Accounting for persistence and volatility of good-level real exchange rates: the role of sticky information*

Mario J. Crucini[†], Mototsugu Shintani[‡] and Takayuki Tsuruga[§]

First Draft: October 2007

This version: March 2009

Abstract

Volatile and persistent real exchange rates are observed not only in aggregate series but also in micro-price data at the retail level. Kehoe and Midrigan (2007) recently showed that, under a standard assumption on nominal price stickiness, empirical frequencies of micro price adjustment cannot replicate the time-series properties of the Law of One Price deviations. We extend their sticky price model by combining good-specific price adjustment with information stickiness in the sense of Mankiw and Reis (2002). Our framework allows for multiple cities within a country. Using a panel of U.S.-Canadian city pairs, we estimate a dynamic price adjustment process for 165 individual goods. Under a reasonable assumption on the money growth process, we show that the model matches the LOP persistence of the median good and accounts for approximately one-half of its volatility when information updates occur every 12 months.

JEL Classification: E31, F31, D40

Keywords: Good-level real exchange rates, Law of One Price, Sticky information, Dynamic panel

*We thank Kenn Ariga, Craig Burnside, Eiji Fujii, Carlos Zarazaga and seminar and conference participants at Duke, Hokkaido, North Carolina State and Osaka University, the University of Kansas, the Bank of Japan, the Federal Reserve Bank of Dallas, the 2008 Midwest Macroeconomics Meetings, and the 2007 TCER Macroeconomic Conference for helpful comments and discussions. Mario J. Crucini and Mototsugu Shintani gratefully acknowledge the financial support of National Science Foundation. Takayuki Tsuruga acknowledges the financial support of Grant-in-aid Scientific Research.

[†]Department of Economics, Vanderbilt University

[‡]Department of Economics, Vanderbilt University

[§]Faculty of Economics, Kansai University

1 Introduction

Aggregate real exchange rates exhibit persistence and volatility much higher than what economists believe is consistent with a plausible degree of price rigidity. The time-dependent pricing model, under local currency pricing, offers a convenient theoretical link between price stickiness and the stochastic properties of real exchange rates. Chari, Kehoe, and McGrattan (2002, CKM) show that to generate the observed persistence of CPI-based aggregate real exchange rates, prices need to be exogenously fixed for at least one year. This degree of price-stickiness, however, appears implausible based on recent evidence by Bils and Klenow (2004) who find a median duration between price changes of only 4.3 months in U.S. micro-data.

Using a broader sample of the Economist Intelligence Unit (EIU) micro-price data than employed here, Crucini and Shintani (2008) find the half-life of deviations from the Law of One Price (LOP) for the median good in the neighborhood of 18 months, considerably lower than the consensus 3-5 year half-lives of aggregate real exchange rates. An important theoretical contribution along similar lines is Kehoe and Midrigan (2007) who allow different price stickiness across individual goods and show that the persistence in LOP deviations is equal to 'the Calvo parameter,' the probability of price non-adjustment at the good level. Their empirical analysis using real exchange rates of 66 individual goods across the U.S. and four European countries shows that the frequency of no price adjustment is higher for goods that exhibit more persistent deviations from the LOP, as suggested by the theoretical model. However, the persistence puzzle is still not resolved in the sense that the observed frequencies of price changes are too high to replicate the persistence of real exchange rates for most goods in the cross-section. In addition, the model does not match the time series variability of LOP deviations observed in the micro-data. Their results point to the need to break the tight link between the frequency of price adjustment and the persistence of LOP deviations predicted by the standard Calvo-type sticky price model.

Our analysis differs from Kehoe and Midrigan (2007) in several ways. First, we break the tight link between the Calvo parameter and LOP persistence by extending the Kehoe-Midrigan model to allow for persistent money growth and information stickiness. Information stickiness is the assumption that only a fraction of firms update their information set each period. Here, LOP persistence arises from the convolution of price adjustment timing and information updating. In the macroeconomic literature, Mankiw and Reis (2002) show that a model of information stickiness, or inattentiveness, is capable of explaining the observed slow response of aggregate inflation to monetary shocks much better than sticky prices alone. When the information stickiness augments the Calvo-type sticky price mechanism, less frequent information updating leads to higher price persistence, at a given frequency of price adjustment (Dupor, Kitamura, and Tsuruga (2008, DKT)). With plausible assumptions on international money growth processes, a similar effect takes place to increase both the persistence and volatility of real exchange rates.

Second, our theoretical model allows for multiple cities in each country and for long-run price deviations between the cross-border city pairs to differ by good and city pair. As such, our model allows us to exploit an international retail price survey at the city level which records local currency prices of individual goods and services spanning most of the CPI basket. Because this survey is conducted by a single agency, the EIU, we expect more comparability of the products among international cities than is true of national CPI surveys, bringing the data more in line with the spirit of the model. Using this survey, from 1990 to 2005, we expand the number of products from 66 used in Kehoe and Midrigan (2007) to 165. We also increase the number of locations from five countries (Austria, Belgium, France, Spain and the U.S.) to 52 U.S.-Canadian city pairs.

Third, we examine the effect of the exclusion of sales on the performance of the model. Recently, Nakamura and Steinsson (2008) claim that the evidence of fast price adjustment reported by Bils and Klenow (2004) may be strongly influenced by the presence of sales, or other temporary price reductions. Focusing on regular price changes, Nakamura and Steinsson (2008) find that the median duration between price changes increases to the range of 8 to 11 months. Since prices are stickier based on this alternative definition of price change, it elevates the Calvo model's ability to account for important features of the data.[1] This improvement is subject to the caveat that we do not explicitly model sales.

The shortcomings of the Calvo model highlighted in Kehoe and Midrigan (2007) persist in the context of the EIU data with or without the adjustment of the frequency of price changes for temporary sales. In contrast, the Calvo model extended to allow for information stickiness fully accounts for LOP persistence of the median good and at least one-half of its volatility when the average duration between information updates is about one year, when we calibrate to the observed frequency of regular price changes.[2]

[1]Our working paper includes results using the Bils and Klenow frequencies, which are omitted here due to space considerations.

[2]When we calibrate to the frequency of price changes including sales the average duration between information updates needed to match LOP persistence rises to between 14 and 20 months.

2 The model

Trade is over a continuum of goods across multiple cities located in two countries. Under monopolistic competition, firms set prices in local currency to satisfy demand for a particular good in a particular city. A representative agent in each country chooses consumption over an infinite horizon subject to a cash-in-advance (CIA) constraint. In what follows, the U.S. and Canada represent the home and foreign country. Due to the symmetry of the model, we mostly focus on the equations of the United States.

The lowest level of aggregation is the brand, z of a particular good. U.S. brands of each good are indexed $z \in [0, 1/2]$, while those of Canada are indexed $z \in (1/2, 1]$. Integrating over brands, we have the CES index for consumption of good j, given by

$$c_t(j,l) = \left[\int_0^1 c_t(j,l,z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{\theta}{\theta-1}} , \tag{1}$$

where $c_t(j,l,z)$ is consumption of a brand z of good j in U.S. city l. CES aggregation across U.S. cities $l \in [0,1]$, gives U.S. national consumption of good j

$$c_t(j) = \left[\int c_t(j,l)^{\frac{\theta-1}{\theta}} dl \right]^{\frac{\theta}{\theta-1}} , \tag{2}$$

and further CES aggregation across goods gives aggregate U.S. consumption, c_t,

$$c_t = \left[\int c_t(j)^{\frac{\theta-1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}} . \tag{3}$$

2.1 Households

As in Kehoe and Midrigan (2007), markets for state-contingent money claims are complete. The asset structure is represented with one period state-contingent bonds. Let $B(s^t, s_{t+1})$ denote the number of units of bonds in U.S. dollars. These bonds are purchased in period t, indexed by the history of events up to period t, s^t, and pay one U.S. dollar per unit purchased, if the state s_{t+1} is realized in period $t+1$. We suppress the state and denote these holdings as B_{t+1} for U.S. households and B^*_{t+1} for Canadian households with associated nominal stochastic discount factor across adjacent time periods denoted, $\Upsilon_{t,t+1}$.[3] Also, $\Upsilon_{t,t+h}$ is used by firms, regardless of their country of origin, to discount profits earned in period $t+h$ back to the period t.

[3] As Kehoe and Midrigan (2007) argue, the choice of currency denomination of bonds is irrelevant when the markets for state-contingent money claims are complete.

Households maximize the discounted sum of utility subject to an intertemporal budget constraint and a CIA constraint. The maximization problem for U.S. households is

$$\max \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (\ln c_t - \chi n_t), \tag{4}$$

$$\text{s.t.} \quad M_t + \mathbb{E}_t(\Upsilon_{t,t+1} B_{t+1}) = R_{t-1} W_{t-1} n_{t-1} + B_t + (M_{t-1} - P_{t-1} c_{t-1}) + T_t + \Pi_t, \tag{5}$$

$$M_t \geq P_t c_t, \tag{6}$$

where n_t is hours of work, and $\mathbb{E}_t(\cdot)$ denotes the expectation operator conditional on the information available in period t. We assume that $\chi > 0$ and β is between 0 and 1. The left hand side of the intertemporal budget constraint (5) represents the U.S. dollar value of household wealth brought into the beginning of period $t+1$. It consists of cash holding, M_t, and bond holdings, B_{t+1}. As shown on the right-hand-side of (5), the household receives nominal labor income $W_{t-1} n_{t-1}$ in period $t-1$ which earns gross nominal interest R_{t-1} until period t.[4] The household carries nominal bonds in amount B_t and cash holding remaining after consumption expenditures $(M_{t-1} - P_{t-1} c_{t-1})$ into period t. Finally, T_t and Π_t are nominal lump sum transfers from the U.S. government and nominal profits of firms operating in the U.S., respectively.

Equation (6) is the CIA constraint. The aggregate price P_t is given by $P_t = [\int P_t(j)^{1-\theta} dj]^{\frac{1}{1-\theta}}$, where $P_t(j)$ is the aggregate price index for good j; it is a CES aggregate over city-specific prices for that good: $P_t(j) = [\int P_t(j,l)^{1-\theta} dl]^{\frac{1}{1-\theta}}$. The price index for good j in a particular city l used in this aggregation is given by

$$P_t(j,l) = \left[\int P_t(j,l,z)^{1-\theta} dz \right]^{\frac{1}{1-\theta}}. \tag{7}$$

Households in Canada solve the analogous optimization problem except we must convert their U.S. dollar bond holdings into Canadian dollars at the spot nominal exchange rate, S_t. Thus the Canadian-dollar intertemporal budget constraint is

$$M_t^* + \frac{\mathbb{E}_t(\Upsilon_{t,t+1} B_{t+1}^*)}{S_t} = \frac{S_{t-1} R_{t-1}}{S_t} W_{t-1}^* n_{t-1}^* + \frac{B_t^*}{S_t} + (M_{t-1}^* - P_{t-1}^* c_{t-1}^*) + T_t^* + \Pi_t^*. \tag{8}$$

The key equations from the consumer choice problem are

$$\frac{W_t}{P_t} = \chi c_t, \tag{9}$$

$$M_t = P_t c_t, \tag{10}$$

[4] We assume that the government pays interest rate $R_t (= 1/\mathbb{E}_t \Upsilon_{t,t+1})$ on labor income in period t. This assumption allows households' intratemporal first-order condition to be undistorted.

$$\Upsilon_{t,t+1} = \beta \left[\left(\frac{c_{t+1}}{c_t} \right)^{-1} \frac{P_t}{P_{t+1}} \right] \quad = \quad \beta \left[\left(\frac{c^*_{t+1}}{c^*_t} \right)^{-1} \frac{S_t P^*_t}{S_{t+1} P^*_{t+1}} \right] . \tag{11}$$

The counterparts of (9) and (10) apply to Canada as well.

