TEXTO PARA DISCUSSÃO

No. 634

The Cross-sectional Distribution of Price Stickiness Implied by Aggregate Data

Carlos Carvalho
Niels Arne Dam
Jae Won Lee

PUC RIO

DEPARTAMENTO DE ECONOMIA
www.econ.puc-rio.br

The Cross-sectional Distribution of Price Stickiness Implied by Aggregate Data*

Carlos Carvalho† Niels Arne Dam‡ Jae Won Lee§

March, 2018

Abstract

This paper provides evidence on three mechanisms that can reconcile frequent individual price changes with sluggish aggregate price dynamics. To that end, we estimate a semi-structural model that allows us to extract information about real rigidities and cross-sectional heterogeneity in price stickiness from aggregate data. Hence the model can also speak to the debate about the aggregate implications of sales and other temporary price changes. Our estimates point to the presence of large real rigidities and a significant degree of heterogeneity in price stickiness. Moreover, the cross-sectional distribution of price stickiness implied by aggregate data is in line with an empirical distribution obtained from micro price data that factors out sales and product substitutions. Our results suggest that all three feaures – i) real rigidities, ii) heterogeneity in price stickiness and iii) exclusion of temporary price changes – help bridge the gap between micro and macro evidence on nominal price rigidity.

JEL classification: E10, E30

Keywords: real rigidities, heterogeneity in price stickiness, sales, regular prices, micro data, macro data, Bayesian estimation

*This paper circulated previously under the title "Real Rigidities and the Cross-Sectional Distribution of Price Stickiness: Evidence from Micro and Macro Data Combined." For helpful comments and suggestions we would like to thank Luis Álvarez, Marco Del Negro, Stefano Eusepi, Henrik Jensen, Ed Knotek, Oleksyi Kryvtsov, John Leahy, Frank Schorfheide, two anonymous referees, and, especially, Giorgio Primiceri. We also thank seminar participants at the ESWC 2010, AEA 2010, Banque de France conference "Understanding price dynamics: recent advances", ESEM 2009, SED 2009, NAMES 2009, Bank of Canada, ECB - WDN Workshop, Rutgers University, LAMES 2008, Danmarks Nationalbank, University of Copenhagen, Riksbank, and NY Fed. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Central Bank of Brazil.

†Central Bank of Brazil and PUC-Rio. E-mail: cvianac@econ.puc-rio.br.

‡Finance Denmark. Email: nad@fida.dk.

§University of Virginia. Email: jl2rb@virginia.edu.

1 Introduction

Our understanding of the real effects of monetary policy hinges, to a large extent, on the existence of some degree of nominal price rigidity. Since the publication of the seminal Bils and Klenow (2004) paper, the availability of large amounts of micro price data has rekindled interest in this area, and allowed us to make progress. Yet, estimates of the extent of nominal price stickiness based on microeconomic data versus those based on aggregate data usually produce a conflicting picture.

According to Klenow and Malin's (2011) survey of the empirical literature based on micro data, prices change, on average, at least once a year – somewhat more often than we thought was the case prior to Bils and Klenow (2004). In contrast, making sense of estimates of the response of the aggregate price level to monetary shocks (from dynamic stochastic general equilibrium – DSGE – models, or vector autoregressions) usually requires much less frequent price adjustments.[1]

If nominal price rigidities are to remain the leading explanation for why monetary policy has large and persistent real effects, it is important that we deepen our understanding of mechanisms that can narrow the gap between the evidence of somewhat flexible individual prices and relatively sluggish aggregate prices – i.e., mechanisms that can produce a large "contract multiplier", as Klenow and Malin (2011) put it. If prices change frequently and each and every price change contributes to fully offset nominal disturbances, then nominal price rigidity cannot be the source of large and persistent monetary non-neutralities. Hence, a large contract multiplier requires that price adjustments, somehow, fail to perfectly neutralize monetary innovations.

In this paper, we contribute to bridging the gap between micro and macro evidence on the extent of nominal price rigidity. To that end, we estimate a standard macroeconomic model of price setting, and use it to speak to three mechanisms that can boost the contract multiplier.[2] The first such mechanism are so-called "real rigidities", in the sense of Ball and Romer (1990). Large real rigidities reduce the sensitivity of individual prices to aggregate demand conditions, and thus serve as a source of endogenous persistence: for any given frequency of price changes, partial adjustment of individual prices makes for a sluggish response of the aggregate price level to monetary shocks.

The other two mechanisms are motivated by empirical evidence uncovered since Bils and Klenow

[1]See, for example, the survey by Maćkowiak and Smets (2008).

[2]Information frictions can also lead to large contract multipliers. Not surprisingly, that literature picked up steam after the empirical literature based on micro price data flourished. Classic contributions include Caballero (1989), Reis (2006), and Maćkowiak and Wiederholt (2009), who obtain large monetary non-neutralities in models with information frictions in which prices change continuously. More recently, Bonomo, Carvalho, Garcia, and Malta (2014) obtain a large contract multiplier in an estimated model with menu costs and partially costly information.

(2004), and subsequent theoretical literature. Cross-sectional heterogeneity in price rigidity, to the extent documented in the micro data, can lead to much larger monetary non-neutralities than the average frequency of price changes would imply (Carvalho 2006, Nakamura and Steinsson 2010). The reason is that, while recurrent price changes by firms in more flexible sectors do not contribute as much to offset monetary shocks, they do count for the frequency of price adjustment.[3] Heterogeneity can become an even more powerful mechanism when coupled with strong real rigidities that lead to strategic complementarities in pricing decisions. In those circumstances, firms in the more sticky sectors become disproportionately important in shaping aggregate dynamics (relative to their sectoral weight), through their influence on pricing decisions of firms that change prices more frequently (Carvalho 2006).[4]

The third mechanism is associated with the presence of sales and other temporary price changes. Guimaraes and Sheedy (2011) and Kehoe and Midrigan (2014) show that such price changes may help reconcile frequent micro adjustments with a sluggish aggregate price response to nominal disturbances.[5] A basic intuition for their results is that temporary price changes fail to offset monetary shocks well, since these shocks tend to induce permanent changes in the level of prices.

Extracting information on the three underlying mechanisms from aggregate data (as opposed to disaggregate data) is a distinct contribution of our paper. To that end, we rely on a relatively standard "semi-structural" model. The price-setting block of the model is a multisector sticky-price economy that allows for heterogeneity in price stickiness, and can feature strategic complementarity or substitutability in pricing decisions. In particular, we consider both Taylor and Calvo pricing schemes. The remaining equations specify exogenous stochastic processes that drive firms' frictionless optimal prices. They provide the model with some flexibility to perform in empirical terms, and thus allow us to focus on the objects of interest in the price-setting block of the economy.[6]

We show that, at least in theory, our model is able to separately identify real and nominal rigidities, and tell apart economies with homogenous from those with heterogeneous price stickiness – based on

[3]Carvalho and Schwartzman (2015) show how this intuition can be formalized in terms of a "selection effect" relative to the timing of price changes, which arises in the class of time-dependent pricing models.

[4]Nakamura and Steinsson (2010) conclude that this interaction is not important in their calibrated menu-cost model.

[5]Coibion et al. (2014) provide evidence that sales are essentially acyclical – which is consistent with the models in Guimaraes and Sheedy (2011) and Kehoe and Midrigan (2014). Kryvtsov and Vincent (2014), on the other hand, argue that sales do not help reconcile micro and macro evidence on price rigidity. They document a large degree of cyclicality of sales in the U.K. micro data, and develop a model that can explain their findings. In bad times, consumers intensify search for bargain prices and firms increase the frequency of sales. This "complementarity" between search effort and sales frequency breaks down the strategic substitutability of sales that would otherwise arise (as in Guimaraes and Sheedy 2011), and leads to cyclical sales.

[6]Several earlier papers in the literature combine structural equations with empirical specifications for other parts of the model (e.g., Sbordone 2002, Guerrieri 2006, and Coenen et al. 2007).

aggregate data only. The model can also discriminate between different (non-degenerate) distributions of price rigidity, providing information on which one helps explain aggregate dynamics better. Hence, our analysis can also speak to the debate about whether price changes due to sales and product substitutions are relevant for macro dynamics.

