Working Papers

FEDERAL RESERVE BANK OF BOSTON™

No. 10-09

In Search of Real Rigidities

Gita Gopinath and Oleg Itskhoki

Abstract:
The closed and open economy literatures both work on evaluating the role of real rigidities, but in parallel. This paper brings the two literatures together. We use international price data and exchange rate shocks to evaluate the importance of real rigidities in price setting. We show that, consistent with the presence of real rigidities, the response of reset-price inflation to exchange rate shocks exhibits significant persistence. Individual import prices, conditional on changing, respond to exchange rate shocks prior to the last price change. At the same time, aggregate reset-price inflation for imports, like that for consumer prices, exhibits little persistence. Competitor prices affect firm pricing, and exchange rate pass-through into import prices is greater in response to trade-weighted, as opposed to bilateral, exchange rate shocks. We quantitatively evaluate sticky-price models (Calvo and menu cost) with variable markups at the wholesale level and constant markups at the retail level, consistent with empirical evidence. Variable markups alone generate price sluggishness at the aggregate level, while they fall short of matching price persistence at the micro level. Finally, variable markups magnify the size of the contract multiplier, but their absolute effects are modest unless they are coupled with exogenous sources of persistence.

JEL Classifications: F1, F3, F4

Gita Gopinath is a professor in the economics department at Harvard University, a visiting scholar in the research department of the Federal Reserve Bank of Boston, and an NBER faculty research fellow. Oleg Itskhoki is an assistant professor of economics and international relations in the department of economics and the Woodrow Wilson School of Public and International Affairs at Princeton University and an NBER faculty research fellow. Their email addresses are: gopinath@harvard.edu and itskhoki@princeton.edu, respectively.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at http://www.bos.frb.org/economic/wp/index.htm.

This paper was originally prepared for the NBER 25th Macroannual Conference, held April 9—10, 2010. We thank the International Price Program of the Bureau of Labor Statistics for access to unpublished micro data.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. We thank Jim Stock, Mark Watson, Ulrich Mueller, Chris Sims, Patrick Kehoe, the editors Daron Acemoglu and Mike Woodford, and our discussants Mike Golosov and Virgiliu Midrigan for comments.

This version: July 23, 2010

1 Introduction

A large literature has recently emerged that documents patterns of nominal price stickiness at the very micro level — the good level. The documented nominal durations are significantly shorter than the estimated real effects of money on output.[1] The long-lasting real effects of monetary shocks can be reconciled with moderate price stickiness if *real rigidities* are an important phenomenon. Real rigidities are mechanisms that dampen price responses of firms because of factors such as strategic complementarities in price setting, real wage rigidity, the dependence of costs on input prices that have yet to adjust, and others.[2] An important empirical literature has emerged recently that evaluates the question: Are quantitatively important real rigidities present in the data? The answer appears to depend on what data one examines.

In international economics there is a large and growing literature that estimates exchange rate pass-through from exchange rate shocks into prices. The estimated exchange rate pass-through is found to be incomplete, that is, if the dollar depreciates by 10 percent relative to the Euro, dollar prices of goods imported from the Euro area increase by less than 10 percent even in the long run. This incompleteness in pass-through is argued to be consistent with the presence of important real rigidities. Exchange rate changes generate relative price movements for the same good across markets despite costs being the same. This destination-specific markup is argued to be consistent with the presence of significant strategic complementarities in price setting.

The closed economy literature, on the other hand, uses indirect tests of real rigidities in the absence of well-identified and sizeable shocks like exchange rate shocks. The recent work based on micro evidence for retail prices argues that real rigidities are not an empirically important phenomenon.

There are many developments in the measurement of real rigidities in the closed and open economy literatures, but these developments have taken place in parallel. In this paper, we bring together the closed economy macro literature, which focuses mainly on indirect tests of

[1] See Bils and Klenow (2004) and Nakamura and Steinsson (2008) for evidence on nominal price durations in the BLS micro price data underlying the construction of the CPI index and Christiano, Eichenbaum, and Evans (1999) for the evidence on real effects of monetary shocks.

[2] For a precise definition, discussion, and examples of real rigidities, see the entry in the New Palgrave Dictionary of Economics by Romer (2008). The ability of real rigidities to magnify the effects of nominal price stickiness is often referred to as the *contract multiplier*, a term coined by Taylor (1980) and commonly used in the macro literature (for example, see Chari, Kehoe, and McGrattan, 2000).

real rigidities, with the international pricing literature, which uses an observable and sizeable shock, namely the exchange rate shock, to evaluate the behavior of prices, and in particular, the behavior of strategic complementarities in pricing. We first review the recent evidence on real rigidities to evaluate whether there is a consensus emerging on the importance of these rigidities in the data. Second, since the two literatures use different metrics to evaluate the importance of real rigidities, we use unpublished international price data collected by the U.S. Bureau of Labor Statistics (BLS) to estimate both metrics, using the same data. Third, we present new evidence on the dynamic response of international prices to exchange rate shocks and the response to competitor prices. Fourth, we calibrate sticky-price macro models (Calvo and menu cost) with a retail and wholesale sector to the evidence on the variable markup channel of real rigidities. We evaluate their ability to match the behavior of prices in the data and to measure the extent of monetary non-neutrality that this channel generates.

In reviewing the literature, we group evidence based on whether the prices studied refer to retail (consumer) prices or wholesale prices. Wholesale prices can alternatively be viewed as intermediate-good prices in business-to-business transactions. The literature on exchange rate pass-through into at-the-dock prices of goods refers to wholesale prices. A review of the existing literature reveals one surprisingly consistent result across several studies— surprising since these studies use different methodologies and data sets. This result is that strategic complementarities, for example, operating through variable markups, play little role for retail prices and appear to be quite important for wholesale prices.

We next use the BLS import price data to perform tests of real rigidity, where we use measures employed in the closed economy literature, namely, the persistence of reset-price inflation (Bils, Klenow, and Malin, 2009, henceforth BKM), and measures employed in the open economy literature, namely, the dynamic response of prices to exchange rate shocks. The actual import-price inflation series has a monthly persistence of 0.56, while the corresponding reset-price inflation series has a persistence of -0.04. In comparison, BKM estimate for retail prices that the inflation series has persistence of -0.05, while the reset-price inflation series has a persistence of -0.41. In comparison to retail prices, import prices have greater persistence, but the magnitude of this persistence suggests very little sluggishness in price adjustment.

However, when we project the aggregate import reset-price inflation on lags of the trade-weighted nominal exchange rate changes, we find that the autocorrelation of the fitted series is substantially higher than that of unconditional reset-price inflation (0.33 versus -0.04). We find similar evidence using micro-level price adjustments. We show that individual import prices, conditional on changing, respond to exchange rate shocks prior to the last time the price was adjusted and that these lagged effects are large and statistically significant. The pass-through, conditional on a price change to the cumulative exchange rate change since the last price adjustment, is 0.11 and the response to the cumulative exchange rate over the previous price duration is 0.08. Both these pieces of evidence that evaluate the response to a specific shock suggest a more important role for real rigidities than for the point estimate of the autocorrelation of reset prices.

Next, we evaluate the importance of strategic complementarities in price setting for incomplete pass-through, using some measures that capture the pricing behavior of competitors and measures that capture the extent of competition in sectors. These measures are not perfect but provide useful information about pricing behavior. We use the prices set by other firms in the same 10-digit or 4-digit harmonized code in the import price sample to control for the behavior of competitor prices, and we find that they have an important positive effect on firms' pricing, reducing the direct pass-through of the exchange rate into prices. The point estimates are consistent with a markup elasticity of 1.5, which implies a 40 percent pass-through for purely idiosyncratic shocks. We also evaluate the sensitivity of firm pricing to shocks to competitors by measuring the response of prices to movements in the U.S. trade-weighted exchange rate that is orthogonal to the bilateral exchange rate for the country. We find the response to be sizeable and significant. In a similar vein, when comparing the response to bilateral exchange rate shocks versus trade-weighted exchange rate shocks, we find that exchange rate pass-through is higher in response to a more aggregate shock than to more idiosyncratic shocks.

We also relate the incompleteness in pass-through to certain sectoral features that proxy for the level of competition among importers. An important distinction between retail prices and wholesale prices is that the latter capture business-to-business transactions. Consequently, the strength of the buyer's bargaining power can impact the extent of pass-through. We use unpublished measures of concentration in the import sector provided to us by the BLS—specifically, the Herfindahl index and the number of importers that make up the top 50 percent of trade—to evaluate this hypothesis. While the point estimates in many cases suggest that sectors dominated by a few large importers have lower pass-through from foreign firms, the estimated standard errors are large.

Lastly, we use estimates from the data to calibrate a closed economy model with different degrees of variable markup elasticity at the wholesale and retail level. In the existing monetary literature there is typically no interesting distinction made between the retail and wholesale sectors. We calibrate the parameters for the wholesale sector, using the evidence from international prices. In the benchmark model, we use Calvo price setting and later evaluate the case of menu cost pricing. First, we show that sluggishness in response to monetary shocks in wholesale prices feeds into slow adjustment of retail prices. However, inflation, as measured by the aggregate inflation and reset-price inflation series, exhibits little persistence, since the movement of these series is dominated by more transitory shocks. Yet, conditional on monetary shocks or exchange-rate-like shocks, inflation series exhibit considerable persistence. Similarly, output series can exhibit significant monetary non-neutralities. Second, while calibrated real rigidities in the form of variable markups increase the size of the contract multiplier, their effects are limited unless they are coupled with exogenous sources of persistence. The model, however, fails to match the slow dynamic in price adjustment that was documented in the empirical data, suggesting that additional sources of persistence are missing from the model.

Why does one observe differences in markup variability at the wholesale and retail level? We do not provide a definitive answer here, but we conjecture that this can be consistent with differences in the competitive environment at the two levels. That is, the retail sector can be described as monopolistically competitive, while the wholesale sector is better described as a bilateral bargaining environment. We present a static bargaining model of wholesale price setting that results in variable markups and incomplete pass-through of shocks into wholesale prices. Specifically, each final good producer bargains with its intermediate good suppliers regarding the price of intermediate goods. Given these bargained prices, the final good producer is free to choose quantities of the intermediate inputs, as well as to set the price of its final good in the monopolistically competitive consumer market. This model results in constant markups at the retail stage, but in variable markups at the wholesale level that depend, among other things, on the relative bargaining power of the final good producer and on the market share of the intermediate good supplier.

An important outstanding question is whether wholesale prices are allocative and also whether contracts specify fixed prices at fixed quantities. While there is no simple way to test this, Gopinath and Rigobon (2008) show that in the case of contracts for international prices they typically involve a fixed price with a quantity range specified, as opposed to a fixed quantity. Moreover, firms export the same good at the same price to multiple destinations and consequently prices behave in many cases like list prices. Further, the behavior of prices is consistent with models of monopolistic price setting, where prices are allocative, as discussed in the papers of Gopinath, Itskhoki, and Rigobon (2010), Gopinath and Itskhoki (2010), and Neiman (2009). Also, as we discuss later, changes in intermediate-good prices affect final-good prices, as they are fully passed through into retail consumer prices. These separate pieces of evidence are consistent with wholesale prices being allocative.

The paper is structured as follows. Section 2 provides a descriptive framework that spells out the sources of real rigidities that can result in sluggish price adjustment. Section 3 reviews the closed and open economy literature on real rigidities that uses micro price data. Section 4 presents new empirical results on price adjustment, using international data. Section 5 presents the closed economy model with differential markup variability in the retail and wholesale sector and sluggish price adjustment. Finally, Section 6 lays out a model of bargaining and variable markups in intermediate-good pricing. All derivations and technical details are relegated to the Appendix.

2 Real Rigidities: A Descriptive Framework

In this section we set up some notation and spell out the sources of real rigidities that can result in sluggish price adjustment. This simple descriptive framework is used to organize the discussion of empirical evidence in the following two sections.

Define the *desired price* of a firm as the price it would set if it could adjust its prices flexibly in a given economic environment. With sticky prices, a forward-looking firm sets its price as a weighted average of future desired prices. The presence of real rigidities slows down the response of the desired price to a shock. Real rigidities, powered by nominal price stickiness and staggered price adjustment (Taylor, 1980; Calvo, 1983), can be either at the aggregate/industry level or at the micro/firm level. Real rigidities at the aggregate level include round-about production structure as in Basu (1995), real wage rigidity as in Blanchard and Galí (2007), and segmented input markets as in Woodford (2003). Real rigidities at the firm level, or *strategic complementarities* in pricing, arise either from non-constant marginal cost (that is, decreasing returns to scale) as in Burstein and Hellwig (2007), or from variable markups. In turn, variable markups can be due to non-CES demand as in Kimball (1995) and Klenow and Willis (2006) or due to strategic complementarities in price setting between large firms as in Atkeson and Burstein (2008).

Empirical studies attempt to determine which of these channels of real rigidities, if any, are present in the data, as well as the plausible magnitudes of their effects. Since the empirical literature has examined evidence at both the retail and wholesale levels, we maintain this distinction in our descriptive framework.

Retail (final-good) pricing: The log desired price for a final good i at time t can be written as a log desired markup over the marginal cost of the firm:

$$\tilde{p}_{it} = \mu_{it}^R + \alpha s_t + (1-\alpha) w_t - z_{it},$$

where μ_{it}^R is the log desired markup, s_t is the log price of the intermediate input, w_t is the log price of other inputs (for example, labor), and z_{it} represents the firm-specific marginal cost or markup shock that may include an aggregate component common across a subset of firms. $\alpha \in [0,1]$ is the share of intermediate inputs in the production cost of the final-good firms.

Throughout the paper we abstract from the non-constant marginal cost channel of strategic complementarities in favor of the variable markup channel. For most purposes the two mechanisms are largely substitutable, however, variable markups can additionally explain *pricing to market*, a phenomenon with strong empirical support, as discussed in the next section. We capture variable markups in the following reduced-form way:

$$\mu_{it}^R \approx \bar{\mu}^R - \Gamma_R \big(p_{it} - p_t\big),$$

where p_t is the industry price index and $\Gamma_R \geq 0$ is the elasticity of the firm's markup with respect to its relative price. That is, a higher relative price of the firm reduces its desired markup. As we discuss in more detail in Section 5, this simple description of the desired markup is consistent with various models of variable markups cited above.

Combining the above two equations, we arrive at

$$\tilde{p}_{it} = \frac{1}{1+\Gamma_R}\big[\bar{\mu}^R + \alpha s_t + (1-\alpha)w_t - z_{it}\big] + \frac{\Gamma_R}{1+\Gamma_R}p_t. \tag{1}$$

When $\Gamma_R > 0$, the desired price of the firm increases in the prices set by the competitors of the firm, which we refer to as strategic complementarities in price setting. Under the same circumstances, a response to an idiosyncratic marginal cost/markup shock (such as z_{it}) is incomplete: the pass-through of an idiosyncratic shock into the desired price equals $1/(1+\Gamma_R) \in [0,1]$.