Combining the intratemporal condition with the CIA constraints, for the U.S., we have

$$W_t \quad = \quad \chi M_t. \tag{12}$$

The nominal wage rate in the U.S. is proportional to the stock of money held by households in the U.S. and an analogous condition holds for Canada.

The aggregate real exchange rate is determined by combining the home and foreign intertemporal conditions:

$$q_t = \frac{S_t P^*_t}{P_t} = \kappa \frac{c_t}{c^*_t}, \tag{13}$$

where $\kappa = q_0 c^*_0 / c_0$. The nominal exchange rate is determined by combining the home and foreign CIA constraints with (13):

$$S_t = \kappa \frac{M_t}{M^*_t}. \tag{14}$$

2.2 Firms

The output of a firm is equal to the number of hours worked:

$$y_t(j, z) = n_t(j, z). \tag{15}$$

Goods are perishable, so the consumption of each good across all cities equals output of that good in the current period:

$$\int c_t(j, l, z) dl + \int [1 + \tau(j, l^*)] c^*_t(j, l^*, z) dl^* = y_t(j, z). \tag{16}$$

We allow for long-run deviations from the LOP across borders through $\tau(j, l^*)$, an iceberg transportation cost of exporting good j from the U.S. to a Canadian city indexed by l^*. A firm must ship $1 + \tau(j, l^*)$ units of good j to city l^* for one unit of that good to arrive at the destination.

An analogous market clearing condition holds for each of the Canadian goods:

$$\int [1 + \tau(j, l)] c_t(j, l, z) dl + \int c^*_t(j, l^*, z) dl^* = y^*_t(j, z). \tag{17}$$

2.3 Price adjustment and information updating

This sub-section begins by reviewing Calvo pricing used by Kehoe and Midrigan (2007) and then presents our extension to allow for information updating as in Mankiw and Reis (2002).

2.3.1 Calvo pricing

We model the nominal price rigidities as in Calvo (1983) and Yun (1996): each month a fraction of firms $1-\lambda_j$ are randomly drawn and allowed to reset their prices. Prices are assumed to be set in the currency of the country in which the goods are sold, local currency pricing. As indicated by the subscript, the frequency of price changes varies by good j, but not by country.

The optimal price for U.S. firms selling good j in city l is the solution to the following maximization problem:

$$\max_{P_{H,t}(j,l)} \mathbb{E}_t \sum_{h=0}^{\infty} \lambda_j^h \Upsilon_{t,t+h}[P_{H,t}(j,l) - W_{t+h}] \left(\frac{P_{H,t}(j,l)}{P_{t+h}}\right)^{-\theta} c_{t+h},. \tag{18}$$

We have used the fact that since the elasticity of substitution is the same across all subaggregators, we can express demand for good j in city l as

$$c_t(j,l) = \left(\frac{P_t(j,l)}{P_t}\right)^{-\theta} c_t. \tag{19}$$

The first-order condition for $P_{H,t}(j,l)$ is

$$\begin{aligned} &\mathbb{E}_t \sum_{h=0}^{\infty} \lambda_j^h \Upsilon_{t,t+h} \left(\frac{P_{H,t}(j,l)}{P_{t+h}}\right)^{-\theta} c_{t+h} \\ &= \frac{\theta}{\theta-1} \mathbb{E}_t \sum_{h=0}^{\infty} \lambda_j^h \Upsilon_{t,t+h} \left(\frac{W_{t+h}}{P_{H,t}(j,l)}\right) \left(\frac{P_{H,t}(j,l)}{P_{t+h}}\right)^{-\theta} c_{t+h}. \end{aligned} \tag{20}$$

Similarly, the first-order condition for Canadian firms' selling in a U.S. city l in U.S. dollars, conditional on adjustment, $P_{F,t}(j,l)$, is

$$\begin{aligned} &\mathbb{E}_t \sum_{h=0}^{\infty} \lambda_j^h \Upsilon_{t,t+h} \left(\frac{P_{F,t}(j,l)}{P_{t+h}}\right)^{-\theta} c_{t+h} \\ =&\frac{\theta}{\theta-1} \mathbb{E}_t \sum_{h=0}^{\infty} \lambda_j^h \Upsilon_{t,t+h} \left(\frac{(1+\tau(j,l))S_{t+h}W^*_{t+h}}{P_{F,t}(j,l)}\right) \left(\frac{P_{F,t}(j,l)}{P_{t+h}}\right)^{-\theta} c_{t+h}. \end{aligned} \tag{21}$$

2.3.2 Calvo pricing with infrequent information updating

We now consider firms facing two nominal rigidities. First, each firm has a constant probability of price resetting $1-\lambda_j$ as before. Second, with probability of $1-\omega$, a firm receives an information update in the current month. The fraction of firms that fail to get updates, ω, use the information available from their most recent update.

DKT develop this combined stickiness structure to explain persistent inflation dynamics. In DKT, infrequent price changes arise due to the Calvo assumption of price changes. However, when firms compute the optimal level at which to reset prices, a fraction of firms use the newest information set and the remaining firms use the stale information set to determine prices. Following DKT, we employ this structure and refer to it as "dual stickiness" pricing.

All U.S. firms that sell good j in city l choose different prices according to the vintage of their information set. When firms are allowed to adjust prices, those with the same vintage of information choose the same price. Let $P_{H,t}^k(j,l)$ be the optimal price reset by U.S. firms conditional on information of vintage k, its age in months. The price $P_{H,t}^k(j,l)$ for these firms solves

$$\max_{P_{H,t}^k(j,l)} \mathbb{E}_{t-k} \sum_{h=0}^{\infty} \lambda_j^h \Upsilon_{t,t+h} [P_{H,t}^k(j,l) - W_{t+h}] \left(\frac{P_{H,t}^k(j,l)}{P_{t+h}} \right)^{-\theta} c_{t+h}, \tag{22}$$

for $k = 0, 1, 2, \cdots$ and for all cities $l \in [0, 1]$.

The first-order condition for $P_{H,t}^k(j,l)$, for $k = 0, 1, 2, \cdots$, is

$$\begin{aligned} &\mathbb{E}_{t-k} \sum_{h=0}^{\infty} \lambda_j^h \Upsilon_{t,t+h} \left(\frac{P_{H,t}^k(j,l)}{P_{t+h}} \right)^{-\theta} c_{t+h} \\ &= \frac{\theta}{\theta - 1} \mathbb{E}_{t-k} \sum_{h=0}^{\infty} \lambda_j^h \Upsilon_{t,t+h} \left(\frac{W_{t+h}}{P_{H,t}^k(j,l)} \right) \left(\frac{P_{H,t}^k(j,l)}{P_{t+h}} \right)^{-\theta} c_{t+h}, \end{aligned} \tag{23}$$

Canadian firms change prices to satisfy an analogous first-order condition to (23):

$$\begin{aligned} &\mathbb{E}_{t-k} \sum_{h=0}^{\infty} \lambda_j^h \Upsilon_{t,t+h} \left(\frac{P_{F,t}^k(j,l)}{P_{t+h}} \right)^{-\theta} c_{t+h} \\ =&\frac{\theta}{\theta - 1} \mathbb{E}_{t-k} \sum_{h=0}^{\infty} \lambda_j^h \Upsilon_{t,t+h} \left(\frac{(1+\tau(j,l)) S_{t+h} W_{t+h}^*}{P_{F,t}^k(j,l)} \right) \left(\frac{P_{F,t}^k(j,l)}{P_{t+h}} \right)^{-\theta} c_{t+h}. \end{aligned} \tag{24}$$

This completes the descriptions of the problems solved by consumers and firms.

2.4 Equilibrium

The monetary authority in each country sets the growth rate of the money stock such that it follows an AR(1):

$$\ln \mu_t = \rho \ln \mu_{t-1} + \varepsilon_t, \tag{25}$$

$$\ln \mu_t^* = \rho \ln \mu_{t-1}^* + \varepsilon_t^*, \tag{26}$$

where ε_t and ε_t^* are mean-zero i.i.d shocks and $\mu_t = M_t/M_{t-1}$ and $\mu_t^* = M_t^*/M_{t-1}^*$. The steady state (log) money growth rates are set to zero and the common persistence parameter satisfies $\rho \in [0, 1)$.

Total transfers from the government to individuals in each country equal domestic money injections minus the lump sum tax from the government paying interest. For the U.S., we have $T_t = M_t - M_{t-1} - (R_{t-1} - 1)W_{t-1}n_{t-1}$. The total transfers in Canada are of the same form up to currency conversions $T_t^* = M_t^* - M_{t-1}^* - (S_{t-1}R_{t-1}/S_t - 1)W_{t-1}^* n_{t-1}^*$.

The profits of U.S. firms accrue exclusively to U.S. households. In other words, $\Pi_t = \int_j \int_{z=0}^{\frac{1}{2}} \Pi_t(j,z)dzdj$, where $\Pi_t(j,z)$ is the profit of a U.S. firm. Similarly, the profits of Canadian firms accrue exclusively to Canadian households: $\Pi_t^* = \int_j \int_{z=\frac{1}{2}}^{1} \Pi_t^*(j,z)dzdj$, where $\Pi_t^*(j,z)$ is the profit of a Canadian firm.

Recall, market clearing conditions for good markets were given by (16) and (17). The labor market clearing conditions are

$$n_t = \int_j \int_{z=0}^{\frac{1}{2}} n_t(j,z)dzdj, \tag{27}$$

$$n_t^* = \int_j \int_{z=\frac{1}{2}}^{1} n_t^*(j,z)dzdj. \tag{28}$$

Last, but not least, the state-contingent bond market clears at each date and state: $B_t + B_t^* = 0$ for all t.

An *equilibrium* of the model is a collection of allocations and prices satisfying household and firm maximization problems and market clearing conditions; as well as the money supply processes and transfers satisfying the specifications noted above.

3 Theoretical implications for LOP dynamics

3.1 Calvo pricing

The model allows money to be a non-stationary random variable. To achieve stationarity, we normalize all nominal variables by the nominal money stock. For example, the normalized U.S. dollar price becomes $p_{H,t}(j,l) = P_{H,t}(j,l)/M_t$. We work with log-linear deviations of these transformed variables from their steady-state levels and these deviations are denoted with ' ˆ 's over them. The reset price based on this normalization, in log-deviations from their steady-state level, for U.S. firms selling good j in city l, is

$$\hat{p}_{H,t}(j,l) = (1-\lambda_j\beta)\sum_{h=0}^{\infty}(\lambda_j\beta)^h \mathbb{E}_t\left(\hat{\mu}_{t+1,t+h}\right) = \left[\frac{\lambda_j\beta\rho}{1-\lambda_j\beta\rho}\right]\hat{\mu}_t \ , \tag{29}$$

where

$$\hat{\mu}_{t+1,t+h} = \begin{cases} 0 & \text{for } h=0 \\ \sum_{d=1}^{h}\hat{\mu}_{t+d} & \text{for } h=1,2,\cdots . \end{cases} \tag{30}$$

Here the proportionality of money and nominal wages, $W_t = \chi M_t$, is used to replace endogenous wages with exogenous money. We see that firms adjusting their prices in period t adjust them in proportion to the present discounted value of future marginal cost changes during periods of price non-adjustment.

