Sectoral price data and models of price setting*

Bartosz Maćkowiak
European Central Bank and CEPR

Emanuel Moench
Federal Reserve Bank of New York

Mirko Wiederholt
Northwestern University

December 2008

Abstract

We use a statistical model to estimate impulse responses of sectoral price indices to aggregate shocks and to sector-specific shocks. In the median sector, 100 percent of the long-run response of the sectoral price index to a sector-specific shock occurs in the month of the shock. The Calvo model and the sticky-information model match this finding only under extreme assumptions concerning the profit-maximizing price. By contrast, the rational inattention model matches this finding without an extreme assumption concerning the profit-maximizing price. Furthermore, we find little variation across sectors in the speed of response of sectoral price indices to sector-specific shocks. The rational inattention model matches this finding, while the Calvo model predicts far too much cross-sectional variation in the speed of response to sector-specific shocks.

JEL: E3, D8, C1.

Keywords: Calvo model, sticky information, rational inattention, Bayesian unobservable index model.

*We thank for comments seminar participants at the Bundesbank. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the European Central Bank or the Federal Reserve. E-mail: bartosz.mackowiak@ecb.int, emanuel.moench@ny.frb.org, and m-wiederholt@northwestern.edu.

1 Introduction

This paper studies sectoral consumer price data in order to evaluate models of price setting. Over the last twenty years, there has been a surge in research on macroeconomic models with price stickiness. In these models, price stickiness arises either from adjustment costs (e.g. the Calvo model and the menu cost model) or from some form of information friction (e.g. the sticky-information model and the rational inattention model). Models with price stickiness are often evaluated by looking at aggregate data. Here, we evaluate models with price stickiness by looking at sectoral data. We estimate a statistical model for sectoral inflation rates. From the statistical model, we compute impulse responses of sectoral price indices to aggregate shocks and to sector-specific shocks. We then ask whether different models of price setting match these empirical impulse responses.

The statistical model that we estimate is the following. The inflation rate in a sector equals the sum of two components, an aggregate component and a sector-specific component. The parameters in the aggregate component and in the sector-specific component may differ across sectors. Each component follows a stationary Gaussian process. An innovation in the aggregate component may affect the inflation rates in all sectors. An innovation in the sector-specific component affects only the inflation rate in this sector. We estimate the statistical model using monthly sectoral consumer price data from the U.S. economy for the period 1985-2005. The data are compiled by the Bureau of Labor Statistics. From the estimated statistical model, we compute impulse responses of the price index for a sector to an innovation in the aggregate component and to an innovation in the sector-specific component.

We find that, in the median sector, after a sector-specific shock 100 percent of the long-run response of the sectoral price index occurs in the month of the shock, and then the response equals the long-run response in all months following the shock. By contrast, after an aggregate shock, only 15 percent of the long-run response of the sectoral price index occurs in the month of the shock, and then the response gradually approaches the long-run response. In other words, in the median sector, the sector-specific component of the sectoral inflation rate is a white noise process, while the aggregate component of the sectoral inflation rate is positively autocorrelated.

We then ask whether different models of price setting match the median impulse response of a sectoral price index to a sector-specific shock. Recall that this impulse response looks like the impulse response function of a random walk: the sectoral price index jumps on impact of a sector-specific shock, and then stays there. We show that the Calvo model matches this impulse response of the sectoral price index to a sector-specific shock only under an extreme assumption concerning the response of the profit-maximizing price to a sector-specific shock. After a sector-specific shock, the profit-maximizing price needs to jump by about $\left(1/\lambda^2\right) x$ in the month of the shock, and then jump back to x in the next month to generate a response equal to x of the sectoral price index on impact and in all months following the shock. Here, λ denotes the fraction of firms that can adjust their prices in a month. We obtain a similar, though less extreme, result for the sticky-information model developed in Mankiw and Reis (2002). After a sector-specific shock, the profit-maximizing price needs to jump by $(1/\lambda)\, x$ in the month of the shock, and then decay slowly to x to generate a response equal to x of the sectoral price index on impact and in all months following the shock. Here, λ denotes the fraction of firms that can update their pricing plans in a month. In contrast, we find that the rational inattention model developed in Maćkowiak and Wiederholt (2008a) matches the empirical impulse response of a sectoral price index to sector-specific shocks without an extreme assumption concerning the response of the profit-maximizing price to a sector-specific shock. The intuition is simple. According to the estimated statistical model, sector-specific shocks are on average much larger than aggregate shocks. In the model, this implies that decision-makers in firms pay a lot more attention to sector-specific conditions than to aggregate conditions. For this reason, prices respond quickly to sector-specific shocks and slowly to aggregate shocks.

We also investigate whether the different models of price setting predict the right amount of variation across sectors in the speed of response of sectoral price indices to sector-specific shocks. According to our statistical model, there is little variation across sectors in the speed of response of sectoral price indices to sector-specific shocks. We study a multi-sector Calvo model calibrated to the sectoral monthly frequencies of price changes reported in Bils and Klenow (2004). We find that the model predicts far too much cross-sectional variation in the speed of response to sector-specific shocks. In contrast, the rational inattention model developed in Maćkowiak and Wiederholt (2008a) correctly predicts little cross-sectional variation in the speed of response to sector-specific shocks. The intuition is the following. In the median sector, decision-makers in firms are already paying so much attention to sector-specific conditions that they track sector-specific conditions almost perfectly. Paying even more attention to sector-specific conditions has little effect on the speed of response of prices to sector-specific shocks.

Our work is related to the recent papers by Boivin, Giannoni, and Mihov (2007) and Reis and Watson (2007a). Boivin, Giannoni, and Mihov (2007) use a factor augmented vector autoregressive model to study sectoral data published by the Bureau of Economic Analysis on personal consumption expenditure. Like us, Boivin, Giannoni, and Mihov (2007) find that prices respond quickly to sector-specific shocks and slowly to aggregate shocks, and that sector-specific shocks account for a dominant share of the variance in prices. Compared to Boivin, Giannoni, and Mihov (2007), we use a different methodology and different data. In addition, we write down the Calvo model, the sticky-information model, and the rational inattention model, and we investigate to what extent the models match what we find in the data. Reis and Watson (2007a) use a dynamic factor model to study sectoral data published by the BEA on personal consumption expenditure. The focus of Reis and Watson (2007a) is on estimating the numeraire, a common component in prices that has an equiproportional effect on all prices.[1]

Our work is also related to the recent paper by Kehoe and Midrigan (2007). They study data from Europe and the United States on sector-level real exchange rates. The Calvo model predicts much more heterogeneity in the persistence of sector-level real exchange rates compared to what Kehoe and Midrigan (2007) find in the data.

The statistical model that we use belongs to the class of unobservable index models. The history of unobservable index models goes back to Geweke (1977) and Sargent and Sims (1977). These authors estimated unobservable index models using aggregate data by max-

[1] See also Reis and Watson (2007b).

imizing the spectral likelihood function. We estimate an unobservable index model using disaggregate data by Bayesian Markov Chain Monte Carlo in the time domain. Unobservable index models and dynamic factor models have had many applications since Geweke (1977) and Sargent and Sims (1977).[2] Otrok and Whiteman (1998) is an influential paper introducing Bayesian methods to estimation of dynamic factor models.

The paper is organized as follows. In Section 2, we present the statistical model. In Section 3, we describe the data. In Section 4, we study the out-of-sample forecasting performance of the statistical model. In Sections 5 and 6, we present the substantive results from the statistical model. In Section 7, we study whether the model of Calvo (1983), the sticky-information model developed in Mankiw and Reis (2002), and the rational inattention model developed in Maćkowiak and Wiederholt (2008a) match what we find in the data. We conclude in Section 8. Appendix A gives econometric details. Appendices B and C contain proofs of theoretical results.

2 Statistical model

We consider the statistical model

$$\pi_{nt} = \mu_n + A_n(L)\, u_t + B_n(L)\, v_{nt}, \tag{1}$$

where π_{nt} is the month-on-month inflation rate in sector n in period t, μ_n are constants, $A_n(L)$ and $B_n(L)$ are square summable polynomials in the lag operator, u_t follows a unit-variance Gaussian white noise process, and each v_{nt} follows a unit-variance Gaussian white noise process. The processes v_{nt} are pairwise independent and independent of the process u_t.

It is straightforward to generalize equation (1) such that u_t follows a vector Gaussian white noise process with covariance matrix identity. We estimate specifications of equation (1) in which u_t follows a scalar process, and we also estimate a specification of equation (1) in which u_t follows a vector process.

[2]See, for example, Stock and Watson (1989), Kim and Nelson (1999), Forni, Hallin, Lippi, and Reichlin (2004), and Bernanke, Boivin, and Eliasz (2005).

Let π_{nt}^{A} denote the aggregate component of the inflation rate in sector n, that is,

$$\pi_{nt}^{A} = A_n(L)\, u_t.$$

We parameterize the aggregate component of the inflation rate in sector n as a finite-order moving average process. We choose the order of the polynomials $A_n(L)$ to be as high as computationally feasible. Specifically, we set the order of the polynomials $A_n(L)$ to twenty four, that is, we let u_t and twenty four lags of u_t enter equation (1).

Let π_{nt}^{S} denote the sector-specific component of the inflation rate in sector n, that is,

$$\pi_{nt}^{S} = B_n(L)\, v_{nt}.$$

To reduce the number of parameters to estimate, we parameterize the sector-specific component of the inflation rate in sector n as an autoregressive process:

$$\pi_{nt}^{S} = C_n(L)\, \pi_{nt}^{S} + B_{n0} v_{nt},$$

where $C_n(L)$ is a polynomial in the lag operator satisfying $C_{n0} = 0$. We estimate specifications in which the order of the polynomials $C_n(L)$ equals six, and we also estimate a specification in which the order of the polynomials $C_n(L)$ equals twelve.

Before estimation, we demean the sectoral inflation rates and we normalize the sectoral inflation rates to have unit variance. This means that we estimate the model

$$\tilde{\pi}_{nt} = a_n(L)\, u_t + b_n(L)\, v_{nt}.$$

Here $\tilde{\pi}_{nt} = \left[\left(\pi_{nt} - \mu_n\right) / \sigma_{\pi_n} \right]$ is the normalized inflation rate in sector n in period t, where σ_{π_n} is the standard deviation of the inflation rate in sector n, and $a_n(L)$ and $b_n(L)$ are square summable polynomials in the lag operator. The following relationships hold: $A_n(L) = \sigma_{\pi_n} a_n(L)$ and $B_n(L) = \sigma_{\pi_n} b_n(L)$. This normalization makes it easier to compare impulse responses across sectors. In what follows, we refer to coefficients appearing in the polynomials $a_n(L)$ and $b_n(L)$ as "normalized impulse responses".

The statistical model (1) is the same as the unobservable index model proposed by Geweke (1977) and Sargent and Sims (1977).[3] The unobservable index in equation (1) is u_t.

[3] See also Geweke and Singleton (1981).

Our approach to inference is Bayesian. We use the Gibbs sampler with a Metropolis-Hastings step to sample from the joint posterior density of the unobservable index and the model's parameters. Taking as given a Monte Carlo draw of the model's parameters, we sample from the conditional posterior density of the unobservable index given the model's parameters. Here we follow Carter and Kohn (1994) and Kim and Nelson (1999). Afterwards, taking as given a Monte Carlo draw of the unobservable index, we sample from the conditional posterior density of the model's parameters given the unobservable index. Here we follow Chib and Greenberg (1994). We employ a prior with zero mean for the parameters appearing in the polynomials $a_n(L)$ and $b_n(L)$, for each n. The prior starts out loose and becomes gradually tighter at more distant lags. See Appendix A for econometric details, including details of the prior.

3 Data

We use the data underlying the consumer price index (CPI) for all urban consumers in the United States. The data are compiled by the Bureau of Labor Statistics (BLS). The data are monthly sectoral price indices. The sectoral price indices are available at four different levels of aggregation: from least disaggregate (8 "major groups") to most disaggregate (205 sectors).[4] We focus on the most disaggregate sectoral price indices. For some sectors, price indices are available for only a short period, often starting as recently as in 1998. We focus on the 79 sectors for which monthly price indices are available from January 1985. These sectors comprise 68.1 percent of the CPI. Each "major group" is represented. The sample used in this paper ends in May 2005.

The median standard deviation of sectoral inflation in the cross-section of sectors in our dataset is 0.0068. For comparison, the standard deviation of the CPI inflation rate in our sample period is 0.0017. In 76 out of 79 sectors, the sectoral inflation rate is more volatile than the CPI inflation rate.

To gain an idea about the degree of comovement in our dataset, we computed principal

[4] The "major groups" are (with the percentage share in the CPI given in brackets): food and beverages (15.4), housing (42.1), apparel (4.0), transportation (16.9), medical care (6.1), recreation (5.9), education and communication (5.9), other goods and services (3.8).

components of the normalized sectoral inflation rates. The first few principal components explain only little of the variation in the normalized sectoral inflation rates. In particular, the first principle component explains 7 percent of the variation, and the first five principle components together explain 20 percent of the variation. This suggests that changes in sectoral price indices are caused mostly by sector-specific shocks.

4 Out-of-sample forecasts

In this section, we compare the out-of-sample forecasting performance of a number of specifications of the unobservable index model to the out-of-sample forecasting performance of simple, autoregressive models for sectoral inflation. We find that the unobservable index model forecasts sectoral inflation better than autoregressive models for sectoral inflation. This gives us confidence in the ability of the unobservable index model to fit the data. Furthermore, the forecast results lead us to focus in Sections 5-6 on one specification of the unobservable index model, which we refer to as the benchmark specification.

We report out-of-sample forecast results based on estimation of the unobservable index model without the last twenty four periods in the dataset. We consider the following three specifications: (i) the unobservable index u_t follows a scalar process and the order of the polynomials $C_n(L)$ equals six; (ii) the unobservable index u_t follows a scalar process and the order of the polynomials $C_n(L)$ equals twelve; (iii) the unobservable index u_t follows a bivariate vector process and the order of the polynomials $C_n(L)$ equals six. For each specification and each sector, we forecast the normalized sectoral inflation rate, $\tilde{\pi}_{nt}$, one-step-ahead in the last twenty four periods in the dataset, and we save the average root mean squared error of the twenty four forecasts. We perform the same out-of-sample forecast exercise using an autoregressive model for the normalized sectoral inflation rate, $\tilde{\pi}_{nt}$. We estimate the autoregressive model separately for each sector by ordinary least squares, setting the number of lags to, alternatively, six and twelve. We compare the AR(6) specification to the specifications of the unobservable index model in which the order of the polynomials $C_n(L)$ equals six. We compare the AR(12) specification to the specification of the unobservable index model in which the order of the polynomials $C_n(L)$ equals twelve.

Table 1 provides information concerning the out-of-sample forecasts from the three specifications of the unobservable index model and from the two specifications of the autoregressive model. Table 1 summarizes the sectoral distribution of the average RMSE of the forecasts for each model and each specification. As an example, 1.059 is the median of the sectoral distribution of the RMSE from the unobservable index model in which the unobservable index follows a scalar process and the order of the polynomials $C_n(L)$ equals six. This is smaller than the median of the sectoral distribution of the RMSE from the AR(6) model, which equals 1.066. As another example, 1.529 is the 95th percentile largest RMSE from the same unobservable index model. This is smaller than the 95th percentile largest RMSE from the AR(6) model, which equals 1.648.

The out-of-sample forecast results show the following. The unobservable index model in which u_t follows a scalar process and the order of the polynomials $C_n(L)$ equals six forecasts better than the AR(6) model. Furthermore, the unobservable index model in which u_t follows a scalar process and the order of the polynomials $C_n(L)$ equals twelve forecasts better than the AR(12) model. The out-of-sample forecast results give us confidence that the unobservable index model fits the data well.

Finally, the specification of the unobservable index model in which u_t follows a scalar process and the order of the polynomials $C_n(L)$ equals six forecasts better than the other two specifications of the unobservable index model. Therefore, we choose this specification as the benchmark specification.

5 Responses of sectoral price indices to sector-specific shocks and to aggregate shocks

In this section, we report results for the estimated unobservable index model (1). We focus on the benchmark specification, in which the unobservable index u_t follows a scalar process and the order of the polynomials $C_n(L)$ equals six. We discuss briefly to what extent the results from the benchmark specification hold in the other specifications of the unobservable index model.

To begin with, we report the variance decomposition of sectoral inflation into aggregate shocks and sector-specific shocks. Sector-specific shocks account for a dominant share of the variance in sectoral inflation. In the median sector, the share of the variance in sectoral inflation due to sector-specific shocks equals 90 percent. The sectoral distribution is tight. In the sector that lies at the 5th percentile of the sectoral distribution, the share of the variance in sectoral inflation due to sector-specific shocks equals 79 percent, and in the sector that lies at the 95th percentile of the sectoral distribution, the share of the variance in sectoral inflation due to sector-specific shocks equals 95 percent.[5]

Next, we report the responses of sectoral price indices to sector-specific shocks and to aggregate shocks. Figure 1 shows the normalized impulse responses of sectoral price indices to sector-specific shocks (top panel) and to aggregate shocks (bottom panel). Each panel shows a posterior density taking into account both variation across sectors and parameter uncertainty. Specifically, we make 7500 draws from the posterior density.[6] We then select at random 1000 draws. Since there are 79 sectors, this gives us a sample of 79000 impulse responses. Each panel in Figure 1 is based on 79000 impulse responses.[7] The median impulse response of a sectoral price index to a sector-specific shock has the following shape. After a sector-specific shock, 100 percent of the long-run response of the sectoral price index occurs in the month of the shock, and then the response equals the long-run response in all months following the shock. The median impulse response of a sectoral price index to an aggregate shock has a very different shape. After an aggregate shock, only 15 percent of the long-run response of the sectoral price index occurs in the month of the shock, and the response gradually approaches the long-run response in the months following the shock. Another way of summarizing the median impulse responses is as follows. The sector-specific component of the sectoral inflation rate is essentially a white noise process, while the aggregate component of the sectoral inflation rate is positively autocorrelated with an autocorrelation coefficient equal to 0.35.[8]

We also compute a simple measure of the speed of the response of a price index to a

[5] The other specifications of the unobservable index model yield a very similar variance decomposition.

