American Economic Review 2012, 102(6): 2798–2825
http://dx.doi.org/10.1257/aer.102.6.2798

Reset Price Inflation and the Impact of Monetary Policy Shocks†

By Mark Bils, Peter J. Klenow, and Benjamin A. Malin*

Many business cycle models use a flat short-run Phillips curve, due to time-dependent pricing and strategic complementarities, to explain fluctuations in real output. But, in doing so, these models predict unrealistically high persistence and stability of US inflation in recent decades. We calculate "reset price inflation"—based on new prices chosen by the subsample of price changers—to dissect this discrepancy. We find that the models generate too much persistence and stability both in reset price inflation and in the way reset price inflation is converted into actual inflation. Our findings present a challenge to existing explanations for business cycles. (*JEL* E31, E52)

Many studies estimate a flat short-run Phillips curve and give it a central role in business cycles.[1] Related, empirical investigations have found it takes several years for permanent aggregate shocks to fully affect prices.[2]

Two key ingredients are used to generate a flat short-run Phillips curve in New Keynesian models: sticky prices and strategic complementarities. Complementarities slow the response of "reset prices"—the new prices chosen by the subsample of price changers—to aggregate shocks.[3] Price stickiness (particularly if time

* Bils: Department of Economics, Harkness Hall, University of Rochester, Rochester, NY 14627 (e-mail: bils@troi.cc.rochester.edu); Klenow: Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305-6072 (e-mail: klenow@stanford.edu); Malin: Federal Reserve Board of Governors, Mail Stop 97, 20th St. and Constitution Ave. NW, Washington, DC 20551 (e-mail: benjamin.a.malin@frb.gov). This research was conducted with restricted access to US Bureau of Labor Statistics (BLS) data. Rob McClelland provided us invaluable assistance and guidance in using BLS data. We thank Jose Mustre Del Rio and Krishna Rao for excellent research assistance. The views expressed here are those of the authors and do not necessarily reflect the views of the BLS or the Federal Reserve System. We are grateful to Carlos Carvalho, Jesus Fernandez-Villaverde, Mark Gertler, Virgiliu Midrigan, Ricardo Reis, Jon Steinsson, Michael Woodford, several anonymous referees, and numerous seminar participants for helpful comments.

† To view additional materials, visit the article page at http://dx.doi.org/10.1257/aer.102.6.2798.

[1] See Smets and Wouters (2003, 2007) for some estimates on European and US data.

[2] Christiano, Eichenbaum, and Evans (1999); Romer and Romer (2004); and Bernanke, Boivin, and Eliasz (2005), each based on US data, are a few of the many examples.

[3] As has been well known since Ball and Romer (1990) and Kimball (1995), strategic complementarities in the pricing decisions of individual sellers can make reset prices sluggish in response to shocks. Examples of such complementarities include sticky wages, sticky intermediate prices, and a kinked demand curve with respect to a firm's relative price. Recent papers using one or more of these complementarities include Carvalho (2006), Blanchard and Galí (2007), Gertler and Leahy (2008), Nakamura and Steinsson (2010), and Altig et al. (2011). The evidence for complementarities is mixed: Klenow and Willis (2006) and Kryvtsov and Midrigan (2010) find it difficult to reconcile firm-level complementarities with, respectively, large idiosyncratic price changes and countercyclical inventories/sales. Gopinath, Itskhoki, and Rigobon (2010) and Gopinath and Itskhoki (2010), in contrast, see strategic complementarities behind the incomplete pass-through of exchange rates to import prices.

dependent rather than state dependent) slows the response of actual prices to changes in reset prices.[4]

By intention, a flat short-run Phillips curve is a machine for dampening inflation volatility and lengthening inflation persistence. In the last two decades, however, US inflation persistence has fallen markedly. Stock and Watson (2007, 2009), Benati (2008), and Cogley, Primiceri, and Sargent (2010) document this phenomenon. As we will illustrate, even a reestimated New Keynesian model has trouble matching this low persistence. This discrepancy could reflect reset price inflation (strategic complementarities), conversion of reset price inflation into actual inflation (time dependence), or both.

To shed light on the conflict, we use micro data on prices from the US Consumer Price Index to construct an empirical measure of reset price inflation from January 1990 through October 2009. We impute to all items, those changing and those not, the reset price changes exhibited by price changers. To arrive at the reset price change for an item changing price, we compare the item's new price to its estimated reset price the previous period—not the item's last new price, set perhaps many periods earlier. We do this separately for over 60 categories of consumption (autos, dental services, etc.) and then aggregate up.

A useful analogy is to home price indices constructed from repeat sales (e.g., Shiller (1991) and Zillow.com). These indices estimate the value of residential homes even when they are not sold. Once a home is sold, the difference between the transacted price and the previous period's estimated value is used to update the estimated value of other homes that were not sold. Our reset price index is the analog for all consumer items.

In the US data, aggregate reset price inflation shows no tendency to build over time. Reset price inflation behaves roughly i.i.d. This finding is not sensitive to excluding food and energy prices, which may face big sector-specific shocks. And it applies to regular prices, not just all posted prices, so it is not driven by transitory price discounts. The pattern holds for sticky items such as services, not just more flexibly priced items. Thus, low inflation persistence reflects, in part, low persistence in reset price inflation.

We next compare the behavior of actual and reset price inflation to that for series generated from a simulated DSGE model. The model is taken from the influential work of Smets and Wouters (2007). Importantly, we remove price indexation, which is not consistent with the observed frequency of price changes observed in the micro data from the CPI. We also allow for the impact of sampling error in measured prices, which not only adds variance to the reset and actual price inflation series, but also contributes a transitory component.

The Smets and Wouters model incorporates time-dependent pricing and strong strategic complementarities. The latter take the form of sticky wages and kinked demand. As a result, the model yields reset price changes that build predictably over time. Model inflation is even more persistent and stable than is reset price inflation. The model falters at both stages vis-à-vis the data. First, empirical reset price

[4] See Dotsey, King, and Wolman (1999), Golosov and Lucas (2007), and Midrigan (2011) for the importance of time- versus state-dependent pricing to the volatility of aggregate real output.

inflation is not persistent. Second, actual inflation is even less persistent than would be predicted from time-dependent pricing *conditional* on actual reset price inflation.

Now, Smets and Wouters include large, transitory price markup shocks precisely to help reconcile a flat short-run Phillips curve with the data on inflation. The required shocks to the average desired price must be large. While such shocks help to reconcile the model with data on actual inflation, we find them at odds with our data on reset inflation. These markup shocks, because they must be so transitory, hit reset inflation for the model with much more force than they hit actual inflation. As a result, they drive variability in reset inflation *relative to that in actual inflation* far above what we see in the data.

De Walque, Smets, and Wouters (2006) suggest that the price markup shocks in the model serve as a reduced form for shocks to a flexible-price sector (e.g., food and energy). When we pursue this extension, we find that allowing large shocks in the flexible-price sector does succeed in pushing down inflation's persistence and pushing up its volatility. But it does so, as with the markup shocks, at the expense of making reset price inflation way too volatile.

To recap, the business cycle literature has coalesced around time-dependent pricing and strategic complementarities to explain business cycles. But these models rely on pricing inertia that is hard to reconcile with the data over the past 20 years. Our micro evidence reveals that this inconsistency is manifested in both the behavior of reset prices and the behavior of actual inflation conditional on reset price inflation.

The rest of the paper proceeds as follows. Section I describes the dataset and how we construct reset price inflation. Section II documents some empirical properties of reset and actual price inflation. Section III compares statistics from a DSGE model to their empirical counterparts. Section IV concludes.

I. An Empirical Measure of Reset Price Inflation

A. *The CPI Research Database*

We use the micro data underlying the nonshelter portion of the CPI to construct our measure of reset price inflation. The BLS surveys about 80,000 items a month in its *Commodities and Services Survey*. Individual prices are collected at approximately 23,000 retail outlets across 87 large urban areas.[5] The survey covers all goods and services except shelter, or about 70 percent of the CPI based on BLS consumer expenditure weights. The *CPI Research Database* (hereafter CPI-RDB) maintained by the BLS Division of Price and Index Number Research contains all prices in the *Commodities and Services Survey* since January 1988. We focus on the sample January 1990 through October 2009 ("1990–2009").[6]

The BLS collects consumer prices monthly for food and fuel items in all areas. The BLS also collects prices monthly for all items in the three largest metropolitan

[5] The BLS selects outlets and items based on household point-of-purchase surveys, which furnish data on where consumers purchase commodities and services. The price collectors have detailed checklists describing each item to be priced—its outlet and unique identifying characteristics. They price each item for up to five years, after which the item is rotated out of the sample.

[6] We use the 1988–1989 data to initialize reset prices, as discussed below.

areas (New York, Los Angeles, and Chicago). The BLS prices items in other categories and other urban areas only bimonthly. To minimize the importance of sampling error, we focus on prices from all areas. We randomly drop the prices from "odd" or "even" month prices for each item priced monthly and report all statistics based on the resulting bimonthly data.

The BLS defines approximately 300 categories of consumption as Entry Level Items (ELIs). Within these categories are prices for particular items (we call a longitudinal series of individual price quotes at the micro level a "quote-line"). The BLS provided us with unpublished ELI weights for each year from 1988–1995 and 1999–2004, based on Consumer Expenditure Surveys in each of those years. We set the 1996 and 1997 ELI weights to the 1995 weights, and the 1998 weights to their 1999 level. We set the 2005 and onward weights to their 2004 level. For each year, we normalize the ELI weights so that they sum to one across the expenditure categories—essentially all nonshelter goods. The CPI-RDB also contains weights for each price within an ELI. (Differential weights within an ELI reflect relative probabilities that an item is purchased versus sampled at an outlet.) We weight individual price quotes within an ELI in each bimonth in proportion to these weights.

The BLS labels each price as either a "sale" price or a "regular" price. Sale prices are temporarily low prices (including clearance prices). Golosov and Lucas (2007), Nakamura and Steinsson (2008), Kehoe and Midrigan (2010), and others filter out such sale prices on the grounds that they are not relevant to aggregate price movements. We use all prices, including sale prices, when constructing our inflation and reset price inflation series. To the extent sales are truly idiosyncratic, their impact on the time series for price inflation, given the large samples of price quotes in each sector, should average close to zero. To the extent sales do affect aggregate inflation, they may not be idiosyncratic.[7] That said, we will show our findings are robust to excluding sales prices when calculating inflation series.

Forced item substitutions occur when an item in the sample has been discontinued from its outlet, and the price collector identifies a similar replacement item (e.g., new model) in the outlet to price going forward. The monthly rate of forced item substitutions is about 3 percent in the sample. The vast majority of item substitutions involve price changes, and we retain these when calculating inflation statistics.[8] Our aggregate statistics are robust, however, to treating all price changes as zero at forced substitutions.