Analogously, the log-deviation of reset price relative to the U.S. nominal money supply for Canadian firms from (21):

$$\hat{p}_{F,t}(j,l) = (1-\lambda_j\beta)\sum_{h=0}^{\infty}(\lambda_j\beta)^h \mathbb{E}_t\left(\hat{\mu}_{t+1,t+h}\right) = \left[\frac{\lambda_j\beta\rho}{1-\lambda_j\beta\rho}\right]\hat{\mu}_t, \tag{31}$$

where we have used the fact that under the monetary approach, the change in the nominal exchange rate is proportional to the change in relative money supplies. As it turns out $\hat{p}_{F,t}(j,l) = \hat{p}_{H,t}(j,l)$ so the short-run dynamics of the optimal prices are the same for home and foreign firms selling the same good at the same location in spite of the transportation costs which drive a wedge between the prices in the long-run.

Thus, the log-deviation of the price index, $p_t(j,l) = P_t(j,l)/M_t$ under Calvo pricing becomes

$$\hat{p}_t(j,l) = \lambda_j\hat{p}_{t-1}(j,l) - \lambda_j\hat{\mu}_t + (1-\lambda_j)\left[\frac{\lambda_j\beta\rho}{1-\lambda_j\beta\rho}\right]\hat{\mu}_t\ . \tag{32}$$

The analogous expression for the Canadian price index for good j and Canadian city l^* is

$$\hat{p}^*_t(j,l^*) = \lambda_j\hat{p}^*_{t-1}(j,l^*) - \lambda_j\hat{\mu}^*_t + (1-\lambda_j)\left[\frac{\lambda_j\beta\rho}{1-\lambda_j\beta\rho}\right]\hat{\mu}^*_t, \tag{33}$$

and the log-deviation of the bilateral real exchange rate for good j across cities l and l^* is $\hat{q}_t(j,l,l^*) = \ln q_t(j,l,l^*) - \ln q(j,l,l^*)$, where $q_t(j,l,l^*)$ is given by

$$q_t(j,l,l^*) = \frac{S_t P^*_t(j,l^*)}{P_t(j,l)}. \tag{34}$$

The next proposition characterizes the short-run, good-level, real exchange rate dynamics under Calvo pricing with a slight generalization of Kehoe and Midrigan (2007).

Proposition 1. *Under the preference assumption* $\ln c - \chi n$*, the CIA constraints, the assumption of money growth* (25) *and* (26) *and good-specific Calvo pricing, the good-level real exchange rate between any cities* l *and* l^* *follows an AR(2) process of the form*

$$\hat{q}_t(j,l,l^*) = (\lambda_j+\rho)\hat{q}_{t-1}(j,l,l^*) - \lambda_j\rho\hat{q}_{t-2}(j,l,l^*) + \theta_j\eta_t, \tag{35}$$

where $\theta_j = \lambda_j - (1-\lambda_j)(\lambda_j\beta\rho)/(1-\lambda_j\beta\rho)$*, and* $\eta_t = \varepsilon_t - \varepsilon^*_t$ *is i.i.d.*$(0,\sigma^2_\eta)$*.*

Proof. From (13) and (14), $\hat{q}_t(j,l,l^*) = \hat{p}^*_t(j,l^*) - \hat{p}_t(j,l)$. Subtracting (32) from (33) yields $\hat{q}_t(j,l,l^*) = \lambda_j \hat{q}_{t-1}(j,l,l^*) + \theta_j(\hat{\mu}_t - \hat{\mu}^*_t)$. Because $\hat{\mu}_t - \hat{\mu}^*_t$ follow an AR(1) from (25) and (26), we obtain (35) and proved Proposition 1. □

Proposition 1 of Kehoe and Midrigan (2007) is a special case of the one above: when money growth rates follow an i.i.d. process ($\rho = 0$) equation (35) reduces to an AR(1) model with persistence λ_j and $\theta_j = \lambda_j$ as Kehoe and Midrigan (2007) prove.[5]

3.2 Calvo pricing with infrequent information updating

Let $\hat{p}^k_{H,t}(j,l)$ be the log deviation of $P^k_{H,t}(j,l)/M_t$ from the steady state. Log-linearizing (23) around the steady state yields

$$\hat{p}^k_{H,t}(j,l) = (1-\lambda_j\beta)\sum_{h=0}^{\infty}(\lambda_j\beta)^h \mathbb{E}_{t-k}\left(\hat{\mu}_{t+1,t+h}\right), \text{ for } k=0 \tag{36}$$

and

$$\begin{aligned}\hat{p}^k_{H,t}(j,l) = (1-\lambda_j\beta)&\sum_{h=0}^{\infty}(\lambda_j\beta)^h \mathbb{E}_{t-k}\left(\hat{\mu}_{t+1,t+h}\right)\\ &+ \mathbb{E}_{t-k}(\hat{\mu}_t + \hat{\mu}_{t-1} + \cdots + \hat{\mu}_{t-k+1}) - (\hat{\mu}_t + \hat{\mu}_{t-1} + \cdots + \hat{\mu}_{t-k+1}), \text{ for } k = 1,2,3,\cdots.\end{aligned} \tag{37}$$

The only differences between this equation and Calvo pricing are that $\hat{p}_{H,t}(j,l)$ and $\mathbb{E}_t$ are replaced with $\hat{p}^k_{H,t}(j,l)$ and $\mathbb{E}_{t-k}$ and that, under $k > 0$, forecast errors are accumulated from the period in which firms last update information.[6]

As in Calvo pricing, it is possible to show that the (normalized) price index for good j in location l can be expressed in terms of lagged prices, reset prices and money growth rates. However, the reset prices are weighted averages of price resets given different vintage of information:

$$\hat{p}_t(j,l) = \lambda_j \hat{p}_{t-1}(j,l) - \lambda_j\hat{\mu}_t + (1-\lambda_j)\hat{x}_t(j,l), \tag{38}$$

where $\hat{x}_t(j,l)$ is the (normalized) weighted average for the newly set prices for good j in city l of the U.S., based upon different information vintages,

$$\begin{aligned}\hat{x}_t(j,l) = (1-\omega)&\sum_{k=0}^{\infty}\omega^k \mathbb{E}_{t-k}\hat{p}_{H,t}(j,l)\\ &+ (1-\omega)\sum_{k=1}^{\infty}\omega^k\left[\mathbb{E}_{t-k}(\hat{\mu}_t + \hat{\mu}_{t-1} + \cdots + \hat{\mu}_{t-k+1}) - (\hat{\mu}_t + \hat{\mu}_{t-1} + \cdots + \hat{\mu}_{t-k+1})\right],\end{aligned} \tag{39}$$

[5] The same proposition was independently derived by Carvalho and Nechio (2008) who focus on the role of sticky prices in producing aggregation bias in the persistence of the aggregate real exchange rate.

[6] Note that $\hat{p}^0_{H,t}(j,l) = \hat{p}_{H,t}(j,l)$ because of the equivalence between (20) and (23) when $k = 0$.

which is similar in mathematical formulation to the price index in Mankiw and Reis (2002, p.1300).

Canadian versions of these expressions are

$$\hat{p}_t^*(j,l^*) = \lambda_j \hat{p}_{t-1}^*(j,l) - \lambda_j \hat{\mu}_t^* + (1-\lambda_j)\hat{x}_t^*(j,l^*) \ , \tag{40}$$

$$\begin{aligned} \hat{x}_t^*(j,l^*) &= (1-\omega)\sum_{k=0}^{\infty}\omega^k \mathbb{E}_{t-k}\hat{p}_{H,t}^*(j,l^*) \\ &+(1-\omega)\sum_{k=1}^{\infty}\omega^k \left[\mathbb{E}_{t-k}(\hat{\mu}_t^* + \hat{\mu}_{t-1}^* + \cdots + \hat{\mu}_{t-k+1}^*) - (\hat{\mu}_t^* + \hat{\mu}_{t-1}^* + \cdots + \hat{\mu}_{t-k+1}^*)\right] . \end{aligned} \tag{41}$$

The next proposition establishes the rich short-run dynamics of the good-level real exchange rate emerging from the extended model.

Proposition 2. *Under the preference assumption* $\ln c - \chi n$*, the CIA constraints, the assumption of money growth* (25) *and* (26)*, along with good-specific Calvo pricing and Mankiw-Reis information updating, the good-level real exchange rate between any cities* l *and* l^* *follows an ARMA(4,2) process of the form*

$$\hat{q}_t(j,l,l^*) = \sum_{r=1}^{4}\phi_{j,r}\hat{q}_{t-r}(j,l,l^*) + \sum_{r=0}^{2}\theta_{j,r}\eta_{t-r} \ , \tag{42}$$

where the coefficients are known functions of β*,* ρ*,* λ_j*, and* ω*.*[7]

When $\omega = 0$ this proposition reduces to Proposition 1.

4 Quantitative theoretical implications for LOP persistence and volatility

The theoretical model places restrictions on the relationship between the structural parameters of the model, ρ, λ_j and ω and the parameters of the univariate ARMA model we estimate in the next section. Here we explore the quantitative implications of the theory by parameterizing the theory to develop intuition for how the structural parameters impact LOP persistence and volatility.

There are only two free parameters in this exercise, ρ and ω, since the λ_j's are pinned down by the estimates of the frequency of price adjustments existing in the literature. We consider two extreme values for each of the free parameters.

The equilibrium conditions of our model imply that the persistence of changes in the nominal exchange rate, changes in the nominal money stock and changes in nominal GDP are equal to

[7] Details are available in the working paper version, Crucini, Shintani, and Tsuruga (2008).

each other. That the data suggest otherwise reflects a combination of model mis-specification and difficulty in precisely estimating the persistence parameter. Our approach is to present the extremes in this section since this enhances intuition, and to consider points in between these extremes in the estimation section.

For the persistence parameter, ρ, two key benchmarks are 0 and 0.83. The first is the value used by Kehoe and Midrigan (2007) and is consistent with the view that the nominal exchange rate follows a random walk (see Meese and Rogoff (1983)). The latter benchmark is the monthly analog to the CKM calibration for M1 growth.[8]

For the information updating parameter, ω, less prior evidence exists to restrict its range. To frame the discussion of its role in accounting for persistence in real exchange rates of individual goods and services, we pick values that encompass the diverse range of LOP persistence estimates found in the micro-price data. In the empirical section we estimate the information updating parameter by fixing the persistence of the money shocks and minimizing the distance between the persistence implied by the restricted model and the estimated persistence in the micro-data.

4.1 Calvo pricing

4.1.1 Persistence

Turning to implications for persistence of the good-level real exchange rates we employ the sum of autoregressive coefficients (SAR) as the persistence metric. This is often the case in applied work when moving beyond the AR(1) model (e.g., Andrews and Chen (1994) and Clark (2006)) because the SAR has a one-to-one relationship to the cumulative long-run impulse response to a shock. We denote the SAR by α_j.

Under Proposition 1, the SAR measure of persistence is $\alpha_j = \lambda_j + \rho(1 - \lambda_j)$; it simplifies to $\alpha_j = \lambda_j$ when $\rho = 0$. Obviously, the SAR is strictly increasing in ρ regardless of the degree of price stickiness under $\lambda_j \in [0, 1)$. The left panel of Figure 1 shows the effect of increasing ρ on the persistence for two hypothetical goods: a good with relatively slow price adjustment ($\lambda_j = 0.95$) and a good with relatively fast price adjustment ($\lambda_j = 0.5$).

The right panel of Figure 1 plots the SAR against λ_j. The figure compares the impact of changing ρ from 0 to 0.83. The impact of introducing persistence in $\hat{\mu}_t$ and $\hat{\mu}_t^*$ on the SAR is clear.