[6] See Appendix A for details of the Gibbs sampler.

[7] We follow this approach to presenting the posterior evidence in the rest of the paper.

[8] Regressing the median impulse response of a sectoral inflation rate on its own lag yields a coefficient of 0.35.

given type of shock. Specifically, we compute the absolute response to the shock in the short run divided by the absolute response to the shock in the long run. We take the short run to be between the impact of the shock and five months after the impact of the shock. We take the long run to be between 19 months and 24 months after the impact of the shock. Formally, let β_{nm} denote the response of the price index for sector n to a sector-specific shock m periods after the shock. We define the speed of response of the price index for sector n to sector-specific shocks as follows:

$$\Lambda_n^S \equiv \frac{\frac{1}{6}\sum_{m=0}^{5} \mid \beta_{nm} \mid}{\frac{1}{6}\sum_{m=19}^{24} \mid \beta_{nm} \mid}.$$

Furthermore, let α_{nm} denote the response of the price index for sector n to an aggregate shock m periods after the shock. We define the speed of response of the price index for sector n to aggregate shocks as follows:

$$\Lambda_n^A \equiv \frac{\frac{1}{6}\sum_{m=0}^{5} \mid \alpha_{nm} \mid}{\frac{1}{6}\sum_{m=19}^{24} \mid \alpha_{nm} \mid}.$$

Figure 2 shows the posterior density of Λ_n^S (top panel) and of Λ_n^A (bottom panel). The posterior density takes into account both variation across sectors and parameter uncertainty. The median speed of response of a sectoral price index to sector-specific shocks equals 1.01. The 68 percent probability interval ranges from 0.89 to 1.05. By contrast, the median speed of response of a sectoral price index to aggregate shocks equals 0.41. The 68 percent probability interval ranges from 0.2 to 1.12. The median speed of response of a sectoral price index to sector-specific shocks is much larger than the median speed of response of a sectoral price index to aggregate shocks.[9]

The finding that sectoral price indices respond quickly to sector-specific shocks and slowly to aggregate shocks is robust. Increasing the order of the polynomials $C_n(L)$ from

[9] We can also look at the speed of response to shocks sector by sector. In 76 out of 79 sectors, the median speed of response of a sectoral price index to sector-specific shocks is larger than the median speed of response of the sectoral price index to aggregate shocks. Furthermore, we can construct, in each sector, a probability interval for the speed of response to sector-specific shocks and a probability interval for the speed of response to aggregate shocks. When we construct 68 percent probability intervals, we find that in 43 out of 79 sectors the probability interval for the speed of response to sector-specific shocks lies strictly above the probability interval for the speed of response to aggregate shocks.

six to twelve does not affect the result. Allowing for a bivariate unobservable index does not affect the speed of response to sector-specific shocks. Allowing for a bivariate unobservable index does raise the speed of response to aggregate shocks somewhat, but the speed of response to each aggregate shock remains much smaller than the speed of response to sector-specific shocks.[10]

The findings reported in this section match the findings of Boivin, Giannoni, and Mihov (2007) obtained using a different methodology and different data. Boivin, Giannoni, and Mihov (2007) use a factor augmented vector autoregressive model to study sectoral data published by the Bureau of Economic Analysis on personal consumption expenditure. Boivin, Giannoni, and Mihov (2007) report 85 percent as the average share of the variance in sectoral inflation due to sector-specific shocks. Furthermore, Boivin, Giannoni, and Mihov (2007) find that sectoral price indices respond quickly to sector-specific shocks and slowly to aggregate shocks.

6 Cross-sectional variation in the speed of response

In this section, we study whether the cross-sectional variation in the speed of response of sectoral price indices to shocks is related to sectoral characteristics that we can measure.

6.1 Frequency of price changes

Bils and Klenow (2004) report the monthly frequency of price changes for 350 categories of consumer goods and services, based on data from the BLS for the period 1995-1997. We can match 75 out of our 79 sectors into the categories studied by Bils and Klenow (2004). Furthermore, Nakamura and Steinsson (2008) report the monthly frequency of price changes for 270 categories of consumer goods and services, based on data from the BLS for the period 1998-2005. We can match 77 out of our 79 sectors into the categories studied by Nakamura and Steinsson (2008).

We consider two regressions. First, we regress the speed of response of a sectoral price index to aggregate shocks (Λ_n^A) on the sectoral monthly frequency of price changes from Bils

[10] We have also experimented with alternative measures of the speed of response to shocks. This yielded the same conclusions.

and Klenow (2004). Note that we do not know the speed of response for certain. Instead, we have a posterior density of the speed of response. To account for uncertainty about the regression line, we consider both the posterior density of the regression coefficient and the posterior density of the associated t-statistic. The median regression coefficient is 1.99. The 90 percent probability interval for the regression coefficient ranges from 0.67 to 4.38. The 90 percent probability interval for the associated t-statistic ranges from 0.9 to 5.1, with the median equal to 2.9. See Table 2. When we use the monthly frequency of regular price changes from Nakamura and Steinsson (2008), the regression changes little.[11] The median regression coefficient becomes 1.49. See Table 3.

Second, we regress the speed of response of a sectoral price index to sector-specific shocks (Λ_n^S) on the sectoral monthly frequency of price changes from Bils and Klenow (2004). The median regression coefficient is 0.14. The 90 percent probability interval for the regression coefficient ranges from 0.1 to 0.19. The 90 percent probability interval for the associated t-statistic ranges from 1.3 to 2.4, with the median equal to 1.9. See Table 4. When we use the monthly frequency of regular price changes from Nakamura and Steinsson (2008), the median regression coefficient becomes negative, -0.09. The associated median t-statistic equals -1.3. See Table 5.

It is important to note that the coefficient in the second regression is an order of magnitude smaller than the coefficient in the first regression. Furthermore, when we use the Nakamura and Steinsson (2008) frequencies, we find that sectoral price indices respond faster to sector-specific shocks in sectors in which the frequency of regular price changes is lower.

6.2 Variance of sectoral inflation due to different shocks

Next, we study whether the speed of response of a sectoral price index to a given shock is related to the variance of sectoral inflation due to this shock.

We consider two regressions. First, we regress the speed of response of a sectoral price index to aggregate shocks (Λ_n^A) on the variance of sectoral inflation due to aggregate shocks.

[11] Regular price changes in Nakamura and Steinsson (2008) exclude price changes related to sales and exclude price changes related to product substitutions.

The median regression coefficient is 357.24. The 90 percent probability interval for the regression coefficient ranges from 36.57 to 1227.48. The 90 percent probability interval for the associated t-statistic ranges from 0.2 to 4.6, with the median equal to 1.4. See Table 6.

Second, we regress the speed of response of a sectoral price index to sector-specific shocks (Λ_n^S) on the variance of sectoral inflation due to sector-specific shocks.[12] The median regression coefficient is 7.08. The 90 percent probability interval for the regression coefficient ranges from 5.32 to 9.22. The 90 percent probability interval for the associated t-statistic ranges from 0.76 to 1.2, with the median equal to 0.96. See Table 7.[13]

It is important to note that the coefficient in the second regression is two orders of magnitude smaller than the coefficient in the first regression. In Section 7, we show that this is what the rational inattention model developed in Maćkowiak and Wiederholt (2008a) predicts.

7 Models of price setting

In this section, we study three models of price setting: the Calvo model, the sticky-information model developed in Mankiw and Reis (2002), and the rational inattention model developed in Maćkowiak and Wiederholt (2008a). For each model, we investigate whether the model matches the empirical findings that we report in Sections 5-6. Since some of the empirical findings concern the response of sectoral price indices to sector-specific shocks, we study a multi-sector Calvo model with sector-specific shocks. Similarly, we study a multi-sector sticky-information model with sector-specific shocks and a multi-sector rational inattention model with sector-specific shocks. Section 7.1 describes the common setup. Most of the theoretical results do not depend on the details of the multi-sector setup, as we show. We mainly specify the details of the multi-sector setup to facilitate discussion.

[12] Note that in each of the regressions reported in Section 6.2 we are uncertain about both the regressand and the regressor.

[13] The regressions reported in Section 6 are based on the benchmark specification of the unobservable index model. Increasing the order of the polynomials $C_n(L)$ from six to twelve has little effect on the regression results. Likewise, allowing for a bivariate unobservable index has little effect on the regression results.

7.1 Common setup

Consider an economy with a continuum of sectors of mass one. Sectors are indexed by n. In each sector, there is a continuum of firms of mass one. Firms within a sector are indexed by i. Each firm supplies a differentiated good and sets the price for the good.

The demand for good i in sector n in period t is given by

$$C_{int} = \left(\frac{P_{int}}{P_{nt}}\right)^{-\theta} \left(\frac{P_{nt}}{P_t}\right)^{-\eta} C_t, \tag{2}$$

where P_{int} is the price of good i in sector n in period t, P_{nt} is the price index for sector n, P_t is the aggregate price index and C_t is aggregate composite consumption. The price index for sector n is given by

$$P_{nt} = \left(\int_0^1 P_{int}^{1-\theta} di\right)^{\frac{1}{1-\theta}}, \tag{3}$$

and the aggregate price index is given by

$$P_t = \left(\int_0^1 P_{nt}^{1-\eta} dn\right)^{\frac{1}{1-\eta}}. \tag{4}$$

The demand function (2) with price indices (3) and (4) can be derived from expenditure minimization by households when households have a CES consumption aggregator, where $\theta > 1$ is the elasticity of substitution between goods from the same sector and $\eta > 1$ is the elasticity of substitution between consumption aggregates from different sectors.

Output of firm i in sector n in period t is given by

$$Y_{int} = Z_{nt} L_{int}^{\alpha}, \tag{5}$$

where Z_{nt} is sector-specific productivity and L_{int} is the firm's labor input in period t. The parameter $\alpha \in (0, 1]$ is the elasticity of output with respect to labor input. In every period, firms produce the output that is required to satisfy demand

$$Y_{int} = C_{int}. \tag{6}$$

The real wage rate in period t is assumed to equal $w(C_t)$, where $w : \mathbb{R}_+ \to \mathbb{R}_+$ is a strictly increasing, twice continuously differentiable function.

The demand function (2), the production function (5) and the requirement that output equals demand (6) yield the following expression for profits of firm i in sector n in period t

$$P_{int}\left(\frac{P_{int}}{P_{nt}}\right)^{-\theta}\left(\frac{P_{nt}}{P_t}\right)^{-\eta}C_t - P_t w\left(C_t\right)\left[\frac{\left(\frac{P_{int}}{P_{nt}}\right)^{-\theta}\left(\frac{P_{nt}}{P_t}\right)^{-\eta}C_t}{Z_{nt}}\right]^{\frac{1}{\alpha}}.$$

Dividing by the aggregate price index yields the real profit function

$$f\left(\hat{P}_{int},\hat{P}_{nt},C_t,Z_{nt}\right)=\hat{P}_{int}^{1-\theta}\hat{P}_{nt}^{1-\eta}C_t - w\left(C_t\right)\left[\frac{\hat{P}_{int}^{-\theta}\hat{P}_{nt}^{-\eta}C_t}{Z_{nt}}\right]^{\frac{1}{\alpha}}, \tag{7}$$

where $\hat{P}_{int}=\left(P_{int}/P_{nt}\right)$ is the relative price of good i in sector n, and $\hat{P}_{nt}=\left(P_{nt}/P_t\right)$ is the relative price index for sector n.

In the following, we work with a log-quadratic approximation of the real profit function around the point $\left(1,1,\bar{C},1\right)$. The value $\bar{C}$ is defined by

$$1=\frac{\theta}{\theta-1}w\left(\bar{C}\right)\frac{1}{\alpha}\bar{C}^{\frac{1}{\alpha}-1}. \tag{8}$$

The point $\left(1,1,\bar{C},1\right)$ has a simple interpretation. It is the solution of the model when all firms set the profit-maximizing price and sector-specific productivity equals one in all sectors and all periods.

After the log-quadratic approximation of the real profit function, the profit-maximizing relative price of good i in sector n in period t is given by

$$\hat{p}_{int}^{\diamondsuit}=\frac{\omega+\frac{1-\alpha}{\alpha}}{1+\frac{1-\alpha}{\alpha}\theta}c_t-\frac{1+\frac{1-\alpha}{\alpha}\eta}{1+\frac{1-\alpha}{\alpha}\theta}\hat{p}_{nt}-\frac{\frac{1}{\alpha}}{1+\frac{1-\alpha}{\alpha}\theta}z_{nt},$$

where $\hat{p}_{int}=\ln\left(\hat{P}_{int}\right)$, $c_t=\ln\left(C_t/\bar{C}\right)$, $\hat{p}_{nt}=\ln\left(\hat{P}_{nt}\right)$ and $z_{nt}=\ln\left(Z_{nt}\right)$. Here ω is the elasticity of the real wage with respect to composite consumption at the point $\bar{C}$. Rearranging the last equation yields the following equation for the profit-maximizing price of good i in sector n in period t

$$p_{int}^{\diamondsuit}=\underbrace{p_t+\frac{\omega+\frac{1-\alpha}{\alpha}}{1+\frac{1-\alpha}{\alpha}\theta}c_t}_{p_{int}^{\diamondsuit A}}+\underbrace{\frac{\frac{1-\alpha}{\alpha}\left(\theta-\eta\right)}{1+\frac{1-\alpha}{\alpha}\theta}\hat{p}_{nt}-\frac{\frac{1}{\alpha}}{1+\frac{1-\alpha}{\alpha}\theta}z_{nt}}_{p_{int}^{\diamondsuit S}}, \tag{9}$$

where $p_{int}=\ln\left(P_{int}\right)$, $p_t=\ln\left(P_t\right)$, $c_t=\ln\left(C_t/\bar{C}\right)$, $\hat{p}_{nt}=\ln\left(\hat{P}_{nt}\right)$ and $z_{nt}=\ln\left(Z_{nt}\right)$. The profit-maximizing price is a log-linear function of the aggregate price index, aggregate composite consumption, the relative price index for the sector, and sector-specific productivity.

Note that the log of the profit-maximizing price equals the sum of two components: an aggregate component, $p_{int}^{\diamondsuit A}$, and a sector-specific component, $p_{int}^{\diamondsuit S}$.[14]

Furthermore, after the log-quadratic approximation of the real profit function, the loss in real profits in period t in the case of a suboptimal price equals

$$\frac{\bar{C}(\theta-1)\left(1+\frac{1-\alpha}{\alpha}\theta\right)}{2}\left(p_{int}-p_{int}^{\diamondsuit}\right)^2. \tag{10}$$

See Maćkowiak and Wiederholt (2008a).

In addition to the log-quadratic approximation of the real profit function, we log-linearize the equations for the price indices. Expressing the equation for the sectoral price index (3) in terms of relative prices yields

$$1=\int_0^1 \hat{P}_{int}^{1-\theta}di.$$

Log-linearizing this equation around the point where all relative prices are equal to one and rearranging yields

$$p_{nt}=\int_0^1 p_{int}di, \tag{11}$$

where $p_{nt}=\ln(P_{nt})$. Similarly, we obtain

$$p_t=\int_0^1 p_{nt}dn. \tag{12}$$

We now study three models of price setting. In each model, the profit-maximizing price is given by equation (9), the loss in real profits in period t in the case of a suboptimal price is given by equation (10), and the sectoral price index and the aggregate price index are given by equations (11) and (12), respectively.

7.2 Calvo model

In the Calvo model, a firm can adjust its price with a constant probability in any given period. We are interested in the implications of the Calvo model for the response of sectoral price indices to aggregate shocks and to sector-specific shocks. For this reason, we study a multi-sector version of the Calvo model with sector-specific shocks: the profit-maximizing

[14] Introducing sector-specific shocks in the form of multiplicative demand shocks in (2) instead of multiplicative productivity shocks in (5) yields an equation for the profit-maximizing price that is similar to equation (9). The only difference is the coefficient in front of z_{nt}.

price of good i in sector n in period t is given by equation (9); the price index for sector n in period t is given by equation (11); a firm in sector n can adjust its price with a constant probability λ_n in any given period; and a firm in sector n that can adjust its price in period t chooses the price that minimizes

$$E_t\left[\sum_{s=t}^{\infty}\left[(1-\lambda_n)\,\beta\right]^{s-t}\frac{\bar{C}\,(\theta-1)\left(1+\frac{1-\alpha}{\alpha}\theta\right)}{2}\left(p_{int}-p_{ins}^{\diamondsuit}\right)^2\right]. \tag{13}$$

In this model, the profit-maximizing price of firm i in sector n in period t equals the sum of an aggregate component and a sector-specific component

$$p_{int}^{\diamondsuit}=p_{nt}^{\diamondsuit A}+p_{nt}^{\diamondsuit S}, \tag{14}$$

where the aggregate component is the same for all firms within sector n and the sector-specific component is the same for all firms within sector n. A firm that can adjust its price in sector n in period t sets the price

$$p_{int}^{*}=\left[1-(1-\lambda_n)\,\beta\right]E_t\left[\sum_{s=t}^{\infty}\left[(1-\lambda_n)\,\beta\right]^{s-t}p_{ins}^{\diamondsuit}\right]. \tag{15}$$

The price equals a weighted average of current and future profit-maximizing prices. Since the firms in a sector that can adjust their prices are drawn randomly and all adjusting firms in a sector set the same price, the price index for sector n in period t is given by

$$p_{nt}=(1-\lambda_n)\,p_{nt-1}+\lambda_n p_{int}^{*}. \tag{16}$$

We now study the implications of the Calvo model for the responses of sectoral price indices to aggregate shocks and to sector-specific shocks. We first ask whether the Calvo model matches the median impulse response of a sectoral price index to sector-specific shocks reported in Figure 1. The following proposition provides a formal answer to this question.