Identification of the distribution of price stickiness based on aggregate data is possible in our framework because sectors that differ in price stickiness have different implications for the response of the macroeconomy to shocks at different frequencies. In particular, sectors where prices are more sticky are relatively more important in determining the low-frequency response to shocks; and vice-versa for more flexible sectors. These differences provide information about the cross-sectional distribution of price stickiness. Separate identification of real and nominal rigidities hinges on the fact that the former induces endogenous persistence in the economy. We show analytically how this introduces a dependence of the aggregate price level on its own lags, which can be exploited to obtain identification of the parameter that governs real rigidities.[7] While we consider both Taylor and Calvo pricing, we find that, in small samples, the former allows for stronger identification of the objects of interest. Hence, in our baseline estimation we rely on Taylor pricing, and use Calvo pricing to assess the robustness of our findings.[8]

Much of the recent literature on price setting pays large attention to empirical facts obtained from micro price data – and so do we. However, treating statistics derived from micro data as the true "population moments" that matter for aggregate dynamics can be misleading, in our view. First, it is possible that some price adjustments do not convey as much information about changes in macroeconomic conditions as others do. While this possibility is at the core of the debate about whether or not to exclude sales from price setting statistics for macro purposes, the argument applies more generally – for example, it also applies to the literature on price setting under information frictions. If so, macro-based estimates should convey useful information about price changes that do matter for aggregate dynamics. Second, and not less importantly, Eichenbaum et al. (2014) show that the BLS micro data underlying the CPI are plagued

[7] Real and nominal rigidities are not separately identified in our model when price setting takes place according to the Calvo model with homogeneous price stickiness. This, however, is the only exception to the class of models that we entertain.

[8] The original Taylor model, in which all prices are fixed for the same number of periods, has been criticized for being at odds with both micro and macro facts. First, that model tends to produce impulse response functions with kinks, which are uncommon in impulse responses produced by estimated VARs. Second, in that model all price spells have the same duration, which is at odds with the microeconomic evidence. The Calvo model is usually more successful along both dimensions. However, as will become clear, our estimated multisector models with Taylor pricing produce smoother dynamics that resemble those of a VAR – especially when pricing decisions are strategic complements. In addition, we later show that our version of the Taylor model can be recast into a model that yields identical aggregate dynamics and yet can be made consistent with a large number of micro price facts. Hence, although it still undeperforms the Calvo model in terms of fit of the aggregate data, our multisector Taylor model can account for both aggregate dynamics and micro price facts significantly better.than the conventional "one-sector" Taylor model.

with measurement problems when it comes to computing statistics based on individual price changes. While Eichenbaum et al. (2014) focus on pitfalls involved in estimating the distribution of the size of price changes, the problems they document certainly add measurement error to available estimates of the cross-sectional distribution of price stickiness that use those data (e.g., Bils and Klenow 2004, Nakamura and Steinsson 2008, Klenow and Kryvtsov 2008). At the same time, we certainly do not want to ignore all the information that micro data can provide.

To strike a balance between extracting information from aggregate data and exploring information contained in the micro data, we employ a full-information Bayesian approach. We use aggregate (time-series) data on nominal and real Gross Domestic Product (GDP) as observables, and incorporate the microeconomic information about the cross-sectional distribution of price stickiness through our prior. In the baseline estimation, we use an uninformative ("flat") prior for the cross-sectional distribution of price stickiness, and hence "let the aggregate data speak." This allows the model to recover the distribution of stickiness that provides the best account of aggregate dynamics. We then supplement the baseline analysis by adopting informative priors based on two empirical distributions of price rigidity: one that takes into account *all* price changes, including sales and product substitutions (derived from Bils and Klenow 2004; henceforth "posted prices"); and another, based on price changes that exclude sales and product substitutions (derived from Nakamura and Steinsson 2008, henceforth, "regular prices").[9]

Our baseline estimation shows that all three mechanisms play an important role in accounting for aggregate dynamics. The estimated model points to heterogeneity in price stickiness and the existence of large real rigidities, which induce strong strategic complementarities in price setting. Importantly, the macro-based estimate of the cross-sectional distribution of price rigidity accords well with the empirical distribution based on regular prices. Moreover, formal statistical model comparisons using informative priors based on the two empirical distributions favor specifications based on these price changes, that exclude sales and product substitutions. This indicates that such price changes are relatively more important for aggregate dynamics than those associated with sales and product substitutions.

Finally, we complement the main findings with additional sets of estimations, in which we effectively shut down each of the three mechanisms at a time by imposing highly informative (or even dogmatic) priors on the cross-sectional distribution of price stickiness or on the degree of real rigidities. Formal statistical model comparisons reveal that our benchmark model (in which all three mechanisms are at

[9]To be clear, the micro information also plays a role in the baseline case, as we choose the support of the price stickiness distribution, which is fixed pre-estimation, to be roughly consistent with the empirical distributions.

play) outperforms alternative models – such as economies with distribution of price stickiness based on posted prices (as opposed to regular prices), with homogeneous price rigidity, or with strategic neutrality (as opposed to strategic complementarity) in price setting.

1.1 Brief literature review

Our work is related to the literature that emphasizes the importance of heterogeneity in price rigidity for aggregate dynamics. However, our focus differs from that of existing papers. Most of the latter focus on the role of heterogeneity in boosting the contract multiplier in calibrated models (e.g., Carvalho 2006, Carvalho and Schwartzman 2008, Nakamura and Steinsson 2010, Carvalho and Nechio 2010, Dixon and Kara 2011). These papers do not address the question of whether such heterogeneity does in fact help sticky-price models fit the data better according to formal statistical criteria.

In terms of empirical work on the importance of heterogeneity in price stickiness, Imbs et al. (2011) study the aggregation of sectoral Phillips curves, and the statistical biases that can arise from using estimation methods that do not account for heterogeneity. They rely on sectoral data for France, and find that the results based on estimators that allow for heterogeneity are more in line with the available microeconomic evidence on price rigidity. Lee (2009) and Bouakez et al. (2009) estimate multisector DSGE models with heterogeneity in price rigidity using aggregate and sectoral data. They also find results that are more in line with the microeconomic evidence than the versions of their models that impose the same degree of price rigidity for all sectors.[10] Taylor (1993) provides estimates of the distribution of the duration of wage contracts in various countries inferred solely from aggregate data, while Guerrieri (2006) provides estimates of the distribution of the duration of price spells in the U.S. based on aggregate data. Both models feature ex-post rather than ex-ante heterogeneity in nominal rigidities, as is the case in our model.[11] Coenen et al. (2007) estimate a model with (limited) ex-ante heterogeneity in *price contracts* using only aggregate data. They focus on the estimate of the Ball-Romer index of real rigidities and on the average duration of contracts implied by their estimates, which they emphasize is in line with the results in Bils and Klenow (2004).[12]

Jadresic (1999) is a precursor to some of the ideas in this paper. He estimates a model with ex-ante

[10] Bouakez et al. (2014) find similar results in an extension of their earlier paper to a larger number of sectors.

[11] Their frameworks are thus closer to the generalized time-dependent model of Dotsey et al. (1997) than to our model with ex-ante heterogeneity.

[12] Their estimated model features indexation to an average of past inflation and a (non-zero) constant inflation objective. Thus, strictly speaking their finding is that the average time between “contract reoptimizations” is comparable to the average duration of price spells documented by Bils and Klenow (2004).

heterogeneous price spells using only aggregate data for the U.S. economy to study the joint dynamics of output and inflation. Similarly to our findings, his statistical results reject the assumption of identical firms. Moreover, he discusses the intuition behind the source of identification of the cross-sectional distribution of price rigidity from aggregate data in his model, which is the same as in our model. Despite these similarities, our paper differs from Jadresic's in several important dimensions. We use a different estimation method, and show the possibility of extracting information about the cross-sectional distribution of price rigidity from aggregate data in a more general context - in particular in the presence of pricing complementarities. Most importantly, the focus of our paper goes beyond assessing the empirical support for heterogeneity in price rigidity from aggregate data. We also investigate the similarities between our macro-based estimates and the available microeconomic evidence, and identify price changes that matter for aggregate dynamics.

Finally, our results speak to the ongoing debate on the role of sales in macroeconomic models. That literature started out as a discussion about whether or not to exclude sales when computing price-setting statistics for macro purposes (Bils and Klenow 2004, Nakamura and Steinsson 2008, Klenow and Kryvtsov 2008). This initial debate was followed by a theoretical literature that provided macroeconomic models with sales and other temporary price changes (Guimaraes and Sheedy 2011, Kehoe and Midrigan 2014). More recently, the literature has focused on the cyclicality of sales and consumer behavior, both in theory and in the micro data (e.g., Coibion et al. 2015, Kryvtsov and Vincent 2014). We provide statistical evidence on the relative performance of macroeconomic models with different distributions of price rigidity that do and do not exclude sales (and product substitutions).

2 The semi-structural model

There is a continuum of monopolistically competitive firms divided into K sectors that differ in the frequency of price changes. Firms are indexed by their sector, $k \in \{1, ..., K\}$, and by $j \in [0, 1]$. The distribution of firms across sectors is summarized by a vector $(\omega_1, ..., \omega_K)$ with $\omega_k > 0$, $\sum_{k=1}^{K} \omega_k = 1$, where ω_k gives the mass of firms in sector k. Each firm produces a unique variety of a consumption good, and faces a demand that depends negatively on its relative price.