[8] The CKM estimate of the autoregressive coefficient is 0.68 using quarterly U.S. data for M1 growth. We transformed this quarterly persistence of M1 growth into the monthly persistence by solving $Cov(\hat{M}_t - \hat{M}_{t-3}, \hat{M}_{t-3} - \hat{M}_{t-6})/Var(\hat{M}_t - \hat{M}_{t-3}) = 0.68$ for ρ and obtained 0.83.

When $\rho = 0$, the model predicts that the SAR equals λ_j, so the two parameters lie on the 45 degree line in the figure. On the other hand, when $\rho > 0$, the model predicts a flatter line. Thus, a high persistence of the money growth rates increases the persistence of LOP deviations, regardless of the frequency of price adjustment, but the quantitative impact is greatest for goods with the highest frequency of price adjustment.

To see the intuition behind the persistent dynamics it is instructive to express the current LOP deviation as a function of its lagged self and the change in the nominal exchange rate

$$\hat{q}_t(j, l, l^*) = \lambda_j \hat{q}_{t-1}(j, l, l^*) + \theta_j \Delta \hat{S}_t \ , \tag{43}$$

where $\Delta \hat{S}_t = \hat{\mu}_t - \hat{\mu}_t^*$ from (14). When $\rho = 0$ as in Kehoe and Midrigan (2007), $\Delta \hat{S}_t$ is an i.i.d shock and the good-level real exchange rate follows an AR(1) with persistence parameter, λ_j. When international money growth differential is positively autocorrelated ($\rho > 0$) the nominal exchange rate change contributes to increased persistence in the real exchange rate.

4.1.2 Volatility

Throughout, real exchange rate volatility will be measured relative to the standard deviation of the change in the nominal exchange rate: $\sigma_j = \text{std}(q_t(j, l, l^*))/\text{std}(\Delta S_t)$. When the nominal exchange rate follows a random walk, $\rho = 0$, the model predicts the normalized standard deviation to be $\sigma_j = \sigma_1(\lambda_j) = \lambda_j / \sqrt{1 - \lambda_j^2}$ and a good with larger λ_j will exhibit more variability. When $\rho > 0$, the normalized standard deviation is predicted to be of the form $\sigma_j = \sigma_2(\lambda_j, \rho, \beta)$ and may be obtained using the variance formula of an AR(2) process along with $\text{std}(\Delta S_t) = \text{std}(\eta_t)/\sqrt{1 - \rho^2}$. Importantly, the volatility function depends not just on λ_j, but also on ρ and β.

An implication of this is that increased persistence in money growth, while helpful in resolving the persistence puzzle, may actually make the volatility puzzle worse because $\sigma_j = \sigma_2(\lambda_j, \rho, \beta)$ is not monotonic in ρ. Moreover the shape of the relationship with ρ depends on the frequency of price adjustment, which we know differs across goods. The practical thrust of this is: changes in money growth persistence will have differential impacts across goods.

The left panel of Figure 2 plots the normalized standard deviations $\sigma_j = \sigma_2(\lambda_j, \rho, \beta)$ against ρ.[9] For a good with relatively infrequent price changes ($\lambda_j = 0.95$), volatility of the real exchange rate rises over the empirically relevant range of money growth persistence. In contrast, for a good with relatively frequent price changes ($\lambda_j = 0.5$), the volatility of the relative price is declining in

[9] We set the discount factor β to 0.99.

the money growth rate persistence throughout. The right panel of Figure 2 shows the ambiguous impact of introducing a positive ρ on the volatility from another dimension. The normalized standard deviation is smaller for $\rho = 0.83$ than for $\rho = 0$ when price adjustment is fast. When the price adjustment is slow, the ranking reverses.

4.2 Calvo pricing with infrequent information updating

4.2.1 Persistence

The SAR under dual stickiness pricing is given by

$$\alpha_j = \sum_{r=1}^{4} \phi_{j,r} = 1 - (1 - \lambda_j)(1 - \omega)(1 - \rho)(1 - \omega\rho). \tag{44}$$

Clearly, the slower the speed of information updating ($\omega \to 1$), the larger the SAR becomes.

For a general ARMA process without parameter restrictions, it is not conventional to use the SAR as a measure of persistence, because of the presence of MA terms. However, if our model is correctly specified, we can show that both the long-run impact of cumulative impulse response of a unit monetary shock on real exchange rates and the SAR is a strictly increasing function of λ_j, ω, and ρ. Furthermore, using the SAR is convenient in computation and for the purpose of making comparison with simpler models introduced in the previous sub-section. For these reasons, we continue to focus on the SAR as an approximate measure of persistence under the assumption that (42) is correctly specified.

The dual stickiness model works well in generating the persistence of real exchange rates. The left panel of Figure 3 shows the SAR among different ω's. The persistence is increasing in ω and is very high regardless of the infrequency of price changes.[10] The right panel of Figure 3 plots the persistence against λ_j. This panel compares cases of two extreme values of ω. One is the case in which firms producing good j updates their information every month. (i.e., $\omega = 0$.) The other is the case in which firms, on average, update information every 50 months (i.e., $\omega = 0.98$). For the former case, the obtained SAR corresponds to the upper line in the right panel of Figure 1 since we set $\rho = 0.83$ in the computation. In the latter case, the persistence measure is very close to one whether prices are sticky or flexible.

[10] Even if $\omega = 0$, $\hat{q}_t(j, l, l^*)$ is already somewhat persistent, because of the AR(1) money growth.

4.2.2 Volatility

Having improved the model in accounting for persistence of real exchange rates, we ask if information updating also helps to account for variability. We calculate the new normalized standard deviation $\sigma_j = \sigma_3(\omega, \lambda_j, \rho, \beta)$, using the fact that the good-level real exchange rates now follow the ARMA(4,2) process according to Proposition 2. The left panel of Figure 4 plots the normalized standard deviations against ω. It shows that the volatility grows exponentially as ω increases. The right panel of Figure 4 shows the effect of increasing λ_j on the normalized standard deviations under the two extreme cases: $\omega = 0$ and 0.98. It shows that real exchange rate volatility becomes substantially greater when the information adjustment is slower. Thus, the introduction of information stickiness enhances the real exchange rate volatility.

The question we pose next is what lengths of information delays are needed to match the persistence and volatility of good-level real exchange rates in the EIU micro-data, conditional on the model and the observed frequency of price changes.

5 Empirical results

5.1 Data

The retail prices come from the *Worldwide Cost of Living Survey* compiled by the Economist Intelligence Unit (EIU). It is an extensive annual survey of international retail prices designed to help managers determine compensation levels of their employees residing in different cities. The coverage of goods and services is broad enough to overlap significantly with what appears in a typical urban consumption basket (see Rogers (2007)). Two advantages of the EIU data are the fact that the prices are in absolute terms and that the survey is conducted by a single agency in a consistent manner over time. Because of this convenient panel data format, a number of recent studies on international price dynamics have used this data, including Crucini and Telmer (2007), Crucini and Shintani (2008), Engel and Rogers (2004), Parsley and Wei (2007) and Rogers (2007).

For a limited number of countries, the EIU data contains observations from multiple cities. In our empirical analysis, we focus on U.S.-Canadian city pairs since the assumption of the common probability of price adjustment for each good seems a reasonable approximation between these two countries.[11]

[11] Amirault, Kwan, and Wilkinson (2006) report frequencies of price adjustment for Canadian firms that match up reasonably well with the U.S. evidence. In other contexts, one may use the average of price change frequencies

After removing missing observations to construct a balanced panel for the period from 1990 to 2005, three of the 16 U.S. cities available in the survey are dropped, while all four Canadian cities remain. This results in a total of 52 cross-border city pairs.

For each good j, the log of $q_t(j, l, l^*)$ for each year t $(= 1, ..., 16)$ is computed using the price level in a U.S. city l $(= 1, ..., 13)$ expressed in U.S. dollars $(P_t(j, l))$, the price level in a Canadian city l^* $(= 1, ..., 4)$ expressed in Canadian dollars $(P_t^*(j, l^*))$, and the spot U.S.–Canadian dollar exchange rate (S_t), all from the EIU data.

Next, for the price stickiness parameter, λ_j, we utilize the frequency of price changes, f_j and transform it with $\lambda_j = 1 - f_j$ for good j. Since the EIU data is annual, it is not useful for constructing estimates of the frequency of price changes. Here we rely on Nakamura and Steinsson (2008) who revisited Bils and Klenow's analysis using more detailed and updated data from the Bureau of Labor Statistics (BLS). Using the CPI Research Database created by the BLS, they re-estimated the frequencies of price change after removing temporary price changes associated with sales. They found that the median duration between regular price changes was 8 - 11 months depending on the treatment of substitutions, considerably higher than 4.3 months for the median good, found by Bils and Klenow (2004). In what follows, we use the Nakamura and Steinsson frequencies as our benchmark.[12]

We take the monthly average frequency of price changes, f_j, and match them with the 165 goods in the EIU sample. Since we require paired persistence and frequency adjustment parameters to evaluate the model, we use only these 165 matched pairs in our analysis. We assume that the frequency of price changes applies to the entire sample period of 1990-2005 in our EIU data set.[13]

For the nominal exchange rate changes required for the theoretical volatility calculation, we use monthly changes in the log of the end-of-month U.S.-Canadian dollar spot rates. While both price stickiness parameter (frequency of no price changes) and nominal exchange rates are available in monthly series, real exchange rates are only observed annually. The small number of time series observations at the annual frequency is the major limitation of the EIU data. In the next sub-

between the two countries, an approach employed in Kehoe and Midrigan (2007), when data from both countries are available.

[12]The working paper version of this paper compared the two.

[13]In some countries which experienced a structural shift in inflation, an assumption of constant frequency of price changes over years may not be satisfied. For example, Ahlin and Shintani (2007) use Mexican price data on 44 goods and report that the average monthly frequency of price changes rose from 28% in 1994 to 50% in 1995 as inflation rose over the same period. We expect that this issue is less serious in our case since both U.S. and Canada had more stable and comparable inflation rates during the period under consideration.

section, we briefly discuss how to reconcile the mixed frequencies of observation in the dynamic panel estimation and describe the procedure to estimate the time series models.

5.2 Estimation

Table 1 shows how monthly ARMA processes predicted by the model are transformed into the ones which have non-zero coefficients for multiples of 12 month lags and finite MA terms. The first row of the table shows the easiest transformation. In Calvo pricing with $\rho = 0$, the equation (35) simplifies to

$$\hat{q}_t(j,l,l^*) = \lambda_j \hat{q}_{t-1}(j,l,l^*) + \lambda_j \eta_t. \tag{45}$$

By repeated substitutions, we get

$$\hat{q}_t(j,l,l^*) = \lambda_j^{12} \hat{q}_{t-12}(j,l,l^*) + \lambda_j \Lambda_j(L)\eta_t \ , \tag{46}$$

where $\Lambda_j(L) = \sum_{r=0}^{11} \lambda_j^r L^r$. In this equation, the AR term is the 12th lag (in months) and the order of the MA term is 11. This ARMA(12,11) process is equivalent to an AR(1) process sampled annually since $\lambda_j \Lambda_j(L)\eta_t$ and $\hat{q}_{t-12}(j,l,l^*)$ are not correlated.

Such a transformation is not necessarily possible with a general ARMA process including AR(2) and ARMA(4,2) processes. However, thanks to a special dynamic feature of the theoretical model, it is possible to make the AR parameters non-zero only if the lags are multiples of 12 and the MA parameters are finite under our extended models (35) and (42).