Proposition 1 *(Calvo model with sector-specific shocks) Suppose that the profit-maximizing price of firm i in sector n in period t is given by equation (14), that is, the profit-maximizing price equals the sum of an aggregate component and a sector-specific component and each component is the same for all firms in the sector. Furthermore, suppose that the sectoral* *price index for sector n is given by equation (16), where p^*_{int} is given by equation (15). Then, the response of the price index for sector n to a sector-specific shock equals x on impact of the shock and in all periods following the shock if and only if the response of the profit-maximizing price equals: (i)*

$$\frac{\frac{1}{\lambda_n} - (1 - \lambda_n)\beta}{1 - (1 - \lambda_n)\beta} x, \tag{17}$$

on impact of the shock, and (ii) x thereafter.

Proof. See Appendix B. ■

The Calvo model matches the median impulse response of a sectoral price index to sector-specific shocks reported in Figure 1 only if the profit-maximizing price jumps by expression (17) on impact of a sector-specific shock, and then jumps to x in the period following the shock. This yields a response of the sectoral price index to the sector-specific shock equal to x on impact and in all periods following the shock. Note that this result is derived from nothing else but equations (14)-(16).

To illustrate Proposition 1, consider the following three examples. In each example, one period equals one month. Therefore, we set $\beta = 0.99^{1/3}$. First, suppose that $\lambda_n = 1/12$. Then the profit-maximizing response on impact has to overshoot the profit-maximizing response in the next month by a factor of 130. In other words, after a sector-specific shock, the profit-maximizing price has to jump up by $130x$ in the month of the shock, and then jump back to x in the next month to generate a response of x of the sectoral price index on impact and in all months following the shock. Second, suppose that $\lambda_n = 0.087$. This is the monthly frequency of regular price changes reported by Nakamura and Steinsson (2008). Then the profit-maximizing response on impact has to overshoot the profit-maximizing response in the next month by a factor of 120. Third, suppose that $\lambda_n = 0.21$. This is the monthly frequency of price changes reported by Bils and Klenow (2004). Then the profit-maximizing response on impact has to overshoot the profit-maximizing response in the next month by a factor of 20. All three examples are depicted in Figure 3. For the sake of clarity, the impulse response of the sectoral price index in Figure 3 is normalized to one.

So far we have shown that there exists a unique impulse response of the profit-maximizing price to a sector-specific shock which implies that, in the Calvo model, all of the response of the sectoral price index to the sector-specific shock occurs on impact of the shock. One can go a step further and derive the impulse response of sector-specific productivity to a sector-specific shock that yields this impulse response of the profit-maximizing price to a sector-specific shock. When the profit-maximizing price is given by equation (9), the sector-specific component of the profit-maximizing price equals

$$p_{nt}^{\diamond S}=\frac{\frac{1-\alpha}{\alpha}\left(\theta-\eta\right)}{1+\frac{1-\alpha}{\alpha}\theta}\hat{p}_{nt}-\frac{\frac{1}{\alpha}}{1+\frac{1-\alpha}{\alpha}\theta}z_{nt}. \tag{18}$$

Given an impulse response function for the profit-maximizing price and given an impulse response function for the sectoral price index, one can compute the impulse response function for sector-specific productivity that satisfies equation (18). Formally, solving the last equation for sector-specific productivity yields

$$z_{nt}=-\frac{1+\frac{1-\alpha}{\alpha}\theta}{\frac{1}{\alpha}}\left[p_{nt}^{\diamond S}-\frac{\frac{1-\alpha}{\alpha}\left(\theta-\eta\right)}{1+\frac{1-\alpha}{\alpha}\theta}\hat{p}_{nt}\right]. \tag{19}$$

Substituting in the desired impulse response function for the profit-maximizing price (e.g. the blue line with diamonds in Figure 3) and the implied impulse response function for the sectoral price index (the red line with circles in Figure 3) yields the impulse response function for sector-specific productivity. To illustrate this result, Figure 4 shows for the parameter values $\alpha=(2/3)$, $\theta=4$ and $\eta=2$ the impulse response functions for sector-specific productivity that yield the impulse response functions for the profit-maximizing price depicted in Figure 3.

Next, we investigate whether the Calvo model matches the cross-sectional variation in the speed of responses of sectoral price indices to shocks. In particular, we redo the regressions reported in Section 6.1 with data simulated from the Calvo model. We set $\beta=0.99^{1/3}$. For simplicity, we suppose that the aggregate component of the profit-maximizing price, $p_{nt}^{\diamond A}$, follows a random walk with a standard deviation of the innovation equal to 0.48, and we suppose that the sector-specific component of the profit-maximizing price, $p_{nt}^{\diamond S}$, follows a random walk with a standard deviation of the innovation equal to 0.88.[15] We tried a variety of different processes for the profit-maximizing price and we always obtained

[15] The median impulse response of a sectoral price index to aggregate shocks reported in Figure 1 equals 0.48 in the long run. The median impulse response of a sectoral price index to sector-specific shocks reported in Figure 1 equals 0.88 in the long run.

the main result that we point out below. We solve equations (14)-(16) for the price index for sector n setting λ_n equal to a number from the cross-section of the frequency of price changes reported by Bils and Klenow (2004). We then compute the speed of response of the price index for sector n to aggregate shocks and to sector-specific shocks (i.e. Λ_n^A and Λ_n^S). We repeat this procedure for 75 sectors, each time setting λ_n equal to a different number from the cross-section of the frequency of price changes reported by Bils and Klenow (2004). Afterwards, we regress the speed of response to aggregate shocks implied by the Calvo model on the sectoral frequency of price changes, and we regress the speed of response to sector-specific shocks implied by the Calvo model on the sectoral frequency of price changes. The regression coefficient in the first regression is 1.11, which falls within the 90 percent probability interval for this coefficient reported in Section 6.1. See Table 2. The regression coefficient in the second regression is again 1.11, which is an order of magnitude larger than the median value of this coefficient reported in Section 6.1. See Table 4. Thus, the Calvo model predicts far more cross-sectional variation in the speed of response to sector-specific shocks compared to what we see in the data.[16]

The upper panel of Figure 2 has two main features: the median speed of response of a sectoral price index to sector-specific shocks is equal to one, and the distribution of the speed of response of sectoral price indices to sector-specific shocks is tight. We think that the first feature together with Proposition 1 casts doubt on the Calvo model. We think that the second feature also casts doubt on the Calvo model.

7.3 Sticky-information model

In the sticky-information model developed in Mankiw and Reis (2002), a firm can update its pricing plan with a constant probability in any given period. A pricing plan specifies a price path (i.e. a price as a function of time). The difference to the Calvo model is that firms choose a price path instead of a price. We are interested in the implications of this model for the impulse responses of sectoral price indices to aggregate shocks and to sector-specific shocks. Therefore, we study a multi-sector version of this model with sector-specific shocks:

[16] This simulation yields essentially the same results when the monthly frequency of regular price changes from Nakamura and Steinsson (2008) is used. See Table 3 and Table 5.

the profit-maximizing price of good i in sector n in period t is given by equation (9); the price index for sector n in period t is given by equation (11); a firm in sector n can update its pricing plan with a constant probability λ_n in any given period; and a firm that can update its pricing plan in period t chooses the price path that minimizes

$$E_t \left[\sum_{s=t}^{\infty} \beta^{s-t} \frac{\bar{C} (\theta - 1) \left(1 + \frac{1-\alpha}{\alpha}\theta\right)}{2} \left(p_{ins} - p_{ins}^{\diamondsuit}\right)^2 \right]. \tag{20}$$

In this model, the profit-maximizing price of firm i in sector n in period t equals the sum of an aggregate component and a sector-specific component

$$p_{int}^{\diamondsuit} = p_{nt}^{\diamondsuit A} + p_{nt}^{\diamondsuit S}, \tag{21}$$

where each component is common to all firms in the sector. A firm that can update its pricing plan in period t chooses a price for period $s \geq t$ that equals the conditional expectation of the profit-maximizing price in period s

$$p_{ins|t} = E_t \left[p_{ins}^{\diamondsuit} \right]. \tag{22}$$

In period t, a fraction $\lambda_n (1 - \lambda_n)^j$ of firms in sector n last updated their pricing plans j periods ago and these firms set a price equal to $E_{t-j}\left[p_{int}^{\diamondsuit}\right]$. Thus, the price index for sector n in period t equals

$$p_{nt} = \sum_{j=0}^{\infty} \lambda_n (1 - \lambda_n)^j E_{t-j} \left[p_{int}^{\diamondsuit} \right]. \tag{23}$$

Comparing equations (15)-(16) and equations (22)-(23) shows two differences between the Calvo model and the sticky-information model. First, in the Calvo model firms front-load expected future changes in the profit-maximizing price, while in the sticky-information model firms wait with the price adjustment until the expected change in the profit-maximizing price actually occurs. Second, in the Calvo model inflation (i.e. a change in the price level) only comes from the fraction λ_n of firms that can adjust their prices in the current period, while in the sticky-information model inflation may also come from the fraction $(1 - \lambda_n)$ of firms that cannot update their pricing plans in the current period. Mankiw and Reis (2002) show that these two differences have interesting implications for the response of inflation and output to nominal shocks and to (anticipated and unanticipated) disinflations.

We now study the implications of the sticky-information model for the response of sectoral price indices to aggregate shocks and for the response of sectoral price indices to sector-specific shocks. In the sticky-information model, it is easy to derive impulse responses. Firms that have updated their pricing plans since a shock occurred respond perfectly to the shock. All other firms do not respond at all to the shock. The fraction of firms in sector n that have updated their pricing plans over the last τ periods equals

$$\sum_{j=0}^{\tau} \lambda_n (1-\lambda_n)^j = 1-(1-\lambda_n)^{\tau+1}. \tag{24}$$

Thus, the response of the price index for sector n in period t to a shock that occurred τ periods ago equals $\left[1-(1-\lambda_n)^{\tau+1}\right]$ times the response of the profit-maximizing price in sector n in period t to the same shock. This is true for any shock.

To illustrate this result, consider the following example. Suppose that the aggregate component of the profit-maximizing price in sector n follows a random walk with a standard deviation of the innovation equal to 0.48, and the sector-specific component of the profit-maximizing price in sector n follows a random walk with a standard deviation of the innovation equal to 0.88. Furthermore, suppose that firms update their pricing plans on average once a year, as assumed in Mankiw and Reis (2002). In a monthly model, this means $\lambda_n = (1/12)$.[17] Figure 5 shows the impulse responses of the sectoral price index to a sector-specific shock and to an aggregate shock. The impulse responses of the sectoral price index to the two shocks have an identical shape, independent of the standard deviation of the two shocks. The reason is that the impulse responses of the profit-maximizing price to the two shocks have an identical shape. For comparison, Figure 5 also reproduces from Figure 1 the median empirical impulse responses of a sectoral price index to the two shocks.

Next, we ask whether the sticky-information model matches the median impulse response of a sectoral price index to sector-specific shocks reported in Figure 1. The following proposition provides a formal answer to this question.

Proposition 2 *(Sticky-information model with sector-specific shocks) Suppose that the profit-*

[17] In the random walk case, the Calvo model and the sticky-information model are observationally equivalent so long as λ_n is the same in the two models, because the optimal pricing plan in the sticky-information model is a constant price.

maximizing price of firm i in sector n in period t is given by equation (21). Furthermore, suppose that the price index for sector n in period t is given by equation (23). Then, the impulse response of the price index for sector n to a sector-specific shock equals x on impact of the shock and in all periods following the shock if and only if, for all $\tau = 0, 1, 2, \ldots$, the impulse response of the profit-maximizing price τ periods after the shock equals

$$\frac{1}{1-(1-\lambda_n)^{\tau+1}}x. \tag{25}$$

Proof. This follows directly from the sentence below equation (24). ■

The sticky-information model matches the median impulse response of a sectoral price index to sector-specific shocks reported in Figure 1 only if the profit-maximizing price jumps by $(1/\lambda_n)\,x$ on impact of a sector-specific shock, and then decays slowly to x in the periods following the shock. This yields a response of the sectoral price index to the sector-specific shock equal to x on impact and in all periods following the shock. This result is derived from equations (21) and (23).

To illustrate Proposition 2, consider the following example. Suppose that firms update their pricing plans on average once a year, as assumed in Mankiw and Reis (2002). In a monthly model, this means $\lambda_n = (1/12)$. Then the profit-maximizing response on impact of a sector-specific shock has to overshoot the long-run profit-maximizing response by a factor of twelve. See Figure 6.

Finally, from equation (19) one can compute the impulse response function for sector-specific productivity that yields the impulse response function for the profit-maximizing price characterized in Proposition 2.

Note that less overshooting is necessary in the sticky-information model than in the Calvo model for the same value of λ_n, but the extent of overshooting is still large.[18]

[18] Reis (2006) shows that a model with a fixed cost of obtaining perfect information can provide a micro-foundation for the kind of slow diffusion of information modeled by Mankiw and Reis (2002). While the model of Reis (2006) does not distinguish between information concerning aggregate conditions and information concerning sector-specific conditions, one can imagine an extension of the model that would make this distinction.

7.4 Rational inattention as in Maćkowiak and Wiederholt (2008a)

In the rational inattention model developed in Maćkowiak and Wiederholt (2008a), decision-makers in firms decide what to pay attention to. Following Sims (2003), agents' inability to attend perfectly to all available information is modeled as a constraint on information flow. Furthermore, in the model, there is a trade-off between paying attention to aggregate conditions and paying attention to disaggregate conditions. Here, we study a version of this model with sector-specific shocks.

In the model, there are multiple sectors. The profit-maximizing price of good i in sector n in period t is given by equation (9). In period $t = 0$, the decision-maker in a firm chooses the stochastic process for the signals that the decision-maker will receive in the following periods. In each period $t \geq 1$, the decision-maker receives signals and the decision-maker sets a price equal to the conditional expectation of the profit-maximizing price, where the expectation is formed given the sequence of all signals that the decision-maker has received up to that point in time. The price index for sector n in period t is given by equation (11).

To make the results for this model as transparent as possible, we solve analytically for the price index for sector n in the case when both the aggregate component and the sector-specific component of the profit-maximizing price in the sector follow a Gaussian random walk. It is straightforward to solve numerically for the price index for sector n in the case when the profit-maximizing price in the sector follows some other Gaussian process.

Formally, in period zero, the decision-maker in a firm chooses the stochastic process for the signals so as to minimize the expected discounted sum of losses in profits due to suboptimal price setting:

$$\min_{(\sigma_\varepsilon,\sigma_\psi)\in\mathcal{R}_+^2} E\left[\sum_{t=1}^{\infty}\beta^t \frac{\bar{C}\left(\theta-1\right)\left(1+\frac{1-\alpha}{\alpha}\theta\right)}{2}\left(p_{int}-p_{int}^{\diamondsuit}\right)^2\right], \tag{26}$$

subject to: (i) the process for the profit-maximizing price

$$p_{int}^{\diamondsuit} = p_{int}^{\diamondsuit A} + p_{int}^{\diamondsuit S}, \tag{27}$$

with

$$p_{int}^{\diamondsuit A} = p_{int-1}^{\diamondsuit A} + \sigma_A u_t, \tag{28}$$

and

$$p^{\diamondsuit S}_{int} = p^{\diamondsuit S}_{int-1} + \sigma_S v_{nt}, \tag{29}$$

where u_t and v_{nt} follow independent, unit-variance Gaussian white noise processes; (ii) the optimal pricing decision in period t given information in period t

$$p_{int} = E\left[p^{\diamondsuit}_{int}|s^t_{in}\right], \tag{30}$$

where $s^t_{in} = \left(s^0_{in}, s_{in1}, s_{in2}, \ldots, s_{int}\right)$ is the sequence of all signals that the decision-maker has received up to period t; (iii) an assumption concerning the set of signal vectors that the decision-maker can choose from

$$s_{int} = \begin{pmatrix} s^A_{int} \\ s^S_{int} \end{pmatrix} = \begin{pmatrix} p^{\diamondsuit A}_{int} \\ p^{\diamondsuit S}_{int} \end{pmatrix} + \begin{pmatrix} \sigma_\varepsilon \varepsilon_{int} \\ \sigma_\psi \psi_{int} \end{pmatrix}, \tag{31}$$

where ε_{int} and ψ_{int} follow idiosyncratic, unit-variance Gaussian white noise processes that are independent of the u process and the v_n process as well as independent of each other; and (iv) the constraint on information flow

$$\forall t = 1, 2, \ldots : \underbrace{H\left(p^{\diamondsuit A}_{int}|s^{t-1}_{in}\right) - H\left(p^{\diamondsuit A}_{int}|s^t_{in}\right)}_{\kappa^A} + \underbrace{H\left(p^{\diamondsuit S}_{int}|s^{t-1}_{in}\right) - H\left(p^{\diamondsuit S}_{int}|s^t_{in}\right)}_{\kappa^S} \leq \kappa, \tag{32}$$

where $H(X|\mathcal{I})$ denotes the conditional entropy of X given the information set $\mathcal{I}$.

Conditional entropy of X given $\mathcal{I}$ is a measure of conditional uncertainty of X given $\mathcal{I}$. The difference $H\left(X_t|s^{t-1}_{in}\right) - H\left(X_t|s^t_{in}\right)$ is a measure of the reduction in uncertainty about X_t that is due to the new information received in period t. Following Sims (2003, Section 5), we use this difference to quantify the amount of information that the signal received in period t contains about X_t. The information flow constraint (32) states that, in every period, the flow of information concerning aggregate conditions, denoted κ^A, plus the flow of information concerning sector-specific conditions, denoted κ^S, cannot exceed the value κ. In other words, the overall flow of information is limited.

