psi^{-1}(\ln p_\delta - \ln p_0) + O(\ln p_\delta - \ln p_0)^2,$$

where as in the text

$$\Psi \equiv \frac{1}{1 + \frac{\varepsilon}{\sigma-1} + \eta\sigma}$$

and we used the fact that $\tilde{\mu}(p_0) + \ln mc(\varphi(p_0)) = 0$ from our normalization assumptions.

⁵³As we show below, the sufficient condition for maximization is $\tilde{\sigma}(p_\delta) > 1$, $\tilde{\varepsilon}(p_\delta) \geq 0$ and $\tilde{\eta}(\varphi(p_\delta)) \geq 0$. We assume that this inequalities are satisfied for all p . This additionally implies that $\tilde{\sigma}(\cdot)$ and $mc(\cdot)$ are non-decreasing, so that both marginal cost and desired mark-up decrease in p . Therefore, the solution to the optimization problem is unique.

Combining the approximation to both sides, we have

$$\Psi^{-1}(\ln p_\delta - \ln p_0) + O(\ln p_\delta - \ln p_0)^2 = \delta + O(\delta^2).$$

The first useful implication of this result is that $\ln p_\delta - \ln p_0 = O(\delta) = o$, i.e. has the same order of magnitude as the cost shocks δ . The second implication is then

$$\ln p_\delta - \ln p_0 = \Psi\delta + O(\delta^2),$$

as stated in the Lemma. Note that it implies, in particular, $\partial \ln p_\delta / \partial \ln(1 + \delta) = \Psi$ and, combining it with $\delta = -a + \phi e + o^2$, proves Corollary 1.

Finally, observe that $p_\delta - p_0 = (p_\delta - p_0)/p_0 = \ln p_\delta - \ln p_0 + O(\delta^2)$, so that all approximations for $(\ln p_\delta - \ln p_0)$ also hold for $(p_\delta - p_0)$. ■

Lemma 2' has immediate implication:

$$\alpha = p_0 - \bar{p}_0 = p(0) - p(\bar{\delta}_\Delta) = O(\bar{\delta}_\Delta),$$

i.e. the order of magnitude of α is the same as that of $\bar{\delta}_\Delta = \mathbb{E}_\Delta \delta$. We need now to characterize region Δ , which is determined by the properties of the profit loss function, $L(\delta|\bar{p}_0)$. To do so, we introduce another intermediate

Lemma 3': *The Taylor approximation to the profit loss function results in*

$$L(\delta|\bar{p}_0) = \frac{1}{2}(\sigma - 1)\Psi\delta^2 + O(\delta^3) + O(\alpha^2),$$

where as before $\alpha = p_0 - \bar{p}_0$ and $\Psi = [1 + \varepsilon/(\sigma - 1) + \eta\sigma]^{-1}$.

Proof: We take the second order Taylor approximation to the profit loss function around the desired price, p_δ :

$$L(\delta|\bar{p}_0) \equiv \Pi(p_\delta|\delta) - \Pi(\bar{p}_0|\delta) = -\frac{1}{2} \frac{\partial^2 \Pi(p_\delta|\delta)}{\partial p^2} (p_\delta - \bar{p}_0)^2 + O(p_\delta - \bar{p}_0)^3,$$

where the first order term is zero due to the first order optimality of profit maximization.

The second derivative of the profit function with respect to price is

$$\frac{\partial^2 \Pi(p|\delta)}{\partial p^2} = \varphi'(p) \frac{\partial \Pi(p|\delta)}{\partial p} - \frac{\tilde{\sigma}(p)\varphi(p)}{p} \left\{ \tilde{\varepsilon}(p) + (1 + \delta) \frac{mc(\varphi(p))}{p} [1 + \tilde{\eta}(\varphi(p)) \cdot \tilde{\sigma}(p) - \tilde{\varepsilon}(p)] \right\}.$$

Evaluating this expression at p_δ , we have:

$$\frac{\partial^2 \Pi(p_\delta|\delta)}{\partial p^2} = -\frac{(\tilde{\sigma}(p_\delta) - 1) \cdot \varphi(p_\delta)}{p_\delta} \left[1 + \frac{\tilde{\varepsilon}(p_\delta)}{\tilde{\sigma}(p_\delta) - 1} + \tilde{\eta}(\varphi(p_\delta)) \cdot \tilde{\sigma}(p_\delta) \right],$$

where we used the first order condition which implies $\partial \Pi(p_\delta|\delta) / \partial p = 0$ and $(1 + \delta)mc(\varphi(p_\delta)) / p_\delta = (\tilde{\sigma}(p_\delta) - 1) / \tilde{\sigma}(p_\delta)$. Note that $\tilde{\sigma}(p_\delta) > 1$, $\tilde{\varepsilon}(p_\delta) \geq 0$ and $\tilde{\eta}(\varphi(p_\delta)) \geq 0$ are indeed sufficient conditions for profit maximization at p_δ .

Assuming that $\tilde{\varepsilon}(\cdot)$ and $\tilde{\eta}(\cdot)$ are smooth functions, we can use the following approximation:

$$\frac{\partial^2 \Pi(p_\delta | \delta)}{\partial p^2} = \frac{\partial^2 \Pi(p_0 | 0)}{\partial p^2} + O(\delta) = -\frac{\sigma - 1}{\Psi} + O(\delta),$$

where second equality evaluates second derivative of the profit function at $\delta = 0$ and $p_0 = 1$ taking into account our normalization of demand and cost functions. In addition, we can use Lemma 2' to obtain

$$p_\delta - \bar{p}_0 = (p_\delta - p_0) + (p_0 - \bar{p}_0) = \Psi \delta + O(\delta^2) + \alpha.$$

Combining these results, we can rewrite the approximation to the profit loss function as

$$\begin{aligned} L(\delta | \bar{p}_0) &= \frac{1}{2} \left[\frac{\sigma - 1}{\Psi} + O(\delta) \right] (\Psi \delta + O(\delta) + \alpha)^2 + O(\delta^3) + O(\alpha^3) \\ &= \frac{1}{2} (\sigma - 1) \Psi \delta^2 + O(\delta^3) + O(\alpha^2). \end{aligned}$$

where we have used the fact that $O(\delta + \alpha)^k = O(\delta^k) + O(\alpha^k)$ for $k > 0$. ■

Recall that $\Delta = \{\delta : L(\delta | \bar{p}_0) \leq \kappa\}$. We will abuse notation by using Δ for a generic set of the form:

$$\Delta \equiv \{\delta : \Theta \delta^2 + O(\delta^3) + O(\alpha^2) \leq \kappa\}, \quad \Theta \equiv \frac{1}{2} (\sigma - 1) \Psi.$$

Lemma 3' implies that $\{\delta : L(\delta | \bar{p}_0) \leq \kappa\}$ is a special case of such set. In what follows, we will prove results for a generic set Δ and, hence, they will hold for a particular set $\{\delta : L(\delta | \bar{p}_0) \leq \kappa\}$. Specifically, we need to evaluate the magnitude of $\bar{\delta}_\Delta = \mathbb{E}_\Delta \delta$; recall that $\alpha = O(\bar{\delta}_\Delta)$.