Previously, l and l^* were used for the U.S. and Canadian cities, respectively. Here, they are replaced by a new single index i $(= 1, ..., 52)$ each representing a city pair spanning the border. In addition, the sampling frequency for the model was assumed to be monthly. With some abuse of notation, our new time subscript now represents the time in annual frequency. With this newly introduced index, we define $q_{i,t}^j$ as the *log* of the real exchange rate for good j between the city pair indexed by i at year t

$$q_{i,t}^j = \ln q_t(j,l,l^*). \tag{47}$$

Thus, the former log deviation from the steady state $\widehat{q}_t(j,l,l^*)$ can be rewritten as $q_{i,t}^j - q_i^j$, where q_i^j is the long-run value which equals:

$$q_i^j = \ln q(j,l,l^*) = \ln \frac{[1+\kappa^{1-\theta}(1+\tau(j,l^*))^{1-\theta}]^{\frac{1}{1-\theta}}}{[1+\kappa^{1-\theta}(1+\tau(j,l))^{1-\theta}]^{\frac{1}{1-\theta}}}. \tag{48}$$

Intuitively, the relative price of a good in the long-run is higher in the destination market with the higher shipping cost from the source. Thus if city l^* is, say, farther from the source of the good than city l, q_i^j is positive. These heterogeneous long-run deviations justify the presence of individual effects (the time invariant city pair-specific effect) in the panel estimation.

Based on the annual transformation shown in Table 1, all of the theoretical cases are nested in the following empirical model of the LOP deviation for good j

$$q_{i,t}^j = \sum_{r=1}^{m} \Phi_{j,r} q_{i,t-r}^j + \zeta_i^j + u_t^j + v_{i,t}^j, \tag{49}$$

where ζ_i^j is the time invariant unobserved city pair-specific effect which allows long-run price difference between two cities, u_t^j is a common time-effect which the theory attributes to a single nominal exchange rate shock having differential impacts across goods because of the good-specific Calvo parameter. The last term, $v_{i,t}^j$, are shocks not captured by the theoretical model, that is, shocks specific to both the good and city-pair at date t.

This empirical model nests all the theoretical models under consideration: (i) Calvo pricing with $\rho = 0$ implies $m = 1$; (ii) Calvo pricing with $\rho \neq 0$ implies $m = 2$; and (iii) dual stickiness pricing implies $m = 4$. For the individual specific effect ζ_i^j, we can easily see its relationship to the long-run mean and the persistence from $q_i^j = \zeta_i^j/(1-\alpha_j)$ where $\alpha_j = \sum_{r=1}^m \Phi_{j,r}$. For the common time effect u_t^j, Calvo pricing with $\rho \neq 0$ predicts a serial correlation of order one, while dual stickiness pricing predicts a serial correlation of order three. However, in a short panel asymptotic with finite T, the common time effects can be treated as unknown parameters to be estimated with time dummies.

Since our main interest is to estimate the SAR, α_j, it is convenient to rewrite the model into the augmented Dickey-Fuller (ADF) form.

$$q_{i,t}^j = \alpha_j q_{i,t-1}^j + \sum_{r=1}^{m-1} \gamma_{j,r} \Delta q_{i,t-r}^j + u_j^\top \widetilde{D}_t + \zeta_i^j + v_{i,t}^j, \tag{50}$$

where $\Delta q_{i,t-r}^j = q_{i,t-r}^j - q_{i,t-r-1}^j$, $\gamma_{j,r} = -\sum_{v=r+1}^m \Phi_{j,v}$ for $r = 1, ..., m-1$, $u_j = (u_{m+1}^j, ..., u_T^j)^\top$ is a vector of constants, $\widetilde{D}_t$ is a $(T-m) \times 1$ time dummy vector with one in the t-th position and zero elsewhere.

To estimate this short dynamic panel model, we employ the generalized method of moments (GMM) estimator in first difference form for the purpose of eliminating the individual effect ζ_i^j. We follow Arellano and Bond (1991) in the choice of instruments and initial weighting matrix, the moment condition is given by

$$\mathbb{E}\left[q_{i,s}^j \left(\Delta q_{i,t}^j - \alpha_j \Delta q_{i,t-1}^j - \sum_{r=1}^{m-1} \gamma_{j,r} \Delta^2 q_{i,t-r}^j - \delta_j^\top D_t\right)\right] = 0 \tag{51}$$

for $s = 1, ..., t-m-1$ and $t = m+2, ..., T$, where $\Delta^2 q^j_{i,t-r} = \Delta q^j_{i,t-r} - \Delta q^j_{i,t-r-1}$, $\delta_j = (\Delta u^j_{m+2}, ..., \Delta u^j_T)^\top$ is a vector of constants, D_t is a $(T-m-1)\times 1$ time dummy vector with one in the t-th position and zero elsewhere. The total number of parameters to be estimated is $T-1$ with the number of moment conditions given by $(T-m)(T-m-1)/2$.[14] This GMM estimator for α_j is consistent under large N fixed T asymptotics.

5.3 Persistence

In this sub-section, we evaluate the Kehoe-Midrigan model and its extension in explaining the observed persistence of the real exchange rate for each good. Following the theoretical analysis, our empirical persistence measure is the good-specific SAR, α_j.

We first revisit the original Kehoe-Midrigan model with an assumption of an i.i.d. money growth ($\rho = 0$). In this case, the theory predicts an AR(1) model and α_j is the AR(1) coefficient. GMM estimation of α_j yields a median of 0.56 with the median standard error equal to 0.03 using annual retail price data for U.S.-Canada city pairs.[15] In terms of monthly frequency, our value corresponds to $0.56^{1/12} = 0.95$, slightly less than 0.98, the median value obtained by Kehoe and Midrigan (2007) based on bilateral real exchange rates of 66 goods between the U.S. and four European countries.

The first panel in Figure 5 plots the estimated persistence measure α_j against the annualized price stickiness parameter $\lambda_j^{12} = (1-f_j)^{12}$ computed from monthly frequency of price adjustment f_j as estimated by Nakamura and Steinsson (2008). A cross-sectional regression of α_j on λ_j^{12} yielded a significantly positive slope coefficient of 0.25 (with a standard error of 0.06) which has the sign consistent with the theoretical prediction: more price stickiness implies higher persistence. However, 84 percent of the goods have persistence levels that lie above the 45 degree line ($\alpha_j = \lambda_j^{12}$) in the scatter plot. If the model performance is evaluated by computing the ratio of the predicted persistence (on the 45 degree line) to the observed persistence for each good, the median good obtains a ratio of 48 percent. This confirms Kehoe and Midrigan's claim that a simple model of price stickiness alone fails to reproduce the observed persistence in good-level real exchange rates.

We next consider the effect of introducing positively serially correlated money growth. On the whole, α_j remains almost unchanged with a median value of 0.57 based on the AR(2) model with

[14] For the model to be (over-) identified, at least $T = 4$ is required for $m = 1$, $T = 6$ is required for $m = 2$, and $T = 9$ is required for $m = 4$. Since $T = 16$ is available in our sample, the number of over-identifying restrictions is 51, 76, and 90, respectively.

[15] This value lies between the medians for OECD city pairs (0.65) and non-OECD city pairs (0.51) obtained by Crucini and Shintani (2008) based on the same data source.

the median standard error rising somewhat to 0.035. The regression slope shown in the second panel of Figure 5 is 0.25 and is again significantly positive. Recall that for a given λ_j, α_j is a monotonically increasing function of ρ (see the left panel of Figure 1). In annual frequency, the predicted SAR from Table 1 is given by

$$\alpha_j = 1 - (1 - \rho^{12})(1 - \lambda_j^{12}). \tag{52}$$

The effect of increasing ρ can be seen in the median value of the ratio of predicted and estimated values provided in the first row of Table 2. In terms of the median, the theoretical persistence equals the estimated persistence when ρ is between 0.90 and 0.95. This value is higher than $\rho = 0.83$, the value estimated by CKM for money growth; it also exceeds persistence estimated using nominal GDP growth or changes in the nominal exchange rate. Using nominal U.S. GDP growth the persistence level is $\rho = 0.75$.[16] Using monthly changes in the U.S.-Canadian nominal exchange rate from 1970:1 to 2008:12, the first-order autocorrelation is 0.25. Taking the model literally, these macroeconomic variables will possess the same persistence level as the money growth rate. This feature is due to the combined effect of particular assumptions we make about functional forms and the stochastic singularity of assuming a single nominal shock drives the model.

Indeed, when $\rho = 0.83$, about 66 percent of the persistence can be explained by the model. Even at this high value, the inability of persistent money growth to account for the persistence of real exchange rates can be seen from the scatter plot. Notice from equation (52), increasing ρ pivots the theoretical line upward and leftward from the 45 degree line such that it becomes flatter with a higher intercept. At $\rho = 0.83$, expressed on the annual frequency basis, we draw the theoretical prediction in the second panel of Figure 5. Relative to zero persistence, the intercept of this line rises from 0 to $\rho^{12} = 0.83^{12} = 0.11$, while the slope flattens from 1 to $1 - \rho^{12} = 0.89$. Yet, about 76 percent of data points remain above the theoretical line. Thus, persistence in money growth helps a bit, but the model with Calvo pricing remains largely unsuccessful in explaining the persistence with a reasonable choice of the money growth process.

Turning to the role of information delay in explaining α_j, the persistence estimates based on the AR(4) model become somewhat lower with a median value of 0.51 (with a median standard error of 0.035), but still are much higher than the level predicted by the standard Calvo pricing without information delay (which corresponds to the $\omega = 0$ line shown in the lower panel of Figure

[16] We estimated the autoregressive coefficient for quarterly U.S. nominal GDP growth from 1947:Q3 to 2008:Q2 and obtained a coefficient of 0.55. We convert this to a monthly persistence level using the formula shown in the footnote 8.

5). Recall that from the left panel of Figure 3, for a fixed value of λ_j and $\rho = 0.83$, α_j is strictly increasing in ω. This pattern is preserved in the SAR expressed in annual frequency:

$$\alpha_j = 1 - (1 - \rho^{12})(1 - \lambda_j^{12})(1 - \omega^{12})(1 - (\omega\rho)^{12}). \tag{53}$$

Notice the symmetry of ρ and ω in the theoretical SAR expression. In effect, the information delay takes some of the burden off the persistence of the money growth rate in accounting for persistence in LOP deviations. In the lower panel of Figure 5, we present a very high information delay case, ($\omega = 0.98$ which corresponds to the 50 month average duration of information updates) and the no delay case ($\omega = 0$), which reduces the model back to the Calvo pricing framework, both with persistence of money growth at $\rho = 0.83$. The line for the very high information delay case has an intercept of 0.82 and a slope of $1 - 0.82 = 0.18$. These are computed from equation (53).

Taking these two polar cases of information delay as upper and lower limits defines the shaded triangular region in Figure 5. The regression line through the scatter of empirical estimates falls within this triangle over a substantial fraction of its range. The fitted line has an intercept of 0.44 and a slope is 0.42. These changes relative to the other panels are brought about by estimating the AR(4) model instead of the AR(1) or AR(2) model in the other panels; the frequencies of price changes (the x-coordinates) are the same across the panels.

Turning to the ability of the model to match the median level of persistence, Table 2 reports the ratio of predicted persistence to estimated persistence of the median good for various parameterizations of the theoretical and empirical models.

We see in this table the symmetry of ω and ρ clearly. For example, the model predicts the same ratio for parameterization $(\omega, \rho) = (0.90, 0.95)$ as it does for $(\omega, \rho) = (0.95, 0.90)$. Focusing on the persistence consistent with the CKM parameterization, namely $\rho = 0.83$, an information delay of $\omega = 0.90$ which corresponds to 10 months of average duration between information updates reconciles the theory with the evidence on median LOP persistence. Less persistent shocks require modest increases in information delay to match median LOP persistence. For example, setting shock persistence between 0.25 and 0.75 requires an increase in information delay to about 12 months to match the micro-data.