To summarize, sluggishness in the response of desired final-good prices may arise either due to staggered price adjustment when $\Gamma_R > 0$, or due to sluggish adjustment in the marginal cost, that is, due to a slow response of s_t and/or w_t to shocks.

Wholesale (intermediate-good) pricing: Let the log desired price of an intermediate variety j be

$$\tilde{s}_{jt} = \mu_{jt} + w_t + \phi_j e_t - a_{jt},$$

where μ_{jt} is the log desired markup, w_t is the price of inputs (for example, labor), and a_{jt} is a firm-specific marginal cost or markup shock that again may contain an aggregate component. Furthermore, we allow the cost of intermediate firms to respond directly to exchange rate fluctuations e_t with various elasticities ϕ_j. For example, some of the intermediate varieties may be produced abroad so that their marginal cost in local currency fluctuates together with the exchange rate.[3] Finally, the log price of the intermediate good s_t is an aggregate of the prices of intermediate varieties s_{jt}, which can be approximated by a geometric average: $s_t \approx \int_j s_{jt}\mathrm{d}j$.

We similarly introduce the possibility of variable markups for the intermediate varieties:

$$\mu_{jt} \approx \bar{\mu} - \Gamma\big(s_{jt} - s_t\big),$$

where Γ is again the price elasticity of markup that may vary across varieties j. We can

[3] Of course, the exchange rate can have additional indirect (general equilibrium) effects via productivity or cost of inputs. We discuss this issue in more detail in Sections 3 and 4.

rewrite the desired price of the intermediate variety as

$$\tilde{s}_{jt} = \frac{1}{1+\Gamma}\left[\bar{\mu} + w_t + \phi_j e_t - a_{jt}\right] + \frac{\Gamma}{1+\Gamma} s_t. \tag{2}$$

Purely idiosyncratic shocks wash out in the aggregate and do not affect s_t. Therefore, they get passed through into desired prices immediately with a pass-through coefficient of $1/(1+\Gamma)$. More aggregate shocks have two channels through which they affect the desired price — directly (w_t or e_t) and via the prices of competitors (s_t). Coupled with nominal price stickiness and staggered price adjustment, variable markups ($\Gamma > 0$) may generate sluggish adjustment that lasts past the periods of nominal stickiness.

Aggregate real rigidities: Various aggregate sources of real rigidities enter through a sluggish adjustment to shocks of the marginal cost component denoted by w_t. Specifically, real wage rigidities, Basu (1995) round-about production, and segmented input markets can all be captured by a slow adjustment of w_t to shocks. For concreteness, one can use the following structure to think about aggregate real rigidities (see Section 5):

$$w_t = \gamma m_t + (1-\gamma) p_t, \tag{3}$$

where m_t is aggregate nominal spending, p_t is the consumer price level, and $\gamma > 0$ is the elasticity of the cost w_t with respect to monetary (nominal spending) shocks. The smaller is γ, the more sluggish is the response of aggregate costs to monetary shocks, that is the stronger are aggregate real rigidities.

Equations (1)–(3) describe our simple framework to account for real rigidities.[4] Sluggish adjustment in the desired final-good prices $\tilde{p}_{it}$ can arise due to one of the three channels: (a) sluggish response of aggregate costs w_t (small γ); (b) sluggish adjustment of intermediate prices s_t ($\Gamma > 0$ provided $\alpha > 0$); and (c) gradual adjustment in final-good desired prices ($\Gamma_R > 0$). We now evaluate the evidence on these different sources of real rigidities.

[4] In Section 5 we combine this framework with specific models of nominal stickiness to study the quantitative predictions of dynamic price-setting models.

3 Evidence on Real Rigidities: A Review

Although appealing at the intuitive level, real rigidities are hard to identify and measure in the data. Aggregate real rigidities imply a sluggish response of the marginal cost of firms to aggregate shocks. Firm-level real rigidities imply a muted response of the firms' prices, conditional on price adjustment, to the marginal cost shocks. Since data on marginal costs are usually unavailable, it is hard to test these mechanisms directly. Therefore, the literature relies either on calibrations or on indirect empirical tests.

In parallel, the international literature has a long tradition of estimating exchange rate pass-through, namely, the response of international prices to changes in exchange rates.[5] A standard assumption in the empirical pass-through literature is that changes in the exchange rate represent a shock to a firm's costs that are orthogonal to other shocks that affect the firm's pricing decision and in reverse are unaffected by the firm's pricing. This assumption is motivated by the empirical finding that exchange rate movements are disconnected from most macro variables at the frequencies studied in the literature (1–2 year horizons).[6] 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. Alternatively, the exchange rate change can be viewed not as a purely exogenous cost shock, but rather as an observable signal about the underlying fundamental macro shocks that differentially impact the costs of domestic and foreign firms. In this case, the coefficients on the exchange rate have a less structural interpretation, but can still provide useful information about the nature of firm pricing.

The estimated exchange rate pass-through has typically been found to be incomplete. This incompleteness in pass-through can be consistent with various sources of firm-level and aggregate real rigidities, as well as with rational inattention and sticky information. The fact that firms export the same good to multiple destinations provides a mechanism to distinguish the source of the real rigidity. The pervasive evidence on *pricing to market*, the practice of charging different prices for the same good in different locations, provides support for the

[5] Goldberg and Knetter (1997) provide a survey of the earlier empirical pass-through literature. The specification of pass-through regressions can either be in real or in nominal terms and accordingly can include the real or the nominal exchange rate. Since the two move closely at most horizons, either specification gives similar results. Typical pass-through regressions also include some controls for the foreign cost level, such as the producer price index or manufacturing wage rate.

[6] See the literature following the seminal observation of Meese and Rogoff (1983).

variable markup channel of incomplete pass-through.

In Sections 3.1 and 3.2, we review the empirical evidence on real rigidities in the closed and open economy literature, respectively. We restrict attention to the most recent literature that uses micro-level price data.

3.1 Closed economy literature

In the closed economy literature, the evidence on real rigidities typically arises from indirect identification strategies. Klenow and Willis (2006) evaluate the variable markup channel of real rigidities by calibrating a menu cost model with Kimball (1995) demand to match the BLS micro retail price evidence on the frequency and size of price adjustment. They conclude that the levels of real rigidity sufficient to generate significant monetary non-neutrality have implausible implications for the required size of menu costs and idiosyncratic productivity shocks. Intuitively, significant real rigidities compress price dispersion so that much larger idiosyncratic shocks are required to match the size of price adjustment. At the same time, large menu costs are required to match the average durations of price rigidity, given large idiosyncratic shocks.

Burstein and Hellwig (2007) replace Kimball demand with increasing marginal costs at the firm level to generate strategic complementarity in price setting. They calibrate the extent of decreasing returns to scale (elasticity of the marginal cost with respect to output) in order to match scanner data from a large chain of supermarkets in the Chicago area (Dominick's) on the co-movement between prices and market shares. This calibration implies a moderate role for strategic complementarities generated via curvature in marginal costs.

BKM develop a test to assess a broad class of models with real rigidities. They examine the persistence of the actual and reset-price inflation series. Reset-price inflation is defined as an average reset price change within the group of goods that adjust prices in a given period, while reset prices of all other goods are indexed by reset-price inflation (for a formal definition see Section 4.1). If real rigidities are important, significant persistence should be evident in both the actual and reset-price inflation series. The paper finds low persistence for actual price inflation and negative persistence for reset-price inflation (in both cases, persistence is measured as the monthly autoregressive coefficient), suggesting that the selection mechanism present in menu costs models (emphasized by Caplin and Spulber, 1987) offsets the effects of real rigidities and that the real effects of money persist for a shorter period than the price effects do.

In a different paper, Klenow and Willis (2007) find that micro price changes respond to old macro information, known prior to their previous price adjustment. They interpret this as evidence of sticky information in the spirit of Mankiw and Reis (2002); nevertheless, this evidence is also consistent with the presence of real rigidities and pricing complementarities that lead to incomplete price adjustment at the micro level.

As previously mentioned, these tests for real rigidities are indirect, relying mainly on calibrations or statistical properties of observable prices, given that marginal costs and markups are not directly observable. The exception to this is a paper by Eichenbaum, Jaimovich, and Rebelo (2009), who use scanner data for a large grocery store chain in the United States. They observe both the wholesale price at which the store obtains the good (specific UPC) and the retail price at which the store sells it. At very short horizons the wholesale price can be viewed as the relevant marginal cost for the firm. They find that, conditional on changing reference prices, there is no evidence of variable markups, that is, all of the reference wholesale cost change is passed through into reference retail prices.[7]

3.2 Open Economy literature

The open economy literature evaluates the response of retail prices and at-the-dock prices to exchange rate shocks. While pass-through is less than one in both cases, pass-through into retail prices is always much lower than into at-the-dock prices. This is not surprising, given that distribution costs, which are mainly for non-traded costs unaffected by the exchange rate, are an important component of retail prices (for example, see Goldberg and Campa, 2010). Goldberg and Knetter (1997) summarize a large body of the earlier empirical literature on pass-through.

A virtue of international price data is that one observes the prices at which the same firm sells its product in different destinations. There is considerable evidence of pricing to market, that is, firms sell the same product at different prices in different destinations. This concept, first proposed in Dornbusch (1987) and Krugman (1987), attributes an important role to strategic complementarities that generate variable markups. In a recent paper, Fitzgerald

[7] Reference price (cost) is defined as the most often quoted price (cost) within a given time period.

and Haller (2008) provide the most direct evidence of this phenomenon. They examine the pricing of Irish manufactures in domestic and export (U.K.) markets. They find a pronounced price differential response to exchange rate movements. Since the goods are manufactured in the same plant, the difference is attributable to variation in markups.[8] Similarly, Burstein and Jaimovich (2008) use supermarket scanner data for stores in the United States and Canada and find that products with the same UPC code produced in the same country sell at different wholesale prices in the two countries. More specifically, relative wholesale prices in the two markets move closely with the exchange rate.

In a recent set of papers, Gopinath and Rigobon (2008), Gopinath, Itskhoki, and Rigobon (2010) and Gopinath and Itskhoki (2010) use micro import price data collected by the U.S. Bureau of Labor Statistics for the period 1994–2005 to provide evidence of incomplete exchange rate pass-through at the dock, even conditional on prices being changed. While incomplete exchange rate pass-through per se is consistent with evidence on variable markups, it can arise from other sources of real rigidities, as previously discussed. Gopinath and Itskhoki (2010) document a positive correlation between the frequency of price adjustment for a good and exchange rate pass-through, conditional on price adjustment. They argue that this positive correlation is consistent with variation in markup elasticity across firms and cannot be consistent with other sources of firm heterogeneity. Neiman (2010) uses the same BLS import price data to document that, consistent with the greater importance of the strategic complementarity channel for arms-length transactions than for intra-firm transactions, for differentiated products, intra-firm prices are characterized by more stickiness, less synchronization, and greater exchange rate pass-through.

Lastly, there are papers that evaluate the response of retail prices and wholesale prices for the same good to an exchange rate shock. Gopinath, Gourinchas, Hsieh, and Li (2009) use scanner data on retail and wholesale prices for stores of the same supermarket chain in the United States and Canada and find that, in response to an exchange rate shock, movements in relative retail prices in a common currency are explained mainly by movements in relative costs and very little by relative markups. Goldberg and Hellerstein (2006) and Nakamura and Zerom (2010) find that for beer and coffee sales, respectively, in a supermarket store

[8] Specifically, Fitzgerald and Haller (2008) find that, conditional on price adjustment in both markets, the common-currency price differential across the two markets moves nearly one-to-one with the nominal exchange rate. According to the descriptive model (2) of Section 2, this is consistent with a $\Gamma \gg 0$, provided that the competitor prices s_t in the two markets move little with the exchange rate, as in the model of Atkeson and Burstein (2008).

the pass-through from exchange rate shocks to wholesale costs is incomplete, but the pass-through from wholesale costs to retail prices is close to complete. In a recent study, Berger, Faust, Rogers, and Steverson (2009) match goods in the BLS import price index to those in the BLS consumer price index for the period 1994–2007 and find that the overall distribution wedge, which is the percentage difference between retail and at-the-dock prices, does not vary systematically with the exchange rate, which implies a nearly complete pass-through from at-the-dock to retail prices.[9]

Summary: Overall, a consistent finding across studies in the closed and open economy literature is that the variable markup channel of real rigidities is an important feature of the wholesale cost data but not of the retail price data. This is a surprising consensus, given the different approaches used in the closed and the open economy literature and the different datasets involved. In the terminology of Section 2, the empirical evidence is consistent with $\Gamma_R \approx 0$ and $\Gamma > 0$. Given the feedback from wholesale prices s_t to retail prices p_t in equation (1), one could expect sluggish adjustment in wholesale reset prices to generate slow adjustment in retail reset prices. The fact that BKM find a negative persistence for overall reset-price inflation could then be consistent with either a small α or a large transitory variance for z_{it}, as long as these shocks are not purely idiosyncratic. In the next section we explore explicitly, using international price data, the dynamics of price adjustment both unconditionally and conditioning on exchange rate shocks.

4 New Evidence on Real Rigidities

As discussed in Section 3.1, BKM find that regular and reset-price inflation for retail price data is consistent with the absence of important real rigidities. In this section, we present evidence on the properties of regular and reset-price inflation for import price data, using the BLS micro data on international prices. We then compare it with the properties of the inflation series, conditional on exchange rate shocks, which allows us to evaluate the

[9]Additionally, a recent study by Berman, Martin, and Mayer (2009) compares the extent of markup variability across French exporters of different sizes. It finds that larger exporters have lower exchange rate pass-through, which is consistent with a model in which both the level and the variability of the markup increases with the productivity and hence the size of the firm. Since exporting firms are typically larger, the implication of this finding is that the extent of markup variability can be greater in the international data than in the full sample of firms producing for the domestic market.

conditional response of prices to a given aggregate cost shock. We then study the dynamic properties of price adjustment at the firm level.

First, we show that reset-price inflation for the import price data is less negatively autocorrelated than that documented for retail prices. This is consistent with the conclusion that international prices depict higher real rigidities than retail prices do. Second, we present evidence regarding the sluggish response of price changes to exchange rate shocks. We do this in two ways. One, we project aggregate regular and reset-price inflation on lags of the exchange rate changes. We find that the autocorrelation of the fitted series is substantially higher than that of the unconditional series. Two, using the micro data, we show that individual prices, conditional on changing, respond to exchange rate shocks that were realized prior to the previous price adjustment. Both pieces of evidence suggest a more important role for real rigidities than do the conclusions one can draw from analyzing unconditional aggregate inflation series.[10]

While this evidence is supportive of the presence of sizeable real rigidities, it does not discriminate between various sources of real rigidities. Therefore, we next try to assess specifically the importance of strategic complementarities in price setting. We do so by studying the response of a firm's prices to the shocks to its competitors. Our central result here is that a firm's prices respond strongly to the prices of its competitors and that this channel explains a significant fraction of exchange rate pass-through into prices.

The data used in this section are the micro import price data underlying the construction of the U.S. import price index. These data cover the period 1994–2005. For details regarding the data see Gopinath and Rigobon (2008).