It is useful to consider another (approximate) set

$$\Delta' \equiv \{\delta : \Theta \delta^2 \leq \kappa\} = \left[-\sqrt{\kappa/\Theta}, \sqrt{\kappa/\Theta} \right].$$

This set is symmetric and, therefore, we have $\int_{\Delta'} \delta dG(\delta) = 0$ due to the symmetry of the distribution of δ . Therefore, we can write

$$|\mathbb{E}_\Delta \delta| = \left| \int_\Delta \delta dG_\Delta(\delta) - \underbrace{\int_{\Delta'} \delta dG_\Delta(\delta)}_{=0} \right| \leq \int_{\Delta \ominus \Delta'} |\delta| dG_\Delta(\delta),$$

where \ominus denotes symmetric difference, or the union of the complement sets, that is

$$\Delta \ominus \Delta' \equiv \{\Delta \setminus \Delta'\} \cup \{\Delta' \setminus \Delta\} \quad \text{and} \quad \Delta \setminus \Delta' \equiv \{\delta : \delta \in \Delta, \delta \notin \Delta'\}.$$

As a result, the order of $\alpha = O(\mathbb{E}_\Delta \delta)$ is bounded above by $O\left(\int_{\Delta \ominus \Delta'} |\delta| dG_\Delta(\delta)\right)$ and our problem reduces to evaluation of the order of magnitude of this integral.⁵⁴

We now introduce a convenient substitution of variables:

$$z \equiv \Theta \cdot \delta / \sqrt{\kappa}.$$

⁵⁴Note, however, that Δ itself depends on α (a fixed point problem). Therefore, we will need to check for consistency of resulting α with the definition of Δ .

Similarly, the counterparts to Δ' and Δ in the z -space are⁵⁵

$$\mathcal{Z}' \equiv \{z : z^2 \leq 1\} = [-1, 1],$$

$$\mathcal{Z} \equiv \{z : z^2 + O(\sqrt{\kappa}) + O(\alpha^2/\kappa) \leq 1\}.$$

Our requirement that $O(\kappa) = O(\sigma_\delta^2)$, where σ_δ^2 is the unconditional variance of δ , ensures that these sets are non-trivial, i.e. $\Pr\{z \in \mathcal{Z}\} = \Pr\{\delta \in \Delta\}$ is separated from both 0 and 1 as $\kappa \rightarrow 0$. It also implies $\delta = O_p(\sqrt{\kappa})$ and $z = O_p(1)$. As a consequence, we replaced $O(\delta^3/\kappa)$ with $O(\sqrt{\kappa})$ in the definition of \mathcal{Z} . Finally, this assumption is equivalent to $G'(z\sqrt{\kappa}/\Theta) = O(\kappa^{-1/2})$ for any given z .⁵⁶

Using the new notation, we can write the following operator:

$$\beta \equiv \beta(\alpha) = \int_{\Delta \ominus \Delta'} |\delta| dG_\Delta(\delta) = \frac{\kappa}{\Theta^2 \Pr\{z \in \mathcal{Z}\}} \int_{\mathcal{Z} \ominus \mathcal{Z}'} |z| G'(z\sqrt{\kappa}/\Theta) dz.$$

β depends on α since α appears in the definition of Δ and, hence, \mathcal{Z} . Recall that $\alpha = O(\beta) = O(\beta(\alpha))$. Therefore, we have a fixed point problem. We now prove the following result:

Lemma 4 *Any fixed point of $\alpha = O(\beta(\alpha))$ satisfies $\alpha = O(\kappa)$.*

Proof: From the discussion above we know that

$$\frac{\sqrt{\kappa} G'(z\sqrt{\kappa}/\Theta)}{\Theta^2 \Pr\{z \in \mathcal{Z}\}} = O(1).$$

Therefore, we can write

$$\beta \sim \sqrt{\kappa} \int_{\mathcal{Z} \ominus \mathcal{Z}'} |z| dz,$$

where \sim denotes the same order of magnitude.

The region of integration has the following form:

$$\mathcal{Z} \ominus \mathcal{Z}' \in U_{\ell_1}(-1) \cup U_{\ell_2}(1),$$

where $U_r(a)$ is a ball of radius r around a . To see this note that $\mathcal{Z}' = [-1; 1]$ and $\mathcal{Z} = [-z_1, z_2]$, where

$$z_{1,2} = \sqrt{1 + O(\sqrt{\kappa}) + O(\alpha^2/\kappa)}$$

⁵⁵Note that this approximations make sense only if $\alpha^2 = o(\kappa)$, that is $\lim_{\kappa \rightarrow 0} \alpha^2/\kappa = 0$. Lemma 4 below asserts that $\alpha = O(\kappa) = o(\sqrt{\kappa})$, which verifies our conjecture.

⁵⁶Note that $\Pr\{\delta \in \Delta'\} = \int_{\Delta'} dG(\delta) = \int_{-1}^1 dG(z\sqrt{\kappa}/\Theta) = \Theta^{-1} \int_{-1}^1 \sqrt{\kappa} G'(z\sqrt{\kappa}/\Theta) dz$ and to ensure that it is separated from both 0 and 1 we need to have $\sqrt{\kappa} G'(z\sqrt{\kappa}/\Theta) = O(1)$. As an example normalize $\Theta = 1$ and consider a normal density: $G'(\delta) = (\sqrt{2\sigma_\delta^2})^{-1/2} \exp\{-\delta^2/(2\sigma_\delta^2)\}$ so that

$$\sqrt{\kappa} G'_\kappa(z\sqrt{\kappa}) = \sqrt{\frac{\kappa}{2\sigma_\delta^2}} \exp\left\{-\frac{\kappa}{\sigma_\delta^2} \frac{z^2}{2}\right\},$$

which is obviously $O(1)$ as long as $O(\kappa/\sigma_\delta^2) = O(1)$.

so that

$$\ell_j = |1 - z_j| = \frac{O(\sqrt{\kappa}) + O(\alpha^2/\kappa)}{1 + \sqrt{1 + O(\sqrt{\kappa})} + O(\alpha^2/\kappa)} = O(\sqrt{\kappa}) + O(\alpha^2/\kappa).$$

As a result, when $z \in \mathcal{Z} \ominus \mathcal{Z}'$, we can approximate it by $|z| = 1 + O(\sqrt{\kappa}) + O(\alpha^2/\kappa)$.

Using these properties, we have:

$$\beta \sim \sqrt{\kappa} \cdot \ell_j = O(\kappa) + O(\alpha^2/\kappa) \quad \Rightarrow \quad \alpha = O(\kappa) + O(\alpha^2/\kappa).$$

It is immediate to check that the only solution to this equation is $\alpha = O(\kappa)$. ■

The implication of this Lemma is that

$$\alpha = O(\kappa) = O(\sigma_\delta^2) = O_p(\delta^2),$$

i.e. α has the same order as δ^2 (in the probabilistic sense). In the space of the original cost shocks (a, e) it translates into

$$\alpha = p(0, 0) - \bar{p}_0 = O_p(\|a, e\|_2)^2,$$

i.e. the difference between \bar{p}_0 and $p(0, 0)$ is second order. This completes the proof of Lemma 1. ■

One important observation is that this prove provides a very crude upper bound on the order of magnitude of α . Nevertheless, this crude evaluation is enough to proof the main results of the Section. We now provide the proofs for Lemmas 2 and 3.