The extent of information delay needed to reconcile LOP persistence and observed price adjustment frequencies are broadly consistent with estimates found in studies using aggregate data. Using the aggregate data on inflation over 1960:Q1 - 2007:Q2, DKT find that information delay, on average, is 7.1 months with 95 percent confidence intervals between 5.0 and 16.1 months. Knotek

(2006) introduces information stickiness into the fixed menu cost model and finds the average duration between information updates to be 20.4 months over 1983:Q1 - 2005:Q4. Therefore, at least in terms of the median, dual stickiness pricing with a reasonable money growth process is capable of replicating the observed persistence.

5.4 Volatility

The second puzzle brought up in Kehoe and Midrigan (2007) is the observation of too much volatility in good-level real exchange rates which is inexplicable by either a simple sticky price model or a model with pricing complementarities. In this sub-section, we evaluate the role of information stickiness in terms of explaining the observed volatility.

We recognize, as other authors have, that the simple monetary model of nominal exchange rates fails to produce the level of nominal exchange rate variability observed in the data when calibrated to match the empirics of the monetary aggregates. Accordingly, we focus on the normalized standard deviation of LOP deviations, as we did in the quantitative theoretical section earlier.

The performance of the model is evaluated by the ratio of 'theoretical' normalized standard deviation to 'observed' normalized standard deviation. The procedure of computing each standard deviation is as follows. First, to compute 'theoretical' normalized standard deviation, note that the standard deviation of real exchange rates predicted by the theory has the same implication for annually sampled data and monthly sampled data. Therefore, unlike the measure of persistence that required transformations shown in Table 1, using annual data poses no complication.

Second, to compute the 'observed' normalized standard deviation, note that using a pooled sample variance as a volatility measure is not appropriate since it includes the variance component due to the dispersion of the long-run real exchange rate q_i^j among city pairs in our panel data. In addition, the theory predicts volatility caused by the nominal exchange rate fluctuation which is common to all the products, but is not designed to incorporate the idiosyncratic variance component such as the one due to time-varying city specific shocks.

For this reason, we regress the level of the LOP deviation on a set of time dummies and city-dummies, good-by-good, pooling all cross-border pairs in a panel. We then use the variance attributed to the time dummies as the component of variance we seek to explain with the model. This procedure is also consistent with the idea of using time dummies in the dynamic panel estimation to incorporate the common time specific shocks in our previous analysis of persistence. We thus use the estimated standard deviation of the time specific component normalized by the sample standard deviation of monthly nominal exchange rate growth as the empirical counterpart to the theoretical normalized standard deviation.

The first row of Table 3 shows the median of the ratios of the theoretical to observed normalized standard deviation. The original Kehoe-Midrigan setting with $\rho = 0$ can explain only 23 percent of the variation in the data. Thus, the evidence of excess volatility discovered by Kehoe and Midrigan (2007) is also confirmed in the EIU panel data of the U.S.-Canadian city pairs. Can we explain this observed volatility with an introduction of serially correlated money growth? Unfortunately, unlike the persistence, the predicted volatility is not a monotonically increasing function of ρ. Examples presented in the left panel of Figure 2 show that the volatility decreases monotonically for goods with a small $\lambda_j = 1 - f_j$ and increases only in some range of ρ for goods with a larger λ_j. As a result of the combination of the two effects for many goods, none of the median ratios presented in the first row of Table 3 are above one, though the maximum value does reach 42 percent at shock persistence in the range of 0.75 to 0.83.

In contrast to the effect of ρ, the left panel of Figure 4 shows that the volatility increases monotonically with ω in dual stickiness pricing for any values of λ_j and ρ. Table 3 also presents the ratio of standard deviations based on dual stickiness pricing with various ω. Introduction of information delay with money persistence set at $\rho = 0.83$, the volatility is now fully explained for the median good when the average duration between information updates is 12 months. When we use $\rho = 0.75$, the dual stickiness pricing can account for the volatility with 15 months of information delay. In this sense, the information delay plays an essential role in explaining volatility even though the model produces too little variability of real exchange rates at low values of the shock persistence.

We believe it is more reasonable to expect the model to match the median and cross-sectional patterns of persistence than volatility. As Crucini and Telmer (2007) show using this data, deviations from the LOP are largely driven by good-specific shocks rather than a common shock to the nominal exchange rate: only 7 percent of the variance of changes in LOP deviations is common across goods. This means that the bulk of the time series variation we see in the micro-data is due to good-specific shocks, which are abstracted from here. It is therefore unreasonable to expect a model driven by only an aggregate shock to account for all of the good-specific variation. With this caveat in mind, we delve into the role of heterogeneous information updating focusing on its role in shaping estimated persistence and not volatility.

5.5 Heterogeneous information updating

Up to this point, good-specific parameter heterogeneity has been restricted to the price adjustment parameter, λ_j while ω has been the same for all goods. This asymmetry is more a reflection of the current state of measurement, than reality. The invention of scanner bar codes has made the direct cost of price changes practically zero in recent decades, weakening the literal interpretation of menu cost models as a rationalization for Calvo pricing. The same technology has facilitated the collection of detailed consumer demand information at most large-scale retail firms and likely complicated the decision theory needed to determine optimal markups. Thus the time-dependent information updating process may capture, in a reduced-form sense, asymmetries across goods in this underlying decision problem.

All of our theoretical propositions and dynamic equations are valid when information updating is good-specific, with ω replaced by ω_j. In this sub-section, we use the theory to infer good-specific information updates by minimizing the difference between estimated persistence at the level of individual goods and what the theory implies given the frequency of price changes for each good.

We estimated good-specific information updating parameters, ω_j for $j = 1, 2, ..., 165$ by solving the problem:

$$\min_{\omega_j \in [0,1)} [\hat{\alpha}_j - \alpha(\omega_j|\lambda_j, \rho)]^2, \tag{54}$$

where $\hat{\alpha}_j$ denotes the SAR estimate of the AR(4) model and $\alpha(\omega_j|\lambda_j, \rho)$ is the theoretical SAR given by, $\alpha_j = 1 - (1 - \rho^{12})(1 - \lambda_j^{12})(1 - \omega_j^{12})(1 - (\omega_j \rho)^{12})$ evaluated at $\rho = 0.83$ and $\lambda_j = 1 - f_j$ from the frequency of regular price changes.

We first ask whether our results are, on the whole, consistent with evidence from micro studies on price reviews. No micro studies provide direct estimates of information delay across goods, but survey results on price reviews by firms may serve our purpose. Fabiani, Druant, Hernando, Kwapil, Laudau, Loupias, Martins, Matha, Sabbatini, Stahl, and Stokman (2005) argue that the frequency of price reviews rather than price changes "could be related to the arrival of information." According to Fabiani et. al. (2005), when additional information on the state of the economy infrequently arrives, it is sensible for firms to review prices infrequently. In this sense, we can exploit survey results for price reviews.

We compare the cross-sectional distribution of information delay parameters to the Blinder, Canetti, Lebow, and Rudd (1998) survey of U.S. firms about price setting behavior in the early 1990s. They ask firms what the customary interval (e.g., daily, weekly, monthly, quarterly, and yearly) was between price reviews for their most important product. Table 4 compares our distribution of durations of information updates with these survey results. Overall, our distribution of durations between information updates seems to match the distribution of price reviews.

An important question to ask is what characteristics of the goods or the markets in which they are traded help us to further understand the patterns of information delay in the cross-section. Recent studies of long-run retail price dispersion and persistence have emphasized the fact that retail prices are jointly determined by the prices of traded and non-traded inputs.[17]

The role of non-traded inputs at the retail level is typically measured by the distribution share: the difference between what consumers pay in retail markets and what manufactures receive, divided by what consumers pay. By construction it includes transportation costs, wholesale costs, retail costs and markups in the movement of goods from the factory door to the retail floor. Since these shares are more aggregated than our price data, we take averages of ω_j across j falling into each distribution sector, s, for which we have a distribution share parameter. A regression of the infrequency of information updating on the distribution share yields the following:

$$\omega_s = \underset{(0.0186)}{0.878} + \underset{(0.0339)}{0.0711}(SHARE_s) + e_s, \tag{55}$$

where $SHARE_s$ is the distribution share for sector s and e_s is the regression residual. The 165 goods in our micro-panel fall into 23 unique sectors, with distribution shares ranging from 0.17 to 0.94. An example of an EIU item from the first category is a low-priced automobile and an example from the latter category is an annual premium for automobile insurance. The implied effect of the distribution share is significant, elevating the information delay from 10 months to 17 months as we move from the sector with the smallest distribution share to the sector with the largest share.

The coefficient on the distribution share suggests firms selling retail goods produced with larger shares of non-traded inputs update their information sets less often. One interpretation of this correlation is the view that information flows are more frequent for traded inputs where globally centralized spot markets exist than for the inputs used in the distribution sector, which often entail confidential wage and rental information.

6 Conclusion

We have confirmed Kehoe and Midrigan's main finding that the standard Calvo-type sticky price model fails to explain the persistence and volatility of good-level real exchange rates using highly

[17]See, for example, Crucini, Telmer, and Zachariadis (2005) and Crucini and Shintani (2008).

disaggregated price data from U.S. and Canadian cities. The robustness of their finding suggests that the baseline model is deficient.

We offer a possible solution to this puzzle by extending the Kehoe-Midrigan model such that only a fraction of firms have up-to-date information when resetting prices.[18] Due to the infrequent arrival of information, real exchange rates become more persistent and track the volatile nominal exchange rate even if price adjustment is fast. Our model can explain estimated LOP persistence and a sizable fraction of volatility with a plausible common duration of information updating.

That the model matches persistence more easily than variability may seem puzzling since this is exactly the opposite of the conventional wisdom. Namely, that monetary models have an easier time matching variability than persistence while the reverse is true of real models. The missing piece of the puzzle is the mapping from LOP deviations to PPP deviations. An emerging literature beginning with Imbs, Mumtaz, Ravn, and Rey (2005) shows that aggregation bias leads PPP persistence to exceed persistence of the median good. A separate literature, including contributions by Crucini and Telmer (2007) and Nakamura (2008), shows substantially more volatility in LOP deviations than PPP deviations, shocks to individual goods mostly average out in the aggregation to PPP. In other words, aggregation increases persistence and reduces variability. This poses a challenge to economic models based on microeconomic foundations since a researcher who matches the persistence and variability of the aggregate may miss the median value of either or both in the micro-data. The best approach may be to use the microeconomic evidence along with tractable heterogeneity at the microeconomic level and aggregate the data in a manner consistent with the practices of national statistical agencies. Such an approach would provide more reliable empirical insights across levels of data aggregation and lead to more focus on both the aggregate and distribution effects of various shocks.

Our results are subject to the caveat that we require a departure from the existing rational expectations approach which presumes continuous information updating and the ability of agents to process the data instantaneously in the sense of knowing the complete model. Our sense is that a model in which some information is released at discrete intervals (e.g., government provision of aggregate data) and managers or analysts differ in how efficiently they process information is what this time-dependent model attempts to emulate. Just as the plausibility of the Calvo framework

[18] Alternative approaches emphasize trade costs, distribution margins and habit formation as exemplified in the work of Eaton and Kortum (2002), Burstein, Eichenbaum, and Rebelo (2005), Crucini and Yilmazkuday (2008) and Johri and Lahiri (2008).

initially required some notion of how sticky prices were, our analysis begs the question regarding information updating and processing. It seems plausible that these tasks are at least as costly and time consuming as altering a menu of prices. We hope to provide more concrete evidence on this dimension in future work.[19] Much remains to be done.