To abstract from transitional dynamics in conditional variances, we make a simplifying assumption. We assume that, after the decision-maker has chosen the stochastic process for the signals in period zero, the decision-maker receives information in period zero such that the conditional variances of $p^{\diamondsuit A}_{in1}$ and $p^{\diamondsuit S}_{in1}$ given information in period zero equal the steady-state values of the conditional variances of $p_{int}^{\diamondsuit A}$ and $p_{int}^{\diamondsuit S}$ given information in period $t-1$. This simplifies computing the solution to problem (26)-(32). In addition, we assume that the decision-maker in a firm can choose the overall attention devoted to the price setting decision, κ, facing the cost function $c(\kappa) = \phi\kappa$ where $\phi > 0$ is the (per-period) real marginal cost of attention.[19] The idea is that the decision-maker in a firm will increase the attention devoted to the price setting decision so long as the marginal value of doing so is high. Formally, we add the term $[\beta/(1-\beta)]\, c(\kappa)$ to objective (26) and we let the decision-maker choose κ. Finally, it is worth pointing out that the assumption that the noise terms in equation (31) are independent captures the idea that paying attention to aggregate conditions and paying attention to sector-specific conditions are separate activities. In Section 8.2 of Maćkowiak and Wiederholt (2008a), we discuss this assumption in detail and we relax the assumption.

Consider first the price setting behavior and the expected loss in profits due to suboptimal price setting for a given allocation of attention (i.e. for a given pair κ^A and κ^S). We show in Appendix C that the price setting behavior for a given allocation of attention is given by

$$\begin{aligned} p_{int}^{\diamondsuit} - p_{int} \;=\; & \sum_{l=0}^{\infty}\left[\left(2^{-2\kappa^A}\right)^{l+1}\sigma_A u_{t-l} - \left(2^{-2\kappa^A}\right)^{l}\left(2^{-\kappa^A}\right)\sigma_A \varepsilon_{int-l}\right] \\ & + \sum_{l=0}^{\infty}\left[\left(2^{-2\kappa^S}\right)^{l+1}\sigma_S v_{nt-l} - \left(2^{-2\kappa^S}\right)^{l}\left(2^{-\kappa^S}\right)\sigma_S \psi_{int-l}\right], \end{aligned} \tag{33}$$

where $p_{int}^{\diamondsuit} - p_{int}$ is the difference between the profit-maximizing price and the actual price. The speed at which the gap $p_{int}^{\diamondsuit} - p_{int}$ closes after an innovation in the aggregate component, u_t, depends on the attention allocated to aggregate conditions, κ^A. The speed at which the gap closes after an innovation in the sector-specific component, v_{nt}, depends on the attention allocated to sector-specific conditions, κ^S. Thus, if the decision-maker pays more attention to sector-specific conditions than to aggregate conditions ($\kappa^S > \kappa^A$), the price set by firm i in sector n responds faster to sector-specific shocks than to aggregate shocks.

[19] The cost $c(\kappa)$ could be an opportunity cost (e.g. spending less time on some other activity at work or at home) or a monetary cost (e.g. a wage payment).

The value of expression (26) for a given allocation of attention equals

$$\frac{\beta}{1-\beta}\frac{\bar{C}(\theta-1)\left(1+\frac{1-\alpha}{\alpha}\theta\right)}{2}\left[\frac{\sigma_A^2}{2^{2\kappa^A}-1}+\frac{\sigma_S^2}{2^{2\kappa^S}-1}\right]. \tag{34}$$

See again Appendix C. It is now straightforward to derive the optimal allocation of attention. The decision-maker in a firm equates the marginal value of attending to aggregate conditions to the marginal cost of attention

$$\frac{\bar{C}(\theta-1)\left(1+\frac{1-\alpha}{\alpha}\theta\right)}{2}\sigma_A^2\frac{2^{2\kappa^A}}{\left(2^{2\kappa^A}-1\right)^2}2\ln(2)=\phi, \tag{35}$$

and the decision-maker equates the marginal value of attending to sector-specific conditions to the marginal cost of attention

$$\frac{\bar{C}(\theta-1)\left(1+\frac{1-\alpha}{\alpha}\theta\right)}{2}\sigma_S^2\frac{2^{2\kappa^S}}{\left(2^{2\kappa^S}-1\right)^2}2\ln(2)=\phi. \tag{36}$$

Thus, the optimal allocation of attention is given by

$$2^{\kappa^A}-\frac{1}{2^{\kappa^A}}=\sigma_A\sqrt{\frac{\bar{C}(\theta-1)\left(1+\frac{1-\alpha}{\alpha}\theta\right)}{\phi}\ln(2)}, \tag{37}$$

and

$$2^{\kappa^S}-\frac{1}{2^{\kappa^S}}=\sigma_S\sqrt{\frac{\bar{C}(\theta-1)\left(1+\frac{1-\alpha}{\alpha}\theta\right)}{\phi}\ln(2)}. \tag{38}$$

Hence, the model predicts that the attention allocated to a given shock depends on the standard deviation of that type of shock. Therefore, the speed at which the gap between the profit-maximizing price and the actual price closes after a given shock depends on the standard deviation of that type of shock. Furthermore, dividing equation (38) by equation (37) yields

$$\frac{2^{\kappa^S}-2^{-\kappa^S}}{2^{\kappa^A}-2^{-\kappa^A}}=\frac{\sigma_S}{\sigma_A}. \tag{39}$$

Hence, the model makes the following prediction. When the sector-specific component of the profit-maximizing price is more volatile than the aggregate component of the profit-maximizing price ($\sigma_S > \sigma_A$), the decision-maker pays more attention to sector-specific conditions than to aggregate conditions ($\kappa^S > \kappa^A$). The price set by the decision-maker then responds faster to sector-specific shocks than to aggregate shocks.

Finally, integrating equation (33) over all i and using equations (11) and (27)-(29) yields an equation for the sectoral price level. This equation for the sectoral price level implies an equation for the sectoral inflation rate that has the form of equation (1), which is the equation that we estimate.[20]

We now ask whether the rational inattention model presented above in which the sector-specific component of the profit-maximizing price follows a random walk can yield an impulse response of the price index for sector n to sector-specific shocks that falls on impact within the 90 percent probability band in the top panel of Figure 1. The lower 90 percent probability band suggests that the sectoral price index underreacts its long-run response by 20 percent on impact. A value of κ^S equal to 1.15 bits yields a 20 percent underreaction on impact. See equation (33). Therefore, we suppose that κ^S equals 1.15 and we consider the implied value for κ^A for plausible values for the ratio (σ_S/σ_A). When the ratio (σ_S/σ_A) equals three, κ^A must equal 0.42. See equation (39). When the ratio (σ_S/σ_A) equals two-and-a-half, κ^A must equal 0.5. When the ratio (σ_S/σ_A) equals two, κ^A must equal 0.62. Maćkowiak and Wiederholt (2008b) show that values for κ^A in this range yield sizable real effects of monetary policy shocks in a DSGE model with rational inattention.[21]

Next, we investigate whether the rational inattention model presented above matches the cross-sectional variation in the speed of response of sectoral price indices to shocks. In particular, we redo the regressions reported in Section 6.2 with data simulated from the rational inattention model. For simplicity, we suppose that the aggregate component of the profit-maximizing price, $p_{nt}^{\diamondsuit A}$, follows a random walk with a standard deviation of the innovation σ_A, and we suppose that the sector-specific component of the profit-maximizing price, $p_{nt}^{\diamondsuit S}$, follows a random walk with a standard deviation of the innovation σ_S. We construct a cross-section of the pair of numbers σ_A and σ_S, using the empirical standard deviation of sectoral inflation due to aggregate shocks as a measure of σ_A and using the empirical standard deviation of sectoral inflation due to sector-specific shocks as a measure of σ_S. Furthermore, we select a value for the marginal cost of information flow, ϕ, such

[20] In the Calvo model, the sectoral price level is given by equations (16) and (15). In the sticky-information model, the sectoral price level is given by equation (23). When the profit-maximizing price (14) follows a Gaussian process with a time-invariant moving-average representation, the equation for the sectoral inflation rate has the form of equation (1).

[21] In the benchmark calibration of Maćkowiak and Wiederholt (2008b, Section 8.3), decision-makers in firms allocate 1.35 bits per quarter to tracking aggregate shocks.

that the median information flow allocated to sector-specific conditions equals 1.15 and the median information flow allocated to aggregate conditions equals 0.4. We solve for the price index for sector n in the rational inattention model repeatedly, each time setting σ_A and σ_S equal to a different pair of numbers from the empirical cross-section. We then compute the speed of response of the price index for sector n to aggregate shocks and to sector-specific shocks (i.e. Λ_n^A and Λ_n^S). Afterwards, we regress the speed of response to aggregate shocks implied by the rational inattention model on the empirical variance of sectoral inflation due to aggregate shocks. The regression coefficient is 263.65, which falls within the 90 percent probability interval for this coefficient reported in Section 6.2. See Table 6. Next, we regress the speed of response to sector-specific shocks implied by the rational inattention model on the empirical variance of sectoral inflation due to sector-specific shocks. The regression coefficient is 9.87, which falls just outside the 90 percent probability interval for this coefficient reported in Section 6.2. See Table 7.

It is straightforward to see why the rational inattention model matches the finding that the coefficient in the second regression is two orders of magnitude smaller than the regression coefficient in the first regression. In the rational inattention model, the speed of response of prices to a given shock is concave in the variance of the shock. When decision-makers in firms already pay close attention to sector-specific shocks, increasing the variance of sector-specific shocks does not make much difference. By contrast, when decision-makers in firms pay a moderate amount of attention to aggregate shocks, increasing the variance of aggregate shocks makes a difference.

Boivin, Giannoni, and Mihov (2007) find a positive relationship between the speed of response of sectoral price indices to aggregate shocks and the variance of sectoral inflation due to sector-specific shocks. We obtain the same finding. This finding is consistent with the rational inattention model described above. In fact, there are two possible explanations. One possible explanation is that there is a positive relationship across sectors between the variance of the profit-maximizing price due to aggregate shocks and the variance of the profit-maximizing price due to sector-specific shocks. Another possible explanation has to do with the fact that parameters in objective (26) affect both the attention allocated to aggregate conditions and the attention allocated to sector-specific conditions. For ease of exposition, suppose that $\alpha = 1$. When $\alpha = 1$, increasing the price elasticity of demand, θ, raises the loss in profits in the case of a price setting mistake without changing the process for the profit-maximizing price. See equations (26) and (9). In sectors with a high price elasticity of demand, firms will then pay more attention to both aggregate conditions and sector-specific conditions to avoid mistakes, implying that prices in those sectors respond faster to both aggregate shocks and sector-specific shocks. This raises the variance of sectoral inflation due to sector-specific shocks.

8 Conclusions

We look at disaggregate data in order to evaluate models that predict a slow response of prices to aggregate shocks. We estimate a statistical model using sectoral price data. In the median sector, 100 percent of the long-run response of the sectoral price index to a sector-specific shock occurs in the month of the shock. The Calvo model and the sticky-information model match this finding only under extreme assumptions concerning the profit-maximizing price. By contrast, the rational inattention model matches this finding without an extreme assumption concerning the profit-maximizing price. Furthermore, we find little variation across sectors in the speed of response of sectoral price indices to sector-specific shocks. The rational inattention model matches this finding, while the Calvo model predicts far too much cross-sectional variation in the speed of response to sector-specific shocks.

We think that the Calvo model is a useful tool for someone who wants an easy-to-solve model for analysis of the effects of slow nominal adjustment on macroeconomic fluctuations. We feel less comfortable when the Calvo model is used for normative analysis of macroeconomic policy. In this respect, we think that this paper's findings provide a warning. We find that the Calvo model does a poor job matching the response of prices to market-specific shocks. Normative analysis with the Calvo model presupposes that the Calvo model, as policy changes, will continue to match macro data well. We are skeptical that this assumption is valid when one contemplates policy changes that will make the macroeconomic environment more volatile.

In the future, it will be interesting to study whether the menu cost model matches the empirical findings of this paper. Furthermore, it will be interesting to see whether the empirical findings of this paper extend to countries other than the United States.

A Econometric details

We consider the unobservable index model

$$\tilde{\pi}_{nt} = a_n' (L) u_t + \tilde{\pi}_{nt}^S, \tag{40}$$

$$\tilde{\pi}_{nt}^S = c_n (L) \tilde{\pi}_{nt}^S + \epsilon_{nt}, \tag{41}$$

where: (i) $\tilde{\pi}_{nt}$ is the zero mean, unit-variance, month-on-month inflation rate in sector n in period t, $n = 1, ..., N$, $t = 1, ..., T$; (ii) $u_t = (u_{1t}; ...; u_{Kt})'$ is a $K \times 1$ unobservable index satisfying $u_{kt} \sim N(0, 1)$ for each $k = 1, ..., K$; (iii) $a_n (L)$ is a $K \times 1$ polynomial in the lag operator of order M; (iv) $c_n (L)$ is a polynomial in the lag operator of order S satisfying $c_{n0} = 0$; and (v) $\epsilon_{nt} \sim N(0, \sigma_n^2)$, for each n. The data, described in Section 3, are monthly sectoral price indices. We perform seasonal adjustment of the log of the price index in each sector, and we construct the month-on-month inflation rate. We subtract the mean from each sector's inflation rate, and we divide each sector's inflation rate by its standard deviation. Since spikes in individual inflation rates were picked up by the unobservable index in preliminary estimation, we eliminated a few outliers from the dataset. For each sector, we replaced the observations falling outside four times the standard deviation of the inflation rate by the mean inflation rate over the rest of the sample period.

When we assemble N equations, with each equation having the form of equation (40), we can write

$$\tilde{\pi}_t = a (L) u_t + \tilde{\pi}_t^S,$$

where $\tilde{\pi}_t$ and $\tilde{\pi}_t^S$ and vectors of length N, and each matrix appearing in the polynomial $a (L)$ has size $N \times K$. Let $\tilde{a}_0$ denote the $K \times K$ matrix consisting of the first K rows of a_0. We assume that $\tilde{a}_0$ is a lower triangular matrix with strictly positive entries on the main diagonal. This is a sufficient condition to estimate u_t uniquely. See Geweke and Zhou (1996). Consider the case when $K = 1$. In order to ensure that the unobservable index is positively correlated with the CPI inflation rate, we add the CPI inflation rate to our dataset, and we order the CPI inflation rate first. As it turns out, the model attributes almost all of the variance in the CPI inflation rate to the unobservable index. In the case when $K = 2$, we order second (after the CPI inflation rate) the variable best explained by the first principle component of the normalized inflation rates.

We use the Gibbs sampler to sample from the joint posterior density of the unobservable index $\{u_t\}$ and the model's parameters ($a_n(L)$, $c_n(L)$, and σ_n, for each n). Given a Monte Carlo draw of the model's parameters, we sample from the conditional posterior density of the unobservable index given the model's parameters. This step is described in Section A.1. Afterwards, given a Monte Carlo draw of the unobservable index, we sample from the conditional posterior density of the model's parameters given the unobservable index. This step is described in Section A.2.

A.1 Sampling unobservable index given parameters

We follow Carter and Kohn (1994) and Kim and Nelson (1999) in sampling from the conditional posterior density of the unobservable index given the model's parameters ($a_n(L)$, $c_n(L)$, and σ_n, for each n). We begin by writing the model (40)-(41) in state space form. Equations (40)-(41) imply that

$$\tilde{\pi}^*_{nt} = g'_n(L)\, u_t + \epsilon_{nt}, \tag{42}$$

where $\tilde{\pi}^*_{nt} = (1 - c_n(L))\, \tilde{\pi}_{nt}$, and $g_n(L) = (1 - c_n(L))\, a_n(L)$, for each n. If we think of the coefficients appearing in the polynomial $g_n(L)$ as forming a vector G_n, and if we define a vector F_t as $F_t = \left(u_t; u_{t-1}; ...; u_{t-(M+S)}\right)$, we can write equation (42) as

$$\tilde{\pi}^*_{nt} = G'_n F_t + \epsilon_{nt}.$$

Note that vectors G_n and F_t have length $l = K(M + S + 1)$. When we assemble N equations of this form, we arrive at the measurement equation:

$$\tilde{\pi}^*_t = G F_t + \epsilon_t, \tag{43}$$

where $\tilde{\pi}^*_t$ and ϵ_t are vectors of length N, and G is a matrix of size $N \times l$. Let R denote the variance-covariance matrix of ϵ_t. Note that R is an $N \times N$ diagonal matrix with diagonal elements $\sigma^2_1, ..., \sigma^2_N$. The state equation is

$$F_{t+1} = J F_t + \tilde{u}_{t+1}, \tag{44}$$

where J is an $l \times l$ matrix defined as

$$J = \begin{bmatrix} 0_{K\times K} & 0_{K\times K(M+S)} \\ \mathcal{I}_{K(M+S)} & 0_{K(M+S)\times K} \end{bmatrix},$$

and $\tilde{u}_{t+1}$ is a vector of length l defined as $\tilde{u}_{t+1} = \left(u_{t+1}; 0_{K(M+S)\times 1}\right)$. Let Q denote the variance-covariance matrix of $\tilde{u}_{t+1}$. Note that Q is an $l \times l$ diagonal matrix, with the first K diagonal elements equal to unity and all other elements equal to zero.