In any given period, profits of firm j from sector k (henceforth referred to as "firm kj") are given by:

$$\Pi_t(k, j) = P_t(k, j) Y_t(k, j) - C(Y_t(k, j), Y_t, \xi_t),$$

where $P_t(k,j)$ is the price *charged* by the firm, $Y_t(k,j)$ is the quantity that it sells at the posted price (determined by demand), and $C(Y_t(k,j), Y_t, \xi_t)$ is the total cost of producing such quantity, which may also depend on aggregate output Y_t, and is subject to shocks (ξ_t). We assume that the demand faced by the firm depends on its relative price $\frac{P_t(k,j)}{P_t}$, where P_t is the aggregate price level in the economy, and on aggregate output. Thus, we write firm kj's profit as:

$$\Pi_t(k,j) = \Pi(P_t(k,j), P_t, Y_t, \xi_t),$$

and make the usual assumption that Π is homogeneous of degree one in the first two arguments, and single-peaked at a strictly positive level of $P_t(k,j)$ for any level of the other arguments.[13]

The aggregate price index combines sectoral price indices, $P_t(k)$'s, according to the sectoral weights, ω_k's:

$$P_t = \Gamma\left(\{P_t(k), \omega_k\}_{k=1,...,K}\right),$$

where Γ is an aggregator that is homogeneous of degree one in the $P_t(k)$'s. In turn, the sectoral price indices are obtained by applying a symmetric, homogeneous-of-degree-one aggregator Λ to prices charged by all firms in each sector:

$$P_t(k) = \Lambda\left(\{P_t(k,j)\}_{j\in[0,1]}\right).$$

We assume the specification of staggered price setting inspired by Taylor (1979, 1980). Firms set prices that remain in place for a fixed number of periods. The latter is sector-specific, and we save on notation by assuming that firms in sector k set prices for k periods. Thus, $\omega = (\omega_1, ..., \omega_K)$ fully characterizes the cross-sectional distribution of price stickiness that we are interested in. Finally, across all sectors, adjustments are staggered uniformly over time.

Before we continue, a brief digression about the Taylor pricing model is in order. As will become clear, this model allows us to tell apart real rigidities from nominal rigidities, and to infer the cross-sectional distribution of price stickiness implied by aggregate data. Hence, it serves our purposes well. However, strictly speaking, that model is at odds with the microeconomic evidence on the duration of price spells. Klenow and Kryvtsov (2008), for example, provide evidence that the duration of individual price spells varies at the *quote line* level. However, this evidence does not invalidate the use of the Taylor model for our purposes. In particular, in Section 6 we provide an alternative model in which the duration

[13] This is analogous to Assumption 3.1 in Woodford (2003).

of price spells varies at the firm level, and yet the aggregate behavior of the model is identical to the one presented here. The alternative model can match additional micro facts documented in the literature. Hence, it provides a cautionary note on attempts to test specific models of price setting by confronting them with descriptive micro price statistics. For ease of exposition, we proceed with the standard Taylor pricing specification. But the reader should keep in mind that the aggregate implications that we are interested survive in models that can match the microeconomic evidence in many dimensions.

When setting its price $X_t(k,j)$ at time t, given that it sets prices for k periods, firm kj solves:

$$\max E_t \sum_{i=0}^{k-1} Q_{t,t+i} \Pi\left(X_t(k,j), P_{t+i}, Y_{t+i}, \xi_{t+i}\right),$$

where $Q_{t,t+i}$ is a (possibly stochastic) nominal discount factor. The first-order condition for the firm's problem can be written as:

$$E_t \sum_{i=0}^{k-1} Q_{t,t+i} \frac{\partial \Pi\left(X_t(k,j), P_{t+i}, Y_{t+i}, \xi_{t+i}\right)}{\partial X_t(k,j)} = 0. \tag{1}$$

Note that all firms from sector k that adjust prices at the same time choose a common price, which we denote $X_t(k)$.[14] Thus, for a firm kj that adjusts at time t and sets $X_t(k)$, the prices charged from t to $t+k-1$ satisfy:

$$P_{t+k-1}(k,j) = P_{t+k-2}(k,j) = ... = P_t(k,j) = X_t(k).$$

Given the assumptions on price setting, and uniform staggering of price adjustments, with an abuse of notation sectoral prices can be expressed as:

$$P_t(k) = \Lambda\left(\{X_{t-i}(k)\}_{i=0,\ldots,k-1}\right).$$

Instead of postulating a fully specified model to obtain the remaining equations to be used in the estimation, we assume exogenous stochastic processes for nominal output ($M_t \equiv P_t Y_t$) and for the unobservable ξ_t process; hence, we refer to our model as semi-structural. Given our focus on estimation of parameters that characterize price-setting behavior in the economy in the presence of heterogeneity, our goal in specifying such exogenous time-series processes is to close the model with a set of equations

[14] In Section 6.3 we discuss how the model can be enriched with idiosyncratic shocks that can help it match some micro facts about the size of price changes without affecting any of its aggregate implications.

that can provide it with flexibility relative to a fully-structural model. Such flexibility is useful because it allows us to draw conclusions about price setting that do not depend on details of structural models that are not the focus of our analysis.[15]

2.1 A loglinear approximation

We assume that the economy has a deterministic zero-inflation steady state characterized by $M_t = \overline{M}, \xi_t = \overline{\xi}, Y_t = \overline{Y}, Q_{t,t+i} = \beta^i$, and for all (k, j), $X_t(k, j) = P_t = \overline{P}$, and loglinearize (1) around it to obtain:[16]

$$x_t(k) = \frac{1-\beta}{1-\beta^k} E_t \sum_{i=0}^{k-1} \beta^i \left(p_{t+i} + \zeta \left(y_{t+i} - y_{t+i}^n\right)\right), \tag{2}$$

where lowercase variables denote log-deviations of the respective uppercase variables from the steady state. The parameter $\zeta > 0$ is the Ball and Romer (1990) index of real rigidities. The new variable Y_t^n is defined implicitly as a function of ξ_t by:

$$\left.\frac{\partial \Pi\left(X_t(k, j), P_t, Y_t^n, \xi_t\right)}{\partial X_t(k, j)}\right|_{X_t(k,j)=P_t} = 0.$$

In the loglinear approximation, y_t^n moves proportionately to $\log\left(\xi_t/\overline{\xi}\right)$. Strictly speaking, it is the level of output that would prevail in a flexible-price economy. In a fully specified model this would tie it down to preference and productivity shocks. Here we do not pursue a structural interpretation of the exogenous processes driving the economy.[17] Nevertheless, for ease of presentation we follow the literature and label y_t^n the "natural level of output." The definition of nominal output yields:

$$m_t = p_t + y_t. \tag{3}$$

Effectively, the two exogenous processes, m_t and y_t^n, serve respectively as aggregate demand and aggregate supply shocks in the model: a positive innovation in m_t (y_t^n) moves output and the price level in the same (opposite) directions. Finally, we postulate that the aggregators that define the overall level of

[15] Needless to say, the results are conditional on the particular model of price setting that we adopt. In Section 6 we discuss the extent to which our conclusions may generalize to alternative price-setting specifications.

[16] We write all such approximations as equalities, ignoring higher-order terms.

[17] We think such an interpretation is unreasonable because we take nominal output to be exogenous. In that context, an interpretation of y_t^n as being driven by preference and technology shocks would imply that these shocks have no effect on nominal output (i.e., that they have exactly offsetting effects on aggregate output and prices).

prices P_t and the sectoral price indices give rise to the following loglinear approximations:[18]

$$p_t = \sum_{k=1}^{K} \omega_k p_t(k), \tag{4}$$

$$p_t(k) = \int_0^1 p_t(k,j)\,dj = \frac{1}{k}\sum_{j=0}^{k-1} x_{t-j}(k). \tag{5}$$

Large real rigidities (small ζ in equation (2)) reduce the sensitivity of prices to aggregate demand conditions, and thus magnify the non-neutralities generated by nominal price rigidity. In fully specified models, the extent of real rigidities depends on primitive parameters such as the intertemporal elasticity of substitution, the elasticity of substitution between varieties of the consumption good, and the labor supply elasticity. It also depends on whether the economy features economy-wide or segmented factor markets, whether there is an explicit input-output structure etc.[19]

In the context of our model, ζ is itself a primitive parameter. Following standard practice in the literature, we refer to economies with $\zeta < 1$ as displaying *strategic complementarities* in price setting. To clarify the meaning of this expression, replace (3) in (2) to obtain:

$$x_t(k) = \frac{1-\beta}{1-\beta^k} E_t \sum_{i=0}^{k-1} \beta^i \left(\zeta\left(m_{t+i} - y_{t+i}^n\right) + (1-\zeta)\, p_{t+i}\right). \tag{6}$$

That is, new prices are set as a discounted weighted average of current and expected future driving variables $\left(m_{t+i} - y_{t+i}^n\right)$ and prices p_{t+i}. $\zeta < 1$ implies that firms choose to set higher prices if the overall level of current and expected future prices is higher, all else equal. On the other hand, $\zeta > 1$ means that prices are *strategic substitutes*, in that under those same circumstances adjusting firms choose relatively lower prices.