Proof of Lemma 2: Follows immediately from Lemmas 1 and 2'. Using $p_0 = 1$, Lemma 2' states that $p_\delta - p_0 = \Psi\delta + O(\delta^2)$. Therefore, we have

$$\frac{p_\delta - \bar{p}_0}{\bar{p}_0} = \frac{p_\delta - p_0 + \alpha}{1 - \alpha} = p_\delta - p_0 + O(\alpha) = \Psi\delta + O(\delta^2) + O(\alpha).$$

Now replacing δ with the original cost shock (a, e) and using the evaluation of α from Lemma 1, we have

$$\frac{p(a, e) - \bar{p}_0}{\bar{p}_0} = \Psi(-a + \phi e) + O_p(\|a, e\|_2)^2,$$

i.e. the expression in the text indeed constitutes a valid first-order approximation. ■

Proof of Lemma 3: Following the same steps as in the proof of Lemma 3', we can write:

$$L(\delta|\bar{p}_0) = \frac{1}{2} \left[\frac{\sigma - 1}{\Psi} + O(\delta) \right] (p_\delta - \bar{p}_0)^2 + O(p_\delta - \bar{p}_0)^3,$$

and $p_\delta - \bar{p}_0 = O(\delta) + \alpha = O(\delta) + O_p(\delta^2) = O_p(\delta)$. Taking into account that $\bar{p}_0 = p_0 + \alpha = 1 + O_p(\delta^2)$, we can rewrite:

$$L(\delta|\bar{p}_0) = \frac{1}{2} \frac{\sigma - 1}{\Psi} \left(\frac{p_\delta - \bar{p}_0}{\bar{p}_0} \right)^2 + O_p(\delta^3),$$

or equivalently, in terms of the cost shocks (a, e) :

$$L(a, e|\bar{p}_0) = \frac{1}{2} \frac{\sigma - 1}{\Psi} \left(\frac{p(a, e) - \bar{p}_0}{\bar{p}_0} \right)^2 + O_p(\|a, e\|_2)^3.$$

Therefore, the expression in the text constitutes a valid second-order approximation. ■

Finally, we provide the exact expression for the approximation to the profit loss function in (11). From Lemma 3' and $\alpha = O_p(\delta^2)$, we have

$$L(\delta|\bar{p}_0) = \frac{1}{2}(\sigma - 1)\Psi\delta^2 + O_p(\delta^3).$$

Substituting in for the original cost shock (a, e) , we obtain

$$L(a, e|\bar{p}_0) = \frac{1}{2}(\sigma - 1)\Psi(-a + \phi e)^2 + O_p(\|a, e\|_2^3),$$

so that the expression in (11) is indeed a valid second-order approximation.

A.2 Size of Price Adjustment

As in the text, denote by

$$X \equiv \frac{-a + \phi e}{\sqrt{\Sigma}}, \quad \Sigma = \sigma_a^2 + \phi^2 \sigma_e^2 = \sigma_a^2 \left(1 + \phi^2 \frac{\sigma_e^2}{\sigma_a^2}\right)$$

the normalized cost shock which has mean zero and unit variance. Then, using Lemma 2, we can write the desired price as

$$p(a, e) \approx \Psi\sqrt{\Sigma} \cdot X.$$

As a result, the mean squared size of the desired price change is approximately $\Psi\sqrt{\Sigma}$, increasing in pass-through, Ψ , and standard deviation of the shocks, $\sqrt{\Sigma}$. However, the firm adjusts its price only if $|X| > \sqrt{2\kappa/[(\sigma - 1)\Psi\Sigma]}$. We denote the expected size of price adjustment (conditional on adjustment taking place) by

$$S = \mathbb{E} \left\{ |p(a, e) - \bar{p}_0| \mid |X| > \sqrt{\frac{2\kappa}{(\sigma - 1)\Psi\Sigma}} \right\} \approx \Psi\sqrt{\Sigma} \cdot \mathbb{E} \left\{ |X| \mid |X| > \sqrt{\frac{2\kappa}{(\sigma - 1)\Psi\Sigma}} \right\}.$$

The conditional expectation on the RHS is increasing in the cutoff.⁵⁷ Therefore, S increases in κ , but may increase or decrease in Ψ and Σ since they induce two counteracting effects: they increase the desired size of price adjustment for a given normalized shock X , but reduce the average size of the normalized shock to which the firm adjusts.⁵⁸ Note that S increases in Ψ whenever it increases in Σ , but not vice versa.

⁵⁷The general result is that $\mathbb{E}\{x|x > a\}$ is increasing in a .

⁵⁸Formally, the condition for $\mathbb{E}\{x|x > a\}/a$ to decrease in a is

$$\frac{af(a)}{1 - F(a)} \frac{\mathbb{E}\{x - a|x > a\}}{\mathbb{E}\{x|x > a\}} > 1,$$

where $f(x)/[1 - F(x)]$ is the hazard rate for the distribution of x . This condition may or may not hold depending on the distribution function and the value of a .

We now consider two special cases for illustration: Pareto distributed $|X|$ with shape parameter χ and exponentially distributed $|X|$ with shape parameter λ . In the Pareto case we have

$$S_P \approx \frac{\chi}{\chi - 1} \sqrt{\frac{2\kappa\Psi}{(\sigma - 1)}},$$

that is the size of price adjustment increase in menu cost, κ , and pass-through elasticity, Ψ , but is constant with respect to the size of the cost shocks, Σ (the two effects exactly cancel each other in this case). As a consequence, if frequency is driven by both variation in κ and Ψ , sorting goods by frequency will result in no particular pattern for the sizes. Recall that in general case, frequency is decreasing in κ and increasing in Ψ and Σ .

In the case of exponentially distributed $|X|$,⁵⁹ we have

$$S_E \approx \frac{\Psi\sqrt{\Sigma}}{\lambda} + \sqrt{\frac{2\kappa\Psi}{(\sigma - 1)}}$$

so that the size of price adjustment is increasing in both Ψ and Σ , as well as in κ . In this case, joint variation in κ and Σ , even when Ψ is constant, can rationalize the absence of pattern for size when the goods are sorted by frequency.

B Calibration and Simulation Details

In this appendix we first discuss in more detail the calibration of parameters and then lay out the simulation procedure for the dynamic model of Section 4.

B.1 Parameter Calibration

We need to calibrate nine parameters:

$$\{\beta, \kappa, \sigma_a, \rho_a, \Delta e, \omega, \phi, \varepsilon, \sigma\}.$$

The calibrated parameters are also summarized in Table 5. The period of the model corresponds to one month and we set the discount factor, β , to be 4% on an annualized basis.

The calibration of the menu cost (κ), volatility and persistence of the idiosyncratic shock process (σ_a and ρ_a), the measure of foreign firms (ω) and their cost sensitivity to the exchange rate (ϕ), as well as the demand parameters (σ and ε) is discussed in the main text. Therefore, it only remains to discuss the calibration of the real exchange rate process. We let the the (log of) the real exchange rate, $e_t \equiv \ln(W_t^*/W_t)$ follow a binomial random walk process within wide boundaries.

⁵⁹Very similar results hold for uniformly distributed $|X|$ as well.

Specifically, its value each period either increases or decreases by Δe with equal probabilities and reflects from the boundaries of a grid which are set to equal $\pm 7\Delta e$.⁶⁰ This procedure generates very persistent exchange rate paths with an auto-regression coefficient of around 0.97, consistent with the properties of the real exchange rate in the data. We set $\Delta e = 2.5\%$ which corresponds to the monthly standard deviation of the innovation to the real exchange rate for developed markets.

B.2 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:

$$D_t = \frac{\sigma - 1}{\sigma} \int_{\Omega} \frac{C_{jt}}{C_t} \exp \left\{ \frac{1}{\varepsilon} \left[1 - \left(\frac{|\Omega| C_{jt}}{C_t} \right)^{\varepsilon/\sigma} \right] \right\} dj, \quad (22)$$

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 $P_j = P = \bar{P}$ and $C_j = \bar{C}/|\Omega|$ for all j , so that $x_j \equiv \bar{D} = (\sigma - 1)/\sigma$. Therefore, in a symmetric steady state the elasticity and super-elasticity of demand equal σ and ε respectively. Equations (14) and (15) in the text together with our demand specification imply the following (implicit) expression for the sectoral price level:

$$P_t = \frac{1}{|\Omega|} \int_{\Omega} P_{jt} \left[1 - \varepsilon \ln \left(\frac{\sigma D_t P_{jt}}{\sigma - 1 P_t} \right) \right]^{\sigma/\varepsilon} dj. \quad (23)$$

We can now prove the following

Lemma 5 (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:*

$$\ln P_t \approx \frac{1}{|\Omega|} \int_{\Omega} \ln P_{jt} dj.$$

In both cases the order of magnitude is $O(\|\{\hat{P}_{jt}\}_{j \in \Omega}\|)$, 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 (22) and (23) in log deviations from the symmetric steady state we obtain two equivalent first-order accurate representations for \hat{D}_t :

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

⁶⁰We use this procedure to numerically generate a highly persistent process for the exchange rate, which is harder to obtain using the Tauchen routine.