About 12 percent of the prices the BLS attempts to collect are unavailable in a given month. The BLS classifies roughly 5 percent of items as out of season in a given month. We put zero weight on out-of-season items when calculating both inflation and the frequency of price changes. The BLS classifies the other 7 percent of missing items as temporarily unavailable. As these items may be only intermittently unavailable during the month, we treat items out of stock as available at the previously collected price. We employ this treatment for calculating both the frequency

[7] See Chevalier and Kashyap (2011) for some evidence that a grocery chain uses the magnitude and frequency of sales to respond to shocks.

[8] For half of forced substitutions, the rate of price change imparted to the CPI reflects a BLS adjustment aimed at capturing quality change. We include these BLS quality adjustments in all price change statistics.

of price changes and time series of inflation rates. Nakamura and Steinsson (2008) follow a similar procedure.

Although the BLS requires its price collectors to explain large price changes in order to minimize measurement error, some price changes in the dataset appear implausibly large. We exclude price changes that exceed a factor of five in either direction (up or down). Such price jumps constitute less than one-tenth of 1 percent of all price changes.

B. *Defining Reset Price Inflation*

Let $p_{i,t}$ denote the log of the price of an individual item i at time t in the CPI-RDB. Let $I_{i,t}$ be a price-change indicator:

$$I_{i,t} = \begin{cases} 1 & \text{if } p_{i,t} \neq p_{i,t-1} \\ 0 & \text{if } p_{i,t} = p_{i,t-1} \end{cases}.$$

Each period we divide items into those that change price $(I_{i,t} = 1)$ and those that do not change price $(I_{i,t} = 0)$. For prices that change, the reset price is simply the current price. For prices that do not change, we index our estimate of the reset price to the rate of reset price inflation among price changers in the current period. Our estimate of the log reset price level for item i in month t is, therefore,

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

Starred variables denote reset values; variables without stars represent actual values. Our estimate of aggregate reset price inflation in period t is then

$$(1) \qquad \pi^*_t = \frac{\sum_i \omega_{i,t}(p_{i,t} - p^*_{i,t-1})I_{i,t}}{\sum_i \omega_{i,t} I_{i,t}},$$

where $\omega_{i,t}$ denotes an item's relative expenditure weight in t.

Although π^*_t employs only time t price changers, price changes from previous months are captured in the base values of $p^*_{i,t-1}$ which are indexed to reflect prior changes. Aggregate reset price inflation can be equivalently defined as the weighted average of micro reset price inflation rates, i.e., $\pi^*_t = \sum_i \omega_{i,t}\pi^*_{i,t}$, where $\pi^*_{it} = p^*_{i,t} - p^*_{i,t-1}$.[9] Actual inflation is $\pi_t = \sum_i \omega_{i,t}\pi_{i,t}$, where $\pi_{i,t} = p_{i,t} - p_{i,t-1}$ and, again, $p_{i,t}$ denotes the log of the actual BLS price of item i at time t.

[9] We considered an alternative measure of reset price inflation based on regressing each price change on bimonthly dummies taking the value 1 for bimonths spanning each price spell. This measure parallels the Case-Shiller Home Price Index (Shiller 1991), which allocates price increases for homes to the months between repeat sales. In our data and model economies, this regression-based measure exhibits very similar statistics to that based on (1).

Table 1—Constructing Reset Price Inflation: A Simple Example

	Period 0	Period 1	Period 2
Price of Good A	1	1.22	1.22
Inflation for Good A		20%	0%
Reset price for Good A	1	1.22	1.22
Reset inflation for Good A		20%	0%
Price of Good B	1	1	1.22
Inflation for Good B		0%	20%
Reset price for Good B	1	1.22	1.22
Reset inflation for Good B		20%	0%
Inflation (π_t)		10%	10%
Inflation for changers ($\tilde{\pi}_t$)		20%	20%
Reset inflation (π_t^*)		20%	0%

Notes: The example assumes expenditure shares of one half for each good. It also assumes that both Good A and Good B exhibited a price change in period 0, establishing the base price for calculating reset price inflation for period 1. The number 1.22 in the table represents exp(0.2) to two decimal places.

In Table 1 we present a stylized example useful for contrasting the rate of reset price inflation (π_t^*) to actual inflation (π_t) and to the average inflation of price changers (call this $\tilde{\pi}_t$). The example has two goods. Both goods change price in period 0, establishing base prices for calculating reset price inflation. Good A's price increases by 20 percent in period 1, with Good B's unchanged. This yields a rate of 20 percent for reset price inflation, same as the average rate of price increase conditional on changing price, while actual inflation is 10 percent. But note that it also kicks up the base price for calculating reset price inflation by 20 percent, not only for Good A, but also for Good B. Thus, when B's price increases by 20 percent in period 2, while A's remains unchanged, B's price just meets its updated reset price from period 1. As a result, reset price inflation for period 2 equals zero, despite the same actual inflation rate and rate of increase for price changers, respectively 10 percent and 20 percent, as in period 1.

It is worth making two more distinctions. First, new prices need not be viewed as frictionless spot prices. If future spot prices are expected to differ from the current spot price, then a newly set price may be influenced by future expected spot prices. Thus, reset price inflation can deviate from spot price inflation. (This forward-looking element is present in state-dependent as well as time-dependent sticky price models.) Second, those items changing price in a given period may be selected based on their idiosyncratic changes in desired prices. This does not happen under Calvo (1983) pricing, in which items changing price are chosen at random, but does occur in state-dependent pricing models such as Dotsey, King, and Wolman (1999) and Golosov and Lucas (2007).

Related, what extra information is contained in π_t^* that cannot be gleaned from π_t alone? With Calvo price setting, π_t^* simply reflects current and lagged inflation

$$\pi_t^* = \frac{\pi_t - (1 - \lambda)\pi_{t-1}}{\lambda}, \tag{2}$$

where λ is the frequency of price change.[10] So, under Calvo pricing, one can infer π_t^* from π_t given the price-change frequency. But, of course, Calvo pricing may poorly reflect reality. Endogenous timing of price changes and selection will break this simple mapping from π_t^* to π_t. By endogenous timing we mean any response in the fraction of items changing price to underlying shocks. By selection of changers we mean that, in contrast to Calvo, the changers may be those with more positive or negative gaps between actual and desired prices.

With data on reset inflation, we can directly test whether the joint behavior of reset and actual inflation is consistent with Calvo time-dependent pricing. For example, from (2), the standard deviation of reset inflation relative to that in actual inflation should be

$$\frac{\sigma_\pi^*}{\sigma_\pi} = \sqrt{1 + \left[2(1 - \rho)(1 - \lambda)/\lambda^2\right]}, \tag{3}$$

where ρ denotes the serial correlation in actual inflation. We will find below that the Smets and Wouters (2007) model requires large and transitory markup shocks to make inflation as volatile and transitory as it is empirically. But we see in (3) that the markup shocks, by bringing down ρ for a given frequency of price change λ, necessarily drive up volatility in reset inflation relative to actual. We can compare the Calvo-predicted ratio of standard deviations $\frac{\sigma_{\pi^*}}{\sigma_\pi}$ from (3) to its counterpart in the data.

Furthermore, even if one rejects a flat short-run Phillips based purely on behavior of actual inflation, information on reset inflation would help diagnose the nature of any rejection. For instance, low persistence of π_t could reflect time-varying frequency of price changes or lack of persistence in π_t^*. Reset price inflation is more directly revealing about strategic complementarities—some forces for low persistence (selection) or high persistence (strategic complementarities) operate on π_t^* directly, whereas their effect on π_t will be clouded by movements in the frequency.

Alternatively, we could focus on the average price change among changers ($\tilde{\pi}_t$) rather than constructing π_t^*. That is, we could simply break π_t into the product of $\tilde{\pi}_t$ and the fraction of prices changing during t. This would speak to the accuracy of the Calvo assumption of a constant frequency of price changes. But it would not shed light on the Calvo assumption that price changers are selected randomly (no selection effect). For this reason we see more power in the statistic π_t^* than in $\tilde{\pi}_t$.

II. Evidence on Inflation and Reset Price Inflation

Here we report statistics on volatility and persistence for both actual inflation and our measure of reset price inflation for the bimonths from January–February 1990 through September–October 2009. Statistics are based on an average of about 80,000 measured prices per month. The CPI-RDB begins in January 1988, but we use the 1988–1989 data to initialize reset prices. We must observe a new (changing) price for a quote-line before it can enter our calculation of reset price inflation. By

[10] This holds in the limit as the number of price setters becomes large.

1990 the fraction of items eligible for calculating reset price inflation is close to its sample mean. All the series are seasonally adjusted by taking out dummies for each bimonth over 1990–2009.

In addition to the aggregate statistics, we examine actual and reset price inflation for four subaggregates: food and energy versus core, and "flexible" items versus "sticky" items. The BLS places individual price quote-lines into about 300 categories (ELIs). The ELIs are easily categorized as either core (68 percent of expenditure weight) or food and energy (32 percent). For the flexible versus sticky categorization, we calculate the average frequency of regular price changes within each ELI, then classify quote-lines as flexible or sticky based on their ELI's frequency. We choose a threshold frequency separating the two groups of one third, not far from the overall weighted mean bimonthly frequency of 31.2 percent. This generates a 66 percent share of spending on the sticky group compared to 34 percent on the flexible group. Having more price quotes in the sticky group mitigates sampling error there, given its smaller number of price changes per period. The flexible items average 22,760 price quotes per month, compared to 57,592 for the sticky items. The mean frequency of price changes is 57 percent in the flexible group, while only 18 percent for the sticky.

While there is overlap between food and energy and the flexible group (and therefore between core items and sticky items), it is far from perfect. Restaurant menu prices, a big category of food expenditures, are in the sticky group. Core items such as vehicles and airfares exhibit frequent price changes, and fall in the flexible group. As a result, the frequency of price changes is not much lower for core items (29 percent) than it is for food and energy items (36 percent). Some of the core goods that are in the flexible group are particularly cyclical—vehicles and airfares are again two examples. Klenow and Malin (2011) show, more generally, that there is a clear positive relationship across goods between the frequency of price change for a good and cyclicality of its output.

We calculate reset price inflation using formula (1) for each of 64 BLS expenditure classes (cereal, computers, medical services, legal services, and so on), before aggregating into one of the four subgroups or overall. We do the same in calculating actual inflation at the aggregate level. Constructing reset price inflation at the EC level first means we do not infer reset prices across disparate items (e.g., basing legal services on computers) and prevents overweighting of ECs with frequent price changes in calculating aggregates.