2.2 Nominal (m_t) and natural (y_t^n) output

We postulate an $AR(p_1)$ process for nominal output, m_t:

$$m_t = \rho_0 + \rho_1 m_{t-1} + ... + \rho_{p1} m_{t-p_1} + \varepsilon_t^m, \tag{7}$$

[18]This is what comes out of a fully-specified multi-sector model with the usual assumption of Dixit-Stiglitz preferences.

[19]For a detailed discussion, see Woodford (2003, chapter 3), Carvalho and Lee (2011) and Carvalho and Nechio (2016).

and an $AR(p_2)$ process for the natural output level, y_t^n:

$$y_t^n = \delta_0 + \delta_1 y_{t-1}^n + ... + \delta_{p_2} y_{t-p_2}^n + \varepsilon_t^n, \tag{8}$$

where $\varepsilon_t = (\varepsilon_t^m, \varepsilon_t^n)$ is i.i.d. $N\left(0_{1\times 2}, \Omega^2\right)$, with $\Omega^2 = \begin{bmatrix} \sigma_m^2 & 0 \\ 0 & \sigma_n^2 \end{bmatrix}$.

2.3 State-space representation and likelihood function

We solve the semi-structural model (3)-(8) with Gensys (Sims, 2002), to obtain:

$$Z_t = \mathcal{C}\left(\theta\right) + G_1\left(\theta\right) Z_{t-1} + B\left(\theta\right) \varepsilon_t. \tag{9}$$

where Z_t is a vector collecting all variables and additional "dummy" variables created to account for leads and lags and ε_t is as defined before. The vector θ collects the primitive parameters of the model:

$$\theta = \left(K, p_1, p_2, \beta, \zeta, \sigma_m, \sigma_n, \omega_1, \cdots, \omega_K, \rho_0, \cdots, \rho_{p_1}, \delta_0, \cdots, \delta_{p_2}\right).$$

In all estimations that follow we make use of the likelihood function $\mathcal{L}\left(\theta | Z^*\right)$, where Z^* is the vector of observed time series (i.e., a subset of Z). Given that our state vector Z_t includes many unobserved variables, such as the natural output level and sectoral prices, the likelihood function is constructed through application of the Kalman filter to the solved loglinear model (9). Letting H denote the matrix that singles out the observed subspace Z_t^* of the state vector Z_t (i.e., $Z_t^* = HZ_t$), our distributional assumptions can be summarized as:

$$\begin{aligned} Z_t | Z_{t-1} &\sim N\left(C\left(\theta\right) + G_1\left(\theta\right) Z_{t-1}, B\left(\theta\right) \Omega B\left(\theta\right)'\right), \\ Z_t^* | \left\{Z_\tau^*\right\}_{\tau=1}^{t-1} &\sim N\left(\mathcal{M}_{t|t-1}\left(\theta\right), V_{t|t-1}\left(\theta\right)\right), \end{aligned}$$

where $\mathcal{M}_{t|t-1}\left(\theta\right) \equiv HC\left(\theta\right) + HG_1\left(\theta\right) \hat{Z}_{t|t-1},\ V_{t|t-1}\left(\theta\right) \equiv HB\left(\theta\right) \hat{\Sigma}_{t|t-1} B\left(\theta\right)' H'$, $\hat{Z}_{t|t-1}$ denotes the expected value of Z_t given $\left\{Z_\tau^*\right\}_{\tau=1}^{t-1}$, and $\hat{\Sigma}_{t|t-1}$ is the associated forecast-error covariance matrix.

2.4 Identification of the cross-sectional distribution from aggregate data

In estimating our multisector model we use only time-series data on aggregate nominal and real output as observables. It is thus natural to ask whether the structure of the model is such that these aggregate data reveal information about the cross-sectional distribution of price stickiness $\omega = (\omega_1, ..., \omega_K)$. As in Jadresic (1999), we start by looking at a simple case where it is easy to show that ω can be inferred from observations of those two aggregate time series. This helps develop the intuition for a more general case for which we also show identification. We then assess the small-sample properties of estimates of ω inferred from aggregate data through a Monte Carlo exercise. As in our estimation, we assume throughout that the discount factor, β, is known.

The key simplifying assumption to show identification in the first case is absence of pricing interactions: $\zeta = 1$. In that case, from (6) new prices $x_t(k)$ are set on the basis of current and expected future values of the two exogenous processes m_t and y_t^n. For simplicity and without loss of generality, assume further that the latter follow the AR(1) processes:

$$m_t = \rho_1 m_{t-1} + \varepsilon_t^m, \text{ and} \tag{10}$$

$$y_t^n = \delta_1 y_{t-1}^n + \varepsilon_t^n. \tag{11}$$

Then, new prices are set according to:

$$x_t(k) = F(\beta, \rho_1, k)\, m_t - F(\beta, \delta_1, k)\, y_t^n,$$

where

$$F(\beta, a, k) \equiv \left(1 + \frac{1-\beta}{1-\beta^k}\frac{\beta a - (\beta a)^k}{1-\beta a}\right).$$

Replacing this expression for newly set prices in the sectoral price equation (5) and aggregating according to (4) produces the following expression for the aggregate price level:

$$p_t = \sum_{j=0}^{K-1}\sum_{k=j+1}^{K} F(\beta, \rho_1, k)\frac{\omega_k}{k} m_{t-j} - \sum_{j=0}^{K-1}\sum_{k=j+1}^{K} F(\beta, \delta_1, k)\frac{\omega_k}{k} y_{t-j}^n. \tag{12}$$

If we observe m_t and y_t - and thus p_t, estimates of the coefficients on m_{t-j} in (12) allow us to infer the sectoral weights ω. The reason is that $F(\beta, \rho_1, k)$ is "known", since ρ_1 can be estimated directly from

(10). Thus, knowledge of the coefficient on the longest lag of m_{t-j} (attained when $j = K-1$) allows us to uncover ω_K. The coefficient on the next longest lag ($m_{t-(K-2)}$) depends on ω_{K-1} and ω_K, which reveals ω_{K-1}. We can thus recursively infer the sectoral weights from the coefficients $F(\beta, \rho_1, k)\frac{\omega_k}{k}$. Moreover, identification obtains with any estimation method that produces consistent estimates of these coefficients.[20]

Checking for identification of ω in the presence of pricing interactions ($\zeta \neq 1$) is slightly more involved. To gain intuition on why this is so, fix the case of pricing complementarities ($\zeta < 1$). Then, because of the delayed response of sticky-price firms to shocks, firms with flexible prices will only react partially to innovations to m_t and y_t^n on impact. They will eventually react fully to the shocks, but also with a delay.

It turns out that the "recursive identification" that applies when $\zeta = 1$ also works in this case. The reason is that, in equilibrium, pricing interactions manifest themselves through a dependence of the aggregate price level on its own lags. This is how they serve as a propagation mechanism. Specifically, the expression for the equilibrium price level becomes:

$$p_t = \sum_{j=1}^{K-1} a_j p_{t-j} + \sum_{j=0}^{K-1} b_j m_{t-j} - \sum_{j=0}^{K-1} b_j y_{t-j}^n, \tag{13}$$

where $a_1, ..., a_{K-1}$, b_0, ..., b_{K-1} are functions of the model parameters. Knowledge of the coefficients on the lags of the aggregate price level and on lagged nominal output again allows us to solve for the sectoral weights – and for ζ.[21]

The intuition behind the identification result in the absence of pricing interactions is clear: the impact of older developments of the exogenous processes on the current price level depends on prices that are sticky enough to have been set when the shocks hit. This provides information on the share of the sector with that given duration of price spells (and sectors with longer durations). More generally, in the presence of pricing interactions, fully forward-looking pricing decisions may also reflect past developments of the exogenous processes. This dependence manifests itself through lags of the aggregate price level. The intuition behind the mechanism that allows for identification extends in a natural way: sectors where prices are more sticky are relatively more important in determining the impact of older

[20] Jadresic (1999) discusses identification in a similar context. The main differences are that he considers a regression based on a first-differenced version of the analogous equation in his model, and assumes $\rho_1 = 1$ and that the term corresponding to $\sum_{j=0}^{K-1} \sum_{k=j+1}^{K} F(\beta, \delta_1, k) \frac{\omega_k}{k} \Delta y_{t-j}^n$ is an i.i.d. disturbance.