Now using the definition of the Kimball aggregator (13) and the fact that $\Psi(\cdot)$ is a smooth function, we have

$$\frac{1}{|\Omega|} \int_{\Omega} (\hat{C}_{jt} - \hat{C}_t) dj = 0$$

up to second-order terms, specifically, the approximation error has the order of $O(\|\{\hat{C}_{jt}\}_{j \in \Omega}\|^2)$. Combining the two results together immediately implies that $\hat{D}_t = 0$ and $\hat{P}_t = |\Omega|^{-1} \int_{\Omega} \hat{P}_{jt} dj$ up to the second-order terms, $O(\|\{\hat{C}_{jt}\}_{j \in \Omega}\|^2)$. Taking the log-differential of the demand equation (19), we have

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

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

$$\ln P_t = \frac{1}{|\Omega|} \int_{\Omega} \ln P_{jt} dj$$

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 5 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.⁶¹ We verify, however, that computing the sectoral price index according to the exact expression (23) 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.

Further, our choice of the domestic wage rate as the numéraire implies that $W_t \equiv 1$ and the log of the real exchange rate is $e_t \equiv \ln(W_t^*/W_t) = \ln W_t^*$. This allows us to reduce the state space for each firm to $\mathbb{S}_{jt} = (P_{j,t-1}, A_{jt}; P_t, W_t^*)$. We iterate the Bellman operator (17) 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). The grid for idiosyncratic shock A_{jt} contains at least 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 foreign wage rate (real exchange rate) has at least 15 grid points with increments equal to Δe .

⁶¹This is also an assumption adopted by Klenow and Willis (2006).

To iterate the Bellman operator (17), a firm needs to form expectations about the future path of the exogenous state variables, (A_{jt}, W_t^*, P_t) . Since A_{jt} and W_t^* follow exogenous stochastic processes specified above, the conditional expectations for these variables are immediate.⁶² 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.⁶³ 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, $(\ln P_t, e_t)$ alone explains over 95% of variation in $\ln P_{t+1}$. 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 $(\gamma_0^{(0)}, \gamma_1^{(0)}, \gamma_2^{(0)})$; (ii) simulate M times the dynamic path of the sectoral price level $\{P_t^{(m)}\}_{t=0}^T$, where in every period we make sure that P_t is consistent with the price setting of the firms; (iii) for each simulation estimate $(\hat{\gamma}_0^{(m)}, \hat{\gamma}_1^{(m)}, \hat{\gamma}_2^{(m)})$ from regressing $\ln P_{t+1}$ on $\ln P_t, e_t$ and a constant; (iv) obtain $(\gamma_0^{(1)}, \gamma_1^{(1)}, \gamma_2^{(1)})$ by taking the median of $(\hat{\gamma}_0^{(m)}, \hat{\gamma}_1^{(m)}, \hat{\gamma}_2^{(m)})$; (v) iterate this procedure till convergence. 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 3.5 price adjustments for each firm), which generates an unbalanced panel of firm price changes, as we observe in the data.⁶⁴ On this simulated data, we estimate the same regressions (1) and (2) as we do on the BLS dataset.

⁶²Recall that A_{jt} follows a first-order autoregressive process; we discretize it using the Tauchen routine.

⁶³In fact, there are two distinct fixed points problem, 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

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

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

⁶⁴Specifically, for each firm we choose a random time interval during which its price is observed by the econometrician, thought 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 (2005): “Trade Costs, Pricing-to-Market, and International Relative Prices,” UCLA.
- 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 Persistence,” *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.
- CAPLIN, A. S., AND D. F. SPULBER (1987): “Menu Costs and the Neutrality of Money,” *Quarterly Journal of Economics*, 102(4), 703–26.
- DORNBUSCH, R. (1987): “Exchange Rate and Prices,” *American Economic Review*, 77(1), 93–106.
- 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.
- 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., O. ITSKHOKI, AND R. RIGOBON (2007): “Currency Choice and Exchange Rate Pass-through,” Harvard University.
- GOPINATH, G., AND R. RIGOBON (2007): “Sticky Borders,” *Quarterly Journal of Economics*, forthcoming.
- GUST, C., S. LEDUC, AND R. J. VIGFUSSON (2006): “Trade Integration, Competition, and the Decline in Exchange-Rate Pass-Through,” Board of Governors of the FRS, International Discussion Paper No. 864.
- 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.
- 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.
- 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.

A Tables and Figures

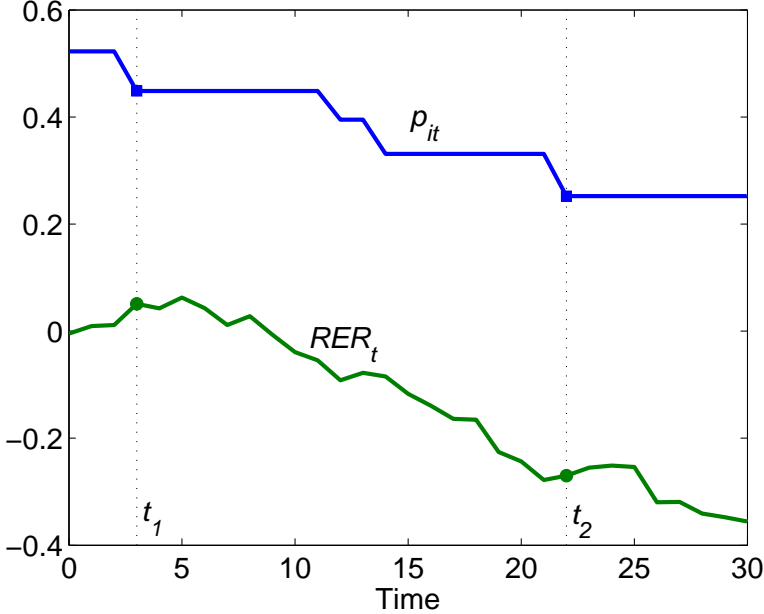


Figure 1: Illustration of Life-long pass-through

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 $\Delta p_{LR}^{i,c} = p_{it_2} - p_{it_1}$ and $\Delta RER_{LR}^{i,c} = RER_{t_2} - RER_{t_1}$. In the baseline specification we additionally adjust $\Delta p_{LR}^{i,c}$ by U.S. CPI inflation over the same period.