The first row of Table 2 shows the standard deviation of bimonthly inflation is 0.52 percent (standard error 0.03 percent). Rows two and three examine inflation persistence. Inflation exhibits a serial correlation of 0.27 (standard error 0.09). This is lower than reported in some other studies, namely because of the time period 1990–2009. Longer time series—extending back to the 1970s or earlier—exhibit much more persistence. Persistence fell markedly by the time our sample begins. See Nason (2006), Stock and Watson (2007, 2009), Benati (2008), and Cogley, Primiceri, and Sargent (2010).

Persistence in inflation is low even at longer horizons. In the third row of Table 2 we report the cumulative impulse response at a one-year horizon. This is the impact of an inflation impulse on the price *level* one-year hence. We estimate it from an ARMA (6,6) process for inflation, with the number of lags chosen

Table 2—Summary Statistics for Reset and Actual Price Inflation

Statistic	All items	Food & energy	Core	Flexible items	Sticky items
Standard deviation of π	0.52%	1.38%	0.22%	1.44%	0.16%
	(0.03)	(0.09)	(0.01)	(0.09)	(0.01)
Serial correlation of π	0.27	0.22	0.33	0.25	0.64
	(0.09)	(0.09)	(0.13)	(0.11)	(0.07)
1-year cumulative π	0.90	0.76	1.29	0.83	2.90
	(0.35)	(0.28)	(0.24)	(0.30)	(0.68)
Standard deviation of π^*	0.66%	1.53%	0.50%	1.59%	0.44%
	(0.04)	(0.10)	(0.03)	(0.10)	(0.03)
Serial correlation of π^*	0.06	0.19	−0.34	0.16	−0.37
	(0.14)	(0.11)	(0.07)	(0.14)	(0.09)
1-year cumulative π^*	0.61	0.71	0.75	−0.18	0.45
	(0.28)	(0.43)	(0.12)	(0.80)	(0.20)

Notes: All data are from the CPI-RDB. Samples run from January–February 1990 through September–October 2009. The threshold frequency of regular price changes is one-third per bimonthly period: quote-lines in ELIs with average frequency higher than one-third are in the flexible group, and those with lower frequency are in the sticky group. All series are seasonally adjusted. Standard errors, with Newey-West correction, are in parentheses.

based on the Akaike criterion.[11] The impact of a 1 percent impulse in inflation on the price level is 0.90 percent (standard error 0.35 percent) after one year. Figure 1 presents the impulse response function going out 15 bimonths (30 months, or 2.5 years). The level response in prices builds modestly in the first couple of periods, to just under 1.4, but then gradually declines back to its initial impact within the 15 periods.

The fourth through sixth rows of Table 2 report statistics for *reset* price inflation. Reset inflation is more volatile than actual inflation, with a standard deviation of 0.66 percent (standard error of 0.04 percent). There is no persistence in reset price inflation as measured by its first-order autocorrelation of 0.06 (standard error 0.14). Row six reports the cumulative impulse response—the impact on the reset price level—at a one-year horizon. (Estimates are based on an ARMA(6,6) process for reset inflation). The impact on the reset price of a 1 percent innovation is less than one-for-one at the one-year horizon, equaling 0.61 (with standard error 0.28). Figure 2 displays the impulse response function up to 2.5 years. The response in reset prices is considerably greater on impact than over time. The impact effect is more than double the long-run response. So reset prices do not build.

As discussed at the end of Section I above, the ratio of the standard deviations for reset versus actual inflation provides a direct test of Calvo pricing, as $\frac{\sigma_{\pi^*}}{\sigma_\pi}$ equals $\sqrt{1 + 2(1 - \rho)(1 - \lambda)/\lambda^2}$ under Calvo, where λ is the frequency of price change and ρ is the serial correlation in actual inflation. For the set of all goods, given $\lambda = 0.31$ and $\rho = 0.27$, the Calvo-implied value for $\frac{\sigma_{\pi^*}}{\sigma_\pi}$ is 3.4. By contrast, this ratio from the data is only 1.3. So the Calvo model exaggerates $\frac{\sigma_{\pi^*}}{\sigma_\pi}$ by a factor of 2.5. Another way to see the magnitude of this discrepancy is to ask what bimonthly frequency of price setting under Calvo actually would be consistent with the observed

[11] The number of significant lags may partly reflect sampling error (see Granger and Morris 1976). We discuss sampling error in greater detail below in contrasting model and data statistics.

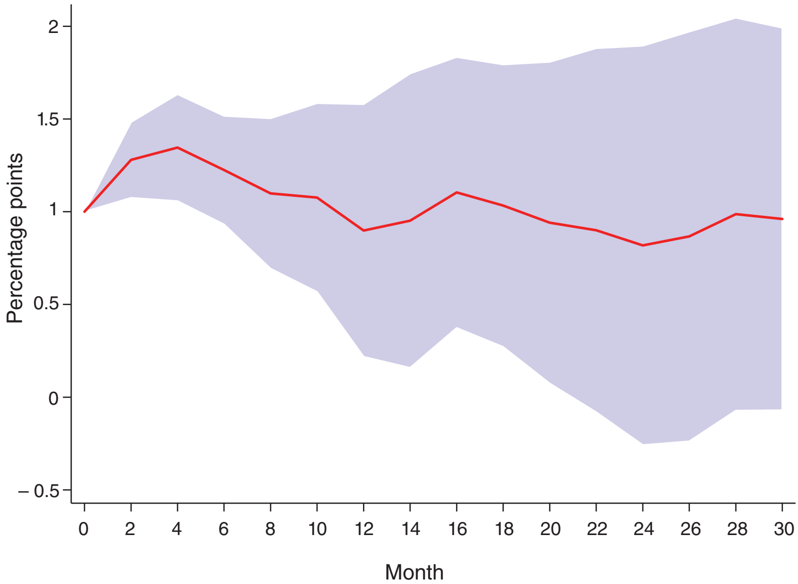

FIGURE 1. EMPIRICAL IMPULSE RESPONSE OF *ACTUAL* PRICES, ALL ITEMS

Notes: Estimates are accumulated responses to an ARMA(6,6) for actual price inflation. Shaded area denotes 95 percent confidence intervals.

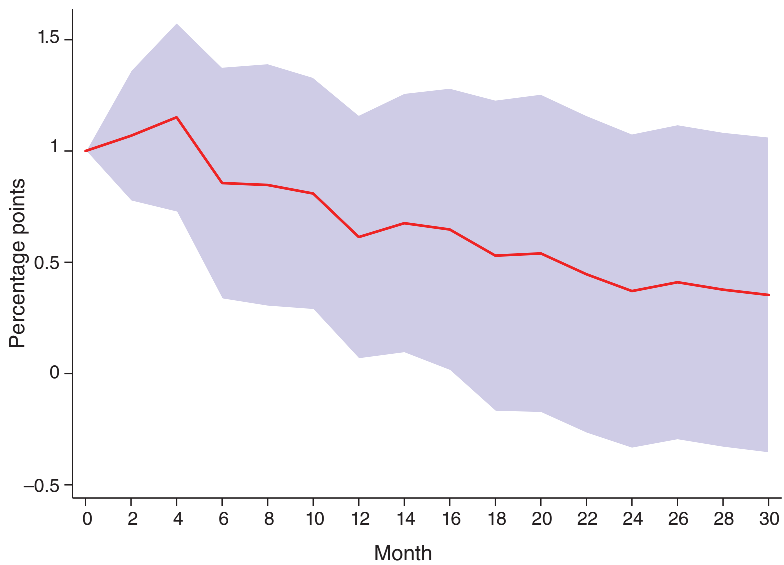

FIGURE 2. EMPIRICAL IMPULSE RESPONSE OF *RESET* PRICES, ALL ITEMS

Notes: Estimates are accumulated responses to an ARMA(6,6) for reset price inflation. Shaded area denotes 95 percent confidence intervals.

values of $\rho = 0.27$ and $\frac{\sigma_{\pi^*}}{\sigma_\pi} = 1.3$. The required value is 76 percent. That is, the joint behavior of actual and reset inflation would be consistent with the Calvo model if sellers change prices more frequently than observed (31 percent).

The remaining four columns of Table 2 repeat the statistics from the first column for the food and energy, core, flexible and sticky groups, respectively. The motivation is that large, transitory shocks may be hitting the food and energy (or flexible) sector, masking the telltale predictions of a flat short-run Phillips curve. By separately examining core (and sticky) inflation, we might isolate a clearer reflection of New Keynesian price dynamics.

As expected, actual inflation is much more volatile for food and energy items, with a standard deviation about six times higher than that for the core items. Reset price inflation is also more volatile for food and energy (standard deviation of 1.5 percent) than in the core sector (standard deviation of 0.5 percent). Core inflation is only somewhat more persistent than food and energy inflation as measured by its serial correlation (0.3 versus 0.2). After one year, the estimated impulse responses are about 70 percent larger for core items as for food and energy. But, even for core items, this impact is only 1.3 times the initial impact, so price responses build only modestly. Perhaps surprisingly, food and energy displays a higher serial correlation in its reset price inflation (around 0.2) than do core items (around -0.3). But after one year this pattern is eliminated. For both groups of items the cumulative impact at one year of an innovation to reset inflation is about 0.7, with the 95 percent confidence interval for this impact on reset price below 1.5 for food and energy goods, and below 1.0 for core.

The last two columns of Table 2 report statistics for items with frequent versus infrequent price changes (i.e., flexible items versus sticky items). Actual inflation is much more volatile for flexible items, with a standard deviation nine times that of sticky items (1.44 percent versus 0.16 percent). This reflects, in part, the important smoothing effect of many unchanging prices in the sticky sector. But even reset price inflation is more volatile in the flexible sector, by more than a factor of three (1.59 percent versus 0.44 percent).[12]

Actual inflation clearly shows more persistence for the sticky group. The serial correlation for sticky items is around 0.6, whereas for flexible items it is 0.25. This fits the prediction of many sticky price models that infrequent price changes act as a force for inflation inertia. This is reflected in the impulse responses as well. For sticky items a 1 percent impulse in actual inflation builds to nearly 3 percent after a year, whereas for flexible items the initial impact has been cut 20 percent by the end of a year. Figures 3 and 4 depict the estimated responses to a 1 percent impulse for flexible and sticky items, respectively.

For reset inflation, the serial correlation is markedly negative at -0.4 for sticky items, but not for flexible items (0.2). Based on the estimated impulse responses, reset prices do not build for either set of items. For the flexible group, the cumulative impact at one year is -0.2, while for sticky items it is 0.5. For both sets of items, the

[12] The correlation between actual inflation rates in the flexible and sticky sectors is -0.02, while the correlation between reset inflation rates is 0.12. The aggregate actual inflation rate is correlated 0.98 with inflation in the flexible sector and 0.19 with inflation in the sticky sector. The aggregate reset inflation rate is correlated 0.89 with reset inflation in the flexible sector and 0.55 with reset inflation in the sticky sector.