[21] In the Appendix we illustrate how the process works in a two-sector model.

shocks to the exogenous processes on the current price level, and vice-versa for sectors where prices are more flexible. Moreover, the relative sizes of the coefficients on past prices and past nominal output in (13) pin down the index of real rigidities ζ.

These results on identification are of little practical use to us if the mechanism highlighted above does not work well in finite samples. To analyze this issue we rely on a Monte Carlo exercise. We generate artificial data on aggregate nominal and real output using parameter values that roughly resemble what we find when we estimate the model with actual data. Then, we estimate the parameters of the model by maximum likelihood. We conduct both a large- and a small-sample exercise. Details and results are reported in the (online) Appendix.

The bottom line is that, for large samples, the estimates are quite close to the true parameter values, and fall within a relatively narrow range. For samples of the same size as our actual sample, we also find the aggregate data to be informative of the distribution of sectoral weights. However, in this case the estimates are less precise and some of them are relatively biased. This finding motivates our supplementary estimation exercise in Section 5.2, where we incorporate prior information from the microeconomic evidence on price-setting.

3 Empirical methodology and data

With the challenges involved in bridging the gap between price-setting statistics based on micro data and time series of aggregate nominal and real output, the choice of empirical methodology is critical. We employ a Bayesian approach, which allows us to integrate those two sources of information and also to compare different models in a formal way.

With some abuse of notation, the Bayesian principle can be shortly stated as:

$$f\left(\theta|Z^{*}\right)=f\left(Z^{*}|\theta\right)f\left(\theta\right)/f\left(Z^{*}\right)\propto\mathcal{L}\left(\theta|Z^{*}\right)f\left(\theta\right),$$

where f denotes density functions, Z^{*} is the vector of observed time series, θ is the vector of primitive parameters, and $\mathcal{L}\left(\theta|Z^{*}\right)$ is the likelihood function.

As observables, we use time series of aggregate nominal and real output. For constructing our prior distribution over the vector of sectoral weights, $f\left(\omega_{1},...,\omega_{K}\right)$, we derive empirical distributions from Bils and Klenow (2004) and Nakamura and Steinsson (2008), as discussed in detail in Subsection 3.1

below. In the ensuing subsections we detail our prior distributions, sources of data, and estimation approach.

3.1 Prior over ω

We specify priors over the set of sectoral weights $\omega = (\omega_1, ..., \omega_K)$, which are then combined with the priors on the remaining parameters to produce the joint prior distribution for the set of all parameters of interest. We impose the combined restrictions of non-negativity and summation to unity of the ω's through a Dirichlet distribution, which is a multivariate generalization of the beta distribution. Notationally, $\omega \sim D(\alpha_1, ..., \alpha_K)$ with density function:

$$f_\omega(\omega|\alpha_1, ..., \alpha_K) \propto \prod_{k=1}^{K} \omega_k^{\alpha_k - 1}, \ \forall \alpha_k > 0, \ \forall \omega_k \geq 0, \ \sum_{k=1}^{K} \omega_k = 1.$$

The Dirichlet distribution is well known in Bayesian econometrics as the conjugate prior for the multinomial distribution, and the hyperparameters $\alpha_1, ..., \alpha_K$ are most easily understood in this context, where they can be interpreted as the "number of occurrences" for each of the K possible outcomes that the econometrician assigns to the prior information.[22] Thus, for given $\alpha_1, ..., \alpha_K$, the parameter $\alpha_0 \equiv \sum_k \alpha_k$ captures, in some sense, the overall level of information in the prior distribution. The information about the cross-sectional distribution of price stickiness comes from the relative sizes of the α_k's. The latter also determine the marginal distributions for the ω_k's. For example, the expected value of ω_k is simply α_k/α_0, whereas its mode equals $(\alpha_0 - K)^{-1}(\alpha_k - 1)$ (provided that $\alpha_i > 1$ for all i).

Whenever we want to estimate a cross-sectional distribution of price rigidity based solely on aggregate data, we can impose an uninformative ("flat") prior, in which all ω vectors in the K-dimensional unit simplex are assigned equal prior density. This corresponds to $\alpha_k = 1$ for all k – and thus $\alpha_0 = K$. This allows us to extract the information that the aggregate data contain about the cross-sectional distribution of price stickiness.

To incorporate microeconomic information in the estimation, we relate the relative sizes of the hyperparameters $(\alpha_1, ..., \alpha_K)$ to the empirical sectoral weights derived from the micro data, and choose the value $\alpha_0 > K$ to determine the tightness of the prior distribution around the empirical distribution. Specifically, let $\widehat{\omega}$ denote the set of sectoral weights from a given empirical distribution. We specify the

[22] Gelman et al. (2003) offers a good introduction to the use of Dirichlet distribution as a prior distribution for the multinomial model.

relative sizes of the hyperparameters $(\alpha_1, ..., \alpha_K)$ so that the mode of the prior distribution for ω coincides with the empirical sectoral weights $\widehat{\omega}$. This requires setting $\alpha_k = 1 + \widehat{\omega}_k (\alpha_0 - K)$. The case of flat priors analyzed previously obtains when $\alpha_0 = K$. Henceforth, we refer to α_0/K as the degree of "prior informativeness".

3.2 Priors on remaining parameters

The remaining parameters of the model fall into three categories that we deal with in turn. Our goal in specifying their prior distributions is to avoid imposing any meaningful penalties on most parameter values – except for those that really seem extreme on an ***a priori*** basis. The first set comprises the ρ's and δ's, parameterizing the exogenous AR processes for nominal and natural output, respectively. These are assigned loose Gaussian priors with mean zero. We choose to fix the lag length at two for both processes, i.e. $p_1 = p_2 = 2$.[23] The second set of parameters consists of the standard deviations of the shocks to nominal (σ_m) and natural output (σ_n). These are strictly positive parameters to which we assign loose Gamma priors. The last parameter is the Ball-Romer index of real rigidity, ζ, which should also be non-negative. This is captured with a very loose Gamma prior distribution, with mode at unity and a 5-95 percentile interval equal to $(0.47, 16.9)$. Hence, any meaningful degree of pricing complementarity or substitutability should be a result of the estimation rather than of our prior assumptions. These priors are summarized in Table 1.[24]

Table 1: Prior distributions for remaining parameters

Parameter	Distribution	Mode	Mean	Std.dev.
ζ	$Gamma\,(1.2, 0.2)$	1.00	6.00	5.48
ρ_j, δ_j	$N\,(0, 5^2)$	0.00	0.00	5.00
σ_n, σ_m	$Gamma\,(1.5, 20)$	0.025	0.075	0.06

Note: The hyper-parameters for the Gamma distribution specify shape and inverse scale, respectively, as in Gelman et al. (2003).

[23] In principle we could specify priors over p_1, p_2 and estimate their posterior distributions as well. However, the computational cost of estimating all the models in the paper is already quite high, and we restrict ourselves to this specification with fixed number of lags. Our conclusions are robust to alternative assumptions about the number of lags (see Section 6).

[24] We do not include β in the estimation, and set $\beta = 0.99$.

3.3 Macroeconomic time series

We estimate the model using quarterly data on nominal and real output for the U.S. economy. These are measured as seasonally-adjusted GDP at, respectively, current and constant prices, from the Bureau of Economic Analysis. We take natural logarithms and remove a linear trend from the data. Whereas the assumptions underlying the model include one of an unchanged economic environment, the U.S. economy has undergone profound changes in the recent decades, including the so-called "Great Moderation" and the Volcker Disinflation. As a consequence, we choose not to confront the model with the full sample of post-war data. We use the period from 1979 to 1982 as a pre-sample, and evaluate the model according to its ability to match business cycle developments in nominal and real output in the period 1983-2007.[25]

3.4 Empirical distributions of price stickiness

We work with the statistics on the frequency of price changes for the 350 categories of goods and services ("entry level items") reported by Bils and Klenow (2004, henceforth BK), and with the 272 entry level items covered by Nakamura and Steinsson (2008, henceforth NS). In the latter case we use the statistics for regular prices (those excluding sales and product substitutions). We refer to the corresponding empirical distributions of price rigidity as distributions *with* (BK) and *without* (NS) *sales*.