Table 1: Life-long pass-through

	Median Freq.	β_{LR}	$\sigma(\beta_{LR})$	N
Panel A				
All Countries				
Manufacturing				
– Low Frequency	0.07	0.20	0.03	5111
– High Frequency	0.39	0.40	0.05	5078
Differentiated				
– Low Frequency	0.07	0.19	0.04	2655
– High Frequency	0.29	0.40	0.06	2573
High-Income OECD				
Manufacturing				
– Low Frequency	0.07	0.27	0.04	3000
– High Frequency	0.40	0.58	0.07	2867
Differentiated				
– Low Frequency	0.07	0.26	0.07	1503
– High Frequency	0.33	0.58	0.08	1461
Panel B				
All Countries				
Manufacturing				
– Low Frequency	0.10	0.18	0.04	3299
– High Frequency	0.50	0.41	0.07	3316
Differentiated				
– Low Frequency	0.09	0.15	0.06	1645
– High Frequency	0.33	0.40	0.08	1573
High-Income OECD				
Manufacturing				
– Low Frequency	0.10	0.18	0.05	2081
– High Frequency	0.50	0.70	0.07	1896
Differentiated				
– Low Frequency	0.09	0.18	0.09	975
– High Frequency	0.40	0.60	0.09	960

Table 2: Life-long pass-through, 3 and more price changes

	Median Freq.	β_{LR}	$\sigma(\beta_{LR})$	N
All Countries				
Manufacturing				
– Low Frequency	0.13	0.22	0.04	2281
– High Frequency	0.58	0.44	0.07	2299
Differentiated				
– Low Frequency	0.11	0.15	0.07	1035
– High Frequency	0.42	0.51	0.09	1095
High-Income OECD				
Manufacturing				
– Low Frequency	0.12	0.30	0.07	1436
– High Frequency	0.60	0.73	0.08	1323
Differentiated				
– Low Frequency	0.11	0.23	0.12	657
– High Frequency	0.50	0.77	0.09	646

Table 3: Life-long pass-through, Regions

	Median Freq.	β_{LR}	$\sigma(\beta_{LR})$	N
Japan				
– Low Frequency	0.07	0.31	0.07	714
– High Frequency	0.27	0.62	0.15	704
Euro Area				
– Low Frequency	0.07	0.28	0.09	972
– High Frequency	0.33	0.49	0.09	980
Canada				
– Low Frequency	0.10	0.36	0.12	621
– High Frequency	0.87	0.74	0.23	529
Non HIOECD				
– Low Frequency	0.07	0.12	0.04	2031
– High Frequency	0.36	0.26	0.06	2291