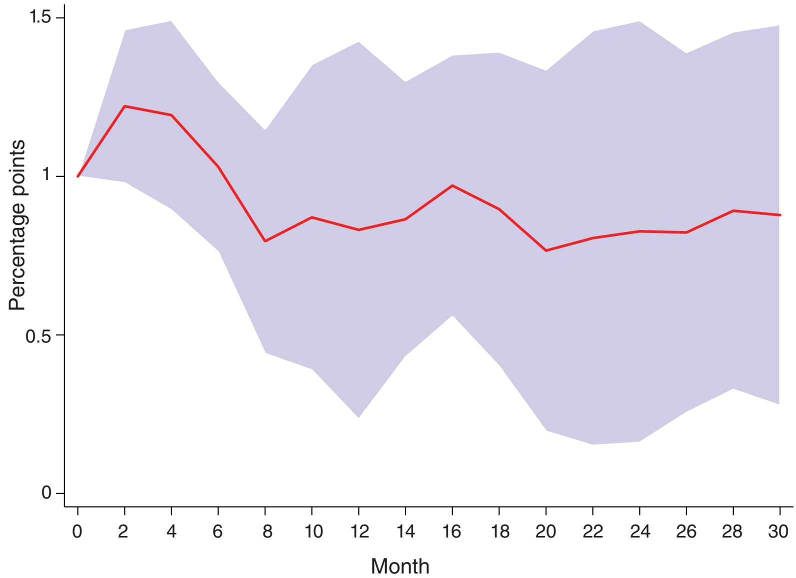

Figure 3. Empirical Impulse Response of *Actual* Prices, *Flexible* Items

Notes: Estimates are accumulated responses to an ARMA(6,6) for actual price inflation. Shaded area denotes 95 percent confidence intervals.

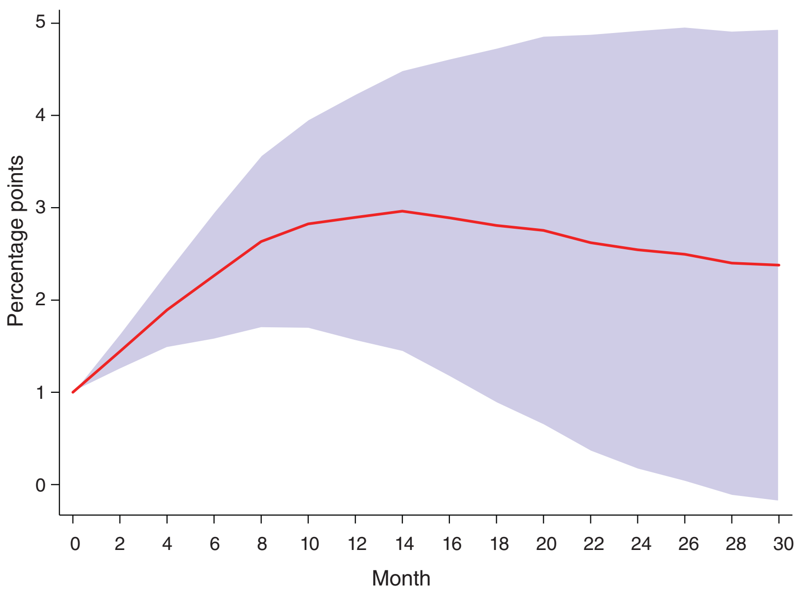

Figure 4. Empirical Impulse Response of *Actual* Prices, *Sticky* Items

Notes: Estimates are accumulated responses to an ARMA(6,6) for actual price inflation. Shaded area denotes 95 percent confidence intervals.

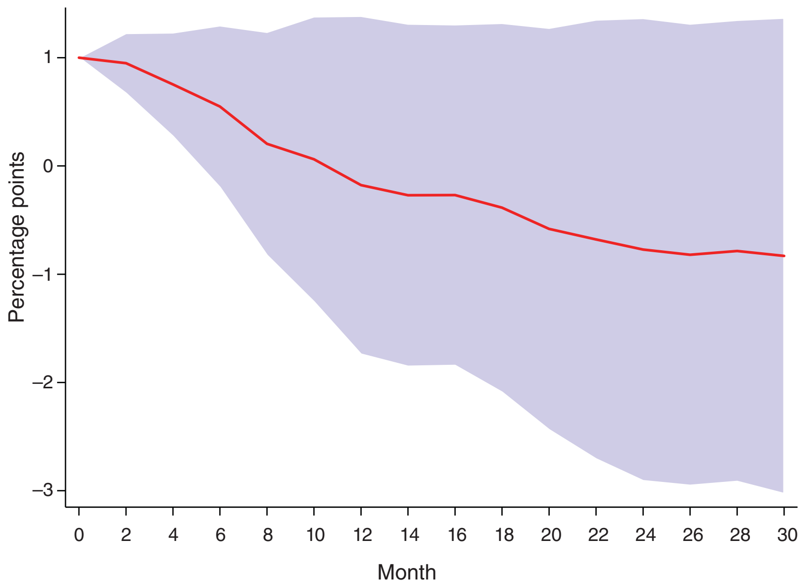

Figure 5. Empirical Impulse Response of *Reset* Prices, *Flexible* Items

Notes: Estimates are accumulated responses to an ARMA(6,6) for reset price inflation. Shaded area denotes 95 percent confidence intervals.

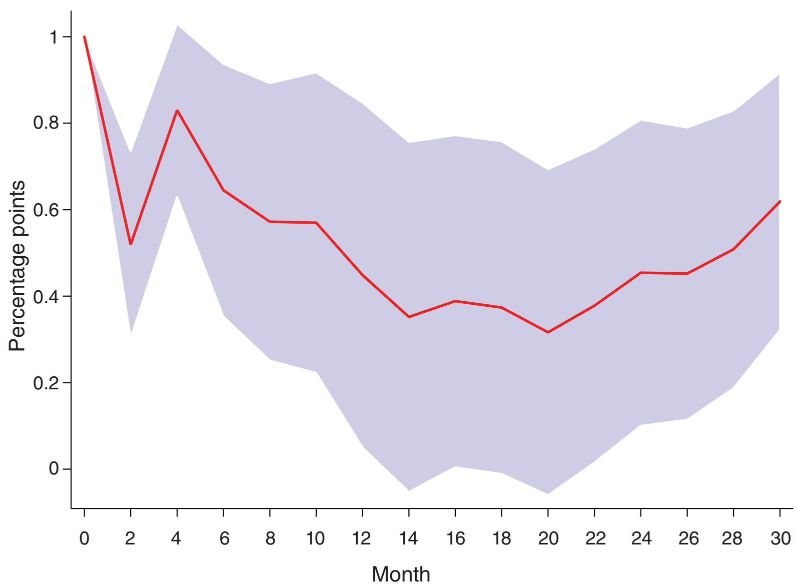

Figure 6. Empirical Impulse Response of *Reset* Prices, *Sticky* Items

Notes: Estimates are accumulated responses to an ARMA(6,6) for reset price inflation. Shaded area denotes 95 percent confidence intervals.

TABLE 3—SUMMARY STATISTICS EXCLUDING SALE PRICES

Statistic	All items	Food & energy	Core	Flexible items	Sticky items
Standard deviation of π	0.51%	1.37%	0.19%	1.43%	0.14%
	(0.03)	(0.09)	(0.01)	(0.09)	(0.01)
Serial correlation of π	0.26	0.21	0.33	0.24	0.82
	(0.09)	(0.09)	(0.13)	(0.11)	(0.06)
1-year cumulative π	0.91	0.92	1.41	0.53	4.82
	(0.36)	(0.20)	(0.18)	(0.34)	(0.81)
Standard deviation of π^*	0.63%	1.56%	0.42%	1.60%	0.38%
	(0.04)	(0.10)	(0.03)	(0.10)	(0.03)
Serial correlation of π^*	0.13	0.18	–0.23	0.15	–0.12
	(0.13)	(0.12)	(0.08)	(0.14)	(0.12)
1-year cumulative π^*	0.79	0.65	0.93	0.45	0.94
	(0.33)	(0.28)	(0.14)	(0.35)	(0.40)

Notes: All data are from the CPI-RDB. Samples run from January 1990 through October 2009. The threshold frequency of regular price changes is one-third per bimonthly period: quote-lines in ELIs with average frequency higher than one-third are in the flexible group, and those with lower frequency are in the sticky group. All series are bimonthly and are seasonally adjusted. Standard errors, with Newey-West correction, are in parentheses.

95 percent confidence interval for the impact after one year lies below 1.4. Figures 5 and 6 depict the estimated responses in reset prices out 2.5 years. For sticky items, as well as flexible, the longer-run impact remains below the initial impact.

We also constructed series solely for services, the sector with the least frequent price changes. (Services display relatively little volatility of spending, making sectoral shocks of less concern.) The results for both reset and overall inflation closely parallel those for the sticky sector reported in Table 2.

The price series (actual and reset) described in Table 2 reflect sale prices as well as regular prices. The results, however, do not hinge on this treatment. Table 3 repeats all statistics from Table 2 but treats sales prices as temporarily missing, carrying forward the most recent regular price as the price for that month à la Nakamura and Steinsson (2008). The patterns highlighted from Table 2 carry over. In particular, the serial correlation of actual inflation is almost unchanged at 0.26 (versus 0.27 in Table 2), and reset price inflation continues to show little serial correlation at 0.13 (versus 0.06 in Table 2). This means that sale prices either wash out in the aggregate or mimic the persistence in regular prices (see Klenow and Malin 2011). The impulse responses for actual and reset prices, reported above, likewise do not reflect temporarily sales. Figures 7 and 8 give, respectively, the own responses to a 1 percent impulse in actual and reset price ignoring sales price changes. The figures are very similar to their counterparts with sales (Figures 1 and 2): the actual price levels off at near one-for-one, while the response in reset price declines with time.

The most notable change in persistence from Table 2 to Table 3 is for actual inflation for sticky-price items. Its serial correlation is 0.82 after dropping sale prices in Table 3 (versus 0.64 in Table 2), and its impulse response builds by a factor of nearly five after a year. But the serial correlation and impulse responses, even for actual inflation, are largely unchanged by dropping sales in the other three groups (food and energy, core, and flexible).