We run Kalman filter iterations from period $t = 1$ to period $t = T$ to obtain

$$F_{t|t} = F_{t|t-1} + P_{t|t-1}G' \left(GP_{t|t-1}G' + R\right)^{-1} \left(\tilde{\pi}_t^* - GF_{t|t-1}\right), \tag{45}$$

$$P_{t|t} = P_{t|t-1} - P_{t|t-1}G' \left(GP_{t|t-1}G' + R\right)^{-1} GP_{t|t-1}, \tag{46}$$

$$F_{t+1|t} = JF_{t|t},$$

and

$$P_{t+1|t} = JP_{t|t}J' + Q.$$

We use the unconditional density of the state vector to initialize the Kalman filter. We set $F_{1|0}$ equal to a zero vector, and we set

$$vec\left(P_{1|0}\right) = \left[\mathcal{I}_{l^2} - (J \otimes J)\right]^{-1} vec\left(Q\right).$$

Next, we sample from the probability density function of F_T given the data until period T. This is a Gaussian density with mean $F_{T|T}$ and variance-covariance matrix $P_{T|T}$. Subsequently, for each $t = T-1, T-2, ..., 1$, we sample from the probability density function of F_t given the data until period t and F_{t+1}. This is a Gaussian density with mean $F_{t|t,F_{t+1}}$ and variance-covariance matrix $P_{t|t,F_{t+1}}$, where

$$F_{t|t,F_{t+1}} = F_{t|t} + P_{t|t}J'P_{t+1|t}^{-1}\left(F_{t+1} - F_{t+1|t}\right),$$

and

$$P_{t|t,F_{t+1}} = P_{t|t} - P_{t|t}J'P_{t+1|t}^{-1}JP_{t|t}.$$

In practice, since in our model Q is a singular matrix, we must modify the densities we sample from slightly. Let Q^* be the matrix of size $l^* \times l^*$ (where $l^* < l$) obtained after removing from Q each row that contains nothing but zeros. Furthermore, let F_{t+1}^* be the corresponding $l^* \times 1$ state vector, and let J^* be the corresponding $l^* \times l$ matrix in the state equation. For each $t = T-1, T-2, ..., 1$, we sample from the probability density function of F_t given the data until period t and F^*_{t+1}. This is a Gaussian density with mean $F_{t|t,F^*_{t+1}}$ and variance-covariance matrix $P_{t|t,F^*_{t+1}}$, where

$$F_{t|t,F^*_{t+1}} = F_{t|t} + P_{t|t}J^{*\prime}\left(J^*P_{t|t}J^{*\prime} + Q^*\right)^{-1}\left(F^*_{t+1} - J^*F_{t|t}\right),$$

and

$$P_{t|t,F^*_{t+1}} = P_{t|t} - P_{t|t}J^{*\prime}\left(J^*P_{t|t}J^{*\prime} + Q^*\right)^{-1}J^*P_{t|t}.$$

In our model, it turns out that $F_{t|t,F^*_{t+1}} = F_{t|t}$ and $P_{t|t,F^*_{t+1}} = P_{t|t}$, so that the Gaussian probability density functions we sample from are characterized fully by expressions (45) and (46). The reason is that the unobservable index in our model follows a vector white noise process.

A.2 Sampling parameters given unobservable index

We follow Chib and Greenberg (1994) in sampling from the conditional posterior density of the model's parameters ($a_n(L)$, $c_n(L)$, and σ_n, for each n) given the unobservable index. Taking the unobservable index as given and collecting N pairs of equations (40)-(41) yields N independent Gaussian regressions, each regression with autoregressive errors of order S. Therefore, given the unobservable index, we can analyze each pair of equations (40)-(41) using the results of Chib and Greenberg (1994) on Gaussian regression with autoregressive errors. Note that, like Chib and Greenberg (1994), we use the full likelihood function without conditioning on initial observations.

In the rest of this section, for simplicity, we drop the subscript n, we think of the coefficients appearing in the polynomial $a(L)$ as forming a vector θ, and we think of the coefficients appearing in the polynomial $c(L)$ as forming a vector ϕ. Furthermore, we define a vector x_t according to $x_t = (u_t; u_{t-1}; ...; u_{t-M})$, and we define a vector $\tilde{x}_t$ according to $\tilde{x}_t = \left(\tilde{\pi}^S_{t-1}; ...; \tilde{\pi}^S_{t-S}\right)$. Then, we can write the model (40)-(41) as

$$\tilde{\pi}_t = x_t'\theta + \tilde{\pi}^S_t,$$

$$\tilde{\pi}^S_t = \tilde{x}_t'\phi + \epsilon_t,$$

and, in matrix notation,

$$\tilde{\Pi} = X\theta + \tilde{\Pi}^S,$$

$$\tilde{\Pi}^S = \tilde{X}\phi + \epsilon.$$

Note that vectors x_t and θ have length $K(M+1)$, vectors $\tilde{x}_t$ and ϕ have length S, vectors $\tilde{\Pi}$, $\tilde{\Pi}^S$, and ϵ have length T, matrix X has size $T \times K(M+1)$, and matrix $\tilde{X}$ has size $T \times S$.

Let Φ be an $S \times S$ matrix satisfying

$$\Phi = \begin{bmatrix} \tilde{\phi}' & \phi_S \\ \mathcal{I}_{S-1} & 0_{(S-1)\times 1} \end{bmatrix},$$

where $\tilde{\phi} = (\phi_1; ...; \phi_{S-1})$ is a vector of length $S-1$. We assume the following joint prior density of θ, ϕ and σ^2:

$$[\theta, \phi, \sigma^2] = [\theta][\phi][\sigma^2] = \mathcal{N}(\theta_0, \Theta_0^{-1})\left[\mathcal{N}(\phi_0, \Phi_0^{-1})\mathcal{I}(\Phi)\right]\mathcal{IG}[\nu_0/2, \delta_0/2],$$

where $\mathcal{N}$ denotes the normal density, $\mathcal{IG}$ denotes the inverse gamma density, and $\mathcal{I}(\Phi)$ is an indicator function equal to one when all eigenvalues of Φ are less than one in modulus.

The following notation is useful. We partition $\tilde{\Pi}$ so that $\tilde{\Pi} = \left(\tilde{\Pi}_1; \tilde{\Pi}_2\right)$, where $\tilde{\Pi}_1$ has length S and $\tilde{\Pi}_2$ has length $T-S$. Analogously, we partition X so that $X = (X_1; X_2)$, where X_1 has size $S \times K(M+1)$ and X_2 has size $T - S \times K(M+1)$. We define Σ to be an $S \times S$ matrix satisfying

$$\Sigma = \Phi\Sigma\Phi' + \left(1; 0_{(S-1)\times 1}\right)\left(1; 0_{(S-1)\times 1}\right)',$$

where $\left(1; 0_{(S-1)\times 1}\right)$ is a vector of length S. That is,

$$vec(\Sigma) = \left[\mathcal{I}_{S^2} - (\Phi \otimes \Phi)\right]^{-1} vec\left[\left(1; 0_{(S-1)\times 1}\right)\left(1; 0_{(S-1)\times 1}\right)'\right].$$

Let $chol(\Sigma)$ denote the lower triangular Choleski square root of Σ. We define $\tilde{\Pi}^* = \left(\tilde{\Pi}_1^*; \tilde{\Pi}_2^*\right)$ and $X^* = (X_1^*; X_2^*)$, where: (i) $\tilde{\Pi}_1^* = [chol(\Sigma)]^{-1}\tilde{\Pi}_1$; (ii) $X_1^* = [chol(\Sigma)]^{-1} X_1$; (iii) $\tilde{\Pi}_2^*$ is a vector of length $T-S$ with t'th row given by $[1 - c(L)]\tilde{\pi}_t$; and (iv) X_2^* is a matrix of size $T - S \times K(M+1)$ with t'th row given by $[1 - c(L)]x_t$. We let $e_t = \tilde{\pi}_t - x_t'\theta$, for each $t = S+1, ..., T$; we let e be a vector of length $T-S$ satisfying $e = (e_{S+1}; ...; e_T)$; and we let E be a matrix of size $T - S \times S$ with t'th row given by $(e_{t-1}, ..., e_{t-S})$, for each $t \geq S+1$.

Chib and Greenberg (1994) derive the following conditional posterior densities:

$$\theta \mid \phi, \sigma^2 \sim \mathcal{N}\left[\Theta_T^{-1}\left(\Theta_0\theta_0 + \sigma^{-2}X^{*\prime}\tilde{\Pi}^*\right), \Theta_T^{-1}\right], \tag{47}$$

$$\sigma^2 \mid \phi, \theta \sim \mathcal{IG}\left[\left(\frac{\nu_0+T}{2}\right), \frac{\delta_0+\left(\tilde{\Pi}^*-X^*\theta\right)'\left(\tilde{\Pi}^*-X^*\theta\right)}{2}\right], \tag{48}$$

$$\phi \mid \theta, \sigma^2 \propto \Psi\left(\phi\right) \times \mathcal{N}\left[\Phi_T^{-1}\left(\Phi_0\phi_0+\sigma^{-2}E'e\right), \Phi_T^{-1}\right]\mathcal{I}\left(\Phi\right), \tag{49}$$

where

$$\Theta_T = \Theta_0 + \sigma^{-2}X^{*\prime}X^*,$$

$$\Phi_T = \Phi_0 + \sigma^{-2}E'E,$$

$$\Psi\left(\phi\right) = \mid \Sigma\left(\phi\right) \mid^{-1/2} \exp\left[-\frac{1}{2\sigma^2}\left(\tilde{\Pi}_1 - X_1\theta\right)'\Sigma\left(\phi\right)^{-1}\left(\tilde{\Pi}_1 - X_1\theta\right)\right],$$

and ϕ in brackets in the last expression reminds us that Σ depends on ϕ.

The conditional density of θ and the conditional density of σ^2 are standard, but the conditional density of ϕ cannot be sampled from directly. Following Chib and Greenberg (1994), we sample from the density of ϕ using the Metropolis-Hastings algorithm. See also Otrok and Whiteman (1998). At each iteration j of the Gibbs sampler, we generate a candidate draw ϕ^* from the density $\mathcal{N}\left[\Phi_T^{-1}\left(\Phi_0\phi_0+\sigma^{-2}E'e\right), \Phi_T^{-1}\right]\mathcal{I}\left(\Phi\right)$. We then set $\phi^j = \phi^*$ with probability

$$\rho = \min\left[\frac{\Psi\left(\phi^*\right)}{\Psi\left(\phi^{j-1}\right)}, 1\right],$$

and we set $\phi^j = \phi^{j-1}$ with probability $1-\rho$.

We choose values for the prior hyperparameters following the Minnesota prior. We set $\theta_0 = 0_{K(M+1)\times 1}$. We assume that Θ_0 is a diagonal matrix of size $K\left(M+1\right) \times K\left(M+1\right)$. Furthermore, we assume that each of the first K entries on the main diagonal of Θ_0 equals 1, each of the subsequent K entries also equals 1, each of the K entries after that equals 4, and so on until the last K entries on the main diagonal of Θ_0 equal M^2. This means that the prior mean on each loading equals zero, the standard deviation of the prior on a contemporaneous loading equals 1, and the standard deviation of the prior on a loading on an m'th lag of the observable index equals $(1/m)$. We set $\phi_0 = 0_{S\times 1}$. We assume that Φ_0 is a diagonal matrix of size $S \times S$ with entries on the main diagonal given by $\left(s/0.2\right)^2$, $s = 1, ..., S$. This means that the prior mean on each autoregressive coefficient equals zero, the standard deviation of the prior on a coefficient on a first lag equals 0.2, and the standard deviation of the prior on a coefficient on an s'th lag equals $(0.2/s)$. We set

$\nu_0 = 4$ and $\delta_0 = 0.1$. In the case when $K = 2$, we use a tighter prior. We set the standard deviation of the prior on a loading on an m'th lag of the observable index equal to $\left(1/m^{1.5}\right)$. In the autoregressive component of the model, we set the standard deviation of the prior on a coefficient on an s'th lag equal to $\left(0.2/s^{1.5}\right)$.

To initialize the Gibbs sampler, we regress the data on, alternatively, the first K principal components of the data or randomly generated K "indices" (current and lagged). We use the ordinary least squares coefficients as initial values for the model's parameters, $\left\{\theta_n, \phi_n, \sigma_n^2\right\}^{j=0}$, $n = 1, ..., N$. Each j'th iteration of the Gibbs sampler proceeds as follows. Given $\left\{\theta_n, \phi_n, \sigma_n^2\right\}^{j-1}$, $n = 1, ..., N$, we make a draw of the unobservable index $\{u_t\}^j$, as described in Section A.1. Next, given $\{u_t\}^j$, we make a draw of the model's parameters $\left\{\theta_n, \phi_n, \sigma_n^2\right\}^j$, $n = 1, ..., N$, as described in this section above. Here, we begin by drawing θ_n^j given $\left\{\phi_n, \sigma_n^2\right\}^{j-1}$ from density (47), afterwards we draw $\left(\sigma_n^2\right)^j$ given θ_n^j and ϕ_n^{j-1} from density (48), and finally we draw ϕ_n^j given $\left\{\theta_n, \sigma_n^2\right\}^j$ from density (49). We make 20000 draws. We discard the initial 5000 draws. We save every second draw out of the remaining 15000 draws. This gives us 7500 draws from the posterior density.

A.3 Convergence of the Gibbs sampler

We use formal and informal diagnostics to assess convergence of the Gibbs sampler. We compute two formal convergence diagnostics: the Raftery-Lewis measure of the number of draws required to achieve a certain precision of the sampler (see Raftery and Lewis, 1992); and the Geweke relative numerical efficiency indicator (see Geweke, 1992). We specify the parameters for the Raftery-Lewis diagnostic as follows: quantile = 0.025; desired accuracy = 0.0125; required probability of attaining the desired accuracy = 0.95. Note that, since we compute the Raftery-Lewis diagnostic and the Geweke diagnostic for each parameter in the model, we obtain a cross-section of the Raftery-Lewis diagnostics and a cross-section of the Geweke diagnostics across the parameters. Table A.1 summarizes the cross-section of the Raftery-Lewis diagnostics across the parameters. Table A.2 summarizes the cross-section of the Geweke diagnostics across the parameters. Both tables refer to the benchmark specification. For ease of exposition, in both tables we divide the parameters into the loadings (the θ's), the autoregressive parameters (the ϕ's), and the standard deviations (the σ's). Consider Table A.1. For 99 percent of the parameters, the Raftery and Lewis diagnostic suggests that one should make 4278 draws or fewer for the sampler to be precise. 4278 is much less than the 20000 draws that we actually make. Only for two parameters the Raftery and Lewis diagnostic suggests that one should make many more draws than 20000. Both of these parameters are autoregressive parameters in a single sector, "Tires". In this sector, the acceptance rate in the Metropolis-Hastings step of the Gibbs sampler turns out to be relatively low. Next, consider Table A.2. With only few exceptions, the Geweke indicator lies well below 20, which is the value of the Geweke indicator considered as small enough to signal good mixing properties of the sampler. See, for example, Primiceri (2005). We also monitor convergence of the Gibbs sampler informally, by plotting the evolution of draws for a set of randomly selected parameters. Furthermore, when we initialized the Gibbs sampler at random points, we obtained very similar results. We also ran the Gibbs sampler on artificial data, and we found that the estimated unobservable index and the estimated parameters were very close to the true ones. All of this give us confidence that the Markov chain that we use has converged to its ergodic distribution.

B Proof of Proposition 1

First, let xv_{nt} denote the long-run response of the price index for sector n to a sector-specific shock of size v_{nt} in period t. Second, the price index for sector n in period t satisfies $p_{nt} = p_{nt-1} + xv_{nt}$ if and only if the price set by adjusting firms in sector n in period t satisfies

$$p^*_{int} = p_{nt-1} + \frac{1}{\lambda_n} xv_{nt}. \tag{50}$$

This follows from equation (16). Third, the sectoral price index satisfies $p_{nt+\tau} = p_{nt}$ for all $\tau \geq 1$ if and only if the price set by adjusting firms satisfies

$$\forall \tau \geq 1 : p^*_{int+\tau} = p_{nt}. \tag{51}$$

This follows from equation (16). Combining these two results yields that $p_{nt+\tau} = p_{nt-1} + xv_{nt}$ for all $\tau \geq 0$ if and only if

$$p^*_{int+\tau} = \begin{cases} p_{nt-1} + \frac{1}{\lambda_n} xv_{nt} & \text{for } \tau = 0 \\ p_{nt-1} + xv_{nt} & \text{for all } \tau \geq 1 \end{cases}. \tag{52}$$

Fourth, the adjustment price in period $t+\tau$ satisfies equation (52) for all $\tau \geq 1$ if and only if the profit-maximizing price in period $t+\tau$ satisfies

$$\forall \tau \geq 1 : p_{int+\tau}^{\diamondsuit} = p_{nt-1} + x v_{nt}. \tag{53}$$

This follows from equation (15). Fifth, given equation (53), the adjustment price in period t satisfies equation (52) if and only if the profit-maximizing price in period t satisfies

$$p_{int}^{\diamondsuit} = p_{nt-1} + \frac{\frac{1}{\lambda_n} - (1-\lambda_n)\beta}{1-(1-\lambda_n)\beta} x v_{nt}. \tag{54}$$

This follows again from equation (15). Collecting results yields that results yields that $p_{nt+\tau} = p_{nt-1} + x v_{nt}$ for all $\tau \geq 0$ if and only if

$$p_{int+\tau}^{\diamondsuit} = \begin{cases} p_{nt-1} + \left(\frac{\frac{1}{\lambda_n} - (1-\lambda_n)\beta}{1-(1-\lambda_n)\beta} \right) x v_{nt} & \text{for } \tau = 0 \\ p_{nt-1} + x v_{nt} & \text{for all } \tau \geq 1 \end{cases}. \tag{55}$$

C Solving the rational inattention model

Kalman filtering: The state-space representation of the dynamics of the signal concerning aggregate conditions is

$$\begin{aligned} p_{int}^{\diamondsuit A} &= p_{int-1}^{\diamondsuit A} + \sigma_A u_t, & (56) \\ s_{int}^{A} &= p_{int}^{\diamondsuit A} + \sigma_\varepsilon \varepsilon_{int}. & (57) \end{aligned}$$

The first equation is the state equation and the second equation is the observation equation. For ease of exposition, we use the following notation in this appendix: $X_t = p_{int}^{\diamondsuit A}$, $S_t = s_{int}^{A}$, $X_{t|t} = E\left[X_t | S^t\right]$, $X_{t|t-1} = E\left[X_t | S^{t-1}\right]$, $\sigma_{t|t}^2 = Var\left[X_t | S^t\right]$ and $\sigma_{t|t-1}^2 = Var\left[X_t | S^{t-1}\right]$. The usual Kalman filter equations yield

$$X_{t|t} = X_{t|t-1} + \frac{\sigma_{t|t-1}^2}{\sigma_{t|t-1}^2 + \sigma_\varepsilon^2} \left(S_t - X_{t|t-1}\right),$$

and

$$\sigma_{t|t}^2 = \sigma_{t|t-1}^2 - \frac{\sigma_{t|t-1}^2}{\sigma_{t|t-1}^2 + \sigma_\varepsilon^2} \sigma_{t|t-1}^2.$$