4.1 Reset-price inflation

We follow BKM in estimating the reset-price inflation series for U.S. imports. The log price of a good i at time t is denoted by p_{it}. The log of the reset price at time t for good i is denoted by p^*_{it} and defined as:

$$p^*_{it} = \begin{cases} p_{it} & \text{if} \quad p_{it} \neq p_{i,t-1}, \\ p^*_{i,t-1} + \pi^*_t & \text{if} \quad p_{it} = p_{i,t-1}, \end{cases} \tag{4}$$

[10] This feature is consistent with the evidence in Boivin, Giannoni, and Mihov (2009), who find that while overall disaggregated prices are volatile, they are sluggish in response to aggregate macro shocks.

where π_t^* is the average reset-price inflation of those goods whose prices change at time t. The inflation of the actual price series is referred to as *regular-price inflation.* In a Calvo pricing environment, where firms differ only in the exogenous frequency of price adjustment, the behavior of reset-price inflation will capture the extent of real rigidities. As discussed in detail in BKM, in the presence of real rigidities reset prices will adjust sluggishly, as firms have multiple price adjustments before they respond fully to a shock. The regular-price inflation series display even greater persistence, as each firm waits for the random arrival of the opportunity to change its price. In the case of state-dependent pricing or variations in desired price responses across sectors, the behavior of measured reset prices is not as direct because there is selection of which firms change prices. Gopinath and Itskhoki (2010) show that sectors that have a higher frequency of price adjustment also have higher long-run pass-through (fewer real rigidities). Consequently, the reset-price series is affected by those goods that change prices more frequently and eventually pass through a lot more.

It is important to point out that the sample sizes in the import price data are not as large as in the CPI data, and given the low frequency of price adjustment, the number of actual price changes used to impute the reset price series is low.[11] This would suggest that the import reset-price inflation series is more subject to noise than is the construct of BKM for the U.S. CPI.

The volatility and persistence of the regular and reset-price inflation series within various subsamples are reported in Table 1.[12] We measure persistence as the AR(1) coefficient for the series and we report the standard deviation of the series as a measure of volatility.[13] The first row of Table 1 reports these moments for consumer-price inflation from BKM. The rest of the rows provide these statistics for the BLS import-price data.[14]

[11]In the sample that excludes petrol classifications, the median number of price observations per month is 6,335 and the median number of price observations whose price is different from the previous month is 770.

[12]The small number of price changes limits the analysis to large groups of goods.

[13]We estimate an AR(1) coefficient so as to compare our results directly to BKM. If, as argued by Stock and Watson (2007), CPI inflation is better modeled as an ARMA(1,1) or an IMA(1,1), then the first-order autocorrelation understates the long-run persistence of the series. We have estimated other measures of persistence, such as the variance ratio to the long-run variance of the series. Although these measures suggest greater persistence for the inflation series, our comparative results for the conditional and unconditional inflation still hold.

[14]In all specifications, we exclude petrol classifications. For each series we use 2002 weights at the 4-digit level to aggregate across prices. More precisely, for actual inflation, we estimate mean price change by 4-digit harmonized code for each month, then we average across the different harmonized codes using weights at the 4-digit level. For a reset-price inflation, we assume that at the start of our sample in January 1994 all prices were reset prices, which is an initial condition assumption. Then, we follow the formula in equation (4) to construct reset-price inflation, π_t^*.

Table 1: Volatility and persistence of regular and reset-price inflation

	Unconditional		Conditional on ER	
	AR(1)	St.D.	AR(1)	St.D.
	Regular-price inflation			
Consumer prices (from BKM)	−0.05	0.14%	—	—
Import prices	0.51	0.33%	0.55	0.25%
– Dollar-priced goods	0.56	0.27%	0.79	0.18%
– Market transactions	0.43	0.30%	0.70	0.20%
	Reset-price inflation			
Consumer prices (from BKM)	−0.41	0.95%	—	—
Import prices	0.02	1.60%	0.31	0.82%
– Dollar-priced goods	−0.04	1.20%	0.33	0.75%
– Market transactions	−0.03	1.70%	0.17	0.91%

Note: Import prices exclude petrol classifications; the rows for market transactions include only dollar-price goods. The last two columns project the inflation series on the current and 24 lags of the log changes of the U.S. trade-weighted exchange rate and compute the moments for the projected series.

Columns 1 and 2 report statistics for the unconditional inflation series. For all subsamples considered, import price inflation is more persistent than consumer-price inflation as calculated by BKM.[15] For example, for the dollar-priced imports, the persistence of inflation is 0.56, while for consumer prices it is close to zero (−0.05). The import-price inflation is also more volatile. A similar comparison holds for reset-price inflation. While BKM find that reset-price inflation for consumer prices is negatively autocorrelated (−0.41), we find essentially zero autocorrelation for import prices (for example, −0.04 for the dollar-priced imports). This difference in persistence of consumer-price and import-price inflation is consistent with the different nature of pricing at the consumer and intermediate-good levels (since most imports are intermediate goods).[16]

An import feature of the international data is that we can examine the response of inflation series to a specific shock, namely, the exchange rate shock. This has advantages over just looking at reset-price inflation that aggregates (imperfectly in small samples) across

[15] BKM sample is longer, from 1989 to 2008. They also exclude energy, fresh fruit and vegetables, and eggs. We report their results for the sample that excludes sales price.

[16] As documented by Stock and Watson (2007) among others, the short-run persistence of consumer-price inflation decreased in the 1990s.

idiosyncratic, sectoral and aggregate shocks. We accordingly project the regular and reset-price inflation series on current and 24 lags of the log changes of the U.S. trade-weighted nominal exchange rate. We use the fitted values from this regression and estimate the AR(1) coefficient and standard deviation for the fitted series. The results are reported in the last two columns of Table 1. In all cases regular and reset-price inflation, conditional (projected) on the exchange rate shocks, exhibits more persistence. For instance, in the case of dollar-priced imports, the conditional regular-price inflation series has an overall persistence of 0.79, while its unconditional persistence is 0.56. Similarly, for the reset-price inflation series the conditional persistence is 0.33 as opposed to the unconditional persistence of -0.04.

This evidence is consistent with the presence of multiple shocks of different degrees of persistence driving the inflation process. Under these circumstances the unconditional persistence of the inflation series might not accurately reflect the underlying sluggishness in the micro-level price adjustment. In the next subsection, we present further evidence of the sluggish response of prices to exchange rate shocks by examining the behavior of individual prices and their response to lagged exchange rates changes.

4.2 Micro-dynamics of price adjustment

At the good level, we estimate pass-through into prices of exchange rate shocks realized during the most recent period of price non-adjustment and of those that were realized prior to the previous price adjustment. In the absence of real rigidities, all adjustment should take place at the first instance of price change and hence the coefficient on the exchange rate change prior to the previous price adjustment should be zero.

Formally, we estimate the following regression:

$$\Delta \bar{p}_{i,t} = \beta_1 \Delta_{\tau_1} e_{i,t} + \beta_2 \Delta_{\tau_2} e_{i,t-\tau_1} + Z'_{i,t}\gamma + \epsilon_{it}, \tag{5}$$

where i indexes the good, $\Delta \bar{p}_{i,t}$ is the change in the log dollar price of the good, *conditional on price adjustment in the currency of pricing.*[17] $\Delta_{\tau_1} e_{i,t} \equiv e_{i,t} - e_{i,t-\tau_1}$ is the cumulative change in the log of the bilateral nominal exchange rate over the duration when the previous price

[17] In the BLS database, the original reported price (in the currency of pricing) and the dollar converted price are both provided. We use the latter, conditional on the original reported price having changed. Since the first price adjustment is censored from the data, we also perform the analysis excluding the first price change and find that the results are not sensitive to this assumption.

was in effect (which we denote τ_1). Similarly, τ_2 denotes the duration of the previous price of the firm, so that $\Delta_{\tau_2} e_{i,t-\tau_1} \equiv e_{i,t-\tau_1} - e_{i,t-\tau_1-\tau_2}$ is the cumulative exchange rate change over the previous (the one prior to the previous price change) period of non-adjustment. Figure 1 illustrates a hypothetical price series: if $\Delta \bar{p}_{i,t}$ is the price change between t_3 and t_{LL}, $\Delta_{\tau_1} e_{i,t}$ is then the exchange rate change between t_3 and t_{LL}, and $\Delta_{\tau_2} e_{i,t-\tau_1}$ is the exchange rate change between t_2 and t_3. Finally, $Z_{i,t}$ includes controls for the cumulative change in the foreign consumer price level, the U.S. consumer price level and fixed effects for every BLS-defined primary stratum (mostly 2–4-digit harmonized codes) and country pair. We allow Z_t to include lagged foreign and domestic inflation. The standard errors are clustered at the level of the fixed effects. We restrict the sample to non-petrol, dollar-priced goods and market transactions. Note that this specification requires a good to have at least two price adjustments during its life. Since there are several goods that have only one price change during their respective lives, we lose about 30 percent of the goods.

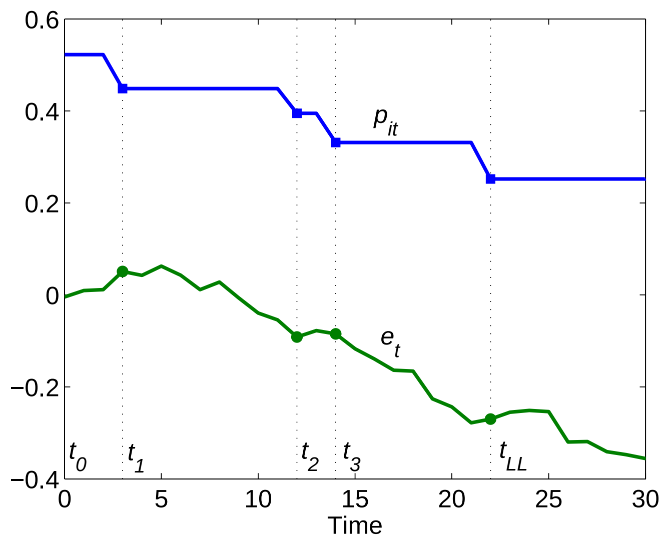

Figure 1: Hypothetical good-level price series and nominal exchange rate

By conditioning on a price change, we get past the period of nominal rigidity, which is essential to understanding the role of real rigidities. In a Calvo pricing environment, since the decision to change prices is exogenous, there are no selection issues to be concerned with. However, in an environment with endogenous frequency of price adjustment, conditioning on a price change will induce a bias in the exchange rate pass-through estimates, as it generates a conditional correlation between the exchange rate and the residual even if the unconditional correlation is zero. This is problematic if one tries to provide a structural interpretation to the coefficient. That is not our purpose here. We use this specification to provide a relation in the data between the response of prices, conditional on adjusting to lagged exchange rate changes. Later, we will estimate these regressions in the model-simulated data, where frequency is chosen endogenously, and infer how well models with real rigidities perform in matching the facts in the data. This exercise is accordingly similar to that of trying to match the behavior of the reset-price inflation series, a series that also is affected by selection issues. Furthermore, when the selection bias is strong, the pass-through coefficient on first adjustment (β_1) is biased upwards, while the pass-through coefficient on second adjustment (β_2) is biased downwards, making it harder to identify the presence of real rigidities.

Table 2: Dynamic response to exchange rate shocks

	β_1	$s.e.(\beta_1)$	β_2	$s.e.(\beta_2)$	N_{obs}	R^2
All countries	0.11	(0.02)	0.08	(0.01)	69,917	0.01
Non-OECD	0.06	(0.02)	0.04	(0.01)	37,108	0.01
High-income OECD	0.23	(0.02)	0.18	(0.02)	32,809	0.02
Euro area	0.22	(0.04)	0.14	(0.03)	5,933	0.02
Japan	0.26	(0.04)	0.24	(0.03)	4,249	0.06
Canada	0.28	(0.16)	0.34	(0.05)	14,620	0.01
Differentiated goods	0.14	(0.02)	0.10	(0.02)	21,360	0.02
No missing prices						
All countries	0.10	(0.02)	0.08	(0.01)	45,765	0.01
High-income OECD	0.19	(0.03)	0.13	(0.02)	22,436	0.01

Note: β_1 and β_2 are the pass-through coefficients at the first and second rounds of price adjustment, respectively, estimated from regression (5). *Standard errors in brackets are clustered at the country $\times$ 4-digit-sector level. N_{obs} is the number of price changes in the sample. Results under "No missing prices" in the lower panel exclude from the sample all price changes that were followed or preceded by a missing price.*

Table 2 reports the results from estimation of regression (5). We provide evidence for various subsamples of the data. Across all specifications we find that exchange rate shocks that took place prior to the current period of non-adjustment have a significant effect on current price adjustments. This is consistent with the existence of real rigidities in pricing. The strength of these lagged effects is much stronger than what would be suggested purely by the reset-price inflation series. The first row of Table 2 points out that the elasticity of current price changes to lagged exchange rate shocks for dollar-priced goods is 0.08, which is only slightly smaller than the response to the contemporaneous exchange rate movement (equal to 0.11). The importance of these lagged effects is consistently present in all sub-samples. For the high-income OECD sample, the contemporaneous and lagged responses are 0.23 and 0.18, respectively. For the non-OECD sample, the pass-through rates are overall lower, but there are still important lagged effects. This is similarly documented for the Euro-area countries, Japan, and Canada.

In the data, there can be spells that have missing prices and where the new price follows or precedes a missing price. In this case, the exact timing of the price change is not known, so lagged effects can arise from getting the timing wrong. The last two rows of Table 2 check for the robustness of the results by including in the sample only those price changes that were not followed or preceded by a missing price. This changes the sample composition, but lagged responses are still strongly evident.

The results in this section are consistent with the evidence in Gopinath, Itskhoki, and Rigobon (2010) that long-run pass-through is much higher than pass-through conditional on the first adjustment to the exchange rate shock. Here, we present explicitly the dynamics and extend the sample to more countries.

We also divide goods into four equally sized bins based on their frequency of price adjustment, and estimate equation 5 within each bin separately. The purpose of this exercise is to evaluate whether the importance of lags varies across goods with different frequencies of price adjustment. One conjecture may be that it is only the very high-frequency goods that have multiple price adjustments to respond to a shock. In fact, this is not the case, as we find that lags are important even for goods that adjust prices very infrequently: for example, the first quartile contains goods that adjust prices less than once a year and in the first round the pass-through is 0.12, while it is 0.08 in the second round. This finding further assuages concerns about measurement issues with the timing of price adjustment.[18]

We also break the sample down by the ‘end use’ of the product. Again, we find the second rounds of price adjustment to be significant for most end-use categories. In the case of ‘Food, feed and beverages’ exported from the non-OECD counties (which dominate the sample of all countries in Panel A), the pass-through is generally very low and insignificantly different from zero. In the case of ‘Consumer goods (non-food and excluding automotive)’ the dynamic is less evident. However, one should be careful about interpreting the results for this subsample

[18]Refer to Gopinath and Itskhoki (2010) for a detailed analysis of pass-through conditional on first adjustment and pass-through conditional on many rounds of adjustment across goods with different frequencies.