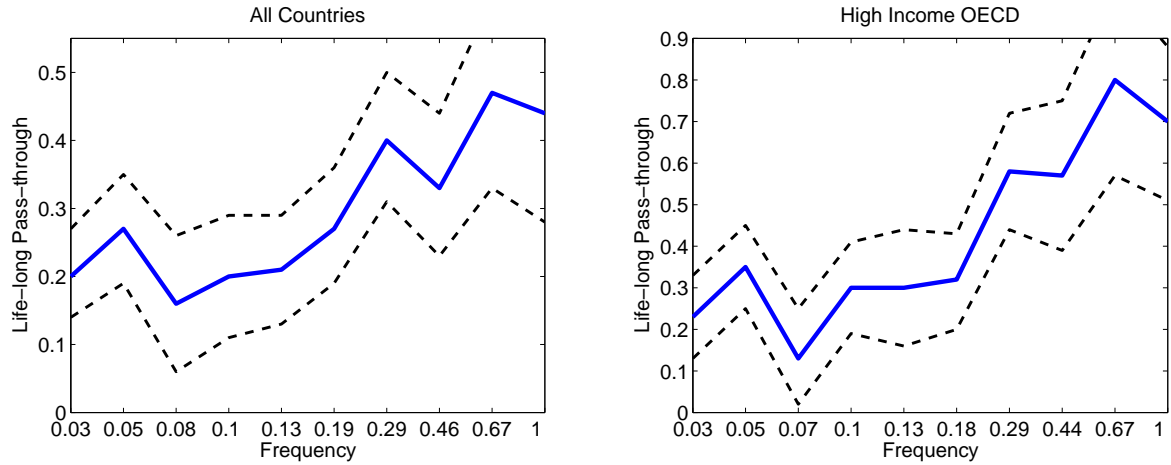


Figure 2: Life-long Pass-through across Frequency Deciles

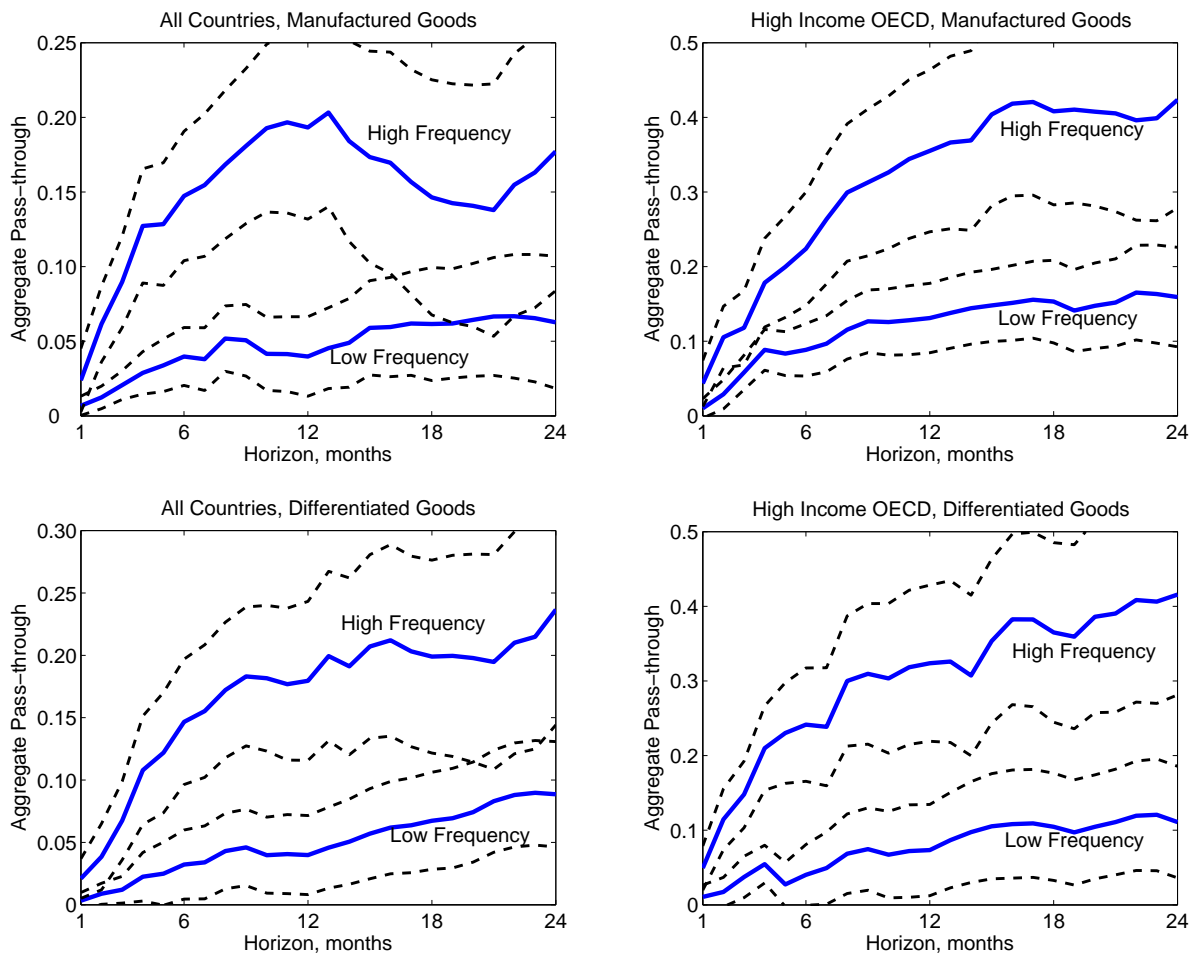


Figure 3: Aggregate Pass-through Regressions

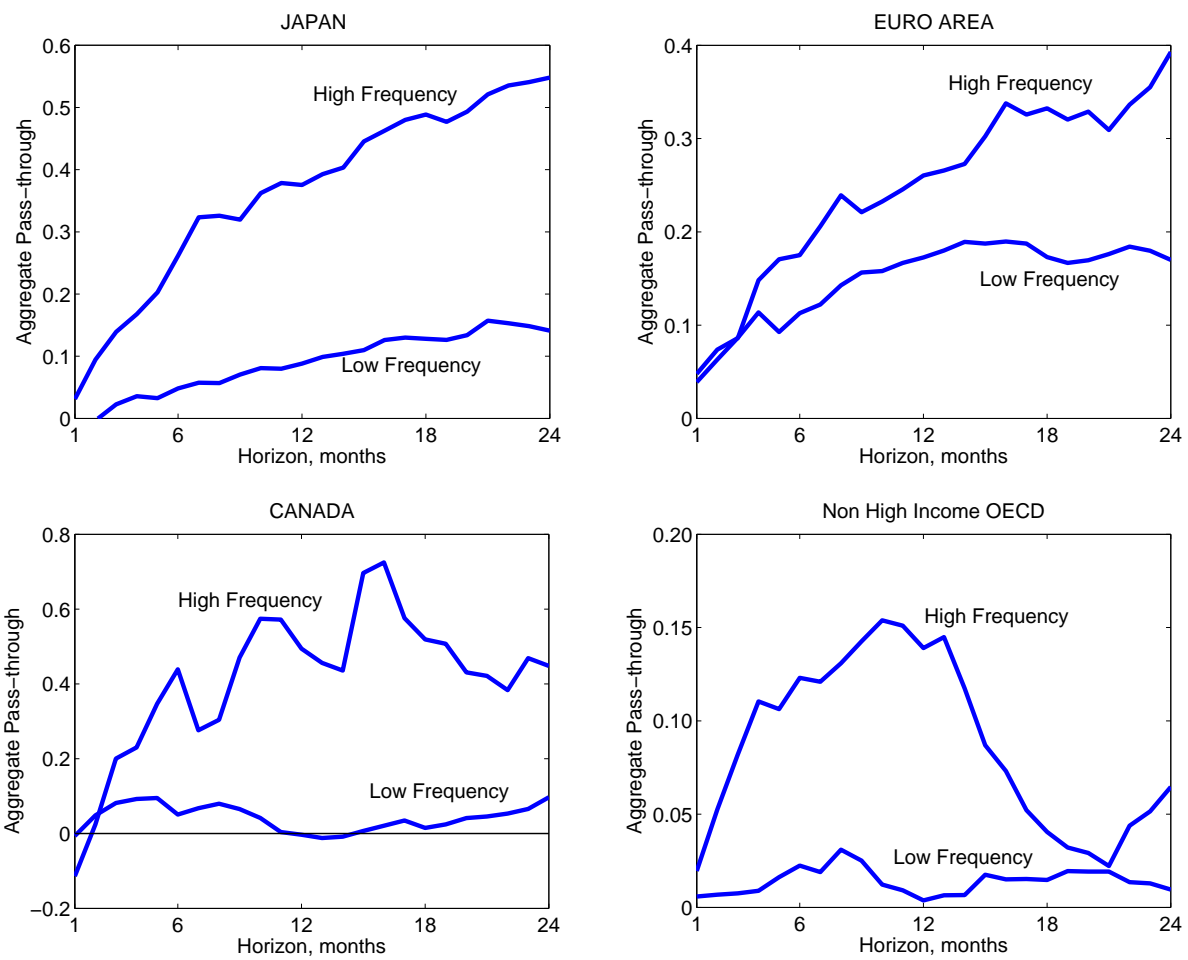


Figure 4: Aggregate Pass-through Regressions: Regions

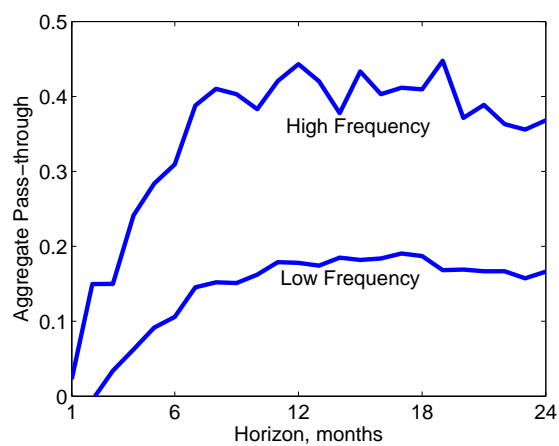


Figure 5: Aggregate Pass-through Regressions: 3 or more price changes

Table 4: Substitutions

Decile	Freq	Life 1	Life 2	Freq sub 1	Freq sub 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