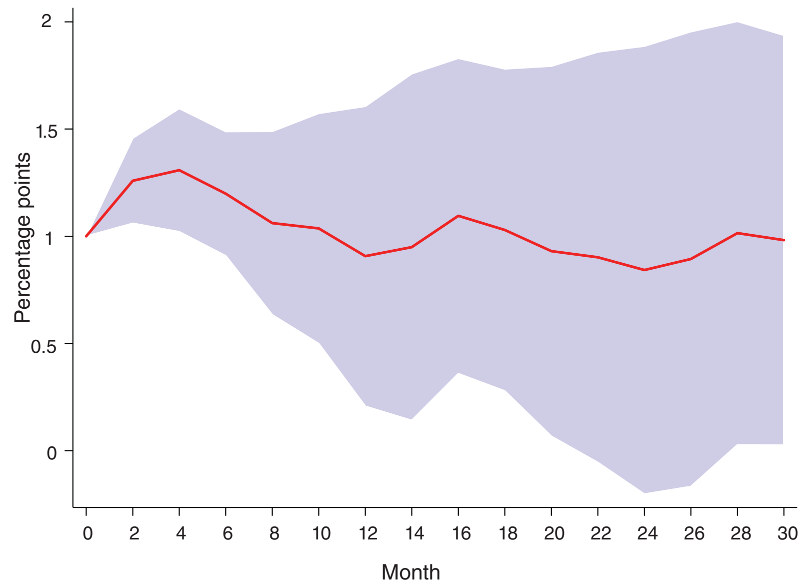

Figure 7. Empirical Impulse Response of *Actual* Prices, All Items Excluding Sale Prices

Notes: Estimates are accumulated responses to an ARMA(6,6) for actual price inflation. Shaded area denotes 95 percent confidence intervals.

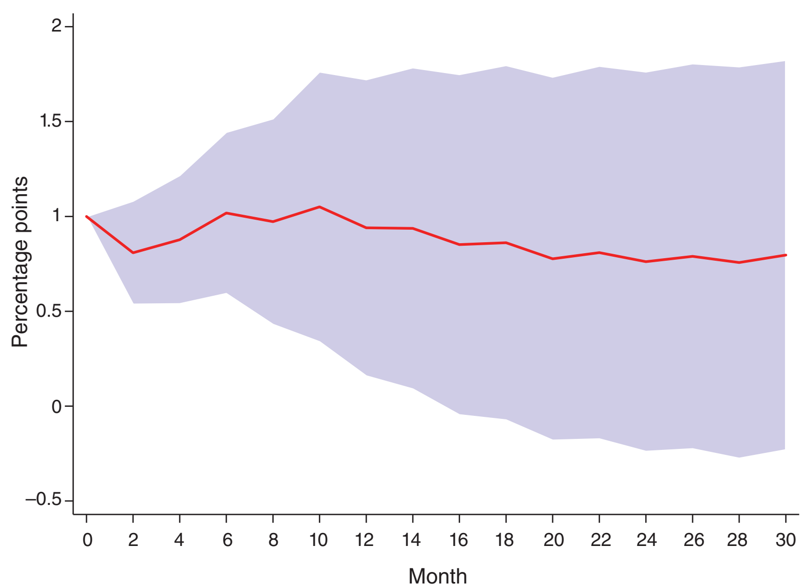

Figure 8. Empirical Impulse Response of *Reset* Prices, All Items Excluding Sale Prices

Notes: Estimates are accumulated responses to an ARMA(6,6) for reset price inflation. Shaded area denotes 95 percent confidence intervals.

Table 4—Summary Statistics for One-Sector Smets-Wouters Models

Statistic	Estimated bimonthly 1990–2009 (1)	Dropping the price markup shocks in (1) (2)	Imposing the BLS frequency (31.2%) (3)	Dropping the price markup shocks in (3) (4)	Data for all items (5)
Standard deviation of π	0.32%	0.09%	0.33%	0.12%	0.52%
	(0.02)	(0.03)	(0.02)	(0.03)	(0.03)
Serial correlation of π	0.18	0.97	0.20	0.91	0.27
	(0.08)	(0.01)	(0.09)	(0.04)	(0.09)
1-year cumulative π	1.08	10.67	1.36	5.59	0.90
	(0.31)	(1.53)	(0.32)	(0.98)	(0.35)
Standard deviation of π^*			1.15%	0.17%	0.66%
			(0.09)	(0.02)	(0.04)
Serial correlation of π^*			–0.38	0.43	0.06
			(0.07)	(0.13)	(0.14)
1-year cumulative π^*			0.45	2.19	0.61
			(0.16)	(0.49)	(0.28)
Slope of the SR Phillips Curve	0.0007	0.0006	0.0042	0.0034	

Notes: In columns 1–4, statistics are averages across 100 model simulations, each of 119 periods. Standard deviations across simulations are in parentheses. Column 5 is based on data from the CPI-RDB for January–February 1990 through September–October 2009, with robust standard errors in parentheses.

III. Sticky Price Models, Inflation, and Reset Price Inflation

Sticky price models have predictions for reset price inflation. In an earlier working paper version (Bils, Klenow, and Malin 2009), we illustrated these predictions using models with one or two aggregate shocks, time-dependent versus state-dependent pricing, and one potential source of strategic complementarities (sticky intermediate prices). Here we adapt the widely used Smets and Wouters (2007) DSGE model, which contains more shocks and more complementarities. The Smets-Wouters (SW) model has the further advantage of including an estimated Taylor Rule for monetary policy as part of Bayesian estimation of parameters to fit US time series.

We adapt the SW model in several ways. First, we reestimated their quarterly model on bimonthly data from 1990–2009 to allow comparison with the bimonthly statistics in the previous section. We construct the natural monthly analogs for most of Smets-Wouters's observable variables and use Stock and Watson's (2010) monthly data for output, investment, and price deflators. We then aggregate these monthly series to bimonthly for our estimation.[13] We also switch from the GDP deflator to the consumption deflator, as the inflation rate in the SW model corresponds to that for consumer goods.

Column 1 in Table 4 contains moments from our reestimated, bimonthly version of the SW model. As there are 119 bimonthly observations in our CPI sample, all model statistics are means (and standard deviations) of each moment across 100 simulated samples of 119 periods. For comparison, column 5 reports corresponding statistics from the 1990–2009 CPI-RDB dataset. The model in column 1 fits

[13] In an online "SW Appendix" we provide more details on how we construct the bimonthly data for our estimation and convert nonestimated parameters from quarterly to bimonthly parameters.

inflation persistence well—its serial correlation is only one standard error from the data (alternatively, about one standard deviation away across simulated samples). But this model falls short on inflation volatility (standard deviation 0.32 percent in the model versus 0.52 percent in the data).[14]

As Smets and Wouters (2007) make clear, much of the inflation variance in their model is driven by price markup shocks, especially at short horizons. To illustrate, column 2 of Table 4 provides model statistics when we drop these price markup shocks entirely. In doing so we reestimate the other parameter values, including shock processes for wage markups, general TFP, investment-specific TFP, government spending, the stochastic discount rate, and monetary policy.[15] The standard deviation of inflation falls from 0.32 percent to 0.09 percent. It is now less than one-fifth of the empirical standard deviation of 0.52 percent. The serial correlation of inflation in the model soars to 0.97 (standard deviation 0.01 across simulations), far above the data's 0.27 (standard error 0.09).

Dropping the markup shocks has little effect on the model's implications for output and other real variables. For example, the standard deviation of output growth for the models shown in columns 1–4 of Table 4 ranges between 1.90 and 1.95 (standard deviation 0.15 across simulations), compared to 1.85 in the data.

For each model column in Table 4, we also estimate the impulse response function (IRF) from a univariate ARMA(6,6) on inflation. For the model without price markup shocks in column 2, the IRF builds sharply. After a 1 percent innovation to the aggregate price level, the aggregate price level is up 10.7 percent one year later (standard deviation 1.5 percent across simulations). This price buildup is in stark contrast to the data, where the one-year cumulative response is no higher than the impact effect of 1 percent (0.90 with a standard error of 0.35).

Why is inflation volatility so low and persistence so high in the SW model without price markup shocks? The bottom of Table 4 provides the slope of the short-run Phillips curve. In columns 1 and 2 the slope is very flat: a 1.0 percent increase in real marginal cost raises current inflation by only 0.0007 percent and 0.0006 percent, respectively. Movements in marginal cost have little impact on current inflation, and take many periods to fully pass through to prices. Of course, this flat short-run Phillips curve is integral to how the SW model produces realistic fluctuations in real output, especially in response to monetary policy shocks.

Again, to mimic the volatile and transitory behavior of actual inflation despite a flat short-run Phillips curve, the Smets-Wouters (2007) model includes a shock to the desired price markup over marginal cost. This price markup shock must be large, transitory, and common to all sellers. Given the minimal pass-through built into the flat short-run Phillips curve, the standard deviation of the shock to the current *desired* price must be 453 percent per bimonth. This is almost 1,000 times bigger

[14] For the model estimation, inflation is defined as the growth rate of the PCE deflator, which had a standard deviation of 0.33 percent and serial correlation of 0.13 over our sample period. We estimate using the PCE deflator because it is the closest counterpart to the model's price index, but we assess the model using the CPI-RDB so we can compare actual to reset price inflation. As we shall see, however, our main conclusions in the remainder of the paper would remain unaffected if we used the PCE deflator rather than the CPI-RDB for our measure of actual inflation in column 5.

[15] For the estimation, we replace the price markup shock with a measurement error shock to the inflation series. A benefit of this approach is that, when we subsequently drop the measurement error shock in the simulations of the estimated model, this alters the model's implications only for inflation and not for other observable variables.

than the standard deviation of actual inflation.[16] De Walque, Smets, and Wouters (2006) suggest smaller, more reasonable shocks might work if the shocks hit only flexible-priced goods such as food and energy. We will pursue this possibility systematically below. In the meantime, we elaborate on the one-sector model.

The entries for reset price inflation are empty in the first two columns of Table 4 for a simple reason: nonadjusting prices are indexed in the Smets-Wouters model, either to steady-state inflation or last period's inflation. The SW model does not feature any nominal price stickiness, only relative price stickiness. As a result, reset price inflation is identical to actual inflation in the SW model. As nominal price stickiness is very much a feature of the CPI data, from here forward we constrain the SW model to feature no nominal price change among nonadjusters. We avoid price indexation altogether because price changes on the order of a few tenths of a percent—as needed to keep up with average or lagged inflation—are uncommon in the CPI micro data (see Klenow and Kryvtsov 2008).

Next in Table 4 (columns 3 and 4) we report moments after reestimating the SW model subject to a price change frequency of 31.2 percent, which is the overall weighted mean bimonthly frequency in our BLS sample. In contrast, the Bayesian estimated frequency of price changes beneath column 1 was 8.2 percent, and that in column 2 was 7.6 percent. Column 3 allows for price markup shocks, while column 4 reestimates dropping these shocks. Because we pin down the frequency from the micro data, we are able to estimate the Kimball-kink pricing parameter. (SW set this exogenously to 10.) Our point estimate is 47.3 allowing markup shocks, and 54.1 without those shocks. The greater price flexibility in columns 3 and 4 makes the Phillips curve steeper, though this is partially offset by the more severe Kimball kink. On net it becomes six times steeper—e.g., a slope of 0.0034 in column 4 versus 0.0006 in column 2.