Table 3: Dynamic response to exchange rate shocks: by end-use sectors

	β_1	$s.e.(\beta_1)$	β_2	$s.e.(\beta_2)$	N_{obs}	R^2
			Panel A: All countries			
Food, feed and beverages	0.05	(0.04)	0.03	(0.02)	17,731	0.01
Industrial supplies and materials	0.13	(0.03)	0.11	(0.02)	22,396	0.02
Capital goods, except automotive	0.16	(0.03)	0.17	(0.03)	4,220	0.05
Consumer goods (non-food)	0.05	(0.04)	0.03	(0.02)	8,222	0.01
			Panel B: High-income OECD			
Food, feed and beverages	0.16	(0.04)	0.15	(0.03)	5,207	0.03
Industrial supplies and materials	0.21	(0.05)	0.25	(0.03)	13,089	0.01
Capital goods, except automotive	0.21	(0.04)	0.22	(0.03)	2,587	0.05
Consumer goods (non-food)	0.17	(0.06)	0.03	(0.04)	2,915	0.02

Note: see notes to Table 2.

because these goods more often have a fixed price during their life and then get discontinued and replaced. Since we do not observe price changes across discontinuations, we might be excluding important adjustments that take place at the time of product replacement.

Overall, the micro evidence is consistent with the aggregate level evidence of sluggish adjustment to exchange rate changes. This sluggishness is consistent with many forms of real rigidities, including variable markups, the Basu (1995) intermediate input channel wherein each firm's output is used as an input in production, and sluggish response of other factor costs (like wages) to the underlying source of exchange rate shocks. Equivalently, it could arise from rational inattention or sticky information. The next subsection evaluates how the extent of product market competition affects the patterns of pass-through, in order to identify the effects of strategic complementarities in price setting.

4.3 Competition and pass-through

In this section, we evaluate the importance of the strategic complementarity in price setting for incomplete pass-through, using some measures that capture the pricing behavior of competitors and measures that capture the extent of competition within sectors. Ideally, one would need to perform this analysis with detailed industry data for each product and information on prices and market shares of different firms. These data, however, do not exist for the large number of products included in our study. Consequently, we use some proxies here and they are necessarily imperfect. Nevertheless, we find evidence that is consistent with the presence of significant strategic complementarities at the firm level.

Trade-weighted versus bilateral exchange rate: First, we evaluate the response of each firm's pricing to its own bilateral exchange rate as compared with its response to the trade-weighted exchange rate. Movements in the trade-weighted exchange rate can be viewed as a more aggregate shock that affects a larger fraction of a firm's competitors than does a shock that affects only the bilateral exchange rate. An alternative interpretation could be that prices of a firm that uses production inputs from the rest of the world are sensitive to the movements in the trade-weighted exchange rate because these prices impact the firm's costs.

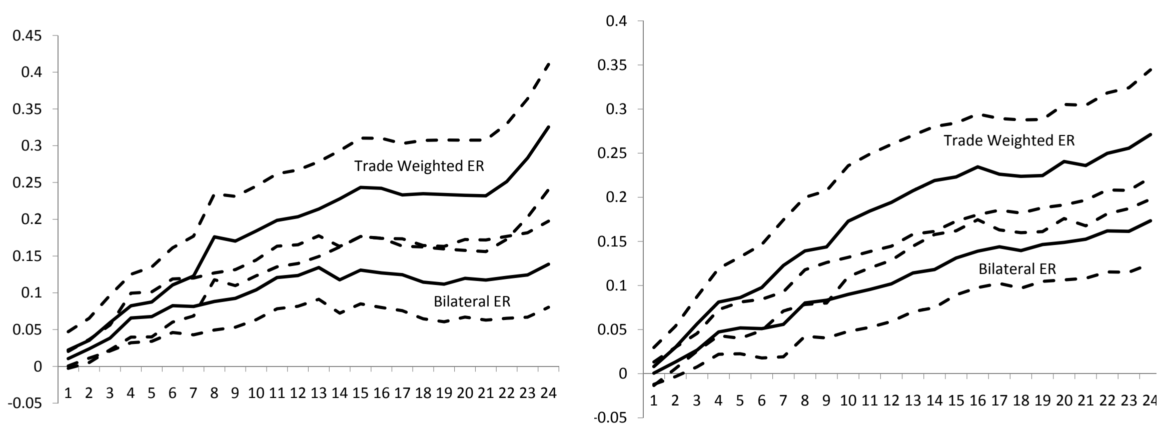

(a) All countries (b) High-income OECD

Figure 2: Impulse responses to bilateral and U.S. trade-weighted exchange rate

More specifically, we estimate the following standard pass-through regression,

$$\Delta p_{k,t} = \alpha_i + \sum_{j=0}^{n} \beta_j \Delta e_{k,t-j} + \sum_{j=0}^{n} \gamma_j \pi_{k,t-j} + \epsilon_{k,t}, \tag{6}$$

where k indexes the country, Δp is the average monthly log price change in dollars, π is the monthly foreign-country inflation using the consumer price index, n is the number of monthly lags that vary from 1 to 24. $\Delta e_{k,t-j}$ is either a bilateral nominal exchange rate or the U.S. trade-weighted nominal exchange rate. Figure 2 plots $\sum_{j=0}^{n} \beta_j$ as a function of n for each case, where we estimate a pooled regression restricting the coefficients β_j to be the same across countries. Panel (a) of the figure plots the results for the all-countries sample, while Panel (b) does it for the high-income OECD subsample. In both figures it is evident that the pass-through from the trade-weighted exchange rate exceeds the bilateral exchange rate pass-through, consistent with the hypothesis that a firm's prices are responsive to cost shocks of a firm's competitors. In further analysis, we find that this pattern is evident for countries in the Euro area as well as in the non-OECD countries, while it is less evident for Japan, Canada, and the United Kingdom.

Table 4: Pass-through of bilateral and U.S. trade-weighted exchange rate

	Bilateral ER		T-W ER		N_{obs}	R^2
All countries	0.11	(0.01)	0.19	(0.02)	83,064	0.01
Non-OECD	0.07	(0.02)	0.18	(0.02)	46,420	0.01
High-income OECD	0.22	(0.02)	0.17	(0.05)	32,809	0.02
Euro area	0.27	(0.03)	0.31	(0.07)	7,856	0.03
Japan	0.21	(0.04)	0.17	(0.06)	5,733	0.02
Canada	0.23	(0.12)	0.12	(0.10)	16,221	0.01
Differentiated	0.12	(0.01)	0.17	(0.03)	21,360	0.02

Note: The first column reports pass-through conditional on price adjustment of the bilateral exchange rate shocks. The second column reports the pass-through of the component of the U.S. trade-weighted exchange rate orthogonal to the bilateral exchange rate (that is, a residual from the projection of the trade-weighted exchange rate on the bilateral exchange rate). Clustered standard errors in brackets.

We also perform the analysis using the individual price data, conditional on a price change. We evaluate the response to the bilateral exchange rate change since the last time the price was adjusted and to movements in the U.S. trade-weighted exchange rate that are orthogonal to the bilateral exchange rate for the country. More specifically, we run a first-stage regression, where we regress the trade-weighted exchange rate on the bilateral exchange rate. We calculate the residual and then estimate a second-stage regression, where we regress the price change, conditional on adjustment, on the cumulative change in the bilateral exchange rate and in the residual.[19] We include a control for the cumulative change in foreign-country CPI inflation since the last price change. The results are reported in Table 4. Consistent with the evidence in Figure 2 using aggregate price changes, the effect of the residual is almost as large as the direct effect of the bilateral exchange rate, and in some cases it is even larger.

[19]The coefficient on the residual will be equivalent to the coefficient on the trade-weighted exchange rate obtained from regressing the price change on the bilateral and the trade-weighted exchange rate, but the coefficient on the bilateral exchange rate will be different across the two specifications.

We should note that this evidence admits a number of interpretations. First, the sensitivity of prices to the trade-weighted exchange rate over and above the bilateral exchange rate is consistent with strategic complementarities in price setting when the competitors of the firm are exporters from other countries. However, it is also consistent with the fact that firms use production inputs imported from the rest of the world, and a movement in the trade-weighted exchange rate could therefore represent a relevant cost shock for the firm. In what follows we test for strategic complementarities in price setting more directly.

Competitor prices: A more direct test of the presence of strategic complementarities is to evaluate whether changes in the competitor prices affect the pricing decisions of the firm. We do so by estimating the following regression:

$$\Delta \bar{p}_{i,k,t} = \beta_e \Delta_{\tau_1} e_{i,k,t} + \beta_I \Delta_{\tau_1} P^I_{k,t} + \gamma Z_{i,t} + \epsilon_{i,t}, \tag{7}$$

where $\Delta \bar{p}_{i,k,t}$ is the change in the log dollar price of good i in sector k, conditional on price adjustment, and $\Delta_{\tau_1} e_{i,k,t}$ is the cumulative change in the log of the bilateral nominal exchange rate over the duration for which the previous price was in effect. Now, $\Delta_{\tau_1} P^I_{k,t}$ is a measure of the cumulative price change by firms other than firm i in sector k.[20] We also estimate the same regression for the life-long change in the price of the good and refer to the coefficients in this case as *long-run pass-through*. In terms of Figure 1, this corresponds to having $p_{t_{LL}} - p_{t_0}$ on the left-hand side of (7) and corresponding cumulative changes in variables on the right-hand side.[21] Finally, $Z_{i,t}$ represents the cumulative change in the consumer price index in the foreign country. We again restrict the sample to non-petrol, dollar-priced goods and market transactions. We include fixed effects for every BLS-defined primary stratum (mostly 2–4-digit harmonized codes) and country pair; the standard errors are clustered at the level of the fixed effects.

[20] For each good i we calculate the average monthly import price change for all goods in the same 10-digit or BLS-defined primary stratum classification, excluding good i. We then cumulate these price changes over time to arrive at an industry competitor price index for each good. In our main specification, we include non-adjacent price changes, that is, if prices are available for January and March, but are missing for February, the price change in March refers to the percentage difference between the March and January prices. We also perform the analysis where we include only price changes across adjacent months and obtain qualitatively the same results.

[21] Results are unaffected if we exclude the first price change that can be censored. For more details on the comparison between life-long pass-through and pass-through conditional on first price adjustment, see Gopinath and Itskhoki (2010).

Table 5: Response to competitor prices

	β_e	$s.e.(\beta_e)$	β_I	$s.e.(\beta_I)$	N_{obs}	R^2
Panel A: Pass-through conditional on first price change						
No $\Delta_{\tau_1} P^I_{k,t}$	0.13	(0.01)	—	—	83,056	0.01
$\Delta_{\tau_1} P^I_{k,t}$ (Primary strata)	0.07	(0.01)	0.61	(0.02)	78,942	0.13
$\Delta_{\tau_1} P^I_{k,t}$ (10-digit HTS)	0.04	(0.01)	0.61	(0.02)	59,972	0.25
Panel B: Long-run pass-through						
No $\Delta_{\tau_1} P^I_{k,t}$	0.31	(0.03)	—	—	16,145	0.06
$\Delta_{\tau_1} P^I_{k,t}$ (Primary strata)	0.13	(0.02)	0.66	(0.03)	15,273	0.24
$\Delta_{\tau_1} P^I_{k,t}$ (10-digit HTS)	0.16	(0.01)	0.62	(0.03)	11,379	0.34

Note: Results from estimation of equation (7). The coefficient γ on the consumer-price inflation in the foreign country also shrinks along with β_e when we include the control for competitor prices.

The results are reported in Table 5. The first row of each panel (labeled 'No $\Delta_{\tau_1} P^I_{k,t}$') presents the results where we exclude any industry competition effect. The next two rows include industry price effects aggregated at the BLS-defined primary stratum level (mostly 2–4-digit harmonized codes) and at the 10-digit harmonized code level, respectively. As is evident in all specifications, the effect of competitor prices is large and highly significant. Moreover, it significantly reduces the direct response of prices to the exchange rate shock. If one were to do a back-of-the-envelope calculation of the extent of strategic complementarities using expression (2) of our accounting framework in Section 2, one would obtain a measure of markup elasticity of $\Gamma \approx 1.5$ (from $\Gamma/(1+\Gamma) \approx 0.6$). This value is consistent with the required markup elasticity to match the evidence of incomplete long-run pass-through, as we discuss in Gopinath, Itskhoki, and Rigobon (2010) and Gopinath and Itskhoki (2010). Furthermore, note that the direct impact of exchange rate changes (β_e) still increases from the specification conditional on one price adjustment (Panel A) to the the life-long specification (Panel B), even when we control for competitor prices. This suggests that, although strategic complementarities are an important feature of price setting, they do not fully explain the delayed pass-through of exchange rate shocks and that there are other sources of real rigidity present in the data.

Sector concentration: Finally, we relate the incompleteness in exchange rate pass-through to certain sectoral features that proxy for the level of competition among importers. An important distinction between retail and wholesale prices is that the latter originate from business-to-business transactions. Consequently, the strength of bargaining power of the buyer can impact the extent of pass-through. To evaluate this hypothesis, we use measures of concentration in the import sectors constructed by the Bureau of Census. This includes a Herfindahl index for importers and a measure of the number of importers that make up the top 50 percent of trade, using census data on all imports entering the United States. We estimate the following long-run pass-through regression where we interact the exchange rate change with a measure of concentration (at the 2–4-digit level):

$$\Delta \bar{p}_{i,k,t} = \beta \Delta_{\tau_1} e_{i,k,t} + \psi\big(\Delta_{\tau_1} e_{i,k,t} \cdot C_k\big) + Z'_{i,k,t}\gamma + \epsilon_{i,t}, \tag{8}$$

where the second regressor is the interaction of the exchange rate change with a given concentration measure. In $Z_{i,k,t}$ we include separate controls for the concentration measure, the change in CPI inflation (both separately and interacted with the concentration measure), and country fixed effects. All standard errors are clustered at the primary-stratum level.

Table 6: Exchange rate pass-through and sectoral characteristics

	β	$s.e.(\beta)$	ψ	$s.e.(\psi)$	N_{obs}	R^2
Herfindahl index	0.30	(0.03)	-0.02	(0.02)	12,432	0.06
No. importers in top 50%	0.23	(0.02)	0.02	(0.02)	12,435	0.06

The results are reported in Table 6 for the two measures of concentration. While the point estimates in both regressions suggest that sectors that are dominated by a few large importers (high Herfindahl index and small number of firms in the top 50 percent) have lower pass-through from foreign firms, the standard errors on these estimates are large. Overall, the evidence is inconclusive. We also performed this exercise for pass-through conditional on first price change, and by restricting the sample to differentiated goods only; in all cases we found no clear relationship in the data.