References

Ahlin, C., and M. Shintani (2007): "Menu Costs and Markov Inflation: A Theoretical Revision with New Evidence," *Journal of Monetary Economics*, 54(3), 753–784.

Amirault, D., C. Kwan, and G. Wilkinson (2006): "Survey of Price-Setting Behaviour of Canadian Companies," Bank of Canada Working Paper 2006-35.

Andrews, D. W. K., and H.-Y. Chen (1994): "Approximately Median-Unbiased Estimation of Autoregressive Models," *Journal of Business and Economic Statistics*, 12(2), 187–204.

Arellano, M., and S. Bond (1991): "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *Review of Economics Studies*, 58(2), 277–297.

Bils, M., and P. J. Klenow (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112(5), 947–985.

Blinder, A. S., E. R. D. Canetti, D. E. Lebow, and J. B. Rudd (1998): *Asking about Prices: A New Approach to Understanding Price Stickiness*. Russell Sage Foundation, New York.

Burstein, A., M. Eichenbaum, and S. Rebelo (2005): "Large Devaluations and the Real Exchange Rate," *Journal of Political Economy*, 113(4), 742–784.

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

Carvalho, C., and F. Nechio (2008): "Aggregation and the PPP Puzzle in a Sticky-Price Model," Federal Reserve Bank of New York Staff Report No. 351.

[19] Examples of models that consider modifications of decision theory along these lines, include Sims (2003), Woodford (2008) and Gorodnichenko (2008).

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2002): "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?," *Review of Economics Studies*, 69(3), 533–563.

Clark, T. E. (2006): "Disaggregate Evidence on the Persistence of Consumer Price Inflation," *Journal of Applied Econometrics*, 21(5), 563–587.

Crucini, M. J., and M. Shintani (2008): "Persistence in Law of One Price Deviations: Evidence from Micro-data," *Journal of Monetary Economics*, 55(3), 629–644.

Crucini, M. J., M. Shintani, and T. Tsuruga (2008): "Accounting for Persistence and Volatility of Good-level Real Exchange Rates: the Role of Sticky Information," NBER Working Paper # 14381.

Crucini, M. J., and C. Telmer (2007): "Microeconomic Sources of Real Exchange Rate Variability," unpublished manuscript.

Crucini, M. J., C. I. Telmer, and M. Zachariadis (2005): "Understanding European Real Exchange Rates," *American Economic Review*, 95(3), 724–738.

Crucini, M. J., and H. Yilmazkuday (2008): "A Model of International Cities: Implications for Real Exchange Rates," unpublished manuscript.

Dupor, B., T. Kitamura, and T. Tsuruga (2008): "Integrating Sticky Prices and Sticky Information," *Review of Economics and Statistics* forthcoming.

Eaton, J., and S. Kortum (2002): "Technology, Geography and Trade," *Econometrica*, 70(5), 1741–1779.

Engel, C., and J. H. Rogers (2004): "European Market Integration after Euro," *Economic Policy*, 19, 347–384.

Fabiani, S., M. Druant, I. Hernando, C. Kwapil, B. Laudau, C. Loupias, F. Martins, T. Y. Matha, R. Sabbatini, H. Stahl, and A. C. J. Stokman (2005): "The Pricing Behaviour of Firms in the Euro Area New Survey Evidence," ECB working paper No 535.

Gorodnichenko, Y. (2008): "Endogenous Information, Menu Costs, and Inflation Persistence," NBER working paper # 14184.

IMBS, J., H. MUMTAZ, M. RAVN, AND H. REY (2005): "PPP Strikes Back: Aggregation and the Real Exchange Rate," *Quarterly Journal of Economics*, 120(1), 1–43.

JOHRI, A., AND A. LAHIRI (2008): "Persistent Real Exchange Rates," *Journal of International Economics*, 76(2), 223–236.

KEHOE, P. J., AND V. MIDRIGAN (2007): "Sticky Prices and Sectoral Real Exchange Rates," Federal Reserve Bank of Minneapolis Working Paper No 656.

KNOTEK, E. S. (2006): "A Tale of Two Rigidities: Sticky Prices in a Sticky Information Environment," Federal Reserve Bank of Kansas City Research Working Paper 06-15.

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), 1295–1328.

MEESE, R. A., AND K. ROGOFF (1983): "Empirican Exchange Rate Models of the 1970's: Do They Fit out of Sample," *Journal of International Economics*, 14(1-2), 3–27.

NAKAMURA, E. (2008): "Pass-Through in Retail and Wholesale," *American Economic Review*, 98(2), 430–437.

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

PARSLEY, D. C., AND S.-J. WEI (2007): "A Prism into the PPP Puzzles: The Micro-foundations of Big Mac Real Exchange Rates," *Economic Journal*, 117, 1336–1356.

ROGERS, J. H. (2007): "Monetary Union, Price Level Convergence, and Inflation: How close is Europe to the USA?," *Journal of Monetary Economics*, 54(3), 785–796.

SIMS, C. A. (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50(3), 665–690.

WOODFORD, M. (2008): "Information-Constrained State-Dependent Pricing," NBER Working paper # 14620.

YUN, T. (1996): "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles," *Journal of Monetary Economics*, 37(2), 345–370.

Figure 1: Persistence without information delay: function of money growth parameter(ρ) and Calvo parameter (λ_j)

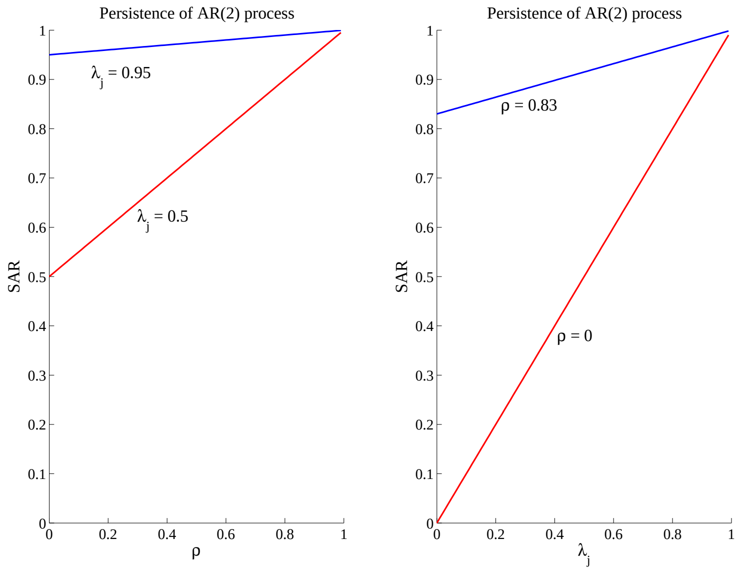

Figure 2: Volatility without information delay: function of money growth parameter(ρ) and Calvo parameter(λ_j)

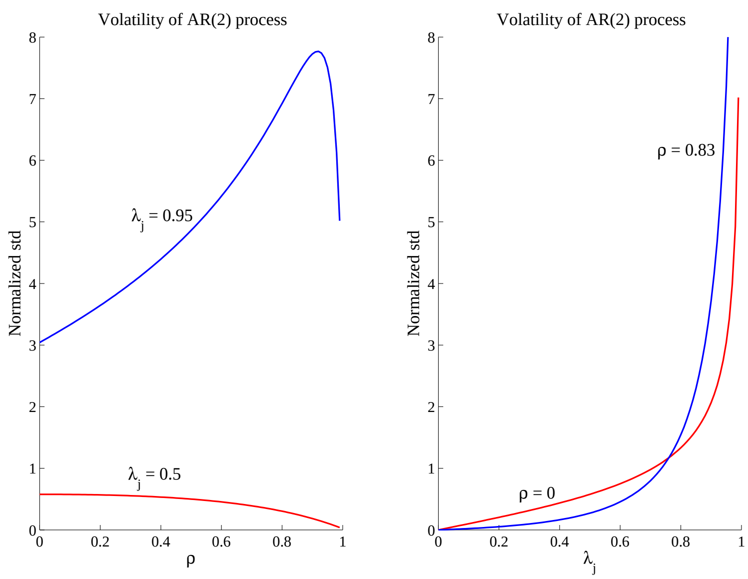

NOTES: The discount factor β is 0.99.

Figure 3: Persistence with information delay: function of information stickiness parameter(ω) and Calvo parameter(λ_j)

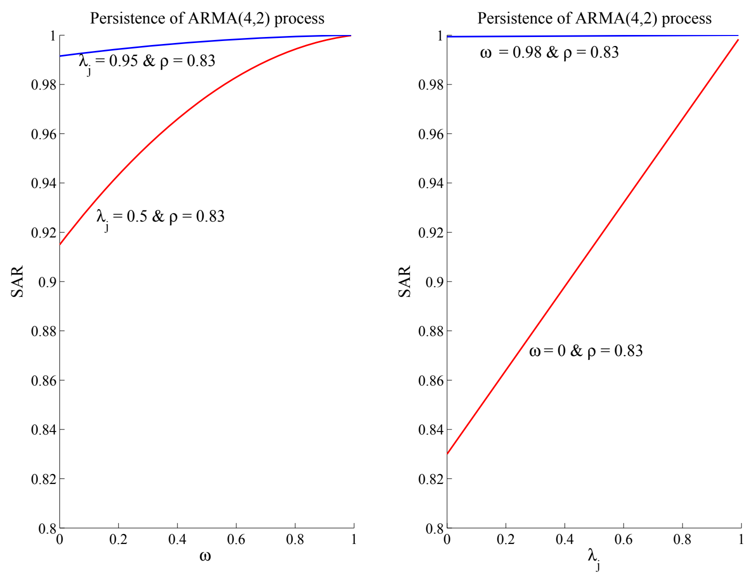

Figure 4: Volatility with information delay: function of information stickiness parameter(ω) and Calvo parameter(λ_j)

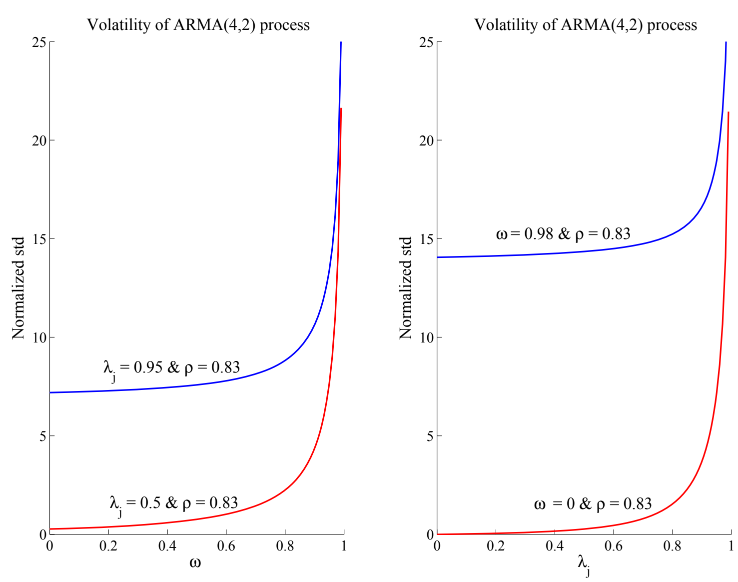

NOTES: The discount factor β is 0.99.