Furthermore,

$$X_{t+1|t} = X_{t|t},$$

and

$$\sigma^2_{t+1|t} = \sigma^2_{t|t} + \sigma^2_A. \tag{58}$$

Substituting the last two equations into the two equations before yields

$$X_{t|t} = X_{t-1|t-1} + \frac{\sigma^2_{t-1|t-1} + \sigma^2_A}{\sigma^2_{t-1|t-1} + \sigma^2_A. + \sigma^2_\varepsilon} \left(S_t - X_{t-1|t-1}\right), \tag{59}$$

and

$$\sigma^2_{t|t} = \frac{\sigma^2_\varepsilon}{\sigma^2_{t-1|t-1} + \sigma^2_A + \sigma^2_\varepsilon} \left(\sigma^2_{t-1|t-1} + \sigma^2_A\right). \tag{60}$$

When $\sigma^2_{t|t} = \sigma^2_{t-1|t-1}$, the unique positive solution to the last equation is

$$\bar{\sigma}^2_{t|t} = \left(-\frac{1}{2} + \sqrt{\frac{1}{4} + \frac{\sigma^2_\varepsilon}{\sigma^2_A}}\right) \sigma^2_A. \tag{61}$$

The information flow constraint: When X_t and $S^t = \left(S^0, S_1, S_2, \ldots, S_t\right)$ have a multivariate Gaussian distribution, the conditional distribution of X_t given S^t is Gaussian. In this case, the conditional entropy of X_t given S^t is a simple function of the conditional variance of X_t given S^t

$$H\left(X_t|S^t\right) = \frac{1}{2}\log_2\left(2\pi e \sigma^2_{t|t}\right).$$

Similarly,

$$H\left(X_t|S^{t-1}\right) = \frac{1}{2}\log_2\left(2\pi e \sigma^2_{t|t-1}\right).$$

The equation

$$H\left(X_t|S^{t-1}\right) - H\left(X_t|S^t\right) = \kappa^A,$$

then reduces to

$$\frac{1}{2}\log_2\left(\frac{\sigma^2_{t|t-1}}{\sigma^2_{t|t}}\right) = \kappa^A.$$

Using equation (58) yields

$$\frac{1}{2}\log_2\left(\frac{\sigma^2_{t-1|t-1} + \sigma^2_A}{\sigma^2_{t|t}}\right) = \kappa^A. \tag{62}$$

When $\sigma^2_{t|t} = \sigma^2_{t-1|t-1}$, the last equation becomes

$$\bar{\sigma}^2_{t|t} = \frac{\sigma^2_A}{2^{2\kappa^A} - 1}. \tag{63}$$

The assumption concerning initial information (s_{in}^0): To abstract from (purely deterministic) transitional dynamics in conditional variances, we assume that, after the decision-maker has chosen the allocation of attention (i.e. a pair κ^A and κ^S) in period zero, the decision-maker receives information in period zero such that the conditional variance of X_1 given information in period zero equals the steady-state value of the conditional variance of X_t given information in period $t-1$. This assumption implies that in period one the conditional variance of X_t given information in period t equals its steady-state value

$$\sigma_{1|1}^2 = \bar{\sigma}_{t|t}^2. \tag{64}$$

This simplifies computing the solution to problem (26)-(32).

Variance of noise, value of the objective and pricing behavior for a given allocation of attention: Equating the right-hand sides of equations (61) and (63) yields the variance of noise in the signal concerning aggregate conditions, σ_ε^2, for given attention allocated to aggregate conditions, κ^A,

$$\begin{aligned} \sigma_\varepsilon^2 &= \frac{1}{4}\left[\left(\frac{2^{2\kappa^A}+1}{2^{2\kappa^A}-1}\right)^2 - 1\right]\sigma_A^2 \\ &= \frac{2^{2\kappa^A}}{\left(2^{2\kappa^A}-1\right)^2}\sigma_A^2. \end{aligned} \tag{65}$$

Similarly, we obtain the variance of noise in the signal concerning sector-specific conditions, σ_ψ^2, for given attention allocated to sector-specific conditions, κ^S,

$$\sigma_\psi^2 = \frac{2^{2\kappa^S}}{\left(2^{2\kappa^S}-1\right)^2}\sigma_S^2. \tag{66}$$

Furthermore, using the fact that

$$\begin{aligned} E\left[\left(p_{int} - p_{int}^\diamondsuit\right)^2\right] &= E\left[\left(p_{int}^\diamondsuit - E\left[p_{int}^\diamondsuit | s_{in}^t\right]\right)^2 | s_{in}^t\right] \\ &= Var\left(p_{int}^\diamondsuit | s_{in}^t\right) \\ &= Var\left(p_{int}^{\diamondsuit A} | s_{in}^t\right) + Var\left(p_{int}^{\diamondsuit S} | s_{in}^t\right) \\ &= \frac{\sigma_A^2}{2^{2\kappa^A}-1} + \frac{\sigma_S^2}{2^{2\kappa^S}-1}, \end{aligned}$$

yields the value of objective (26) for a given allocation of attention:

$$\frac{\beta}{1-\beta}\frac{\bar{C}\left(\theta-1\right)\left(1+\frac{1-\alpha}{\alpha}\theta\right)}{2}\left[\frac{\sigma_A^2}{2^{2\kappa^A}-1} + \frac{\sigma_S^2}{2^{2\kappa^S}-1}\right]. \tag{67}$$

Next, we derive the pricing behavior for a given allocation of attention. Substituting equations (63) and (65) into equation (59) yields

$$X_{t|t} = \left(2^{-2\kappa^A}\right) X_{t-1|t-1} + \left(1 - 2^{-2\kappa^A}\right) S_t.$$

Furthermore, using equations (56) and (57) yields

$$\left(X_t - X_{t|t}\right) = \left(2^{-2\kappa^A}\right)\left(X_{t-1} - X_{t-1|t-1}\right) + \left(2^{-2\kappa^A}\right)\sigma_A u_t - \left(1 - 2^{-2\kappa^A}\right)\sigma_\varepsilon \varepsilon_{int}.$$

In addition, using equation (65) to substitute for σ_ε yields

$$\left(X_t - X_{t|t}\right) = \left(2^{-2\kappa^A}\right)\left(X_{t-1} - X_{t-1|t-1}\right) + \left(2^{-2\kappa^A}\right)\sigma_A u_t - \left(2^{-\kappa^A}\right)\sigma_A \varepsilon_{int}.$$

Solving this difference equation by repeated substitution, we arrive at

$$\begin{aligned} X_t - X_{t|t} &= \sum_{l=0}^{t-2}\left(2^{-2\kappa^A}\right)^l\left[\left(2^{-2\kappa^A}\right)\sigma_A u_{t-l} - \left(2^{-\kappa^A}\right)\sigma_A \varepsilon_{int-l}\right] \\ &\quad + \left(2^{-2\kappa^A}\right)^{t-1}\left(X_1 - X_{1|1}\right). \end{aligned}$$

Thus,

$$\begin{aligned} p_{int}^{\diamondsuit A} - E\left[p_{int}^{\diamondsuit A}|s_{in}^t\right] &= \sum_{l=0}^{t-2}\left(2^{-2\kappa^A}\right)^l\left[\left(2^{-2\kappa^A}\right)\sigma_A u_{t-l} - \left(2^{-\kappa^A}\right)\sigma_A \varepsilon_{int-l}\right] \\ &\quad + \left(2^{-2\kappa^A}\right)^{t-1}\left(p_{in1}^{\diamondsuit A} - E\left[p_{in1}^{\diamondsuit A}|s_{in}^1\right]\right). \end{aligned} \tag{68}$$

It follows from equations (27), (30), (68) and the corresponding equation for $p_{int}^{\diamondsuit S} - E\left[p_{int}^{\diamondsuit S}|s_{in}^t\right]$ that

$$\begin{aligned} p_{int}^{\diamondsuit} - p_{int} &= \sum_{l=0}^{t-2}\left(2^{-2\kappa^A}\right)^l\left[\left(2^{-2\kappa^A}\right)\sigma_A u_{t-l} - \left(2^{-\kappa^A}\right)\sigma_A \varepsilon_{int-l}\right] \\ &\quad + \sum_{l=0}^{t-2}\left(2^{-2\kappa^S}\right)^l\left[\left(2^{-2\kappa^S}\right)\sigma_S v_{nt-l} - \left(2^{-\kappa^S}\right)\sigma_S \psi_{int-l}\right] \\ &\quad + \left(2^{-2\kappa^A}\right)^{t-1}\left(p_{in1}^{\diamondsuit A} - E\left[p_{in1}^{\diamondsuit A}|s_{in}^1\right]\right) + \left(2^{-2\kappa^S}\right)^{t-1}\left(p_{in1}^{\diamondsuit S} - E\left[p_{in1}^{\diamondsuit S}|s_{in}^1\right]\right) \end{aligned} \tag{69}$$

Equation (69) and equations (27)-(29) already pin down the response of p_{int} to various shocks. We derive one more equation because that equation is easier to read than equation

(69). As $t \to \infty$ (or for the right values of $p_{in1}^{\diamondsuit A} - E\left[p_{in1}^{\diamondsuit A} | s_{in}^1\right]$ and $p_{in1}^{\diamondsuit S} - E\left[p_{in1}^{\diamondsuit S} | s_{in}^1\right]$), equation (69) becomes

$$\begin{aligned} p_{int}^{\diamondsuit} - p_{int} &= \sum_{l=0}^{\infty} \left(2^{-2\kappa^A}\right)^l \left[\left(2^{-2\kappa^A}\right) \sigma_A u_{t-l} - \left(2^{-\kappa^A}\right) \sigma_A \varepsilon_{int-l}\right] \\ &\quad + \sum_{l=0}^{\infty} \left(2^{-2\kappa^S}\right)^l \left[\left(2^{-2\kappa^S}\right) \sigma_S v_{nt-l} - \left(2^{-\kappa^S}\right) \sigma_S \psi_{int-l}\right]. \end{aligned} \tag{70}$$

The optimal allocation of attention: The optimal allocation of attention is derived in the main text from equation (67).

References

[1] 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, 387-422.

[2] Bils, Mark, and Peter J. Klenow (2004): "Some Evidence on the Importance of Sticky Prices." Journal of Political Economy, 112, 947-985.

[3] Boivin, Jean, Marc Giannoni, and Ilian Mihov (2007): "Sticky Prices and Monetary Policy: Evidence from Disaggregated U.S. Data." American Economic Review, forthcoming.

[4] Geweke, John, and Guofu Zhou (1996): "Measuring the Pricing Error of the Arbitrage Pricing Theory." Review of Financial Studies, 9, 557-587.

[5] Calvo, Guillermo (1983): "Staggered Prices in a Utility-Maximizing Framework." Journal of Monetary Economics, 12, 383-398.

[6] Carter, C. K., and R. Kohn (1994): "On Gibbs Sampling for State Space Models." Biometrika, 81, 541-553.

[7] Chib, Siddartha, and Edward Greenberg (1994): "Bayes Inference in Regression Models with ARMA(p,q) errors." Journal of Econometrics, 64, 183-206.

[8] Forni, Mario, Marc Hallin, Marco Lippi, and Lucrezia Reichlin (2004): "The Generalized Factor Model: Indetification and Estimation." Review of Economics and Statistics, 82, 540-554.

[9] Geweke, John (1977): "The Dynamic Factor Analysis of Economic Time Series." In "Latent Variables in Socioeconomic Models", D. Aigner and A. Goldberger, eds., Amsterdam, North Holland.

[10] Geweke, John (1992): "Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments." In "Bayesian Statistics", J.M. Bernardo et al., eds., Oxford, Oxford University Press.

[11] Geweke, John, and Kenneth J. Singleton (1981): "Maximum Likelihood 'Confirmatory' Factor Analysis of Economic Time Series." International Economic Review, 22, 37-54.

[12] Kehoe, Patrick J., and Virgiliu Midrigan (2007): "Can Heterogeneity in Price Stickiness Account for the Persistence and Volatility of Good-Level Real Exchange Rates?" Federal Reserve Bank of Minneapolis Research Department Staff Report.

[13] Kim, Chang-Jin, and Charles R. Nelson (1999): "State-Space Models with Regime Switching." The MIT Press, Cambridge, Massachusetts.

[14] Maćkowiak, Bartosz, and Mirko Wiederholt (2008a): "Optimal Sticky Prices under Rational Inattention." American Economic Review, forthcoming.

[15] Maćkowiak, Bartosz, and Mirko Wiederholt (2008b): "Business Cycle Dynamics under Rational Inattention." Discussion paper, European Central Bank and Northwestern University.

[16] Mankiw, N. Gregory, and Ricardo Reis. (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." Quarterly Journal of Economics, 117, 1295-1328.

[17] Nakamura, Emi, and Jón Steinsson (2008): "Five Facts About Prices: A Reevaluation of Menu Cost Models." Quarterly Journal of Economics, forthcoming.

[18] Otrok, Christopher, and Charles H. Whiteman (1998): "Bayesian Leading Indicators: Measuring and Predicting Economic Conditions in Iowa." International Economic Review, 39, 997-1014.

[19] Primiceri, Giorgio E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy." Review of Economic Studies, 72, 821-852.

[20] Raftery, Adrian E., and Stephen M. Lewis (1992): "How Many Iterations in the Gibbs Sampler?" In "Bayesian Statistics", J.M. Bernardo et al., eds., Oxford, Oxford University Press.

[21] Reis, Ricardo (2006): "Inattentive Producers." Review of Economic Studies, 73, 793-821.

[22] Reis, Ricardo, and Mark W. Watson (2007a): "Measuring Changes in the Value of the Numeraire." Discussion paper, Princeton University.

[23] Reis, Ricardo, and Mark W. Watson (2007b): "Relative Goods' Prices and Pure Inflation." NBER working paper 13615.

[24] Sargent, Thomas J., and Christopher A. Sims (1977): "Business Cycle Modeling Without Pretending to Have Too Much A Priori Economic Theory." In "New Methods of Business Cycle Research", C.A. Sims, ed., Minneapolis, Federal Reserve Bank.

[25] Sims, Christopher A. (2003): "Implications of Rational Inattention." Journal of Monetary Economics, 50, 665-690.

[26] Stock, James H., and Mark W. Watson (1989): "New Indexes of Coincident and Leading Economic Indicators." In "NBER Macroeconomics Annual 1989", O. Blanchard and S. Fischer, eds., Cambridge, MIT Press.

Table 1: Out-of-Sample Forecast Results: the Sectoral Distribution of the Average RMSE

		5th percentile	**Median**	**95th percentile**
Unobservable index model	u_t follows a scalar process and the order of $C_n(L)$ equals 6	0.615	1.059	1.529
	u_t follows a scalar process and the order of $C_n(L)$ equals 12	0.649	1.069	1.596
	u_t follows a bivariate vector process and the order of $C_n(L)$ equals 6	0.611	1.078	1.585
AR(6) model for sectoral inflation		0.621	1.066	1.648
AR(12) model for sectoral inflation		0.658	1.098	1.683

Note: See Section 4 for a desciption of the out-of-sample forecast exercise.

Table 2: Regression of the Speed of Response of a Sectoral Price Index to Aggregate Shocks on the Sectoral Monthly Frequency of Price Changes (Bils and Klenow, 2004)

	Empirical	Calvo model
Regression coefficient	**1.99**	**1.11**
90% probability interval for the empirical regression coefficient	(0.67, 4.38)	
t-statistic for the empirical regression coefficient	2.9	

Note: The empirical regression coefficient is the median of the posterior density.
The empirical t-statistic is the median of the posterior density.

Table 3: Regression of the Speed of Response of a Sectoral Price Index to Aggregate Shocks on the Sectoral Monthly Frequency of Regular Price Changes (Nakamura and Steinsson, 2008)

	Empirical	Calvo model
Regression coefficient	**1.49**	**1.06**
90% probability interval for the empirical regression coefficient	(0.47, 4.32)	
t-statistic for the empirical regression coefficient	2.6	

Note: The empirical regression coefficient is the median of the posterior density.
The empirical t-statistic is the median of the posterior density.

Table 4: Regression of the Speed of Response of a Sectoral Price Index to Sector-Specific Shocks on the Sectoral Monthly Frequency of Price Changes (Bils and Klenow, 2004)

	Empirical	Calvo model
Regression coefficient	**0.14**	**1.11**
90% probability interval for the empirical regression coefficient	(0.1, 0.19)	
t -statistic for the empirical regression coefficient	1.9	

Note: The empirical regression coefficient is the median of the posterior density.
The empirical *t* -statistic is the median of the posterior density.

Table 5: Regression of the Speed of Response of a Sectoral Price Index to Sector-Specific Shocks on the Sectoral Monthly Frequency of Regular Price Changes (Nakamura and Steinsson, 2008)

	Empirical	Calvo model
Regression coefficient	**-0.09**	**1.06**
90% probability interval for the empirical regression coefficient	(-0.14, -0.05)	
t -statistic for the empirical regression coefficient	-1.3	

Note: The empirical regression coefficient is the median of the posterior density.
The empirical *t* -statistic is the median of the posterior density.

Table 6: Regression of the Speed of Response of a Sectoral Price Index to Aggregate Shocks on the Variance of Sectoral Inflation due to Aggregate Shocks

	Empirical	RI model
Regression coefficient	**357.24**	**263.65**
90% probability interval for the empirical regression coefficient	(36.57, 1227.48)	
t-statistic for the empirical regression coefficient	1.4	

Note: The empirical regression coefficient is the median of the posterior density.
The empirical *t*-statistic is the median of the posterior density.

Table 7: Regression of the Speed of Response of a Sectoral Price Index to Sector-Specific Shocks on the Variance of Sectoral Inflation due to Sector-Specific Shocks

	Empirical	RI model
Regression coefficient	**7.08**	**9.87**
90% probability interval for the empirical regression coefficient	(5.32, 9.22)	
t-statistic for the empirical regression coefficient	0.96	

Note: The empirical regression coefficient is the median of the posterior density.
The empirical *t*-statistic is the median of the posterior density.