Summary: The presented evidence is consistent with real rigidities having an important role in the data, reflected in a sluggish response of prices to exchange rate shocks, even conditional on prices changing. The reset-price inflation series, when conditioned on exchange rate shocks, also exhibits persistence. We further find that the response of prices is greater to trade-weighted exchange rate shocks than to bilateral shocks, and the response to competitor prices is significant. These last two pieces of evidence are consistent with the importance of strategic complementarities in price setting; however, they are also possibly consistent with other explanations. For example, the trade-weighted exchange rate can reflect changes in the prices of inputs imported from the rest of the world, while competitor prices may proxy for sectoral shocks that are not reflected in the exchange rate and are common across firms from a number of countries. Presumably, the most direct evidence of strategic complementarity arises from studies of pricing to market that control for costs, as discussed earlier in Section 3. The best description of our evidence is as important facts regarding the dynamics of price adjustment that should be matched by models of price setting. We turn to one such model in the next section.

5 Model

In this section, we quantitatively evaluate a reduced-form sticky price model with a retail and a wholesale sector. Consistent with the data, we allow for variable markups at the wholesale level and constant markups at the retail level. In the next section, we discuss a bargaining-based micro-foundation for this reduced-form assumption. In the existing monetary literature there is typically no interesting distinction made between the retail and wholesale sectors. The goal is to evaluate the behavior of regular and reset-price inflation, both unconditional and conditional on aggregate shocks, as well as the dynamic response of good-level prices conditional on changing, so as to compare it with the evidence in Section 4. We also evaluate the extent of monetary non-neutrality generated by this source of real rigidities.

The model generates sluggishness in response to monetary shocks in wholesale prices and this feeds into the slow adjustment of retail prices. However, aggregate inflation and reset-price inflation exhibit little persistence, since their movements are dominated by more transitory shocks. On the other hand, conditional on monetary shocks or exchange-rate-like shocks, inflation series exhibit considerable persistence.[22] This is consistent with the data. The model however fails to match the slow dynamics in price adjustment documented in the empirical data, suggesting a need for an additional source of persistence in prices.

[22]Note that our simple model does not rely on any mechanisms of differential response to aggregate versus idiosyncratic shocks as, for example, in Maćkowiak and Wiederholt (2009).

Similarly, the output series can exhibit significant monetary non-neutralities if the money growth process is sufficiently persistent. While calibrated real rigidities in the form of variable markups increase the size of the contract multiplier, their effects are modest unless they are coupled with exogenous sources of persistence.

We begin by laying out the familiar equilibrium conditions of the model while the details are relegated to the appendix. We first consider the model with Calvo price setting and later evaluate the robustness of the predictions in a menu cost model of price setting in the wholesale sector. We should clarify that this is only a numerical exercise where the empirical evidence disciplines some parameters of the model.

5.1 Setup of the model

Wholesale sector: Wholesale firms use labor and a constant returns to scale production function to produce intermediate goods. Therefore, a wholesale firm j faces a constant marginal cost mc_{jt} (all variables in logs):

$$mc_{jt} = w_t + \phi_j e_t - a_{jt},$$

where w_t is the nominal wage rate and e_t captures an exogenous exchange-rate-like shock that affects the wholesale firm with elasticity ϕ_j that varies across firms. Further, $a_{jt} \equiv \bar{a}_t + \tilde{a}_{jt}$ is the sum of an aggregate (wholesale-sector-wide) and idiosyncratic (firm-specific) shock to the firm; a_{jt} represents some combination of shocks to the marginal cost and the markup that affects the firm's desired price.

The desired log-price of a wholesale producer equals a log desired markup over the marginal cost:

$$\tilde{s}_{jt} = \mu_{jt} + mc_{jt}.$$

By desired prices we mean prices that a firm would set if it could adjust prices every period in a given general equilibrium environment; desired price is not the same as a reset price, which is set for a number of future periods.

We assume variable markups that depend on the firm's relative price:

$$\mu_{jt} = \bar{\mu} - \Gamma(s_{jt} - s_t),$$

where $\bar{\mu}$ is the steady state level of markup, Γ is the elasticity of markup with respect to price, and

$$s_t \equiv \int s_{jt} \mathrm{d}j$$

is (an approximation to) the price index in the wholesale sector. This specification is a first-order approximation to a more general model of variable markups. For example, it can be obtained from the Kimball demand with non-constant elasticity (see, for example, Klenow and Willis, 2006; Gopinath and Itskhoki, 2010), or a model of strategic interactions between large firms (see, for example, Atkeson and Burstein, 2008, and the bargaining model of Section 6).

Retail sector: In the retail sector, firms combine labor and intermediate goods supplied by the wholesale sector to produce a final good. Specifically, firm i's marginal cost is given by:

$$mc_{it}^R = \alpha s_t + (1-\alpha) w_t - z_{it},$$

where α is the production-cost share of intermediate goods, and $z_{it} \equiv \bar{z}_t + \tilde{z}_{it}$ is the sum of aggregate (retail-sector-wide) and idiosyncratic (firm-specific) marginal cost and/or markup shocks that affect the firm's desired price. Note that we assume that the exchange rate shock e_t does not affect the retail sector directly and that each retail firm uses a full bundle of intermediate goods as input in production.

We assume constant-markup pricing in the retail sector (for example, monopolistic competition and CES demand) so that the desired price of firm i is given by

$$\tilde{p}_{it} = \bar{\mu}^R + mc_{it}^R.$$

In the notation of Section 2, it is equivalent to assuming that $\Gamma_R = 0$. This assumption, along with $\Gamma > 0$, is consistent with the evidence discussed in Section 3. In Section 6, we provide one economic explanation that can rationalize this difference.

Wage rate and real output: We assume that the nominal wage rate depends on the consumer (final-good) price level p_t and aggregate nominal spending, m_t:

$$w_t = \gamma m_t + (1-\gamma) p_t, \qquad \gamma > 0. \tag{9}$$

This reduced form model of wages is common in macroeconomics (for example, see Chari, Kehoe, and McGrattan, 2000; Burstein and Hellwig, 2007) and can be derived, for example, from a cash-in-advance model of money demand and the intratemporal optimality condition for consumption-leisure choice.[23] Smaller values of γ imply a more gradual response of wages to aggregate nominal spending shocks and hence are a stand-in for various unmodeled aggregate real rigidities, such as real wage rigidity (Blanchard and Galí, 2007), segmented labor markets (Woodford, 2003), and round-about production structure (Basu, 1995).

With our definition of m_t as aggregate nominal spending, real output is given by

$$y_t = m_t - p_t.$$

Therefore, the extent of monetary non-neutrality can be measured as the persistence of y_t in response to nominal spending shocks, since in a flexible-price world exogenous m_t-shocks have no effect on real output.

Exogenous shock processes: As commonly assumed in the literature (for example, see Chari, Kehoe, and McGrattan, 2000, BKM), nominal spending m_t follows an exogenous AR(1) process in first differences:

$$\Delta m_t = \rho_m \Delta m_{t-1} + \sigma_m \epsilon_t^m,$$

where $\rho_m \geq 0$ is the measure of *exogenous persistence* in the model. All other exogenous shocks follow persistent but stationary AR(1) processes:

$$x_t = \rho_x x_{t-1} + \sigma_x \epsilon_t^x,$$

where $x_t \in \{e_t, \bar{z}_t, \bar{a}_t, \tilde{z}_{it}, \tilde{a}_{jt}\}$. All innovations ($\epsilon_t^m$ and ϵ_t^x's) are mean-zero unit-variance i.i.d. random variables.

[23] Specifically, a model with log-utility of consumption and linear disutility of labor results in $\gamma = 1$, provided there is no additional source of aggregate real rigidities. Golosov and Lucas (2007) derive the same specification in a money-in-the-utility model. Ball and Romer (1990) refer to this benchmark as the case of strategic neutrality. Aggregate real rigidities work to reduce the value of γ.

5.1.1 Calvo price-setting

In the case of Calvo price setting, a given wholesale firm j adjusts prices with probability $(1-\theta)$ each period, while for any retail firm i the adjustment probability equals $(1-\theta_R)$. Up to a first-order approximation, at the instances of adjustment the firms set their prices to the discounted expectation of their future desired prices (see the appendix for a formal derivation):

$$\begin{aligned}\bar{s}_{jt} &= (1-\beta\theta)\textstyle\sum_{\ell=0}^{\infty}(\beta\theta)^{\ell}\mathbb{E}_t\tilde{s}_{j,t+\ell},\\ \bar{p}_{it} &= (1-\beta\theta_R)\textstyle\sum_{\ell=0}^{\infty}(\beta\theta_R)^{\ell}\mathbb{E}_t\tilde{p}_{i,t+\ell},\end{aligned}$$

where β is the discount factor, $\bar{s}_{jt}$ and $\bar{p}_{it}$ are the (theoretical) reset prices, and $\tilde{s}_{j,t+\ell}$ and $\tilde{p}_{i,t+\ell}$ are the future desired prices (derived above) for the wholesale and retail firms, respectively.

Under Calvo-pricing assumptions, the dynamics of aggregate wholesale and retail prices are given, respectively, by $s_t = \theta s_{t-1} + (1-\theta)\mathbb{E}_j\bar{s}_{jt}$ and $p_t = \theta_R p_{t-1} + (1-\theta_R)\mathbb{E}_i\bar{p}_{it}$, where $\mathbb{E}_j$ and $\mathbb{E}_i$ denote the cross-sectional means. Combining these equations with the expressions for reset prices and substituting in the expressions for desired prices, we arrive at the familiar forward-looking Phillips curves—dynamic equations for aggregate wholesale and retail inflation (see the appendix):

$$\Delta s_t = \beta\mathbb{E}_t\Delta s_{t+1} + \frac{\lambda}{1+\Gamma}\left[\gamma(m_t-p_t) - (s_t-p_t) + \bar{\phi}e_t - \bar{a}_t\right], \tag{10}$$

$$\Delta p_t = \beta\mathbb{E}_t\Delta p_{t+1} + \lambda_R\left[\alpha(s_t-p_t) + (1-\alpha)\gamma(m_t-p_t) - \bar{z}_t\right], \tag{11}$$

where the expressions in square brackets are the average marginal costs of retail and wholesale firms with w_t substituted in from (9). Note that all idiosyncratic shocks wash out from the aggregate price dynamic equations. The slopes of the Phillips curves equal $\lambda \equiv (1-\beta\theta)(1-\theta)/\theta$ and analogously for λ_R. Finally, $\bar{\phi} \equiv \int\phi_j \mathrm{d}j$ is the sensitivity of the average wholesale marginal cost to exchange rate shock, e_t.

Dynamic equations (10)–(11) together with the specifications for the exogenous shock processes, fully describe equilibrium dynamics in the case of Calvo pricing. The solution to this dynamic system can be obtained numerically using a conventional Blanchard and Kahn (1980) method.[24] Finally, in the Calvo case, reset-price inflation can be measured simply as

[24] When final-good prices are flexible ($\theta_R = 0$), this dynamic system has a simple closed-form solution,

$\Delta s_t^* = \big(\Delta s_t - \theta \Delta s_{t-1}\big)/(1-\theta)$ and $\Delta p_t^* = \big(\Delta p_t - \theta_R \Delta p_{t-1}\big)/(1-\theta_R)$, since the adjusting firms are selected randomly. Given the equilibrium dynamics of the aggregate variables, we can simulate firm-level prices by using the expressions for optimal reset prices provided above. Aggregating firm-level prices, we arrive at the sample measures of regular and reset-price inflation, the counterparts to the empirical measures studied in Section 4.

Nominal and real rigidities in the model: We now discuss the sources of nominal and real rigidity in the model. First, nominal stickiness enters through the Calvo parameters θ and θ_R, which reduce the slopes of the Phillips curves (λ and λ_R) and increase the persistence of inflation. Real rigidities in the form of variable markups as measured by Γ further reduce the slope of the wholesale inflation Phillips curve and contribute to the sluggish adjustment of wholesale prices. Furthermore, aggregate real rigidities measured inversely by γ slow down the pass-through of monetary shocks into the marginal costs of both types of firms and reduce the slopes of the Phillips curves. Finally, the share of intermediate inputs in the final good production costs α links retail marginal costs to wholesale prices and constitutes a channel through which sluggish adjustment in wholesale prices translates into persistence in retail prices.

5.1.2 Calibration

We calibrate the model to monthly data and summarize the benchmark parameters in Table 7. We set the discount rate to 4 percent annually, which implies a monthly discount factor $\beta = 0.96^{1/12}$. We calibrate the money growth process and exchange rate process to the data. Specifically, we use the monthly BEA data on the M2 supply to calibrate $\rho_m = 0.5$ and $\sigma_m = 0.25$ percent. Other papers in the literature use different numbers for the persistence of money growth. For example, Chari, Kehoe, and McGrattan (2000) use $\rho_m = 0.57^{1/3} \approx 0.83$, while BKM use $\rho_m = 0$. Therefore, for robustness we also simulate the model for $\rho_m = 0$ and 0.8. Next, we let the exchange rate follow a very persistent AR(1) process with the standard deviation of innovation equal to $\sigma_e = 2$ percent and autocorrelation parameter $\rho_e = 0.995$, consistent with the data on bilateral nominal exchange rates for developed countries.

We select the parameters for the idiosyncratic shock processes ($\tilde{\sigma}_a, \tilde{\sigma}_z, \tilde{\rho}_a$ and $\tilde{\rho}_z$) to match

which we discuss in the appendix.

Table 7: Benchmark parameters

Parameter	Symbol	Value	Source
Discount factor	β	$0.96^{1/12}$	Monthly data
Money growth process, Δm_t			BEA data on M2
volatility	σ_m	0.25%	
persistence	ρ_m	0 or 0.5	
Exchange rate process, Δe_t			OECD exchange rates
volatility	σ_e	2%	
persistence	ρ_e	0.995	
Retail idiosyncratic shocks, $\tilde{z}_{it}$			BLS CPI data
volatility	$\tilde{\sigma}_z$	8%	Size of price adjustment of 8.5%
persistence	$\tilde{\rho}_z$	0.90	
Wholesale idiosyncratic shocks, $\tilde{a}_{jt}$			BLS IPP and PPI data
volatility	$\tilde{\sigma}_a$	10%	Size of price adjustment of 7.5%
persistence	$\tilde{\rho}_a$	0.95	Persistence of new prices of 0.77
Retail aggregate shocks, $\bar{z}_t$			Volatility and persistence of CPI
volatility	$\bar{\sigma}_z$	5%	regular and reset-price inflation
persistence	$\bar{\rho}_z$	0.50	from BKM
Wholesale aggregate shocks, $\bar{a}_t$			Volatility and persistence of IPP
volatility	$\bar{\sigma}_a$	4%	regular and reset-price inflation
persistence	$\bar{\rho}_a$	0.75	
Calvo parameters			
Retail	θ_R	0.75	Duration of 4 months, CPI data
Wholesale	θ	0.90	Duration of 10 months, IPP data
Share of intermediate inputs	α	0.5	Nakamura and Steinsson (2010)
Wholesale markup elasticity	Γ	1.5	Evidence on pass-through
Aggregate real rigidities	γ	0.75	
Sensitivity to the ER shock	$\bar{\phi}$	0.225	Gopinath and Itskhoki (2010)

the micro-data on price adjustment. Specifically, we set the persistence of idiosyncratic shocks to match the high autocorrelation of new prices in the BLS IPP data (0.77 for import prices), and we set the standard deviation of idiosyncratic shocks to match the absolute size of price adjustment (7.5 percent for import prices and 8.5 percent for consumer prices).[25] This results in the standard deviation of idiosyncratic shocks being equal to 10 percent and 8 percent for wholesale and retail prices, respectively, while the persistence is set to 0.95 and 0.9, respectively.