Figure 5: Real exchange rate persistence and price stickiness

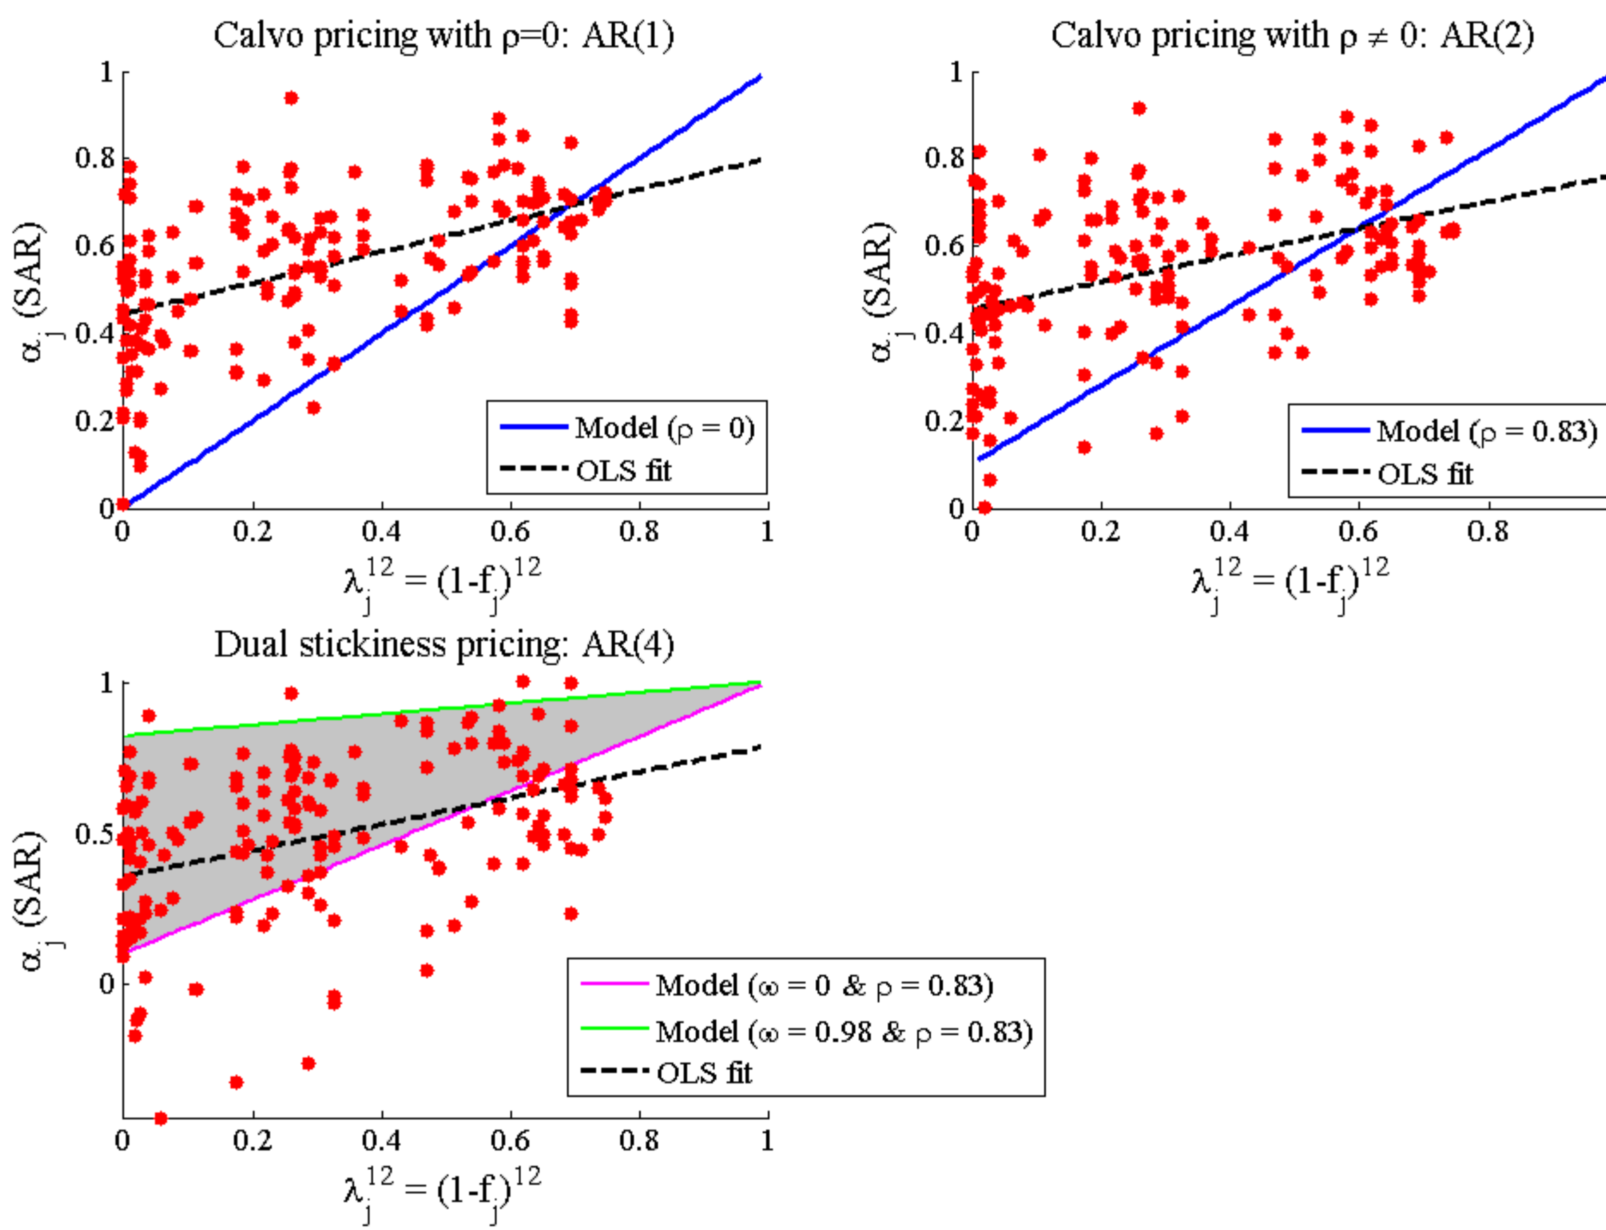

Table 1: Summary of transformations from monthly to annual specification

	Monthly specification	Annual specification
Calvo (ρ=0)	$\hat{q}_t(j,l,l^*) = \lambda_j \hat{q}_{t-1}(j,l,l^*) + \lambda_j \eta_t$	$\hat{q}_t(j,l,l^*) = \lambda_j^{12} \hat{q}_{t-12}(j,l,l^*) + \lambda_j \Lambda_j(L)\eta_t$
Calvo ($\rho > 0$)	$\hat{q}_t(j,l,l^*) = (\lambda_j + \rho)\hat{q}_{t-1}(j,l,l^*)$	$\hat{q}_t(j,l,l^*) = (\lambda_j^{12} + \rho^{12})\hat{q}_{t-12}(j,l,l^*)$
	$-\lambda_j \rho \hat{q}_{t-2}(j,l,l^*) + \theta_j \eta_t$	$-\lambda_j^{12}\rho^{12}\hat{q}_{t-24}(j,l,l^*) + \lambda_j \Lambda_j(L) R(L) \eta_t$
Dual stickiness	$\hat{q}_t(j,l,l^*) = \sum_{r=1}^{4} \phi_{j,r} \hat{q}_{t-r}(j,l,l^*)$	$\hat{q}_t(j,l,l^*) = \sum_{r=1}^{4} \Phi_{j,r} \hat{q}_{t-12r}(j,l,l^*)$
	$+\sum_{r=0}^{2} \theta_{j,r} \eta_{t-r}$	$+\Theta_j(L)\eta_t$

NOTES: The left panel shows the original monthly ARMA processes which are in the main text. The right panel shows corresponding conversions such that autoregressive coefficients are non-zero only if the lags are multiples of 12 and that moving average terms are finite. These conversions allow us to estimate the original monthly ARMA process with annually sampled data. The autoregressive parameters $\Phi_{j,r}$ and moving average polynomials, $\Lambda_j(L)$, $R(L)$ and $\Theta_j(L)$ are given in the appendix of Crucini, Shintani, and Tsuruga (2008).

Table 2: Proportions of explained persistence of good-level real exchange rates

	ρ							
Information delay	0	0.25	0.50	0.75	0.83	0.90	0.95	0.98
No delay ($\omega = 0$)	0.48	0.51	0.51	0.57	0.66	0.92	1.23	1.52
10 months ($\omega = 0.90$)	0.92	0.91	0.91	0.93	1.02	1.20	1.48	1.68
12 months ($\omega = 0.92$)	0.98	1.00	1.00	1.02	1.11	1.29	1.52	1.70
15 months ($\omega = 0.93$)	1.09	1.10	1.10	1.14	1.20	1.35	1.59	1.72
17 months ($\omega = 0.94$)	1.13	1.16	1.16	1.18	1.26	1.42	1.62	1.73
20 months ($\omega = 0.95$)	1.23	1.21	1.21	1.26	1.33	1.48	1.66	1.74

NOTES: Numbers are median ratios of the theoretical persistence, predicted by Nakamura and Steinsson (2008), to observed persistence measured by the SAR estimated from real exchange rate data. Theoretical persistence for the first row is the SAR for various ρ when Calvo pricing is used. Theoretical persistence from the second to the bottom row is the SAR for various ρ with information delay from 10 to 20 months when dual stickiness pricing is used. Median SAR estimates for AR(1), AR(2) and AR(4) models are 0.563, 0.568, and 0.508, respectively.

Table 3: Proportions of explained volatility of good-level real exchange rates

Information delay	ρ = 0	0.25	0.50	0.75	0.83	0.90	0.95	0.98
No delay ($\omega = 0$)	0.23	0.29	0.35	0.42	0.42	0.40	0.31	0.21
10 months ($\omega = 0.90$)	0.41	0.51	0.64	0.84	0.88	0.91	0.88	0.65
12 months ($\omega = 0.92$)	0.43	0.54	0.69	0.90	1.01	1.03	0.98	0.78
15 months ($\omega = 0.93$)	0.46	0.58	0.75	1.02	1.12	1.24	1.19	0.95
17 months ($\omega = 0.94$)	0.48	0.61	0.78	1.07	1.21	1.33	1.33	1.09
20 months ($\omega = 0.95$)	0.51	0.64	0.83	1.15	1.31	1.46	1.51	1.29

NOTES: Numbers are median ratios of the theoretical volatility, predicted by Nakamura and Steinsson (2008), to observed volatility measured by normalized standard deviation of real exchange rate data. Theoretical volatility for the first row is the normalized standard deviation for various ρ when Calvo pricing is used. Theoretical volatility from the second to bottom row is the normalized standard deviation for various ρ with information delay from 10 to 20 months when dual stickiness pricing is used. The normalized sample standard deviation of real exchange rate is the extracted standard deviation component due to time specific shocks in the one-way error component model.

Table 4: Intervals between information update

	one month or less	1.01-5.99 months	6-11.99 months	12 months or above
Blinder et. al.'s survey	25.6	13.2	16.5	44.6
Our estimates $1/(1-\omega_j)$	33.3	12.7	18.2	35.8

NOTES: The numbers in the first row represent the distribution, in percentages, of the frequency of price reviews reported in Blinder, Canetti, Lebow, and Rudd (1998, Table 4.7 in p. 90). The second row shows the distribution of information delay implied by the observed persistence of real exchange rates based on Nakamura and Steinsson's (2008) data on the frequency of regular price changes is used.