Our goal is to map those statistics into an empirical distribution of sectoral weights over spells of price rigidity with different durations. We work at a quarterly frequency, and for computational reasons consider economies with at most 8 quarters of price stickiness. Sectors correspond to price spells which are multiples of one quarter. We denote an empirical cross-sectional distribution of price rigidity by $\{\widehat{\omega}_k\}_{k=1}^{8}$, where $\widehat{\omega}_1$ denotes the fraction of firms that change prices every quarter, $\widehat{\omega}_2$ the fraction with an expected duration of price spells between one (exclusive) and two quarters (inclusive), and so on. The sectoral weights are aggregated accordingly by adding up the corresponding CPI expenditure weights. We proceed in this fashion until the sector with 7-quarter price spells. Finally, we aggregate all the remaining categories, which have mean durations of price rigidity of 8 quarters and beyond, into a sector with 2-year price spells. This gives rise to the empirical cross-sectional distributions of price stickiness that we use in our estimation, which are summarized in Table 2. We denote the sectoral weight for

[25]We make use of the pre-sample 1979-1982 by initializing the Kalman filter in the estimation stage with the estimate of Z_t and corresponding covariance matrix obtained from running a Kalman filter in the pre-sample. We use the parameter values in each draw. For the initial condition for the pre-sample, we use the unconditional mean and a large variance-covariance matrix.

sector k obtained from this procedure by $\widehat{\omega}_k$. For each of the BK and NS distributions, we also compute the average duration of price spells, $\widehat{\overline{k}} = \sum_{k=1}^{8} \widehat{\omega}_k k$, and the cross-sectional standard deviation of the underlying distribution, $\widehat{\sigma}_k$=$\sqrt{\sum_{k=1}^{8} \widehat{\omega}_k \left(k - \widehat{\overline{k}}\right)^2}$.

Table 2: Empirical cross-sectional distributions of price stickiness

Parameter	With sales (BK)	Without sales (NS)
$\widehat{\omega}_1$	0.395	0.273
$\widehat{\omega}_2$	0.240	0.071
$\widehat{\omega}_3$	0.116	0.098
$\widehat{\omega}_4$	0.118	0.110
$\widehat{\omega}_5$	0.037	0.060
$\widehat{\omega}_6$	0.033	0.129
$\widehat{\omega}_7$	0.030	0.061
$\widehat{\omega}_8$	0.032	0.198
$\widehat{\overline{k}}^{(*)}$	2.54	4.23
$\widehat{\sigma}_k^{(*)}$	1.86	2.66

(*) In quarters. $\sum \widehat{\omega}_k$ might differ from unity due to rounding.

3.5 Simulating the posterior distribution

The joint posterior distribution of the model parameters is obtained through application of a Markov-chain Monte Carlo (MCMC) Metropolis algorithm. The algorithm produces a simulation sample of the parameters that converges to the joint posterior distribution under certain conditions.[26] We provide details of our specific estimation process in the Appendix. The outcome is a sample of one million draws from the joint posterior distribution of the parameters of interest, based on which we draw the conclusions that we start to report in the next section.

Having obtained a sample of the posterior distribution of parameters from any given model, we can estimate the marginal likelihood (henceforth ML) of the data given the model as:

$$\text{ML}_j = f\left(Z^*|\mathcal{M}_j\right) = \int \mathcal{L}\left(\theta|Z^*, \mathcal{M}_j\right) f\left(\theta|\mathcal{M}_j\right) d\theta, \tag{14}$$

and use it for model-comparison purposes. In (14), $\mathcal{M}_j$ refers to a specific configuration of the model and prior distribution, and $f\left(\theta|\mathcal{M}_j\right)$ denotes the corresponding joint prior distribution. Specifically, we

[26]These conditions are discussed in Gelman et al. (2003, part III).

approximate $\log(\text{ML}_j)$ for each model using Geweke's (1999) modified harmonic mean. We use these estimates to evaluate the empirical fit of the models relative to one another. The ML ratio of two model configurations yields the *Bayes factor*, which, when neither configuration is *a priori* considered more likely, constitutes the *posterior odds*. It indicates how likely the two models are relative to one another given the observed data Z^*.

4 Results based on macro data only (flat prior over ω)

This section reports our main findings. We first show that the aggregate data point to the existence of strong strategic complementarities in price setting and heterogeneity in price stickiness. Importantly, the estimated distribution of price stickiness resembles the empirical distribution based on regular price changes. We then present exercises that help develop intuition for the role of each mechanism through simple counterfactual exercises in which, departing from the estimated model, we "shut down" each mechanism at a time.

4.1 Estimates

Table 3 and Figure 1 report the results for the case of uninformative priors, in terms of marginal distributions for the parameters.[27] The empirical distributions of price rigidity from Table 2 are reproduced in the last columns, for ease of comparison. In what follows, we use the posterior means as the point estimates for the sectoral weights, reported in the third column of the table.[28]

The cross-sectional distributions that we infer from aggregate data conform quite well with the empirical ones. The macro-based estimates imply that approximately 28% of firms change prices every quarter; 43% change prices at least once a year; 13% change prices once every two years. The average duration of price spells is 13 months, and the standard deviation of the duration of price spells is approximately 8 months. These numbers are quite close to the empirical distribution based on regular price changes (last column of the table). The correlation between our macro-based estimates and those empirical weights is 0.63. The correlation of the estimates with the empirical distribution based on posted prices is somewhat lower, at 0.43. This is a first, informal indication that the distribution that excludes

[27] We use a Gaussian kernel density estimator to graph the marginal posterior density for each parameter. The priors on $\bar{k}$ and σ_k are based on 100,000 draws from the prior Dirichlet distribution.

[28] The results are almost insensitive to using alternative point estimates, such as the values at the joint posterior mode, or taking medians or modes from the marginal ditributions and renormalizing so that the weights sum to unity.

Table 3: Parameter estimates under a flat prior

	$\alpha_k = 1$ *for all* k $(\alpha_0 = 8)$			*Empirical distributions*	
				With sales	*W/o sales*
ζ	4.440 (0.466;16.863)	0.042 (0.015;0.111)	0.050	–	–
ω_1	0.094 (0.007;0.348)	0.264 (0.099;0.493)	0.276	0.395	0.273
ω_2	0.094 (0.007;0.348)	0.072 (0.007;0.212)	0.086	0.240	0.071
ω_3	0.094 (0.007;0.348)	0.020 (0.002;0.078)	0.027	0.116	0.098
ω_4	0.094 (0.007;0.348)	0.027 (0.002;0.107)	0.037	0.118	0.110
ω_5	0.094 (0.007;0.348)	0.144 (0.017;0.337)	0.156	0.037	0.059
ω_6	0.094 (0.007;0.348)	0.123 (0.011;0.345)	0.144	0.033	0.129
ω_7	0.094 (0.007;0.348)	0.120 (0.010;0.353)	0.143	0.030	0.061
ω_8	0.094 (0.007;0.348)	0.112 (0.010;0.323)	0.132	0.032	0.198
$\bar{k}$	4.501 (3.245;5.760)	4.394 (3.214;5.462)	4.37	2.54	4.25
σ_k	2.139 (1.584;2.678)	2.523 (2.112;2.893)	2.62	1.86	2.66
ρ_0	0.000 (−8.224;8.224)	0.000 (−0.001;0.001)	0.000	–	–
ρ_1	0.000 (−8.224;8.224)	1.426 (1.273;1.576)	1.426	–	–
ρ_2	0.000 (−8.224;8.224)	−0.446 (−0.593;−0.296)	−0.446	–	–
σ_m	0.059 (0.009;0.195)	0.005 (0.005;0.006)	0.005	–	–
δ_0	0.000 (−8.224;8.224)	0.002 (−0.002;0.007)	0.003	–	–
δ_1	0.000 (−8.224;8.224)	0.541 (0.270;0.763)	0.532	–	–
δ_2	0.000 (−8.224;8.224)	0.146 (−0.027;0.331)	0.149	–	–
σ_n	0.059 (0.009;0.195)	0.069 (0.030;0.172)	0.081	–	–

Note: The first two columns report the medians of, respectively, the marginal prior and posterior distributions; the third column gives the mean of the marginal posterior distribution; numbers in parentheses correspond to the 5$^{\text{th}}$ and 95$^{\text{th}}$ percentiles; the last columns reproduce the empirical distributions from Table 2.

sales and product substitutions helps the model fit aggregate dynamics better. Below we investigate this possibility by performing formal model comparisons using a standard measure of fit.

The index of real rigidities implies strong pricing complementarities. The posterior mean of ζ is 0.05 and the 95$^{\text{th}}$ percentile equals 0.11, which falls within the 0.10-0.15 range that Woodford (2003) argues can be made consistent with fully specified models. As highlighted by Carvalho (2006), such complementarities interact with heterogeneity in price stickiness to amplify the aggregate effects of nominal rigidities in this type of sticky-price model.