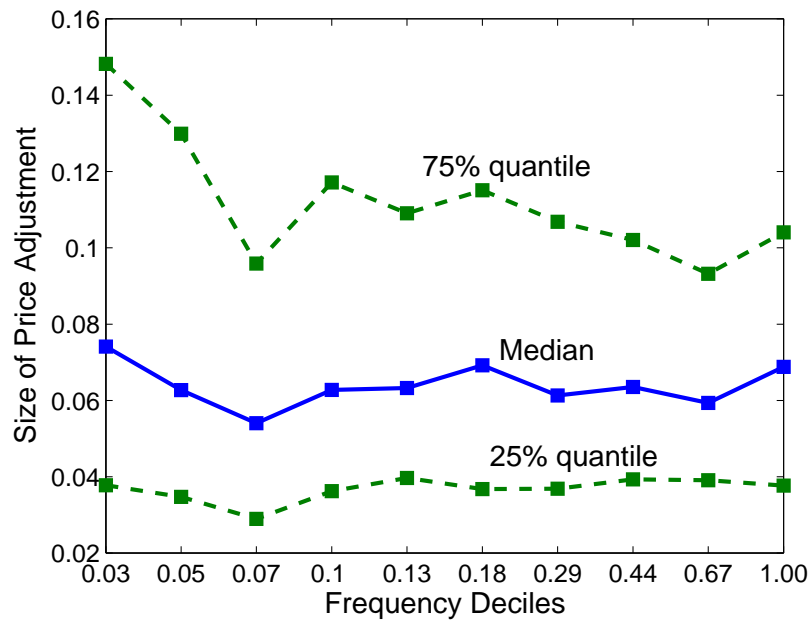


Figure 6: Size of Price Adjustment

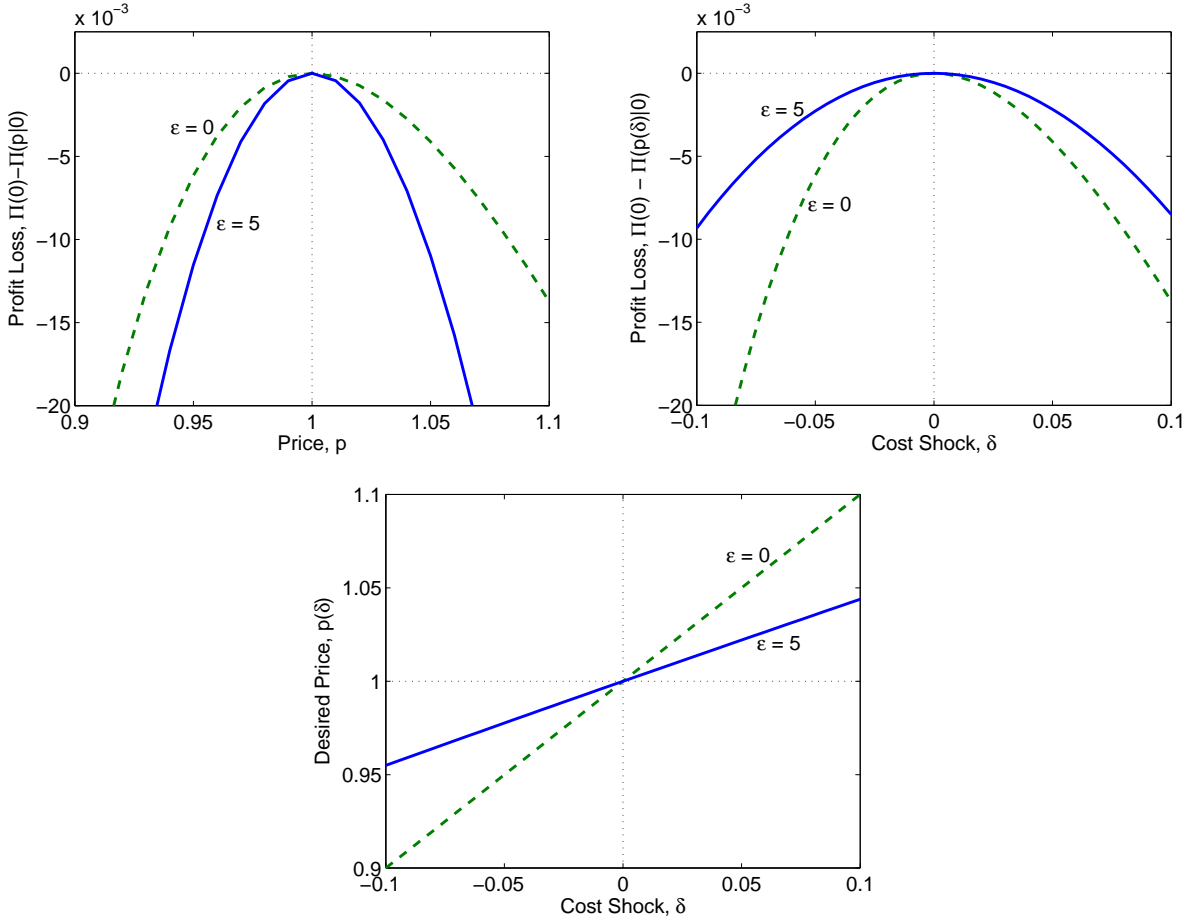


Figure 7: Pass-through and Curvature of the profit function

For this figures we use constant marginal cost and demand specification $\varphi(p) = (1 - \varepsilon \ln p)^{\sigma/\varepsilon}$ for $\sigma = 5$ and $\varepsilon = 0$ and 5 . We denote by δ the overall cost shock such that $(1 + \delta) = (1 - a)(1 + \phi e)$. Similarly, $\Pi(p|\delta) = \varphi(p)[p - (1 + \delta)c]$, where $c = (\sigma - 1)/\sigma$ is the constant marginal cost. Finally, we denote $p(\delta) \equiv \arg \max_p \Pi(p|\delta)$ and $\Pi(\delta) \equiv \Pi(p(\delta)|\delta)$. The top left panel plots $\Pi(0) - \Pi(p|0)$ as a function of p , that is profit loss when current shock is $\delta = 0$ and the price is preset at some exogenous level p . In this space, the curvature of the profit function is larger the greater is ε , as suggested by Lemma 3. The top right panel then plots $\Pi(0) - \Pi(p(\delta)|0)$, that is the profit loss when current shock is $\delta = 0$ and the price was optimally preset to level $p(\delta)$ when the previous shock was δ . In other words, this figure endogenizes the movement in the desired price in response to a cost shock. Note that in this space, the curvature of the profit function is smaller the greater is ε , as suggested by (11). Finally, the bottom panel plots the response of the desired price to the cost shock, $p(\delta)$. Note that the sensitivity of the desired price to cost shock decreases with greater ε ; specifically, pass-through for $\varepsilon = 0$ is complete (one-to-one), while for $\varepsilon = 5$ it is less than complete (slightly less than 50%). This panel describes the mechanism behind the difference in the two top panels.

Table 5: Calibrated Parameters

Parameter	Symbol	Values
Discount factor	β	$0.96^{1/12}$
Menu Cost	κ	2.5%
Exchange Rate Shock	Δe	2.5%
Idiosyncratic Shock	σ_a	8.5%
	ρ_a	0.95
Fraction of Imports	$\omega/(1 + \omega)$	16.7%
Cost Sensitivity to W^*	ϕ	0.75

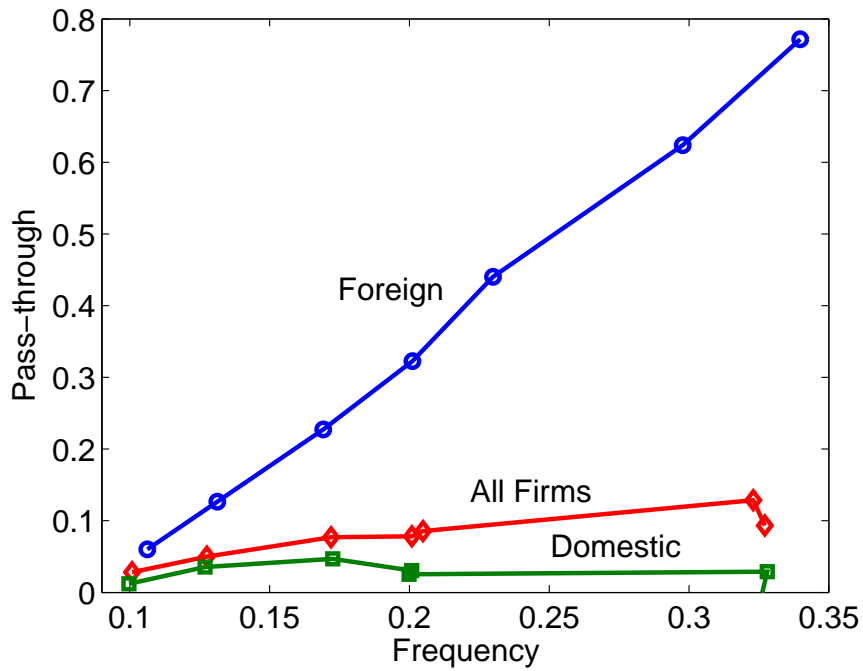


Figure 8: Frequency and Pass-through in the model (variation in ε)