Comparing the models in columns 2 and 4, both without markup shocks, we see that the steeper Phillips curve makes inflation a bit more volatile and transitory. Still, model inflation in column 4 is far too stable (0.12 percent model standard deviation versus 0.52 percent in the data) and far too persistent (serial correlation 0.91 versus 0.27 in the data, and IRF after one year of 5.6 versus 0.9 in the data). As seen in comparing 1 and 2, if we allow sufficiently large and transitory markup shocks, then the model can do better in generating volatility in inflation. For the model in 3, the markup shocks nearly triple inflation's standard deviation while cutting its persistence considerably. (The one-year impulse for inflation is cut from 5.0 to 1.4, much closer to the data's level of 0.9.)

Columns 3 and 4 of Table 4 also provide our first look at reset price inflation in the SW model, as distinct from actual inflation. Without markup shocks (column 4) the serial correlation of reset price inflation is 0.43 (above the data's 0.06); and the IRF accumulates to 2.2 percent after one year versus only 0.6 percent in the data.

With the markup shocks (column 3), the behavior of reset price inflation is very different. In order to cut persistence in actual inflation, the markup shocks in the SW model must hit reset price inflation more dramatically. Comparing columns 3 and 4,

[16] See Chari, Kehoe, and McGrattan (2009) for a related critique of the SW model's wage markup shocks. Justiniano, Primiceri, and Tambalotti (2011) contend these shocks can be largely replaced by measurement error in wages. In contrast, we will find only a small role for sampling error in aggregate inflation.

Table 5—One-Sector Smets-Wouters with Sampling Error

Statistic	(4) from Table 4 (1)	Adding cross-sectional sampling error (2)	Data for all items (3)
Standard deviation of π	0.12%	0.12%	0.52%
	(0.03)	(0.03)	(0.03)
Serial correlation of π	0.91	0.85	0.27
	(0.04)	(0.07)	(0.09)
1-year cumulative π	5.59	4.42	0.90
	(0.98)	(0.86)	(0.35)
Standard deviation of π^*	0.17%	0.20%	0.66%
	(0.02)	(0.02)	(0.04)
Serial correlation of π^*	0.43	0.19	0.06
	(0.13)	(0.16)	(0.14)
1-year cumulative π^*	2.19	1.77	0.61
	(0.49)	(0.39)	(0.28)
Slope of the SR Phillips Curve	0.0034	0.0034	

Notes: In columns 1 and 2, statistics are averages across 100 model simulations, each of 119 periods. Standard deviations across simulations are in parentheses. In column 2 statistics take into account sampling error from a finite number of prices in the cross-section. Column 3 is based on data from the CPI-RDB for January–February 1990 through September–October 2009, with robust standard errors in parentheses.

the markup shocks increase the standard deviation of reset inflation by a factor of nearly seven. With the markup shocks, the standard deviation of reset inflation is predicted to be 3.5 times larger than for actual inflation, whereas in the data it is only modestly larger (0.66 percent versus 0.52 percent). As discussed in Section II, the bigger impact on reset inflation reflects that, under Calvo pricing, $\frac{\sigma_{\pi^*}}{\sigma_\pi}$ equals $\sqrt{1 + 2(1 - \rho)(1 - \lambda)/\lambda^2}$.

The markup shocks also generate a mismatch in the relative persistence of reset and actual inflation. With the markup shocks, the model's one-year impulse response in actual prices is three times that in reset inflation, whereas in the data that ratio is only one and a half. So we see that, although sufficiently large markup shocks can generate realistic volatility and persistence of actual inflation, this creates unrealistic behavior of reset inflation *relative to actual inflation*.

Could sampling error close the gap between the SW model and the inflation facts? Columns 1–4 of Table 4 already incorporate time series sampling error by taking means and standard deviations across short (119 month) simulated time series. In Table 5 we add an additional, *cross-sectional* source of sampling error. The CPI is an aggregate of a finite sample of individual prices. The SW model, in contrast, contains a continuum of prices and, hence, no idiosyncratic noise. Such sampling error may add volatility and pull down the persistence of actual and reset price inflation.

Based on looking at replicate subsamples, Shoemaker (2004 to 2010) estimates the standard deviation of cross-sectional sampling error in the CPI. He provides estimates for each year from 2003 to 2009, for inflation at horizons of one, two, and six months, and for subaggregates as well as the overall CPI. He estimates a standard deviation from sampling error for bimonthly inflation for the entire CPI that

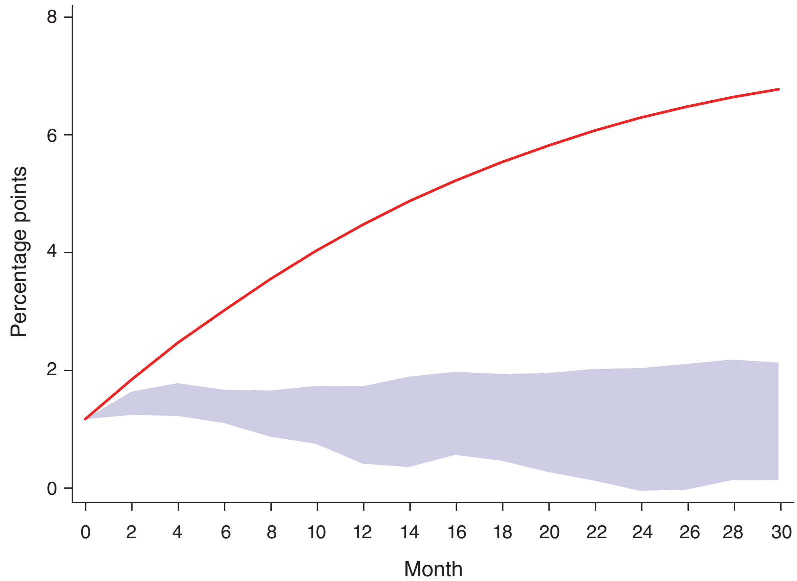

FIGURE 9. SW MODEL IMPULSE RESPONSE OF *ACTUAL* PRICES

Notes: Estimates are accumulated responses to an ARMA(6,6) for actual price inflation. Shaded area denotes 95 percent confidence intervals for estimates based on CPI-RDB data for all items.

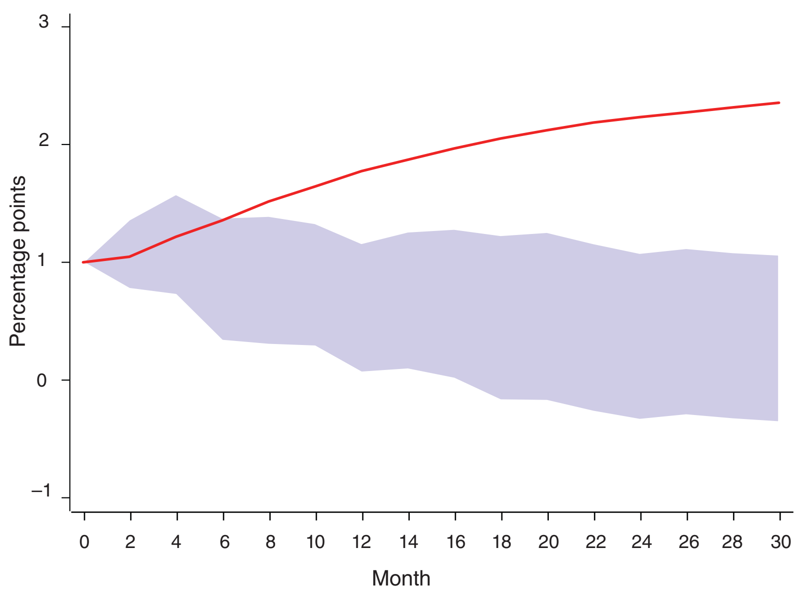

FIGURE 10. SW MODEL IMPULSE RESPONSE OF *RESET* PRICES

Notes: Estimates are accumulated responses to an ARMA(6,6) for reset price inflation. Shaded area denotes 95 percent confidence intervals for estimates based on CPI-RDB data for all items.

Table 6—One-Sector Smets-Wouters without Complementarities

Statistic	(2) from Table 5 (1)	(1) minus Kimball Kink for final goods and for labor (2)	(2) minus sticky wages and wage indexation (3)	Data for all items (4)
Standard deviation of π	0.12%	0.91	1.20	0.52%
	(0.03)	(0.16)	(0.18)	(0.03)
Serial correlation of π	0.85	0.78	0.71	0.27
	(0.07)	(0.07)	(0.08)	(0.09)
1-year cumulative π	4.42	3.79	3.20	0.90
	(0.86)	(0.89)	(0.75)	(0.35)
Standard deviation of π^*	0.20%	1.77	2.63	0.66%
	(0.02)	(0.13)	(0.19)	(0.04)
Serial correlation of π^*	0.19	0.11	0.03	0.06
	(0.16)	(0.11)	(0.10)	(0.14)
1-year cumulative π^*	1.77	1.41	1.13	0.61
	(0.39)	(0.38)	(0.34)	(0.28)
Slope of the SR Phillips curve	0.0034	0.137	0.137	

Notes: In columns 1–3, statistics are averages across 100 model simulations, each of 119 periods. Standard deviations across simulations are in parentheses. These statistics take into account sampling error from a finite number of prices in the cross-section. Column 4 is based on data from the CPI-RDB for January–February 1990 through September–October 2009, with robust standard errors in parentheses.

averages 0.077 percent for the seven years. For our sample of goods, which excludes housing services, his estimates imply a bimonthly sampling error of 0.094 percent. As the standard deviation of bimonthly inflation is 0.52 percent in our 1990–2009 sample, Shoemaker's estimates imply that 3.3 percent of the variance of observed bimonthly inflation stems from sampling a finite set of prices. For inflation at the six-month horizon, Shoemaker's estimates yield a standard deviation from sampling error of 0.142 percent for our sample of goods. Based on the higher variance for sampling error at six versus two months, we can infer that the idiosyncratic price movements underlying this sampling error are reasonably persistent.

To capture sampling error we simulate Calvo time-dependent pricing, as in Table 4, but allow for persistent, idiosyncratic cost shocks to price setters. We perform 100 simulations of the model each for a sample of 80,000 price setters, matching the size of our BLS bimonthly sample. For each simulation, we calculate bimonthly and six-month inflation rates. We calibrate the standard deviation of the idiosyncratic shocks to replicate the share of sampling error in the variance of measured inflation implied by Shoemaker's estimates. We calibrate the persistence of the idiosyncratic shocks so that the relative standard deviation for its six-month versus bimonthly inflation also matches Shoemaker. This requires an AR(1) parameter for these shocks of 0.73. We sum the inflation series drawn from one of 100 simulations of the SW model with the inflation series drawn from one of 100 simulations of the Calvo model with idiosyncratic shocks. We then recalculate statistics based on the sum of these two series. We repeat this process 100 times, sampling with replacement.