4.2 The role of the three mechanisms

Our baseline estimation reveals that all three mechanisms are useful to account for the dynamics of the aggregate price level and output. To further highlight this result, we contrast the estimated benchmark model with counterfactual economies in which we shut down each of the mechanisms at a time. In addition, we compare model-implied dynamics for inflation and output to those of a restricted bivariate VAR including nominal and real output.

In estimating the VAR we impose the same assumption used in the models, that nominal output is exogenous and follows an AR(2) process. We allow real output to depend on four lags of both itself and nominal output, and to be contemporaneously affected by innovations to nominal output. Estimation is by ordinary least squares.

Figure 2 shows the (mean) impulse response functions of inflation (π_t, first row) and output (y_t, second row) to positive innovations ε_t^m (left column) and ε_t^n (right column) of one standard deviation in size.[29] In addition, we report the responses of the price level (p_t, third row), constructed from those of inflation.

The impulse responses of the benchmark model are fairly close to the data. As expected, the counterfactual models do a worse job at approximating the impulse response functions produced by the VAR. This happens at both short and medium horizons, and in response to both shocks. The impulse response function of inflation is more volatile and switches sign at an earlier date after then shocks in the counterfactual economies, when compared to the benchmark or the VAR. This implies swifter adjustments of the price level in the former: it responds more – especially at short horizons – and peaks earlier (shown in the third row,) which in turn causes the output *gap* $(y_t - y_t^n)$ to respond less and peak earlier (shown

[29]Following the notation of the semi-structural model, in the VAR ε_t^m denotes innovations to nominal output, and ε_t^n denotes the other (orthogonal) innovations.

in the second row).

Let us first consider the homogeneous-firm economy, where all firms change prices every four quarters (i.e. roughly the mean duration of price spells in the benchmark economy) and is otherwise identical to the benchmark. The response of inflation features a *kink* after four quarters. The aggregate price level and output peak earlier, compared to the benchmark model, in which low-frequency firms still play an importanty role after the forth quarter. Furthermore, in the benchmark, pricing interactions between low-frequency and high-frequency firms generate more persistent dynamics than what is implied by the mean frequency of price adjustments, which helps bring the model closer to the data. Introducing multiple sectors that differ in price rigidity smoothes out such kinks, thereby producing empirically more plausible responses.

Turning to the counterfactual economy with the distribution of price stickiness based on posted prices, the overall degree of nominal rigidities turns out to be too low to fit the data well. According to that distribution (forth column of Table 3 – i.e. the BK distribution), the majority of price adjustments are done within two quarters after a shock. In addition, while the aforementioned pricing interactions are still operative, their impact is muted, because the mass of firms with a large degree of price rigidity is small. These properties cause the aggregate price level to respond more and peak earlier.

Finally, real rigidities are especially important to generate persistent aggregate dynamics. In particular, the right column of the Figure 2 shows that the model with strategic neutrality in price setting ($\zeta = 1$) fails to generate a hump-shaped response of the price level and output when shocks are relatively transient (y_t^n). In the absence of pricing interactions, optimizing firms are not held back by non-adjusting firms and thus change their prices by a *full* amount. This property induces a fast adjustment of the aggregate price level, which is at odds with the data.

Overall, we find that the three mechanisms, i) real rigidities, ii) heterogeneity in price stickiness and iii) temporary price changes, are all quantitatively important. In particular, our exercise suggests that real rigidities that generate pricing interactions across heterogeneous firms play a key role to generate persistent aggregate dynamics.

5 Results based on informative priors

Results in the previous section show the importance of the three mechanisms under consideration in our estimated model. But they do not speak to whether models estimated without these features can still provide a good account of the aggregate data. For instance, a homogeneous-firm economy with a greater degree of nominal rigidities, a heterogeneous-firm economy with the distribution of price stickiness based on posted prices, yet with stronger complementarities in price setting, and an economy with no real rigidities, yet with a greater number of low-frequency firms, all have the ability to generate more muted and persistent responses of the aggregate price level. In other words, while the previous counterfactual analysis allows us to understand the role of each mechanism in the estimated model, it does not provide a formal way to assess whether those features are crucial.

We therefore complement our main findings with three additional sets of estimation exercises, which feature tighter (or even dogmatic) priors on the cross-sectional distribution of price stickiness, $\omega = (\omega_1, ..., \omega_K)$, or on the extent of real rigidities, ζ. The main goal is to see whether and how the aggregate data allow us to distinguish between different configurations that have the potential to produce sluggish aggregate dynamics. Overall, we find that the insights from the counterfactual analysis continue to apply in our estimation exercises.

5.1 Homogeneous firms

We first ask how sharply the data allow us to discriminate between multisector models with heterogeneity in price stickiness and one-sector models with homogeneous firms. To that end, we estimate one-sector models with price spells ranging from two to eight quarters. We keep the same prior distributions for all parameters besides the sectoral weights. A one-sector model with price spells of length k, say, can be seen as a restriction of the multisector model, with a degenerate distribution of sectoral weights ($\omega_k = 1$, $\omega_{k'} = 0$ for all $k' \neq k$).

We pick the best-fitting one-sector model according to the marginal likelihood of the data given the models. Results are reported in Table 4.[30] The best-fitting model is the one in which all price spells last for seven quarters. This seems unreasonable in light of the microeconomic evidence. Given the extent of nominal rigidity, not surprisingly the degree of pricing complementarity is smaller. The posterior distributions for the parameters of the nominal output process are quite similar to the ones obtained in the multisector models. Perhaps this should be expected given that this variable is one of the observables used in the estimation. In contrast, the distributions of the parameters of the unobserved driving process are different under the restriction of homogeneous firms. We defer a discussion of what might drive this

[30] Figures that show prior and posterior distributions from the estimation exercises in section 5.1-5.3 are provided in the online appendix.

result to the end of this subsection.

Table 4: Best-fitting homogeneous economy

	Prior	$K=7, \omega_7=1$	
ζ	4.440 (0.466;16.863)	0.362 (0.193;0.830)	0.419
ρ_0	0.000 (−8.224;8.224)	0.000 (−0.001;0.001)	0.000
ρ_1	0.000 (−8.224;8.224)	1.430 (1.284;1.568)	1.428
ρ_2	0.000 (−8.224;8.224)	−0.454 (−0.590;−0.310)	−0.452
σ_m	0.059 (0.009;0.195)	0.005 (0.005;0.006)	0.005
δ_0	0.000 (−8.224;8.224)	0.003 (−0.003;0.011)	0.004
δ_1	0.000 (−8.224;8.224)	0.064 (−0.154;0.319)	0.071
δ_2	0.000 (−8.224;8.224)	0.135 (−0.027;0.327)	0.141
σ_n	0.059 (0.009;0.195)	0.216 (0.087;0.421)	0.230

Note: The first two columns report the medians of, respectively, the marginal prior and posterior distributions; the third column gives the mean of the marginal posterior distribution; numbers in parentheses correspond to the 5^{th} and 95^{th} percentiles.

The multisector model with $K = 8$ nests the best-fitting homogeneous-firms model. Thus, under measures of fit that do not "correct" for the number of parameters, the former model will necessarily perform at least as well as the latter model. To circumvent that problem we base our comparison on the marginal likelihood, which already accounts for the fact that the multisector model has more parameters than the homogeneous-firms model.[31]

Table 5 reports the results for the multisector model with the flat prior for ω, and the best-fitting one-sector model. The fit of the multisector model is much better than that of the best-fitting one-sector model: the posterior odds in favor of the former model is of the order of $10^{11} : 1$.

Table 5: Model comparison - heterogeneous versus homogeneous economy

	multisector model	Best-fitting 1-sector model
log ML	808.03	781.33

Note: The logarithm of the marginal likelihood of the data given the models (log ML) is approximated with Geweke's (1999) modified harmonic mean.

Our model-comparison criterion has the disadvantage that it does not provide information on what drives the improved empirical fit of the multisector model. To shed some light on this question we com-

[31] The reason is that the vector of parameters is "integrated out" in (14).

pare model-implied dynamics for inflation and output to those of the bivariate VAR, as in the previous section. The impulse response functions from the (benchmark) multisector model and those from the VAR in Figure 3 are the same as those in Figure 2. We include the impulse responses implied by the best-fitting homogeneous-firms model. Since the impulse response functions are conditional on specific parameter values (the posterior means) the comparison does not correct for the larger number of parameters in the multisector model. Thus, it is only meant to provide some indication of the sources of the large differences in the posterior odds of the models.[32]

Figure 3 indicates that, relative to the one-sector model, the estimated multisector model does a better job at approximating the impulse response functions produced by the VAR at both short and medium horizons, in response to both shocks. The overwhelming statistical support for heterogeneity does not seem to depend on any single feature of the dynamic response of macroeconomic variables to the shocks. However, one noticeable feature is once again the *kink* (or abrupt change) in the response of inflation (and thus the price level) around the time when all prices have responded to the shock – now that corresponds to the seventh quarter after a shock. Such feature of the model, which is now more pronounced because the exogenous processe y_t^n is more transient, is clearly at odds with the data. Compared to the results for the homogeneous-firms model in the previous section, the overall size of the responses of inflation and the price level, are now closer to the data (and the benchmark model), thanks to the greater degree of nominal rigidities (i.e. prices are now fixed for 7 quarters instead of 4 quarters).