Next, we set the parameters for the aggregate shock processes ($\bar{\sigma}_a, \bar{\sigma}_z, \bar{\rho}_a$ and $\bar{\rho}_z$) to match the standard deviation and autocorrelation of regular and reset-price inflation series, reported previously in Table 1. This requires fairly large and transitory aggregate shocks at both the wholesale and retail levels (standard deviation of 4 percent and 5 percent and persistence of 0.75 and 0.5, respectively). These processes are a stand-in for all unmodeled shocks that hit the economy, including various economy-wide and industry-level marginal cost and markup shocks, as well as measurement error in prices.

We set the Calvo probabilities of non-adjusting prices (θ and θ_R) to match the micro-data on nominal price durations. Specifically, we choose parameters to produce 10-month durations in the wholesale sector (consistent with the evidence in Nakamura and Steinsson, 2008; Gopinath and Rigobon, 2008) and 4-month durations in the retail sector (consistent with Bils and Klenow, 2004). Next, we calibrate γ, the slope of the wage equation (9), and the aggregate real rigidity parameter of the model. The literature uses a wide variety of values for γ, ranging between 0.1 in models with segmented labor markets and round-about production and 4 in models with no real rigidities and strong concavity in the utility function. We set the benchmark value for γ to be 0.75 and for robustness we also use a greater value of 1.5, so that these parameters lie on both sides of the strategic neutrality case of $\gamma = 1$ and depart from it only moderately. We view this as a conservative choice for an aggregate parameter for which we have little direct information.

We choose the benchmark value for markup elasticity to be $\Gamma = 1.5$. This number implies a 40 percent pass-through of idiosyncratic shocks and is consistent with the evidence in Gopinath, Itskhoki, and Rigobon (2010) and Gopinath and Itskhoki (2010) on long-run exchange rate pass-through of about 50 percent. Moreover, it is consistent with the coefficients on the competitor prices reported in Table 5 of Section 4. We additionally

[25] For details see Gopinath and Itskhoki (2010).

evaluate the robustness of our results using the values for Γ of 0 and 4, the former being the case of constant markups and the latter being the case of strong strategic complementarities at the firm level.[26]

Finally, we set the share of intermediate inputs in the final good production α to equal 50 percent, at the conservative end of the spectrum of calibrations considered in Nakamura and Steinsson (2010). The sensitivity of the aggregate marginal cost to exchange rate shocks is set to $\bar{\phi} = 0.225$, which is consistent with most domestic firms being unaffected by this shock directly, while a small fraction of importers in the industry (for example, 30 percent) being affected strongly by this shock (for example, $\phi_j = 0.75$).[27]

5.2 Simulation results

First, we study the persistence of regular and reset-price inflation generated by our model. We compute all inflation series as sample averages of the simulated firm prices, using a procedure close to the one used on the BLS data in Section 4.[28] The results are reported in Table 8 for different values of parameters ρ_m, Γ, and γ. The first two columns report the results for the final-good (retail) inflation series, the next two columns provide the results for the unconditional wholesale inflation series, while the last two columns provide the results for the projection of wholesale inflation series onto lags of exchange rate changes. For wholesale prices we use only the subsample of "foreign firms," that is, those that are affected directly by the exchange rate shock (with $\phi_j > 0$), to make this exercise as close as possible to our empirical evidence in Section 4, which uses data on import prices.

The first pattern that emerges from Table 8 is that aggregate consumer-price inflation may not be very persistent even when wholesale inflation is significantly more persistent. Secondly, wholesale inflation is significantly less persistent than the wholesale inflation projected on the exchange rate. Next we examine reset-price inflation. Both for retail and wholesale prices, reset-price inflation is negatively autocorrelated, while when projected on the exchange rate the autocorrelation becomes positive. These patterns are consistent with

[26]Our benchmark number of $\Gamma = 1.5$ is considerably smaller than the markup elasticity of 2.5 implied by the Klenow and Willis (2006) calibration.

[27]For more details, refer to the calibration in Gopinath and Itskhoki (2010).

[28]We use a sample of 5,000 retail firms and 6,000 wholesale firms, 1,800 of which are affected directly by the exchange rate shock. With this sample size, low frequency of price adjustment, and large idiosyncratic price movements, the sampling error in average inflation series is non-trivial.

Table 8: Persistence of regular and reset-price inflation

	Unconditional				Conditional on ER	
	Δp_t	Δp_t^*	Δs_t	Δs_t^*	$\widehat{\Delta s_t}$	$\widehat{\Delta s_t^*}$
$\Gamma = 0$	0.41	−0.25	0.64	−0.23	0.85	−0.03
$\Gamma = 1.5$	0.32	−0.25	0.82	−0.14	0.92	0.10
$\Gamma = 4$	0.37	−0.28	0.87	−0.01	0.91	0.18
$\rho_m = 0$	0.39	−0.23	0.76	−0.19	0.91	0.17
$\rho_m = 0.8$	0.48	−0.14	0.79	−0.07	0.86	0.16
$\gamma = 1.5$	0.37	−0.27	0.76	−0.09	0.87	0.07
Menu cost	—	—	0.29	−0.89	0.38	−0.66

Note: the entries are AR(1) coefficients for each series (Δp refers to final-good inflation and Δs refers to wholesale inflation for the subsample of foreign firms with $\phi_j > 0$; * indicate reset-price inflation); in the last two columns the series are the projections on the current and 24 lags of the exchange rate changes. All inflation series are sample averages of the simulated firm prices, approximating the procedure used in the data. The default parameters are the benchmark parameters from Table 7 (that is, $\Gamma = 1.5$, $\gamma = 0.75$, $\rho_m = 0.5$).

the empirical findings in Section 4.1. In our calibration, negative autocorrelation of the reset-price inflation series arises due to sampling error combined with transitory semi-aggregate shocks affecting wholesale and retail pricing.

The results in Table 8 are largely similar across the different parameter values considered. Higher values of markup elasticity Γ result in more persistent (wholesale) inflation as shown in the series, particularly when projected on the exchange rate shock, while variation in γ and ρ_m has relatively little effect on the persistence of inflation. There are no clear patterns in variation of retail price persistence for different parameter values. This is because in our calibration, monetary shocks are not the key drivers of inflation in the short run (at the monthly frequency), which appears to be a reasonable description of reality.[29]

We next evaluate the extent of monetary non-neutrality produced by the model. Table 9 reports the half-lives of output in response to a monetary shock for different values of the parameters. Specifically, we calculate the AR(1) coefficient of the real output series when all shocks other than monetary shocks are shut down; based on this AR(1) coefficient, we back out a measure of half-life that we report. From Table 9 it is evident that the model

[29] This feature is consistent with aggregate dataset where inflation is well approximated by an ARMA(1,1) process with both large AR and MA roots (see, for example, Stock and Watson, 2007). While monetary shocks are likely to be responsible for the AR component (long memory, low-frequency movements), there needs to be a source of relatively large transitory shocks to explain the MA component.

Table 9: Half-life of output in response to a monetary shock (in months)

	$\rho_m = 0$	$\rho_m = 0.5$	$\rho_m = 0.8$
	Panel A: $\gamma = 0.75$		
$\Gamma = 0$	5.3	17.2	56.1
$\Gamma = 1.5$	7.0	23.7	83.0
$\Gamma = 4$	8.9	31.3	114.8
	Panel B: $\gamma = 1.5$		
$\Gamma = 0$	3.6	11.8	40.0
$\Gamma = 1.5$	4.4	15.5	58.1
$\Gamma = 4$	5.4	19.8	80.0

Note: half-life is defined as $\log(0.5)/\log(\varrho_{y|m})$*, where* $\varrho_{y|m}$ *equals the AR(1) coefficient of output* $y_t \equiv m_t - p_t$*, conditional on monetary shocks* ϵ_t^m *(i.e., when other sources of shocks are shut down).*

can produce a wide range for the extent of non-neutrality, with half-lives of output ranging from about one quarter to over 20 quarters. However, this variation is largely driven by ρ_m, the exogenous persistence introduced through the autocorrelation of the money growth rate. On the other hand, variation in the amount of real rigidity, Γ and γ, has a relatively modest effect on the extent of non-neutrality: an increase in Γ from 0 to 4 nearly doubles the half-life, while a decrease in γ from 1.5 to 0.75 increases the half-life by around 50 percent.

When we fix parameters at their benchmark values, the model produces a fairly large half-life of slightly below eight quarters, while shutting down the variable markup channel drops the half-life to less than six quarters. Without exogenous persistence (that is, $\rho_m = 0$), however, the model produces very little monetary non-neutrality (a half-life of around one quarter).[30] We conclude that the empirically calibrated variable markup channel of real rigidities goes a fair way in amplifying the real effects on output, however, without exogenous persistence its absolute effect is modest.

Note that the variation in the persistence of output deviation in the model is not linked very tightly to the persistence of inflation, which does not vary much with the amount of real rigidities or the extent of exogenous persistence. Therefore, fairly long periods of monetary non-neutrality can be consistent with transitory inflation dynamics and negatively autocorrelated reset-price inflation. Again, this is because monetary shocks are not the main drivers of inflation at very high frequencies.

[30]This finding is consistent with the results in Carvalho and Nechio (2008) on the persistence of real exchange rates.

Table 10: Micro-level exchange rate pass-through

	$\Gamma = 0$	$\Gamma = 1.5$	$\Gamma = 4$	Menu cost
First adjustment	61%	44%	35%	55%
Second adjustment	−5%	4%	7%	−8%

Note: The entries are coefficients from the regression at the firm level (for wholesale foreign firms with $\phi_j > 0$) of price change, conditional on adjustment on cumulative exchange rate change during the most recent and the previous price duration, respectively (i.e., a counterpart to regression (5) *of Section 4). The menu cost uses* $\Gamma = 1.5$.

Our final results evaluate the success of the model at capturing the slow response of firm-level prices to exchange rate shocks at the micro level, conditional on price adjustment. This is done for the wholesale foreign price with $\phi_j > 0$ to approximate the empirical analysis of Section 4.2. The results are reported in Table 10. When there are no strategic complementarities, the second adjustment is negative (due to some mean-reversion that we assumed in the exchange rate process). When strategic complementarities are present, pass-through at the second round of price adjustment is positive; however, it is much smaller than it is in the data (see Table 2), where the second rounds of price adjustment are almost as large as the first. This failure of the model persists even when we assume strong strategic complementarities ($\Gamma = 4$) or shorter nominal price durations (not reported). Although the model with strategic complementarities captures incomplete pass-through in the long run, it predicts very fast dynamics of pass-through relative to the data. This leads us to conclude that our model misses some important sources of persistence that may further contribute to the extent of monetary non-neutrality produced by the model. Matching the very slow adjustment of prices to aggregate shocks at the micro level that we document in Section 4 is an important challenge that we leave for future work.

5.3 A menu cost model

In this section, we briefly describe the setup and provide the simulation results of a menu cost model of price setting. The details are relegated to the Appendix and further discussion of the estimation procedure can be found in Gopinath and Itskhoki (2010). This exercise is important in order to evaluate the severity of the selection effects present in menu cost models and absent in the time-dependent pricing models.

We adopt a two-sector model (wholesale and retail, as above) with three types of shocks: idiosyncratic marginal cost shocks, semi-aggregate marginal cost shocks (a stand-in for exchange rate shocks) in the wholesale sector, and aggregate monetary (nominal spending) shocks. In order to maintain computational feasibility, we assume that retail prices are completely flexible. We introduce the variable markup channel, using the Klenow and Willis (2006) specification of the Kimball (1995) preferences. The rest of the setup is similar to the one discussed above in Section 5.1.

With flexible prices, no strategic complementarities (for example, CES demand) and no aggregate shocks in the retail sector, the final-good price level is given by $p_t = \bar{\mu}^R + \alpha s_t + (1-\alpha) w_t$. The wage rate is still assumed to satisfy (9). These two equations allow one to solve for the final-good price level p_t and the nominal wage rate w_t as functions of aggregate nominal spending m_t and the wholesale price level s_t.

A wholesale firm j faces a marginal cost $mc_{jt} = w_t + \phi_j e_t - a_{jt}$, where a_{jt} is an idiosyncratic shock and e_t is the semi-aggregate shock that affects the firm with elasticity ϕ_j distributed on $[0, 1]$. The firm also faces a demand schedule with elasticity σ and the elasticity of elasticity ε, evaluated when firm's relative price equals 1 (for details see Gopinath and Itskhoki, 2010, and the appendix). This implies a desired markup of $\sigma/(\sigma-1)$ with markup elasticity equal to $\Gamma = \varepsilon/(\sigma-1)$. A firm maximizes its discounted present value by optimally choosing the instances of price adjustment at a menu cost κ and optimally resetting prices at these instances. This problem can be formalized with a standard Bellman equation (see the appendix) which we solve numerically. We then use the derived policy functions to simulate a panel of prices on which we conduct similar empirical tests to those in Section 4.

For calibration, we use the same benchmark parameters as in Table 7. Additionally, we set $\sigma = 5$ and $\varepsilon = 6$ so that the level of wholesale markup is 25 percent and the elasticity of markup is equal, as before, to $\Gamma = 1.5$. The menu cost is chosen to match the duration of 10 months of wholesale prices, which implies a $\kappa = 3.5$ percent of the steady state revenue, conditional on adjustment (corresponding to total annual menu cost paid equal to 0.35 percent of annual revenues), a number consistent with the literature. Finally, in order to match the absolute size of price adjustment of 7.5 percent, we set the standard deviation of the idiosyncratic shocks to $\sigma_a = 6$ percent.

To keep this section brief, we report the results only for the benchmark values of parameters, $\Gamma = 1.5$, $\gamma = 0.75$, and $\rho_m = 0.5$. The unconditional autocorrelation of the wholesale-price inflation is 0.29, and it is 0.38, conditional on the exchange rate shock, both numbers being substantially smaller than in the case of the Calvo model (see Table 8). The corresponding reset-price inflation series is strongly negatively autocorrelated—with autocorrelation of -0.89 unconditionally and -0.66 conditional on the exchange rate shock—emphasizing the powerful selection effects of the menu cost models. This negative autocorrelation of reset-price inflation, conditional on the exchange rate shock goes against our empirical findings in Section 4.