Column 2 of Table 5 shows the outcome of adding this cross-sectional sampling error to the SW model. Column 1, for comparison, repeats model results without sampling error from column 4 of Table 4. The impact is modest. The standard deviations increase only slightly. For actual inflation, the impact on serial correlation is small, going from 0.91 to 0.85. Sampling error does reduce the serial correlation of reset price inflation in the model considerably, from 0.43 to 0.19. But the one-year impulse responses are only modestly reduced (from 2.2 to 1.8) and remain far above those reported for the data in the last column of the table.

Figures 9 and 10 demonstrate the model IRFs for inflation and reset inflation over 30 months once we allow for cross-sectional sampling error. The initial response in reset prices accumulates to a factor of 2.4; the response in actual prices ascends to a factor of almost seven. These model IRFs differ starkly from the flat or declining empirical IRFs for reset and actual inflation. The large disparities between model and data are statistically significant, judging by either the standard error in the data IRFs or by the standard deviation of the SW model IRFs. In short, from Table 5 and these figures, we see sampling error does little to diminish the excess smoothness and inertia in model inflation and reset price inflation.

To recap, actual inflation for the model, absent enormous markup shocks, is too stable and too persistent compared to the data. Much of this discrepancy is mirrored in reset prices. The model's standard deviations in reset and actual inflation are each less than 30 percent of that in the data. Consider next the IRF after 30 months in the model relative to the data. For reset prices this ratio is very large at 5.5, while for actual prices it is even larger at 6.5. The upshot is that the discrepancy comes largely from reset prices, reinforced by how the model translates reset price inflation into actual inflation.

A. *Complementarities and Reset Inflation*

Why does the SW model imply too much persistence in reset and actual inflation? Table 6 provides evidence on the role played by complementarities. Column 1 of Table 6 repeats the statistics for Table 5, column 2—the estimated model with sampling error and no markup shocks. The subsequent columns sequentially shut down complementarities without reestimating. First we get rid of the Kimball "kink" in the demand curve facing individual price setters (firms) and wage setters (workers) in column 2. When the elasticity of demand is no longer increasing in an item's relative price, price setters more freely pass marginal costs into prices. They do not wait so much for other prices to adjust despite the lack of synchronization endemic to Calvo pricing. The slope of the short-run Phillips curve increases from 0.34 percent to 13.7 percent. As a result, the standard deviation of inflation rises by a factor greater than seven, and the standard deviation of reset price inflation by almost a factor of nine.[17] In addition, all persistence measures fall. The IRF for reset prices becomes 1.4 after one year, compared to 1.8 in column 1. After 30 months the response is only 1.6, compared to 2.4 (see Figure 10) for the economy with all complementarities in column 1.

In column 3 of Table 6 we eliminate both nominal wage stickiness and wage indexation. Inflation—and especially reset price inflation—becomes even more

[17] This understates the case. By holding fixed the size of the reduced-form shocks as we peeled away complementarities, we implicitly scaled down the standard deviation of the desired wage markup from 236 percent in column 1 to 39 percent in column 2.

Table 7—Two-Sector Smets-Wouters Models

Statistic	1-sector model aggregate (1)	2-sector aggregate (2)	Flexible sector (3)	Sticky sector (4)
Standard deviation of π	0.34%	0.79%	1.84%	0.45%
	(0.02)	(0.07)	(0.13)	(0.05)
Serial correlation of π	0.19	0.36	0.25	0.47
	(0.09)	(0.10)	(0.10)	(0.12)
1-year cumulative π	1.34	1.63	1.02	2.32
	(0.38)	(0.40)	(0.35)	(0.47)
Standard deviation of π^*	1.18%	2.47%	3.17%	1.52%
	(0.09)	(0.19)	(0.22)	(0.11)
Serial correlation of π^*	–0.38	–0.26	–0.15	–0.21
	(0.06)	(0.08)	(0.09)	(0.10)
1-year cumulative π^*	0.44	0.54	0.56	0.79
	(0.16)	(0.21)	(0.23)	(0.22)
Slope of the SR Phillips curve	0.0042		0.426	0.0011

Notes: Statistics are averages across 100 model simulations, each of 119 periods. Standard deviations across simulations are in parentheses. These statistics take into account sampling error from a finite number of prices in the cross-section.

volatile. They end up *too* volatile relative to the data, by a factor greater than two for inflation and equal to four for reset price inflation. This drives home the role strategic complementarities play in suppressing inflation volatility.

Table 6 held the size of model shocks fixed. The resulting volatility of real output growth is fairly stable as we move from columns 1 to 3. The standard deviation of real output growth is 1.90 percent in model column 3, close to the data's 1.85 percent. If we instead scale down all shock standard deviations equally to hit the volatility of actual inflation in column 3, the model standard deviation of output growth falls to 0.82 percent, less than half the data figure. If we scale down shocks sufficiently to match the volatility of reset price inflation, then output variability sinks further to 0.50 percent. Thus, the model without complementarities cannot simultaneously generate realistic volatility for output growth and inflation.

In the absence of all of these complementarities, reset price inflation builds little over time. The serial correlation of reset price inflation, 0.03, is essentially zero as in the data. The IRF for reset inflation is 1.1 after one year, and 1.2 after 30 months. While above the corresponding data responses, 0.59 and 0.43, this at least stays within the data's confidence interval for the first two years.

Even in this stripped-down SW model, inflation is too persistent relative to the data. Its serial correlation is 0.71 (versus 0.27 in the data). Prices build to 3.2 percent higher one year after a 1 percent impulse, whereas they do not build at all in the data. Thus stripping away the complementarities largely eliminates excess persistence in reset price inflation, but not in actual inflation. Time-dependent pricing still adds too much persistence to inflation conditional on reset price inflation.[18]

[18] In Bils, Klenow, and Malin (2009), we illustrate that dropping time-dependent pricing in favor of state-dependent pricing dramatically reduces persistence in both reset and actual inflation. Of course, dropping time-dependent

B. *Sectoral Shocks*

The one-sector models analyzed so far neglect the possibility of large, temporary sectoral shocks that do not wash out in the aggregate. Boivin, Giannoni, and Mihov (2009) provide evidence that disaggregated inflation rates are much more volatile and transitory than aggregate inflation, consistent with such shocks. And De Walque, Smets, and Wouters (2006) illustrate how one might replace price markup shocks with a productivity shock to a flexible-price sector. We now pursue this possibility.

Our two-sector version of the Smets-Wouters model has the following features. Aggregate consumption is a Cobb-Douglas composite of flexible-price consumption and sticky-price consumption. Flexible-price consumption is a CES aggregate of individual flexible items, whereas sticky-price consumption is a Kimball aggregate of individual sticky items. We omit the Kimball kink for flexible goods because strong strategic complementarities there would mute the impact of any large transitory shocks in that sector, undermining the point of entertaining such shocks. As in De Walque, Smets, and Wouters (2006), instead of an aggregate price markup shock there is a productivity shock that hits only the flexible sector (in addition to the common aggregate shock that hits both sectors). The two sectors compete in the same market for labor, which continues to feature sticky wages, indexation, a Kimball labor-market kink, and a wage markup shock. Finally, in any given quarter, sector capital stocks are predetermined.

We calibrate the two-sector model by exploiting the separate series we constructed on inflation (and reset price inflation) for flexible and sticky items. We choose the Cobb-Douglas exponents of the consumption aggregator to match consumption shares. We impose the BLS pricing frequencies seen for flexible and sticky items, with no price indexation for nonadjusters. Finally, we calibrate the two parameters of an AR(1) process for the flexible sector productivity shock to match the serial correlation and volatility of the (log) of the flexible price relative to the sticky price. The flexible sector productivity shock is less persistent (serial correlation of 0.50) than all model shocks except the wage markup shock and is more volatile (innovation standard deviation of 0.60) than all shocks except the aggregate productivity shock.

Table 7 presents simulated statistics from this two-sector model. For comparison, column 1 starts with a *one-sector* model containing markup shocks and sampling error. Column 2 provides statistics on the two-sector aggregate, including sampling error.[19] In terms of persistence, the two-sector model does a fairly good job of mimicking the one-sector model with extreme markup shocks. The serial correlation in aggregate inflation is 0.36 in the two-sector model versus 0.19 in

pricing, as with dropping strategic complementarities, undercuts the ability of the sticky price model to generate large or persistent output volatility, especially in response to monetary shocks.

[19] We calibrate the sampling error for the flexible and sticky sectors based on Shoemaker's (2004 to 2010) estimates of sampling error in inflation for the sectors we classify as flexible (price-change frequency greater than one-third) and sticky. We find that sampling error constitutes only 2 percent of the time-series variance of inflation for the flexible sector, but 30 percent for the sticky sector. Sampling error also breaks the direct connection between inflation persistence and the relative variability in reset versus actual inflation under Calvo pricing, $\frac{\sigma_{\pi^*}}{\sigma_\pi} = \sqrt{1 + 2(1-\rho)(1-\lambda)/\lambda^2}$, as in a finite sample there is randomness in whether those sellers drawing a price-change exhibit above or below average desired price changes. We find this effect of sampling error is unimportant for the aggregate model economy or the flexible sector; but, for the sticky sector, it lowers inflation persistence without driving up the relative variability of reset to actual inflation so drastically.

the one-sector model (and 0.27 in the data); and the one-year IRF is 1.6 percent in the two-sector model versus 1.3 percent in the one-sector model. Even so, that remains nearly double that for the data (0.9 percent). For reset inflation the serial correlation, at -0.26, is slightly higher than for the one-sector model and lower than the value of zero seen in the data. The one-year IRF for reset inflation nearly matches that for the data, as does the one-sector model with aggregate markup shocks.

But the two-sector model also inherits, and even magnifies, the empirical problems of the one-sector model with markup shocks. In particular, reset inflation for the model is far too volatile, both in absolute terms and relative to actual inflation. Its standard deviation for reset inflation, 2.5 percent, is 3.8 times that in the data. This 2.5 percent is also more than three times that predicted for actual inflation. By contrast, the data show a standard deviation for reset inflation that is only modestly greater (by 26 percent) than for actual. We point out again: this failure of the model is a consequence of hitting low inflation persistence, given its time-dependent pricing.

Columns 3 and 4 of Table 7 report model moments for flexible and sticky sector inflation separately. Not surprisingly, inflation is more volatile in the flexible sector. This reflects not only the productivity shock hitting the flexible sector, but also the quicker pass-through of marginal cost into flexible prices. The slope of the short-run Phillips curve is 400 times larger in the flexible sector (43 percent) than in the sticky sector (0.11 percent). Nevertheless, the two-sector model doesn't generate enough volatility in the flexible sector relative to the sticky sector, as the ratio of their standard deviations is 4 in the model versus 9 in the data (see data Table 2 above).