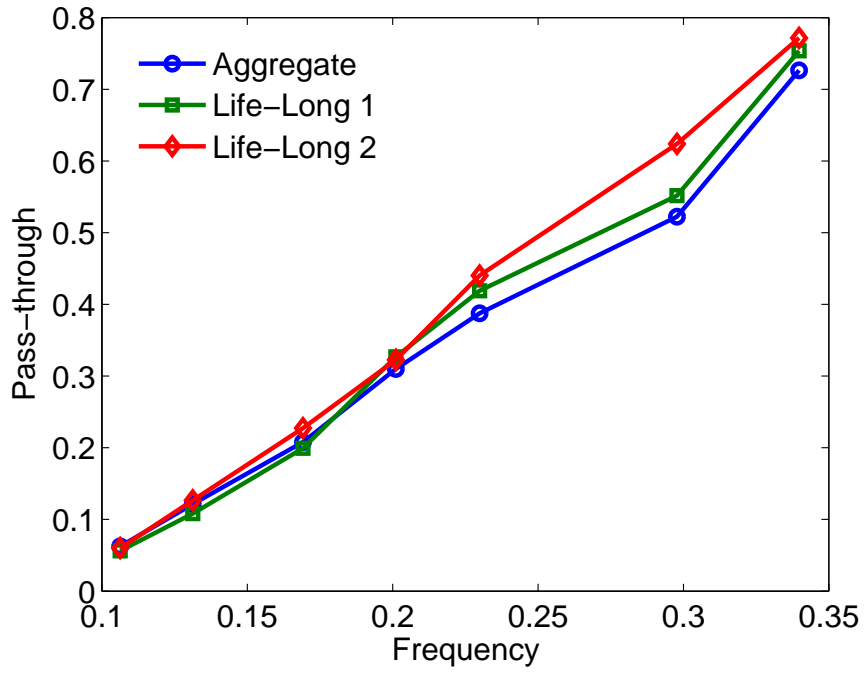


Figure 9: Measures of Long-run Pass-through

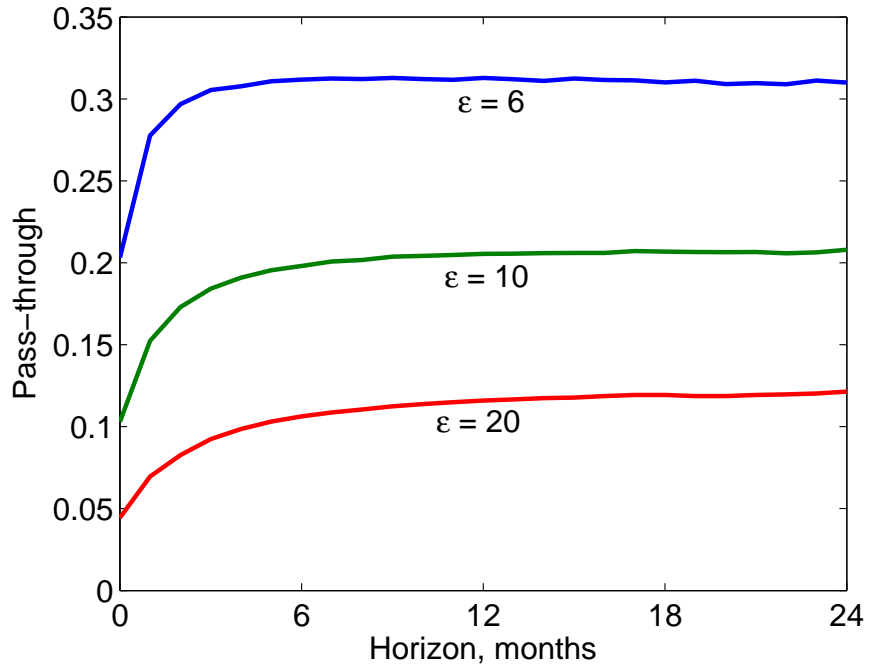


Figure 10: Aggregate Pass-through Regressions (variation in ϵ)

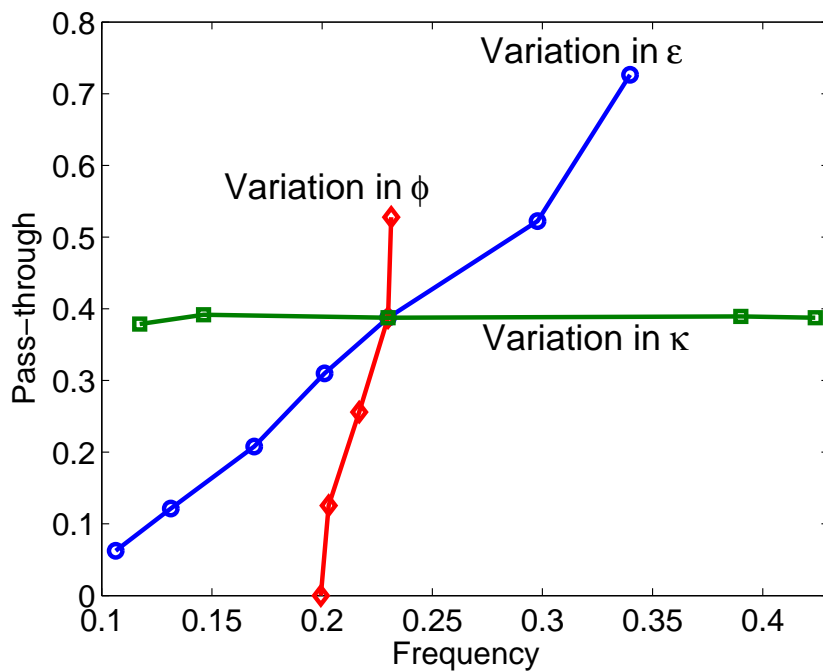


Figure 11: Frequency and Long-run Pass-through: variation in ϕ and κ

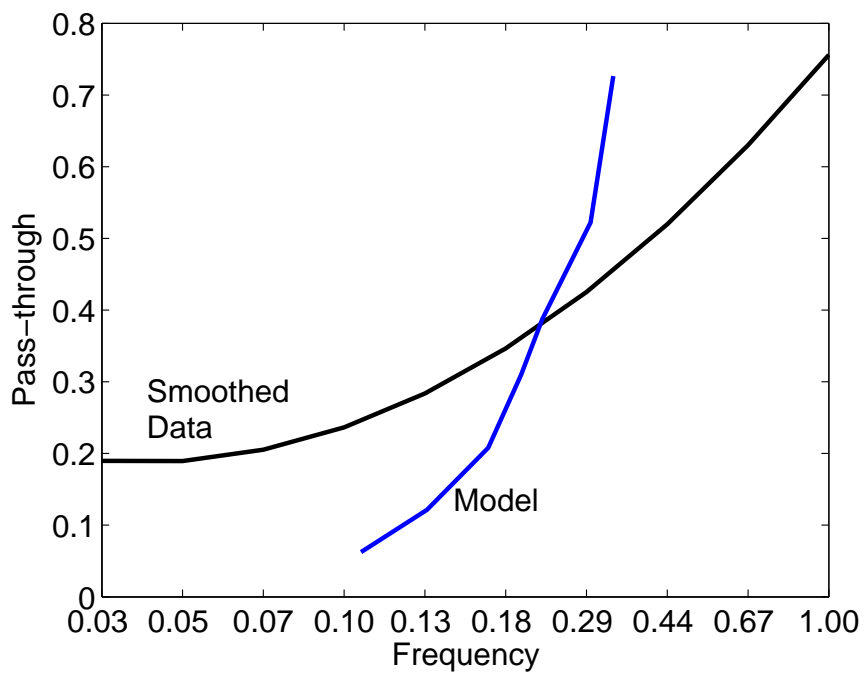


Figure 12: Model against the Data

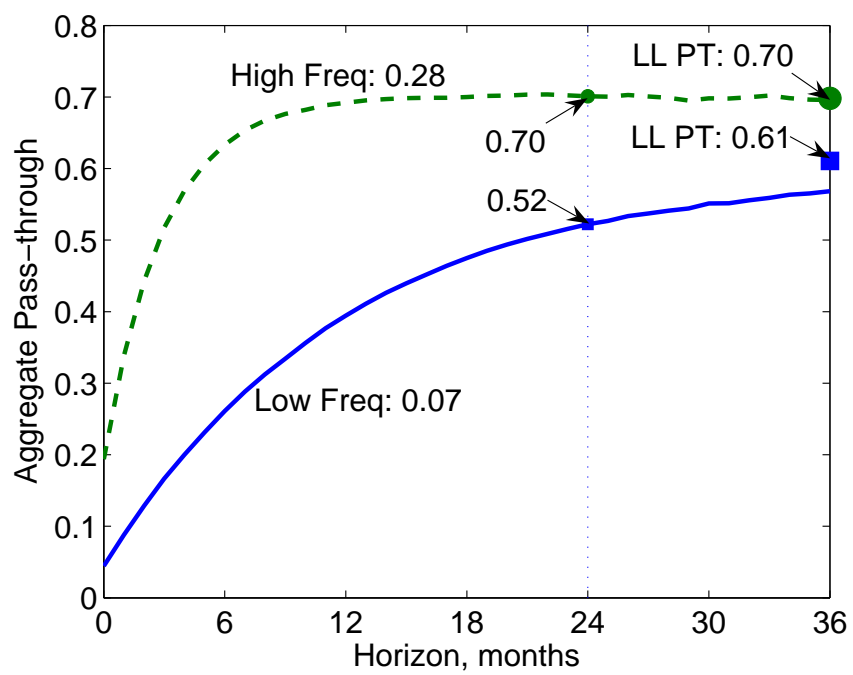


Figure 13: Frequency and Pass-through in a Calvo model