Similarly, in the type of micro-level pass-through regressions run in Section 4.2, the menu cost model generates a pass-through coefficient conditional on the first round of price adjustment equal to 55 percent, while the coefficient for the second round of price adjustment is -8 percent. This is also due to the strong selection effect of the menu cost models that dominated the persistence introduced through strategic complementarities in pricing. Recall that in the Calvo model these two pass-through coefficients were 44 percent and 4 percent, respectively (see Table 10), and this is in contrast to our empirical findings that the pass-through at the second adjustment is nearly as large as the first one. Finally, the half-life of output in response to monetary shocks is 15.1 months in the menu cost model as opposed to 23.7 months in the Calvo model.[31] This illustrates a well-known fact: that the selection effect of the menu cost models reduces substantially the contract multiplier relative to the time-dependent models of price setting.

Summary: A number of insights come out of our simulation exercise. First, transitory aggregate inflation series are consistent with a persistent response of prices to certain aggregate shocks, including exchange rate and monetary shocks. Second, properties of the aggregate inflation series may be largely disconnected from the size of the contract multiplier for monetary shocks. Third, quantitatively exogenous persistence appears to be more important than strategic complementarities in generating long half-lives and large contract multipliers. Although strategic complementarities work to magnify the size of the contract multiplier, their absolute effects are modest unless coupled with exogenous sources of persistence. Fourth, the analyzed models (in particular, the menu cost model, but also the

[31] The approach to computing the conditional half-life of output in the menu cost model is different: Since we cannot simply shut down other sources of shocks in the menu cost model (as its dynamics are non-linear in shocks), we estimate econometrically the impulse response of output y_t to current and lagged monetary shocks ε^m_{t-j} and then use it to compute the projection of output series onto these shocks. The reported number is based on the AR(1) coefficient for the projected series.

Calvo model) cannot match the very sluggish dynamics of prices in response to shocks at the micro level conditional on adjusting prices. This suggests a need for additional sources of persistence lacking in the model.

6 Bargaining Model of Price Setting

We now consider a setup in which final good producers bargain with a number of intermediate good suppliers. This is a more realistic characterization of business-to-business transactions. We show that in a static bargaining model there are strategic complementarities in wholesale price setting, so that markups are variable and wholesale prices exhibit incomplete pass-through. On the other hand, final good producers compete monopolistically, and subject to a CES consumer demand, charge constant markups. This provides a micro-foundation for the reduced-form assumptions on wholesale and retail markups imposed in the dynamic model of Section 5. We leave the extension of this bargaining model to a dynamic price-setting setup for further research.

Although it is quite natural to think that intermediate good prices are set via bargaining, we are unaware of any macroeconomic models of intermediate-good price-setting via bargaining. In a recent paper, Goldberg and Tille (2009) propose a bargaining model of currency choice in international transactions. Our model here is most closely related to Atkeson and Burstein (2008), which is a special case of our model when the intermediate-good suppliers have full bargaining power and hence act as price setters.[32]

Consider a final good producer i. In what follows, we omit the final good producer's identifier i where the omission leads to no confusion. The final good producer uses intermediate inputs to assemble the final good according to a CES technology

$$y = \left[\sum_{j=1}^{N} q_j^{\eta}\right]^{1/\eta}, \tag{12}$$

where $\eta \in (0,1)$ controls the elasticity of substitution between intermediate varieties, N is the number of intermediate varieties used for assembly, and q_j is the input quantity of intermediate variety j.[33] Note that labor is not used for production of the final good, which

[32] Atkeson and Burstein (2008) choose not to interpret their two-tier demand structure as a sequence of the wholesale and retail sectors, but this is an equally coherent interpretation.

[33] Note that in this section, as opposed to the rest of the paper, lower case letters denote the levels of

is produced using intermediates only. This corresponds to the special case of $\alpha = 1$ in the terminology of Sections 2 and 5, and we adopt it for simplicity.

The revenue of the final good producer is given by

$$R = py = Ay^{\zeta},$$

where p is the final good price, A is a demand shifter and $\zeta \in (0, 1)$ is a parameter that controls the elasticity of demand. This revenue function specification arises from CES preferences over the final good with the elasticity of substitution equal to $1/(1-\zeta)$. Finally, the intermediate good producer j has a constant marginal cost of c_j. Therefore, the total surplus to be shared between the final and intermediate good producers is

$$Ay^{\zeta} - \sum_{j=1}^{N} c_j q_j, \tag{13}$$

where y is given in (12).

The surplus in (13) is divided according to bilateral Nash bargaining between the final good producer and each of the intermediate good suppliers. Specifically, we assume that prices are determined through bargaining, while, given prices, the final good producer is free to choose any quantity of the intermediate good supply.[34]

Formally, denote by s_j the price of intermediate good j determined via bargaining. The final good supplier maximizes his revenues minus costs, given by $\sum_{j=1}^{N} s_j q_j$, when choosing quantities. As a result, quantities lie on the demand curve given by

$$\zeta A y^{\zeta-\eta} q_j^{\eta-1} = s_j. \tag{14}$$

This implies (upon aggregation over j) that the final good's price is a constant markup over

variables rather than logs.

[34]If the parties could also bargain over quantities, the bargaining game would result in an efficient supply of the intermediate goods, with prices playing a role of transfers without any allocative role. However, if, for example, the value of the demand shifter A were unknown at the bargaining stage, it could be optimal to set prices without restricting quantities and let the quantities adjust *ex post* in response to the movements in A. Empirically, intermediate-good prices appear to play an allocative role, since changes in intermediate-good prices get fully passed-through into retail consumer prices (see Section 3).

the cost index of the intermediate goods:

$$p = Ay^{\zeta-1} = \frac{s}{\zeta}, \qquad \text{where} \qquad s \equiv \left[\sum_{j=1}^{N} s_j^{-\frac{\eta}{1-\eta}}\right]^{-\frac{1-\eta}{\eta}}.$$

Therefore, we arrive at the first assumption of Section 5: that retail prices are set as a constant markup over marginal cost (that is, $\Gamma_R = 0$). Note that the final good quantities and prices respond to the intermediate good prices s_j, which play an allocative role in this bargaining model.

The Nash bargaining between the final good producer and supplier j determines the price s_j. We assume that the bargaining power of the final good producer is $\phi \in (0, 1)$ and we denote his relative bargaining power by $\lambda \equiv \phi/(1-\phi)$. We do not provide here a micro-foundation for the source of variation in bargaining power. It may come from the differential patience of the final good producer and suppliers, which in turn may be related to the extent of liquidity constraints that different firms face or from the tightness of the supplier market (that is, how easy it is to replace a given supplier).

If bargaining breaks down, the supplier receives zero, while the final good producer has to assemble the final good without the input of this supplier. Therefore, the surplus of the supplier is $(s_j - c_j)q_j$. From (14) it follows that the profit of the final good producer equals $(1-\zeta)Ay^{\zeta}$, where $y = \big(\zeta A/s\big)^{1/(1-\zeta)}$ is the optimal output of the final good, given the intermediate price index s. If bargaining with supplier j breaks down, the cost index of the intermediates becomes

$$s_{-j} = \left[\sum_{k\neq j} s_k^{-\frac{\eta}{1-\eta}}\right]^{-\frac{1-\eta}{\eta}} > s.$$

Consequently, the optimal output of the final good will be $y_{-j} = \big(\zeta A/s_{-j}\big)^{1/(1-\zeta)} < y$, resulting in the profit of the final good producer of $(1-\zeta)Ay_{-j}^{\zeta}$. As a result, the incremental surplus of the final good producer from supplier j is given by

$$(1-\zeta)A\big[y^{\zeta} - y_{-j}^{\zeta}\big] = (1-\zeta)\zeta^{\frac{\zeta}{1-\zeta}}A^{\frac{1}{1-\zeta}}\left[s^{\zeta} - s_{-j}^{\zeta}\right].$$

Under these circumstances, we can write the Nash bargaining problem formally as:

$$\max_{s_j}\left\{\Big[(s_j - c_j)q_j\Big]^{1-\phi}\left[(1-\zeta)\zeta^{\frac{\zeta}{1-\zeta}}A^{\frac{1}{1-\zeta}}\big(s^{\zeta} - s_{-j}^{\zeta}\big)\right]^{\phi}\right\},$$

where q_j is subject to (14) and s and s_{-j} are the cost indexes defined above. Taking the first-order condition, one can demonstrate that the bargained price needs to satisfy the following condition:

$$\frac{s_j}{s_j - c_j} = \frac{1}{1-\eta} + \frac{\zeta}{1-\zeta}\theta_j \left[\frac{\lambda}{1-(1-\theta_j)^{\chi}} - \frac{\eta-\zeta}{1-\eta} \right], \tag{15}$$

where $\chi \equiv \frac{\zeta}{1-\zeta}\frac{1-\eta}{\eta}$ and θ_j is the cost (market) share of supplier j equal to

$$\theta_j \equiv \frac{s_j q_j}{sy} = \left(\frac{s_j}{s}\right)^{-\frac{\eta}{1-\eta}}.$$

Note that (15) implies a markup pricing rule in which the markup depends on the market share θ_j, the relative bargaining power λ, and the parameters of the model ζ and η. Moreover, since prices affect the market shares of firms, in general markups are not constant and there is strategic non-neutrality in pricing (that is, $\Gamma \neq 0$ in the terminology of Section 2).

To keep this section brief, we state here only the main results on the properties of the bargained prices, which follow directly from (15):

1. When $\theta_j = 0$ (infinitesimal supplier), the markup is constant and given by $\mu_j \equiv s_j/c_j = \phi + (1-\phi)/\eta$, which is a convex combination of 1 and $1/\eta$ weighted by the bargaining power.

2. When all bargaining power is with the final good producer (that is, $\phi = 1$ or $\lambda = \infty$), bargaining results in marginal-cost pricing $s_j = c_j$.

3. In the opposite case, when full bargaining power is with intermediate goods producers ($\phi = \lambda = 0$), markup depends on the market share:

$$\frac{\mu_j}{\mu_j - 1} \equiv \frac{s_j}{s_j - c_j} = \frac{1}{1-\eta} - \theta_j \frac{1}{1-\eta}\frac{\eta-\zeta}{1-\zeta}.$$

   Note that when $\eta > \zeta$ (the baseline case), the markup is increasing in the market share and hence decreasing in the relative price of the firm. This corresponds to the case $\Gamma > 0$.[35]

[35] In fact, this case corresponds exactly to the Atkeson and Burstein (2008) model, since $1/(1-\eta)$ equals the elasticity of substitution between intermediate varieties, and $1/(1-\zeta)$ equals the elasticity of substitution between final good varieties.

4. In general, when $\theta_j \in (0,1)$ and $\lambda < \infty$, the markup is variable and pass-through is incomplete. Moreover, pass-through is not necessarily monotonic in the market share.

To summarize, this bargaining model results in variable-markup pricing at the wholesale level and constant-markup pricing at the retail level, with wholesale markups depending on the market shares of the suppliers and on their relative bargaining power, as well as on the parameters of the model. Therefore, the model appears to be consistent with the broad features of the data discussed in Sections 3–4. We leave to future research testing the qualitative success of this model in capturing the dynamics of wholesale and retail prices.

APPENDIX

A Calvo Price Setting

In this section, we derive a general log-linear approximation for the price setting equation and aggregate inflation dynamics (Phillips curve) in a Calvo model. Since these derivations are well known, we keep the exposition brief.

Consider a firm j with a real profit function $\Pi^j(x_j|S)$, where x_j is the firm's log-price and S is the state of the economy.[36] The desired price of the firm is $\tilde{x}_j(S) \equiv \arg\max_{x_j} \Pi^j(x_j|S)$ with the necessary condition $\Pi^j_x(\tilde{x}_j(S)|S) = 0$, where the subscript denotes a partial derivative. We assume that the marginal cost of the firm does not depend on the price of the firm, that is, a firm faces a constant returns to scale in production, where productivity depends on the state of the world. Then, we can decompose the desired price as

$$\tilde{x}_j(S) = \mu_j(\tilde{x}_j - x, S) + mc_j(S),$$

where $mc_j(S)$ is the log nominal marginal cost of the firm and $\mu_j(x_j - x, S)$ is the log (desired) markup, which we allow to depend on the relative price of the firm, with x denoting the log of the relevant price index.

A general first-order approximation to the markup can be written as

$$\mu_j(x_j, S) \approx \bar{\mu} - \Gamma(x_j - x) + \epsilon_j(S),$$

where $\bar{\mu}$ and Γ are some constants (assumed to be common across all firms at the point of approximation) and $\epsilon_j(S)$ is some linear function of the state S. It is natural to assume that $\epsilon_j(S)$ is stationary, while $mc_j(S)$ is co-integrated with the nominal variables of the model. With this approximation, we can solve explicitly for the desired price of the firm:

$$\tilde{x}_j(S) = \frac{1}{1+\Gamma}\big[\bar{\mu} + mc_j(S) + \epsilon_j(S)\big] + \frac{\Gamma}{1+\Gamma}x.$$

[36]Since we look at the real profit function, that is, a profit function normalized by the price level in the economy, it is without loss of generality to assume that the nominal variables enter the sufficient state vector S only normalized by the price level. Therefore, we can treat S as having a stationary distribution even though monetary variables may be trending. For an example see Appendix C.

A given firm sets prices to maximize its discounted expected value. In a Calvo pricing environment the firm may adjust its price in every period with exogenous probability $(1-\theta)$. Therefore, we can write the problem of the firm recursively as:

$$\bar{x}_{jt}(\mathbb{S}) = \arg\max_{x_j} \mathbb{E}\left\{ \Pi^j(x_j|S_t) + \sum_{\ell=1}^{\infty} Q_{t,t+\ell}(\mathbb{S})\theta^{\ell-1}\big[\theta\Pi^j(x_j|S_{t+\ell}) + (1-\theta)\Pi^j(\bar{x}_{j,t+\ell}(\mathbb{S})|S_{t+\ell})\Big| S^t \right\},$$

where $Q_{t,t+\ell}(\mathbb{S})$ is the stochastic discount factor for real variables, $S^t \equiv (S_0, \ldots, S_t)$ is the history of the states, and $\mathbb{S} \equiv (S^t, S_{t+1}, \ldots)$. The first-order condition for the price setting can be written as:

$$\sum_{\ell=0}^{\infty} \theta^\ell \mathbb{E}\left\{ Q_{t,t+\ell}\Pi_x^j(\bar{x}_{jt}|S_{t+\ell})\big| S^t \right\} = 0,$$

where we omit the explicit dependence on $\mathbb{S}$. Taking a first-order approximation of this optimality condition around a non-stochastic steady state with zero inflation, we obtain

$$\sum_{\ell=0}^{\infty} (\beta\theta)^\ell \mathbb{E}\left\{ \bar{x}_{jt} - \tilde{x}_j(S_{t+\ell})\big| S^t \right\} = 0,$$

where β^ℓ is the non-stochastic steady-state value of $Q_{t,t+\ell}$. The price setting formulas in Section 5.1 are direct implications of this linearized optimality condition. Now, using the expression for the desired price, we have:

$$\bar{x}_{jt} = (1-\beta\theta)\sum_{\ell=0}^{\infty}(\beta\theta)^\ell \mathbb{E}_t\left\{ \frac{1}{1+\Gamma}\big(mc_{j,t+\ell} + \epsilon_{j,t+\ell}\big) + \frac{\Gamma}{1+\Gamma}x_{t+\ell} \right\},$$

where we switched notation for conditional expectation, suppressed the explicit dependence on the state of the economy, and omitted the constant by implicitly relabeling the variables to denote the deviations from the non-stochastic steady state.