Table A.1: The Cross-Section of the Raftery-Lewis Diagnostics Across the Parameters

	Median	95th percentile	99th percentile	Max
θ's	623	695	1352	3736
φ's	697	1208	4278	478784
σ's	623	664	879	951

Note: The diagnostics are for the benchmark specification. See Section A.3 for the details concerning the diagnostics.

Table A.2: The Cross-Section of the Geweke Diagnostics Across the Parameters

	Median	95th percentile	99th percentile	Max
θ's	1.8	4.2	6.7	27.9
φ's	1.3	2.3	40.7	79.1
σ's	0.7	1.2	1.4	1.5

Note: The diagnostics are for the benchmark specification. See Section A.3 for the details concerning the diagnostics.

Figure 1: The Cross-Section of the Normalized Impulse Responses of Sectoral Price Indices

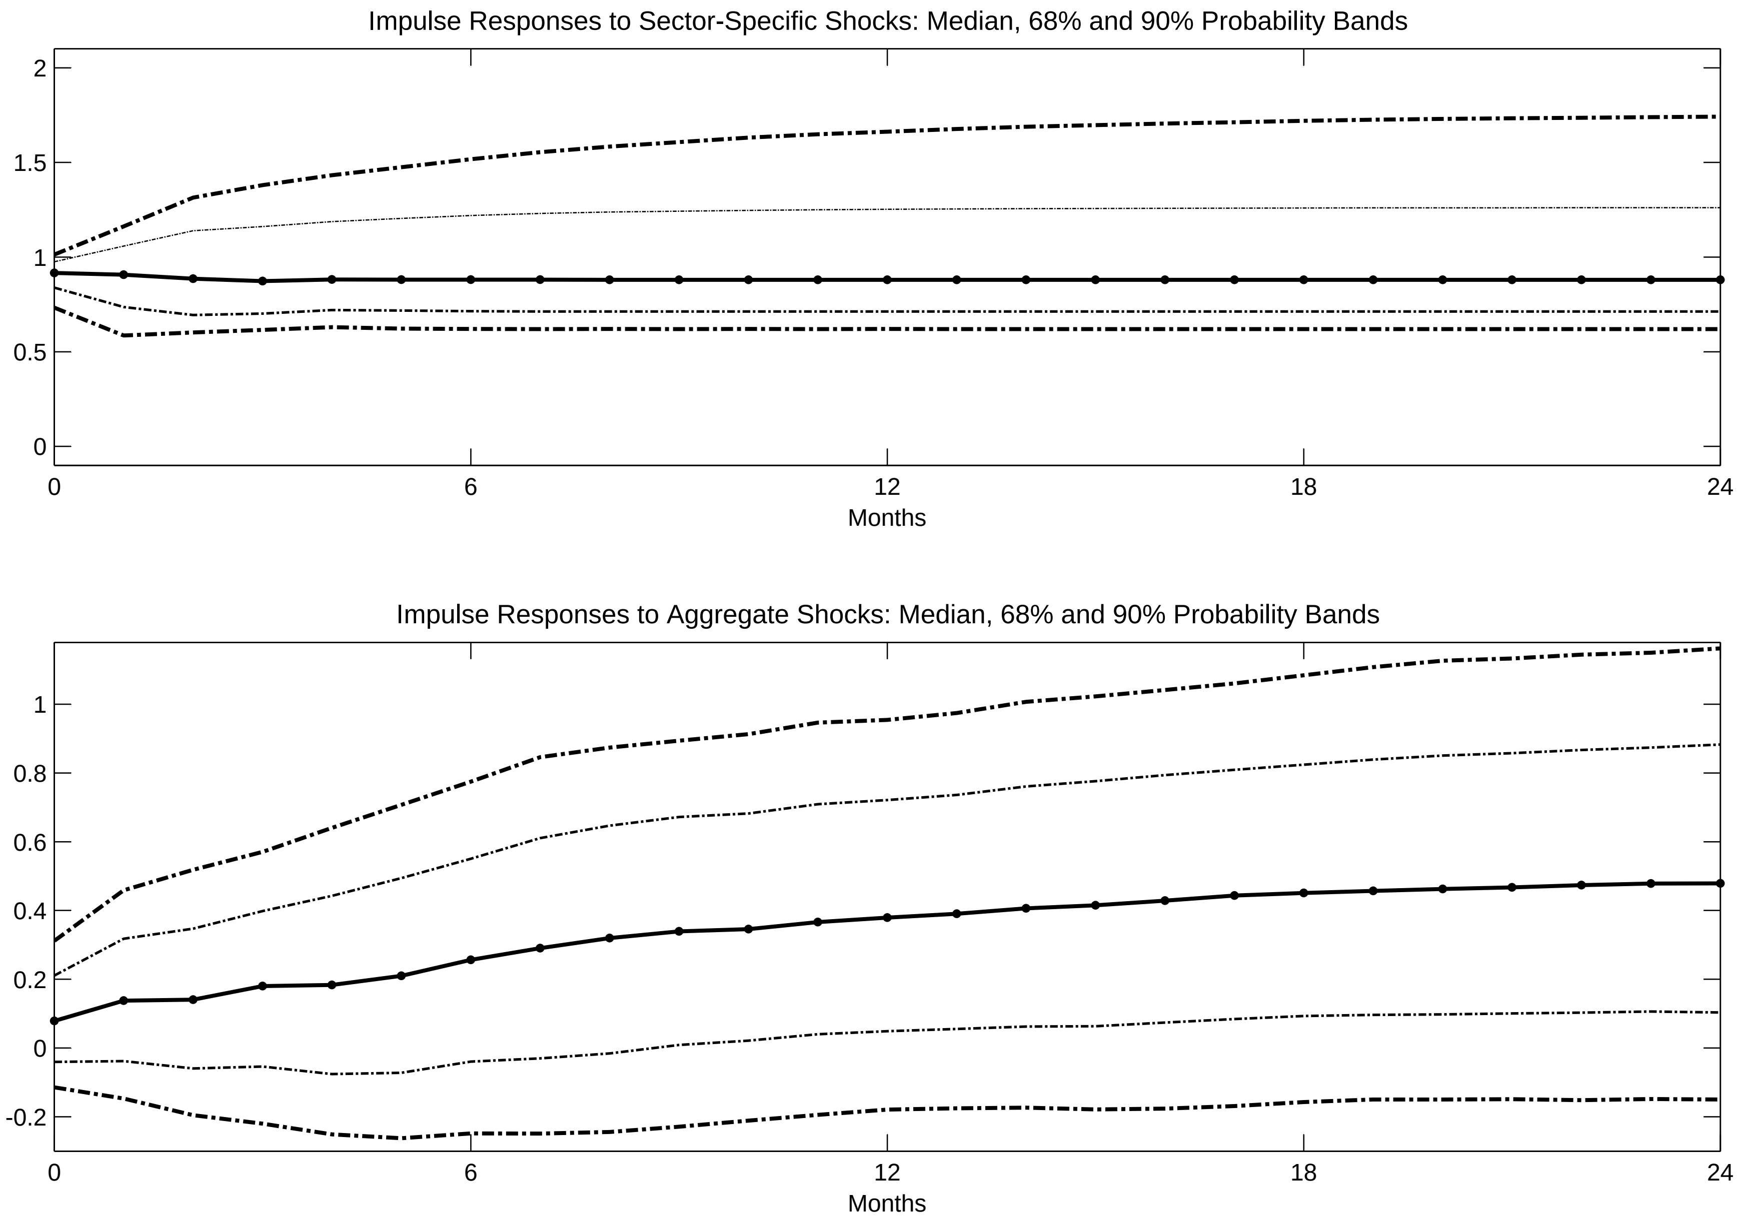

Figure 2: The Cross-Section of the Speed of Response of Sectoral Price Indices to Shocks

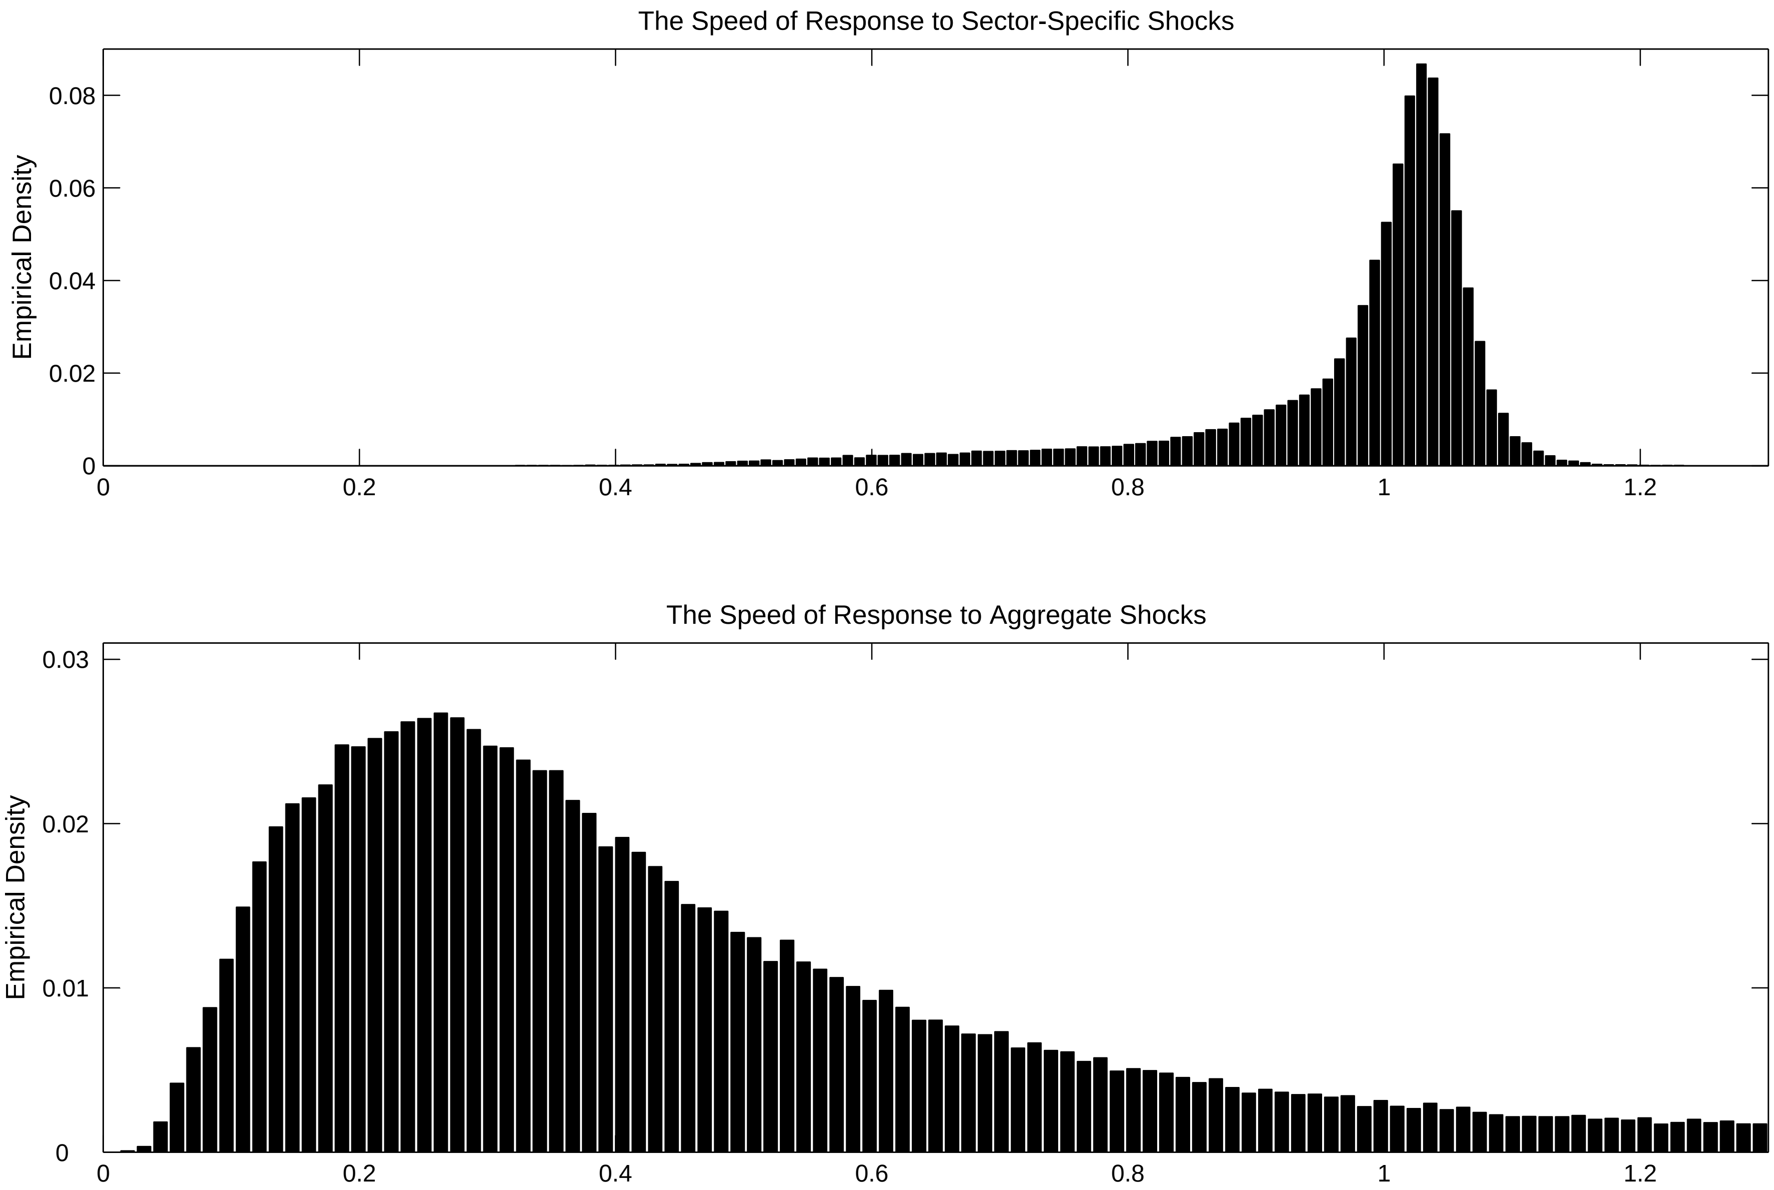

Figure 3: Impulse Responses of the Sectoral Price Index to Sector-Specific Shocks: Profit-Maximizing and Calvo