The two-sector model does well in matching inflation persistence in the two sectors. The one-year response in the sticky sector in reset prices is 0.79 percent, not far above that in the data. But the model requires far too much volatility in reset price inflation to achieve this. For flexible goods, reset price inflation is twice as volatile as in the data; for sticky goods, it is 3.5 times as volatile.

The model generates volatility in sticky-sector inflation, despite its very flat short-run Phillips curve, because the flexible sector shocks create large transitory fluctuations in wages facing the sticky sector. As a result, reset inflation rates in the two sectors are highly correlated at 0.72, despite the large relative technology shocks. But the data do not show this, displaying a correlation of only 0.12. (For actual inflation rates the model correlation is 0.51, versus -0.02 for the data.) We take this as evidence that the two-sector model overstates the importance of flexible-sector shocks to sticky-sector prices. Because the model greatly overstates the correlation in inflation rates, when the two sectors are aggregated the discrepancy between the model and data is more striking. Most notably, the model overstates the standard deviation for reset inflation by a factor of 3.8.

IV. Conclusion

A large empirical literature has estimated that monetary policy shocks affect real variables for several years, much longer than the duration of nominal prices. Popular machinery to explain these findings combines sticky prices and strategic complementarities. The complementarities make reset prices build slowly after permanent shocks, prolonging the real effects beyond the duration of nominal prices. That is, strategic complementarities impart persistence to reset price inflation relative to the persistence in the underlying shocks. Price stickiness imparts further persistence in actual inflation relative to reset inflation. The combination succeeds in flattening the short-run Phillips curve.

But we do not see persistence in reset price inflation using data underlying the US CPI from 1990–2009. In particular, impulse responses for reset prices are flat or declining over 30 months. This statement holds for core goods, as well as food and energy goods; it holds for goods with frequent as well as infrequent price changes; it holds with or without sale prices (i.e., for regular prices). Over the same period persistence in actual inflation has also been modest—impulse responses for actual prices build modestly after 30 months for sticky-priced goods, and little, or not at all, for all goods taken together.

In theory, low persistence in actual inflation can be reconciled with a flat short-run Phillips curve provided that desired spot prices—the prices firms would charge absent sticky prices—are sufficiently volatile and transitory. Smets and Wouters (2007) include large, transitory price markup shocks to accomplish that end. (We find that a two-sector model, with large shocks to the flexible-price sector, can achieve much the same outcome.) While such shocks help to reconcile the model with data on actual inflation, we find they cannot at the same time reconcile the model with our data on reset inflation. Because these shocks must be so transitory, their impact on reset price inflation is exaggerated compared to that on actual inflation— creating variability in reset price inflation well above that seen in the data. This conflict is not specific to the Smets-Wouters model. The modest persistence in actual inflation, together with the low variability of reset versus actual inflation, is difficult to match with time-dependent Calvo pricing.

Our results suggest the high inflation persistence over longer samples might reflect monetary policy rather than a flat short-run Phillips curve. The low inflation persistence of recent decades could be because the Fed stopped adding persistence, revealing low endogenous persistence. Cogley and Sargent (2002), Primiceri (2006), Cogley and Sbordone (2008), and Benati (2008) all argue that US inflation persistence over long samples stems, in part, from monetary regime changes.

Our findings also suggest considering alternatives to Calvo time-dependent pricing. State-dependent pricing models are one possibility, but can bring problems of their own (e.g., too few small price changes as emphasized by Midrigan 2011). And state-dependent pricing models with weak complementarities can generate too little output response to monetary shocks. As we see it, an open challenge is to reconcile the low persistence of inflation (and low volatility of reset price inflation) in recent decades with the VAR evidence that monetary variables exhibit large, sustained real effects.

REFERENCES

Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Linde. 2011. “Firm-Specific Capital, Nominal Rigidities and the Business Cycle.” *Review of Economic Dynamics* 14 (2): 225–47.

Ball, Laurence, and David Romer. 1990. “Real Rigidities and the Non-neutrality of Money.” *Review of Economic Studies* 57 (2): 183–203.

Benati, Luca. 2008. “Investigating Inflation Persistence across Monetary Regimes.” *Quarterly Journal of Economics* 123 (3): 1005–60.

Bernanke, Ben S., Jean Boivin, and Piotr Eliasz. 2005. "Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach." *Quarterly Journal of Economics* 120 (1): 387–422.

Bils, Mark, Peter J. Klenow, and Benjamin A. Malin. 2009. "Reset Price Inflation and the Impact of Monetary Policy Shocks." National Bureau of Economic Research Working Paper 14787.

Bils, Mark, Peter J. Klenow, and Benjamin A. Malin. 2012. "Reset Price Inflation and the Impact of Monetary Policy Shocks: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.102.6.2798.

Blanchard, Olivier, and Jordi Galí. 2007. "Real Wage Rigidities and the New Keynesian Model." *Journal of Money, Credit, and Banking* 39 (1): 35–65.

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

Bureau of Economic Analysis. 1990–2009. "National Data. GDP and Personal Income." United States Department of Commerce. http://www.bea.gov/iTable/index_nipa.cfm (accessed December 8, 2010).

Bureau of Labor Statistics. 1988–2009a. "CPI Research Database." United States Department of Labor. http://stats.bls.gov/bls/blsresda.htm (accessed January 13, 2010).

Bureau of Labor Statistics. 1988–2009b. "Multifactor Productivity." United States Department of Labor. http://www.bls.gov/mfp/ (accessed December 8, 2010).

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

Carvalho, Carlos. 2006. "Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks." *B. E. Journal of Macroeconomics: Frontiers of Macroeconomics* 2 (1): 1–56.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. 2009. "New Keynesian Models: Not Yet Useful for Policy Analysis." *American Economic Journal: Macroeconomics* 1(1): 242–66.

Chevalier, Judith A., and Anil K. Kashyap. 2011. "Best Prices." National Bureau of Economic Research Working Paper 16680.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" In *Handbook of Macroeconomics Volume 1A*, edited by J. B. Taylor and M. Woodford, 65–148. Amsterdam: Elsevier Science, North-Holland.

Cogley, Timothy, Giorgio E. Primiceri, and Thomas J. Sargent. 2010. "Inflation-Gap Persistence in the US." *American Economic Journal: Macroeconomics* 2 (1): 43–69.

Cogley, Timothy, and Thomas J. Sargent. 2002. "Evolving Post-World War II U.S. Inflation Dynamics." In *National Bureau of Economic Research Macroeconomics Annual 2001*, edited by B. S. Bernanke and K. Rogoff, 331–73. Cambridge, MA: MIT Press.

Cogley, Timothy, and Argia M. Sbordone. 2008. "Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve." *American Economic Review* 98 (5): 2101–26.

de Walque, Gregory, Frank Smets, and Raf Wouters. 2006. "Price Shocks in General Equilibrium: Alternative Specifications." *CESifo Economic Studies* 52 (1): 153–76.

Dotsey, Michael, Robert G. King, and Alexander L. Wolman. 1999. "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output." *Quarterly Journal of Economics* 114 (2): 655–90.

Gertler, Mark, and John Leahy. 2008. "A Phillips Curve with an Ss Foundation." *Journal of Political Economy* 116 (3): 533–72.

Golosov, Mikhail, and Robert E. Lucas, Jr. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115 (2): 171–99.

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

Gopinath, Gita, Oleg Itskhoki, and Roberto Rigobon. 2010. "Currency Choice and Exchange Rate Pass-Through." *American Economic Review* 100 (1): 304–36.

Granger, Clive W. J., and M. J. Morris. 1976. "Time Series Modelling and Interpretation." *Journal of the Royal Statistical Society, Series A* 139 (2): 246–57.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2011. "Is There a Trade-off between Inflation and Output Stabilization?" National Bureau of Economic Research Working Paper 17071.

Kehoe, Patrick J., and Virgiliu Midrigan. 2010. "Prices are Sticky After All." Federal Reserve Bank of Minneapolis Research Department Staff Report 413.

Kimball, Miles S. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model." *Journal of Money, Credit, and Banking* 27 (4): 1241–77.

Klenow, Peter J., and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123 (3): 863–904.

Klenow, Peter J., and Benjamin A. Malin. 2011. "Microeconomic Evidence on Price-Setting." In *Handbook of Monetary Economics Volume 3A*, edited by Benjamin Friedman and Michael Woodford, 231–84. Amsterdam: Elsevier Science, North-Holland.

Klenow, Peter J., and Jonathan L. Willis. 2006. "Real Rigidities and Nominal Price Changes." Unpublished.

Kryvtsov, Oleksiy, and Virgiliu Midrigan. 2010. "Inventories, Markups, and Real Rigidities in Menu Cost Models." Unpublished.

Midrigan, Virgiliu. 2011. "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations." *Econometrica* 79 (4): 1139–80.

Nakamura, Emi, and Jon Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123 (4): 1415–64.

Nakamura, Emi, and Jon Steinsson. 2010. "Monetary Non-neutrality in a Multisector Menu Cost Model." *Quarterly Journal of Economics* 125 (3): 961–1013.

Nason, James M. 2006. "Instability in U.S. Inflation: 1967–2005." *Federal Reserve Bank of Atlanta Economic Review* 91 (2): 39–59.

Primiceri, Giorgio E. 2006. "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U.S. Postwar Stabilization Policy." *Quarterly Journal of Economics* 121 (3): 867–901.

Romer, Christina D., and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94 (4): 1055–84.

Shiller, Robert J. 1991. "Arithmetic Repeat Sales Price Estimators." *Journal of Housing Economics* 1 (1): 110–26.

Shoemaker, Owen J. 2004–2010. "Variance Estimates for Price Changes in the Consumer Price Index January–December 2007." http://www.bls.gov/cpi/cpivar2007.pdf (accessed November 15, 2011).

Smets, Frank, and Raf Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1 (5): 1123–75.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.

Stock, James H., and Mark W. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit, and Banking* 39 (Supplement): 3–33.

Stock, James H., and Mark W. Watson. 2009. "Phillips Curve Inflation Forecasts." In *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*, edited by J. Fuhrer, Y. K. Kodrzycki, J. S. Little and G. P. Olivei, 101–86. Foreword by Paul A. Samuelson. Cambridge, MA: MIT Press.

Stock, James H. and Mark W. Watson. 2010. "Research Memorandum on the Distribution of Quarterly values of GDP/GDI across Months within the Quarter: Dataset." http://www.princeton.edu/~mwatson/mgdp_gdi.html (accessed December 8, 2010).