Finally, since the nominal marginal cost is possibly integrated, we need to scale this expression by some monetary variable co-integrated with the marginal cost. A natural candidate is the competitor price index x_t or a sector price level (in a number of models, including the Kimball demand model, the two variables coincide). We therefore define $x_t = \int_j x_{jt}\mathrm{d}j$. With some manipulation, we rewrite the deflated price-setting equation as:

$$\bar{x}_{jt} - x_{t-1} = \frac{1-\beta\theta}{1+\Gamma}\sum_{\ell=0}^{\infty}(\beta\theta)^\ell \mathbb{E}_t\big\{ mc_{j,t+\ell} - x_{t+\ell} + \epsilon_{j,t+\ell}\big\} + \sum_{\ell=0}^{\infty}(\beta\theta)^\ell \mathbb{E}_t \Delta x_{t+\ell},$$

or in a recursive form:[37]

$$\big(\bar{x}_{jt} - x_{t-1}\big) - \beta\theta\mathbb{E}_t\big(\bar{x}_{j,t+1} - x_t\big) = \frac{1-\beta\theta}{1+\Gamma}\big(mc_{jt} - x_t + \epsilon_{jt}\big) + \Delta x_t.$$

Next, with Calvo pricing, the dynamics of the aggregate price level can be written as

$$x_t = \theta x_{t-1} + (1-\theta)\mathbb{E}_j\bar{x}_{jt} \qquad \Rightarrow \qquad \Delta x_t = (1-\theta)\mathbb{E}_j\big\{\bar{x}_{jt} - x_{t-1}\big\},$$

where $\mathbb{E}_j$ is the cross-sectional average across all firms. Combining the above two equations and rearranging, we arrive at the traditional New-Keynesian Phillips curve:

$$\Delta x_t - \beta\mathbb{E}_t\Delta x_{t+1} = \frac{\lambda}{1+\Gamma}\mathbb{E}_j\big\{mc_{jt} - x_t + \epsilon_{jt}\big\}, \qquad \lambda \equiv \frac{(1-\beta\theta)(1-\theta)}{\theta}.$$

Equations (10)–(11) in the text are special cases of this Phillips curve with the expressions for marginal costs substituted in (note that the cross-sectional expectation averages out all purely idiosyncratic shocks).

B Aggregate Dynamics under Calvo Pricing

The aggregate dynamic system contains three equations — the two Phillips curves for the wholesale and retail prices and the aggregate wage equation — for three variables (s_t, p_t, w_t):

$$\begin{aligned}
\Delta s_t &= \beta\mathbb{E}_t\Delta s_{t+1} + \frac{\lambda}{1+\Gamma}\big\{w_t - s_t + \bar{\phi}e_t - \bar{a}_t\big\},\\
\Delta p_t &= \beta\mathbb{E}_t\Delta p_{t+1} + \lambda_R\big\{\alpha s_t + (1-\alpha)w_t - p_t - \bar{z}_t\big\},\\
w_t &= \gamma m_t + (1-\gamma)p_t,
\end{aligned}$$

where Δm_t, e_t, $\bar{z}_t$, and $\bar{a}_t$ follow exogenous stationary processes. This system can be solved using the conventional Blanchard and Kahn (1980) method, which results in expressions for the endogenous variables (s_t, p_t, w_t) as functions of the shocks to the exogenous variables. This solution allows one to study the statistical properties of the endogenous variable time series, including their volatility and persistence.

When the retail prices are set flexibly (that is, $\theta_R = 0$ or $\lambda_R = \infty$), there exists a tractable

[37] This is the step that requires stationarity of the right-hand-side variables.

analytical solution for the aggregate dynamics. We discuss it briefly here. In this case, the expression for the consumer price level becomes static:

$$p_t = \alpha s_t + (1-\alpha) w_t - \bar{z}_t.$$

Together with the wage equation, it allows one to solve for p_t and w_t as linear functions of m_t, s_t, and $\bar{z}_t$. Substituting these expressions into the wholesale-price Phillips curve, we obtain a second-order difference equation in $s_t - m_t$:[38]

$$\Delta s_t - \beta \mathbb{E}_t \Delta s_{t+1} = \kappa (s_t - m_t) + \xi_t,$$

where

$$\xi_t \equiv \frac{\lambda}{1+\Gamma}\big(\bar{\phi} e_t - \bar{a}_t\big) - \frac{\lambda}{1+\Gamma}\frac{1-\gamma}{\alpha(1-\gamma)+\gamma}\bar{z}_t$$

is the summary measure of all shocks other than m_t, and $\kappa \equiv \frac{\lambda}{1+\Gamma}\frac{\gamma}{\alpha(1-\gamma)+\gamma}$ is the summary measure of nominal and real rigidities in the model. This difference equation can be solved forward. Assuming for simplicity that ξ_t follows an AR(1), the process for $s_t - m_t$ is an ARMA(3,1). Therefore, s_t is co-integrated with m_t and movements in m_t dominate the low-frequency movements in s_t; however, the short-run dynamics of s_t (around slow moving m_t) may be dominated by transitory shocks to ξ_t. In particular, the MA component may reduce significantly the short-run persistence, while it does not affect the long-run persistence. This logic is consistent with the empirical findings of Stock and Watson (2007).[39] Finally, one can show that one of the AR-roots is given by ρ_m, while the other root is decreasing in κ and converging to 1 as $\kappa \to 0$. Furthermore, one can show that these roots also drive the persistence of the output response to monetary shocks.

C Menu Cost Model

Kimball demand: To simulate the menu cost model, we first need to specify the explicit source of variable markups. We generate variable markups by introducing the Klenow and Willis (2006) specification of Kimball (1995) demand. The demand function for firm j in

[38] This is under the assumption of a unit root in m_t; otherwise, the difference equation is second order in s_t.

[39] If ξ_t is absent from the model, the process for $s_t - m_t$ is an AR(2), and therefore s_t exhibits both high short-run and long-run persistence.

this case is given by:

$$\psi(s_{jt}-s_t)=\big[1-\varepsilon(s_{jt}-s_t)\big]^{\sigma/\varepsilon},\qquad \sigma\geq 1,\quad \varepsilon>0,$$

where s_{jt} is the log price of the firm and $s_t=\int_j s_{jt}\mathrm{d}j$ is the sectoral log price index. In Gopinath and Itskhoki (2010), we show that this price index is a valid second-order approximation to the ideal price index with this demand system. The price elasticity of this demand is given by

$$\tilde{\sigma}=\frac{\sigma}{1-\varepsilon(s_{jt}-s_t)},$$

which equals σ when the relative price of the firm is 1. With this demand, the desired price is equal to a markup $\tilde{\sigma}/(\tilde{\sigma}-1)$ over the marginal cost. The elasticity of the markup is given by

$$\tilde{\Gamma}=\frac{\tilde{\varepsilon}}{\tilde{\sigma}-1},$$

where

$$\tilde{\varepsilon}\equiv\frac{\tilde{\varepsilon}}{1-\varepsilon(s_{jt}-s_t)}$$

is the elasticity of the elasticity of demand. Therefore, the markup elasticity evaluated at the relative price of 1 is given by $\Gamma=\varepsilon/(\sigma-1)>0$.

Problem of the firm: The real profit of the firm is given by

$$\Pi(s_{jt}|\mathbb{S}_{jt})=\psi(s_{jt}-s_t)\big[\exp\{s_{jt}-p_t\}-\exp\{mc_{jt}-p_t\}\big],$$

where $\mathbb{S}_{jt}$ is the state vector for the firm, mc_{jt} is the log nominal marginal cost of the firm, and p_t is the final-good price level. As discussed in the text, the marginal cost of the firm equals

$$mc_{jt}=w_t+\phi_j e_t-a_{jt}.$$

Therefore, the state vector for the firm includes (p_t,s_t,w_t,e_t,a_{jt}).

We can write the firm's problem recursively as:

$$\begin{cases} V^N(\mathbb{S}_{jt}) &=& \Pi(s_{j,t-1}|\mathbb{S}_{jt})+\mathbb{E}\left\{Q(\mathbb{S}_{j,t+1})V(\mathbb{S}_{j,t+1})\big|\mathbb{S}_{jt}\right\},\\ V^A(\mathbb{S}_{jt}) &=& \max_{\bar{s}_{jt}}\Big\{\Pi(\bar{s}_{jt}|\mathbb{S}_{jt})+\mathbb{E}\left\{Q(\mathbb{S}_{j,t+1})V(\mathbb{S}_{j,t+1})\big|\mathbb{S}_{jt}\right\}\Big\},\\ V(\mathbb{S}_{jt}) &=& \max\big\{V^N(\mathbb{S}_{jt}),V^A(\mathbb{S}_{jt})-\kappa\big\}, \end{cases}$$

where V is the value of the firm, V^N is the value of the firm if it does not adjust its price, and V^A is the value of the firm if it adjusts its price; κ is the menu cost, Q is the stochastic discount factor for real variables (which we set to equal β in the simulation), and $\mathbb{S}_{jt}$ includes, in addition, the previous price of the firm $s_{j,t-1}$.

General equilibrium: We assume flexible prices for the final good and no aggregate productivity shocks in the final-good sector. This implies that $p_t = \bar{\mu}_R + \alpha s_t + (1-\alpha)w_t$, where $\bar{\mu}_R$ is the constant markup in the final-good sector. In turn, the wage is given by $w_t = \gamma m_t + (1-\gamma)p_t$. This allows us to solve for p_t and w_t as functions of m_t and s_t, respectively, and to reduce the aggregate state space to (s_t, m_t, e_t). The state vector for an individual wholesale firm additionally includes $(s_{j,t-1}, a_{jt})$.

In the general equilibrium of the model, firms optimally decide to adjust prices given their current state vector and rational expectations about the evolution of the state vector, while aggregated individual firm pricing decisions are consistent with the aggregate dynamics of the wholesale price level s_t.

Simulation procedure: We iterate the Bellman equation for the firm pricing problem on the grid, given a forecasting rule for the evolution of the state vector.[40] This produces a policy function for firm pricing decisions, which allows us to simulate a panel of firm prices. In each period of the simulation, we make sure that the wholesale price index is consistent with the firm pricing decisions (which constitutes a static fixed point problem). As a result, we obtain a time series for the equilibrium wholesale price level. Given this time series, we update the forecasting equation. We iterate this procedure until the forecasting equation converges. With the equilibrium forecasting rule, we simulate a panel of firm prices and we use it to estimate various statistical moments as in Section 4. Additional details of the simulation procedure can be found in Gopinath and Itskhoki (2010).

[40] To ensure stationarity of the grids, we normalize all nominal variables in the model by m_{t-1}.

References

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

Basu, S. (1995): "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare," *American Economic Review*, 85(3), 512–31.

Berger, D., J. Faust, J. H. Rogers, and K. Steverson (2009): "Border Prices and Retail Prices," Board of Governors of the Federal Reserve, International Finance Discussion Paper No. 972.

Berman, N., P. Martin, and T. Mayer (2009): "How Do Different Exporters React to Exchange Rate Changes? Theory, Empirics and Aggregate implications," CEPR Discussion Paper 7493.

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

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

Blanchard, O. J., and J. Galí (2007): "Real Wage Rigidities and the New Keynesian Model," *Journal of Money, Credit and Banking*, 39, 35–65.

Blanchard, O. J., and C. M. Kahn (1980): "The Solution of Linear Difference Models under Rational Expectations," *Econometrica*, 48(5), 1305–1311.

Boivin, J., M. Giannoni, and I. Mihov (2009): "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data," *American Economic Review*, 99(1), 350–84.

Burstein, A., and C. Hellwig (2007): "Prices and Market Shares in a Menu Cost Model," UCLA Working Paper.

Burstein, A., and N. Jaimovich (2008): "Understanding Movements in Aggregate and Product-level Real Exchange rates," UCLA Working Paper.

Calvo, G. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12(3), 383–398.

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

Carvalho, C., and F. Nechio (2008): "Aggregation and the PPP Puzzle in a Sticky-Price Model," Federal Reserve Bank of New York Working Paper.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2000): "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?," *Econometrica*, 68(5), 1151–1180.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999): "Monetary Policy Shocks: What Have We Learned and to What End?," in *Handbook of Macroeconomics*, ed. by J. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 2, pp. 65–148. Elsevier.

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

Eichenbaum, M., N. Jaimovich, and S. Rebelo (2009): "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review*, forthcoming.

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

Goldberg, P., and R. Hellerstein (2006): "A Framework for Identifying the Sources of Local-Currency Price Stability with an Empirical Application," Federal Reserve Bank of New York Working Paper.

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

Golosov, M., and R. E. Lucas, Jr. (2007): "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115, 171–99.

Gopinath, G., P.-O. Gourinchas, C.-T. Hsieh, and N. Li (2009): "Estimating the Border Effect: Some New Evidence," NBER Working Paper 14938.

Gopinath, G., and O. Itskhoki (2010): "Frequency of Price Adjustment and Pass-through," *Quarterly Journal of Economics*, 125(2), 675–727.

GOPINATH, G., O. ITSKHOKI, AND R. RIGOBON (2010): "Currency Choice and Exchange Rate Pass-through," *American Economic Review*, 100(1), 306–336.

GOPINATH, G., AND R. RIGOBON (2008): "Sticky Borders," *Quarterly Journal of Economics*, 123(2), 531–575.

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 Working Paper.

——— (2007): "Sticky Information and Sticky Prices," *Journal of Monetary Economics*, 54(1), 79–99.

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.

MAĆKOWIAK, B., AND M. WIEDERHOLT (2009): "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 99(3), 769–803.

MANKIW, N. G., AND R. REIS (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 117(4), 1295–328.

MEESE, R., AND K. ROGOFF (1983): "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?," *Journal of International Economics*, 14(1), 3–24.

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

——— (2010): "Monetary Non-Neutrality in a Multi-Sector Menu Cost Model," *Quarterly Journal of Economics*, forthcoming.

NAKAMURA, E., AND D. ZEROM (2010): "Accounting for Incomplete Pass-through," *Review of Economic Studies*, 77(3), 1192–1230.

NEIMAN, B. (2009): "A State-Dependent Model of Intermediate Goods Pricing," Chicago Booth Working Paper.

——— (2010): "Stickiness, Synchronization, and Passthrough in Intrafirm Trade Prices," *Journal of Monetary Economics*, 57(3), 295–308.

Romer, D. (2008): "Real Rigidities," in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf, and L. E. Blume. Palgrave Macmillan, Basingstoke.

Stock, J., and M. Watson (2007): "Why Has U.S. Inflation Become Harder to Forecast?," *Journal of Money, Credit and Banking*, 39(1).

Taylor, J. B. (1980): "Aggregate Dynamics and Staggered Constracts," *Journal of Political Economy*, 88(5), 1–24.

Woodford, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.