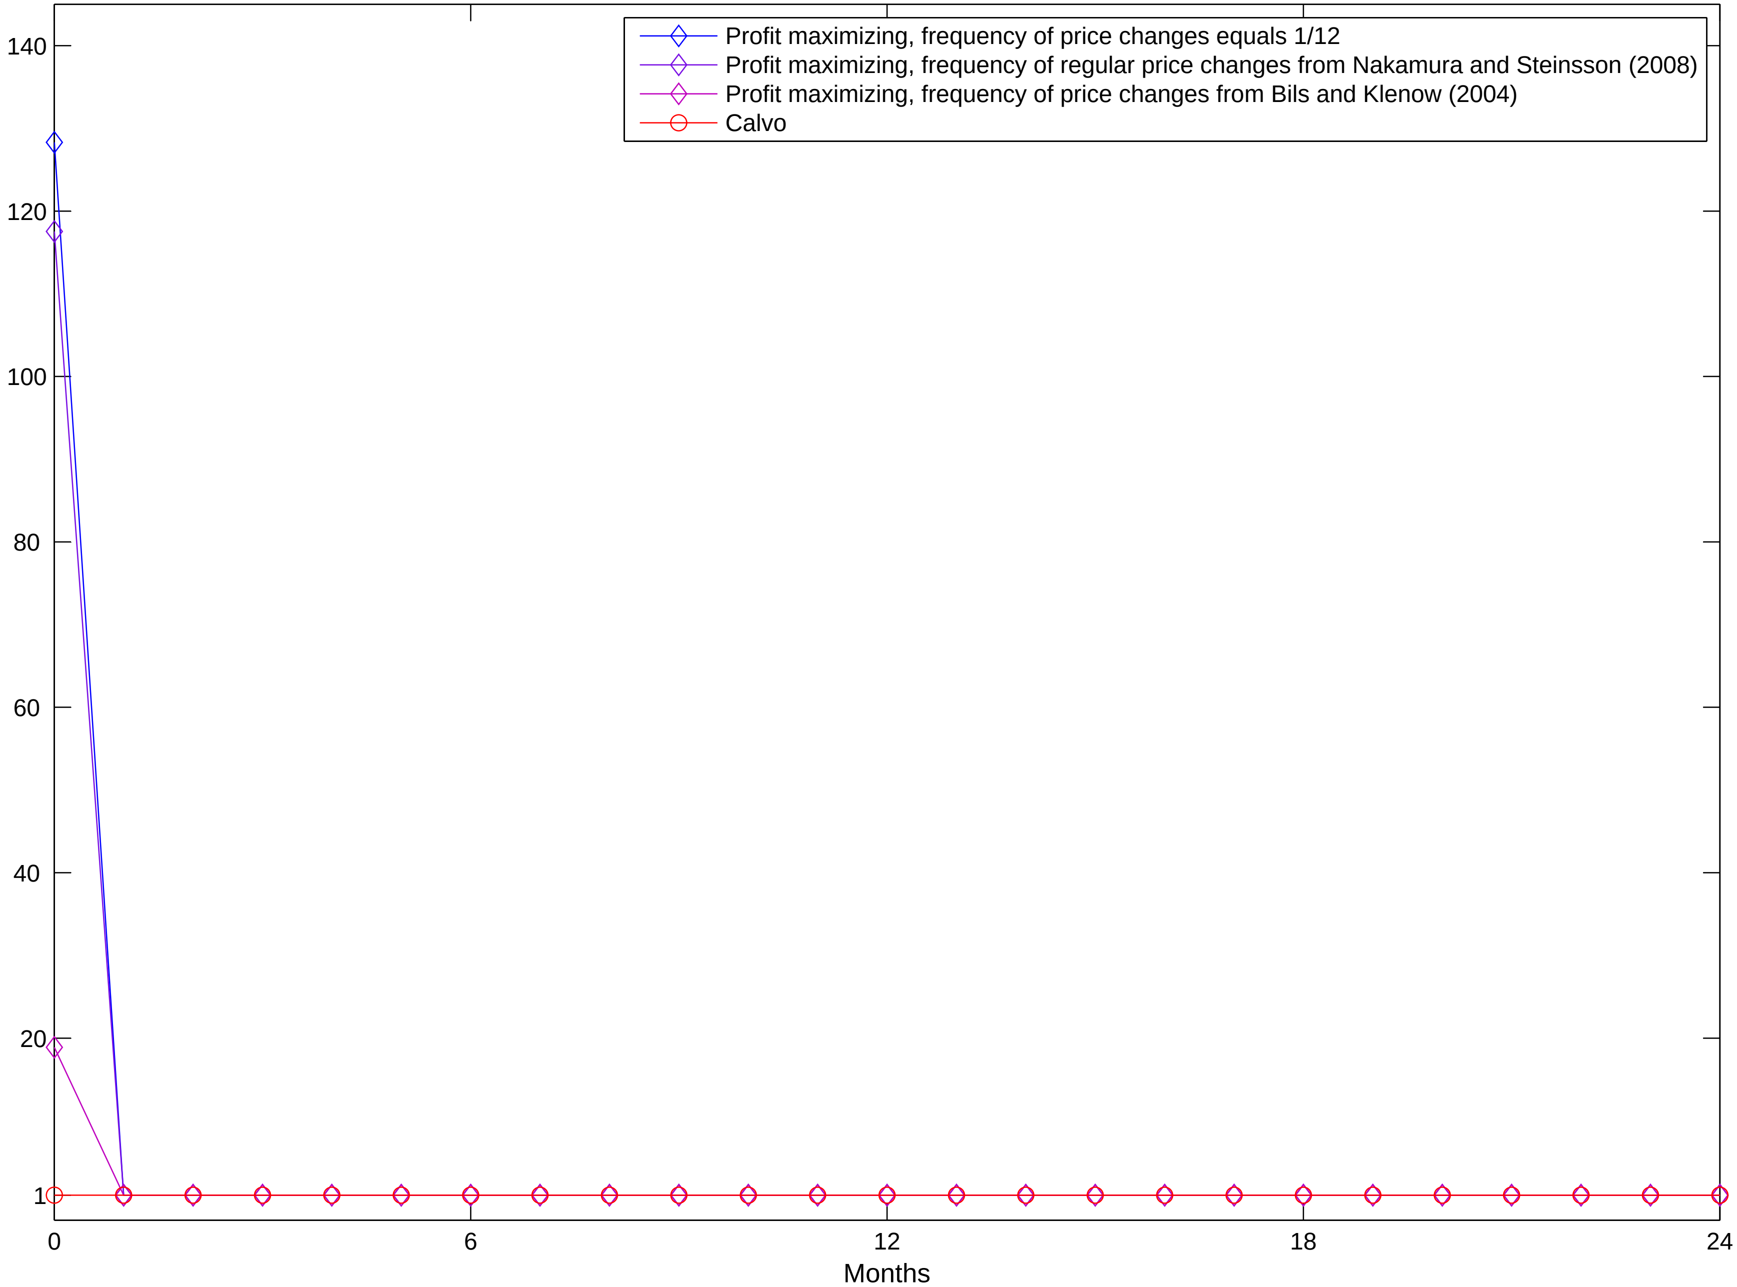

Figure 4: Impulse Responses of Sector-Specific Productivity that Yield the Profit-Maximizing Impulse Responses in Figure 3

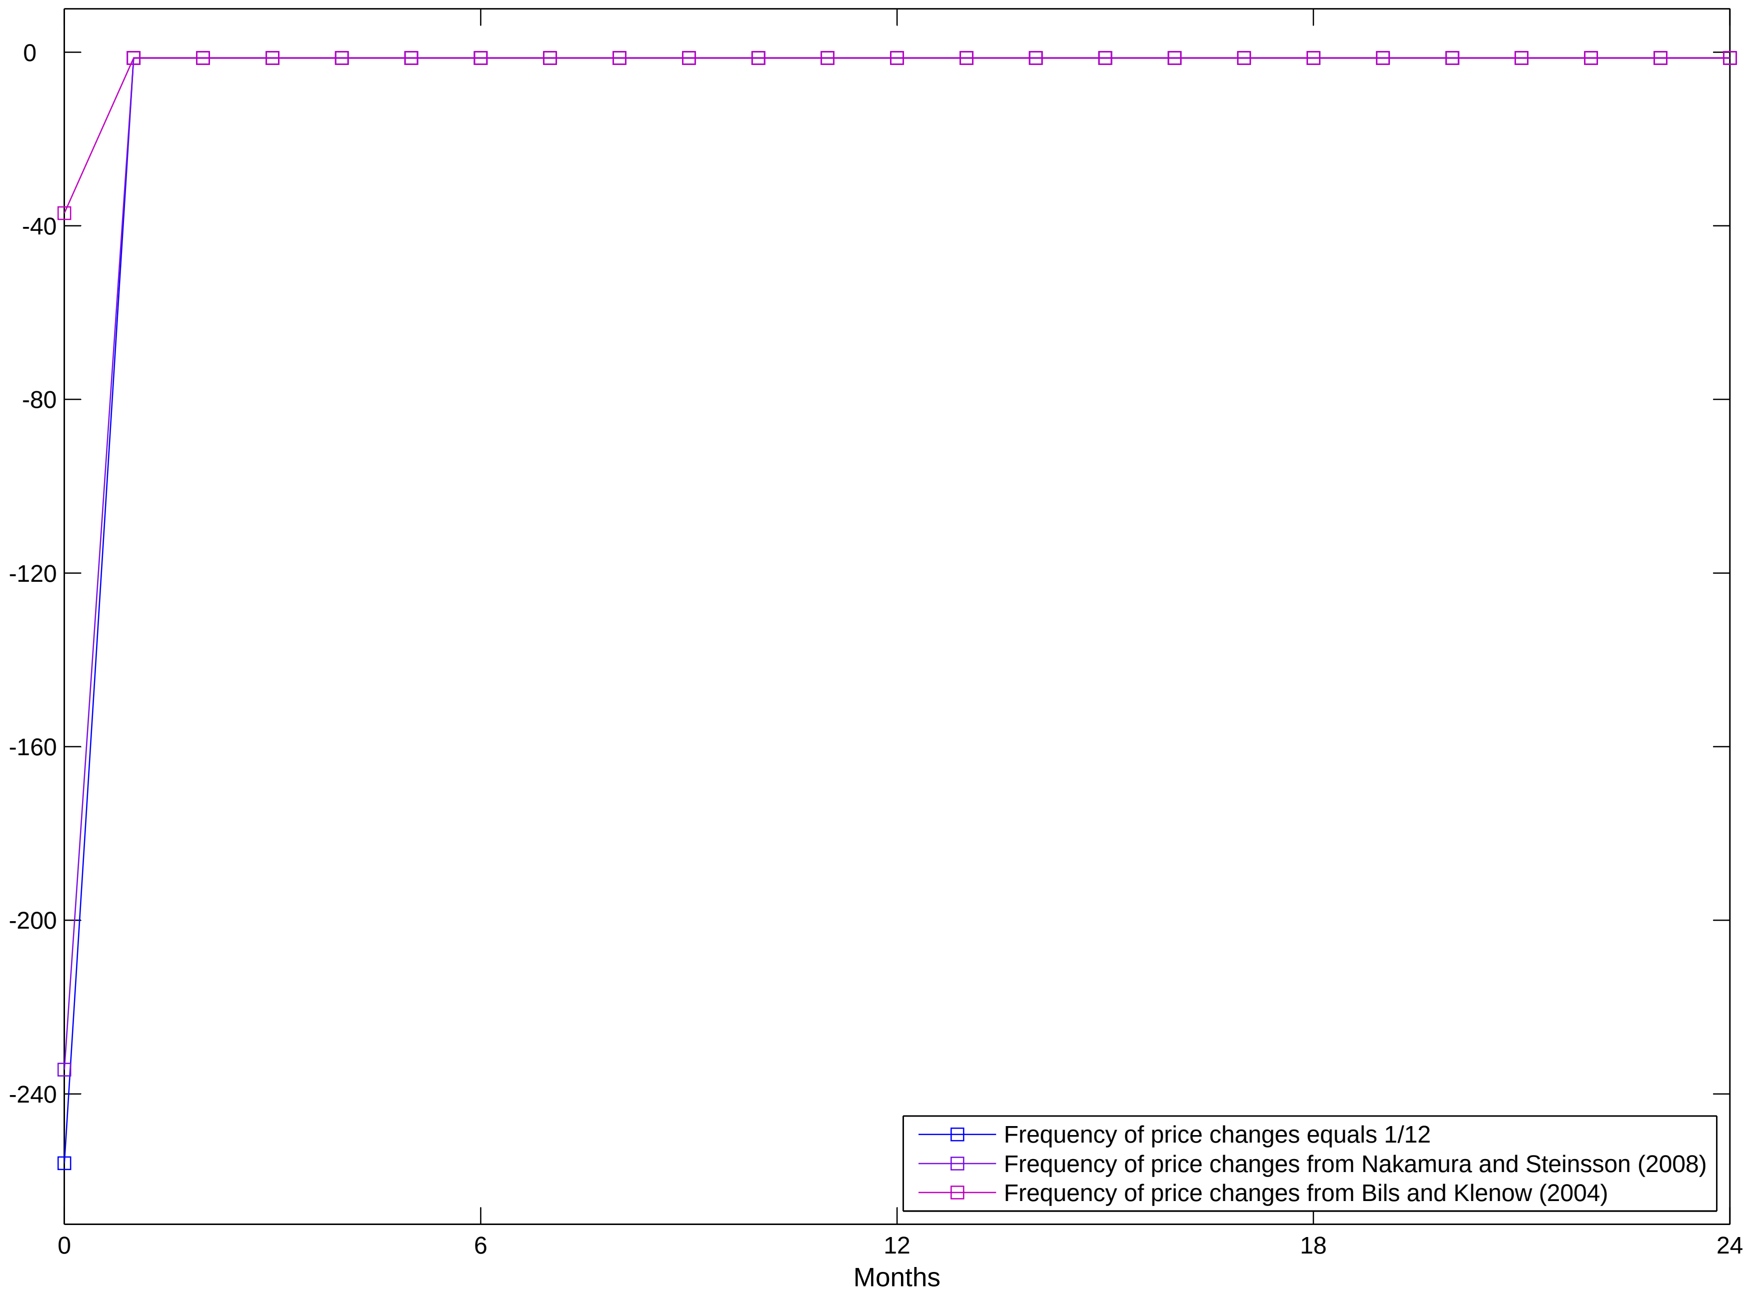

Figure 5: Impulse Responses of the Sectoral Price Index to Shocks: Empirical, Profit-Maximizing, and Sticky Information

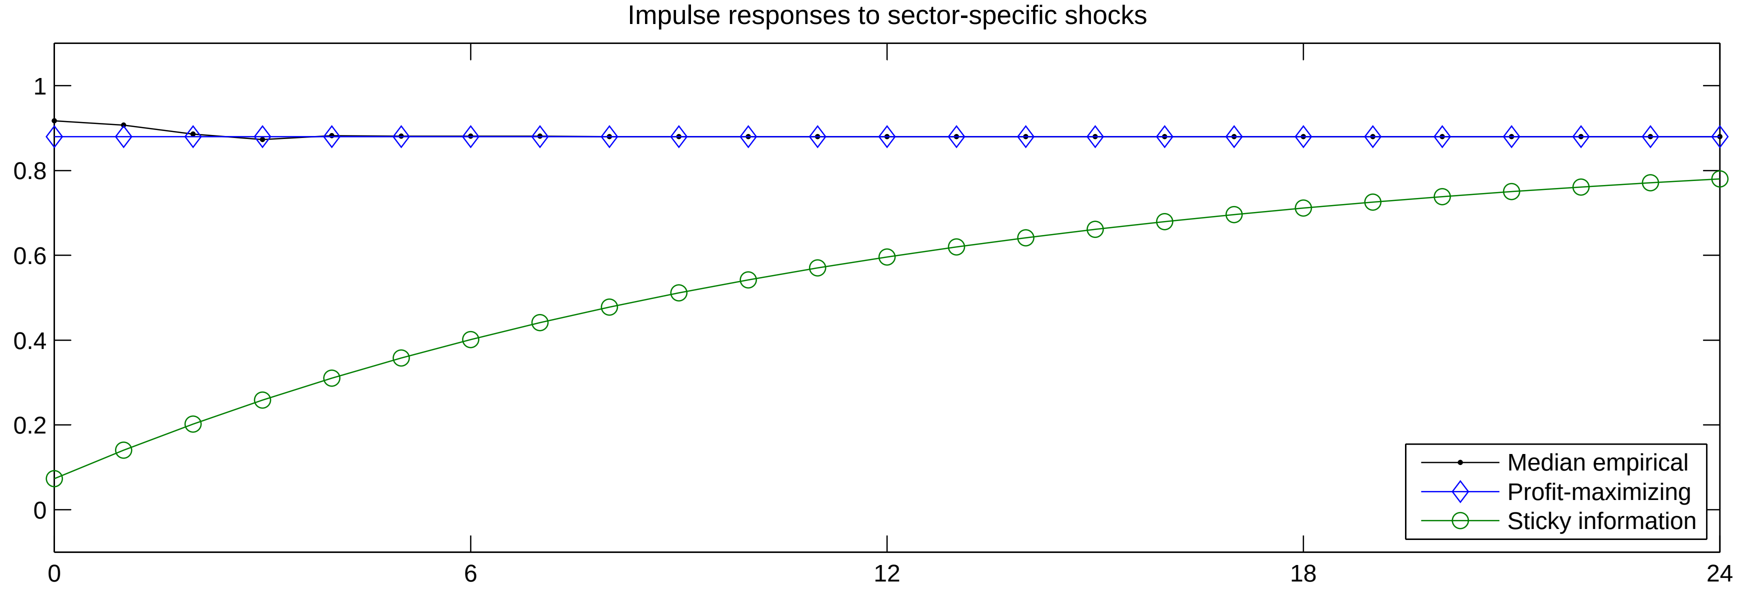

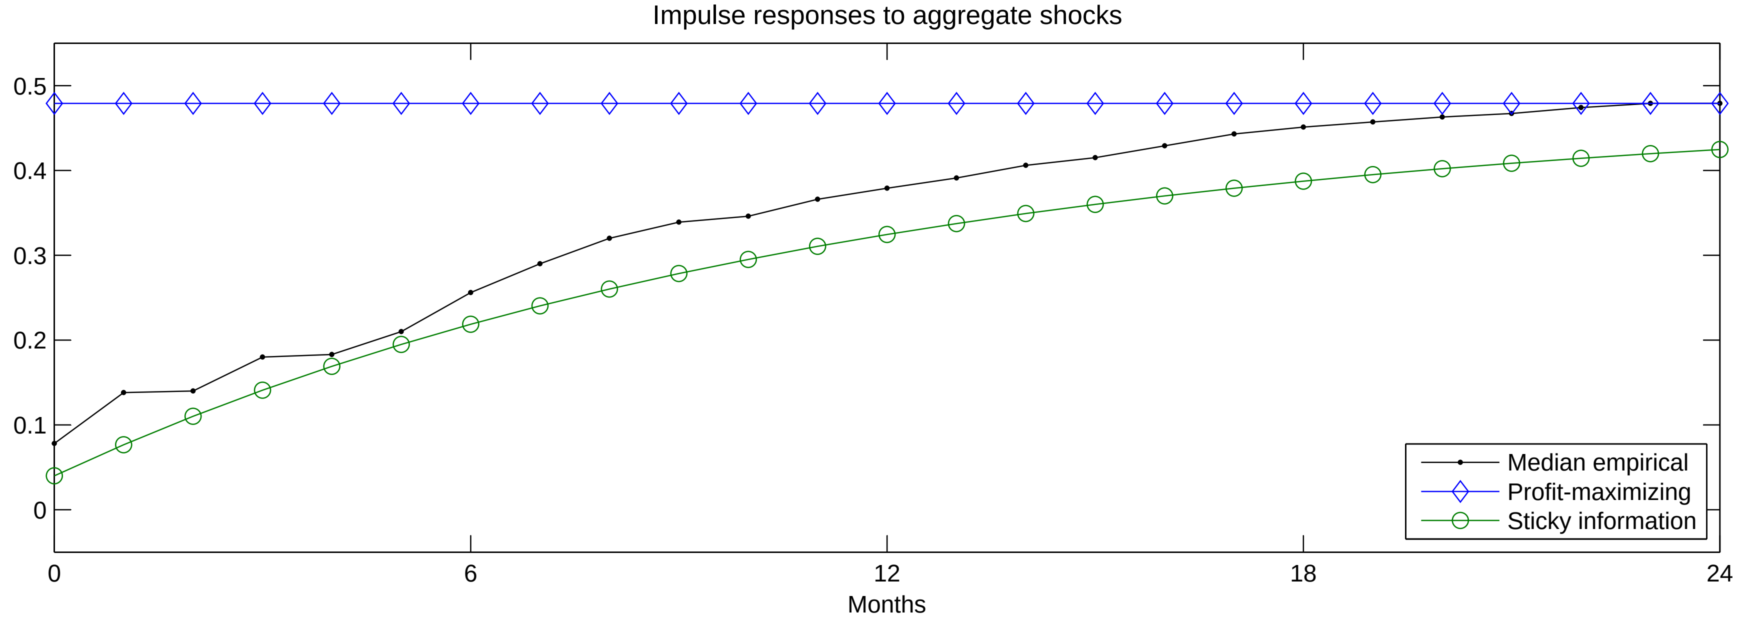

Figure 6: Impulse Responses of the Sectoral Price Index to Sector-Specific Shocks: Profit-Maximizing and Sticky Information