These results suggest one explanation for why the estimated parameters associated with the unobserved driving process (y_t^n) are different in the one-sector economy. While the multisector model can rely on the distribution of sectoral weights to balance the response of the economy to shocks at different horizons, the one-sector model lacks this mechanism. Given the facts that nominal output is observed and that its parameter estimates imply quite persistent dynamics in both economies, perhaps the one-sector economy needs to rely on the unobserved process as a more transient and volatile component that can help it do a better job at matching higher-frequency features of the data.

5.2 Sales and product substitutions

We now turn to estimations that incorporate information from price-setting statistics derived from micro data. This exercise serves two purposes. First, it produces a formal statistical comparison of two model

[32]In addition, while looking at the impulse response functions is clearly instructive, the usual caveat applies because our likelihood-based estimation method fit the model to the entire autocovariance function of the data, not any specific moments.

economies, with distributions of price stickiness based on posted or regular prices. Second, while identification is possible solely based on aggregate data, the estimates can possibly be less precise in small samples, as suggested by our Monte Carlo exercises. Incorporating what we know based on the micro price data through priors may help the model better identify the role of real rigidities (ζ) separately from that of nominal rigidities implied by $\omega = (\omega_1, ..., \omega_K)$. This is useful *per se*, and also provides some information about the validity of our baseline estimation in Section 4.1.

Table 6 presents the results for two sets of informative priors ($\alpha_0/K = 2,\ 5$) for each empirical distribution. The bottom row of Table 6 reports the log of the marginal likelihood of the various models. For the less informative set of priors ($\alpha_0/K = 2$), the two empirical distributions that inform the prior lead to models that perform similarly in terms of fit – and close to the model estimated under a flat prior. The aggregate data appear to be informative of the parameters of interest. The estimates of ω, even under the prior distribution that includes sales, move toward the empirical distribution of price rigidity obtained when discarding sales. On the other hand, for the more informative set of priors ($\alpha_0/K = 5$), the model with prior based on the empirical distribution without sales fits the data better according to the marginal likelihood criterion – the difference of 4.4 log-points implies a posterior odds ratio of roughly $80 : 1$ in favor of that model.

We can compare marginal likelihoods of various estimated models to assess the relative merits of the two sets of priors for the purpose of helping the model explain aggregate dynamics. To that end, we estimate a series of additional models with informative priors based on the two empirical distributions of price rigidity (with and without sales), progressively increasing the degree of prior informativeness (i.e., increasing α_0/K). Specifically, we estimate additional models with $\alpha_0/K = 10, 20, 100$, and 1000. In addition, we estimate models in which the distribution of price stickiness that forms the prior has equal weights in all sectors ("uniform prior"). We summarize the results in Figure 4. It shows clearly that the difference between the fit of estimated models increases as the priors become more informative. While the difference in fit between the models based on the prior distribution without sales and the uniform prior is not that large (it tends to approximately 3 log-points for very informative priors), the difference between models based on prior distributions with and without sales is more substantial. As the degree of prior informativeness increases, that difference approaches 6 log-points – which translates into a posterior odds ratio of roughly $400 : 1$ in favor of the model with prior distribution that excludes sales and product substitutions.[33]

[33]Figure 4 also shows that, as we tighten the priors on the sectoral weights, the fit of models estimated under priors with

Table 6: Parameter estimates with informative priors

	Inform. prior, $\alpha_0/K=2$		*Inform. prior,* $\alpha_0/K=5$		*Flat prior*	*Empirical distributions*	
	With sales	*W/o sales*	*With sales*	*W/o sales*	*Benchmark*	*With sales*	*W/o sales*
ζ	0.032 (0.01;0.08)	0.042 (0.02;0.11)	0.018 (0.01;0.05)	0.041 (0.02;0.10)	0.042 (0.02;0.11)	–	–
ω_1	0.324 (0.17;0.51)	0.277 (0.14;0.45)	0.425 (0.31;0.55)	0.309 (0.21;0.43)	0.264 (0.099;0.49)	0.395	0.273
ω_2	0.123 (0.04;0.24)	0.069 (0.01;0.18)	0.190 (0.11;0.29)	0.059 (0.02;0.13)	0.072 (0.01;0.212)	0.240	0.071
ω_3	0.035 (0.01;0.09)	0.033 (0.01;0.09)	0.059 (0.02;0.11)	0.051 (0.02;0.10)	0.020 (0.00;0.08)	0.116	0.098
ω_4	0.049 (0.01;0.13)	0.047 (0.01;0.12)	0.081 (0.03;0.15)	0.072 (0.03;0.14)	0.027 (0.00;0.11)	0.118	0.110
ω_5	0.106 (0.02;0.26)	0.109 (0.02;0.26)	0.056 (0.01;0.15)	0.066 (0.02;0.15)	0.144 (0.01;0.34)	0.037	0.059
ω_6	0.100 (0.01;0.27)	0.142 (0.04;0.31)	0.052 (0.01;0.15)	0.144 (0.07;0.25)	0.123 (0.01;0.35)	0.033	0.129
ω_7	0.090 (0.01;0.25)	0.086 (0.01;0.24)	0.042 (0.00;0.13)	0.058 (0.02;0.14)	0.120 (0.01;0.35)	0.030	0.061
ω_8	0.088 (0.01;0.24)	0.160 (0.05;0.32)	0.044 (0.01;0.13)	0.200 (0.11;0.31)	0.112 (0.01;0.32)	0.032	0.198
$\bar{k}$	3.776 (2.91;4.69)	4.367 (3.45;5.25)	2.811 (2.31;3.40)	4.262 (3.60;4.91)	4.394 (3.21;5.46)	2.54	4.25
σ_k	2.515 (2.17;2.85)	2.612 (2.28;2.91)	2.184 (1.81;2.56)	2.725 (2.50;2.93)	2.523 (2.11;2.89)	1.86	2.66
ρ_0	0.000 ($-$0.00;0.00)	0.000 ($-$0.00;0.00)	0.000 ($-$0.00;0.00)	0.000 ($-$0.00;0.00)	0.000 ($-$0.00;0.00)	–	–
ρ_1	1.425 (1.27;1.57)	1.427 (1.27;1.58)	1.424 (1.27;1.57)	1.429 (1.28;1.58)	1.426 (1.27;1.58)	–	–
ρ_2	$-$0.445 ($-$0.59;$-$0.30)	$-$0.447 ($-$0.59;$-$0.30)	$-$0.444 ($-$0.59;$-$0.29)	$-$0.449 ($-$0.60;$-$0.30)	$-$0.446 ($-$0.59;$-$0.30)	–	–
σ_m	0.005 (0.00;0.01)	0.005 (0.00;0.01)	0.005 (0.00;0.01)	0.005 (0.00;0.01)	0.005 (0.00;0.01)	–	–
δ_0	0.002 ($-$0.00;0.01)	0.002 ($-$0.00;0.01)	0.002 ($-$0.00;0.01)	0.002 ($-$0.00;0.01)	0.002 ($-$0.00;0.01)	–	–
δ_1	0.514 (0.30;0.72)	0.545 (0.32;0.75)	0.465 (0.28;0.65)	0.563 (0.38;0.75)	0.541 (0.27;0.76)	–	–
δ_2	0.176 (0.01;0.34)	0.151 ($-$0.01;0.32)	0.201 (0.06;0.34)	0.146 ($-$0.01;0.30)	0.146 ($-$0.03;0.33)	–	–
σ_n	0.068 (0.03;0.16)	0.066 (0.03;0.16)	0.072 (0.03;0.17)	0.062 (0.03;0.15)	0.069 (0.03;0.17)	–	–
log ML	807.56	808.27	803.768	808.16	808.03		

Note: The first four columns report the posterior medians under informative priors, and the fifth column reproduces the posterior medians under a flat prior from Table 3; numbers in parentheses correspond to the 5th and 95th percentiles; the last two columns reproduce the empirical distributions from Table 2.

For intuition, we finally report the impulse response functions as in the previous subsections. Figure 5 shows that, in contrast to the counterfactual analysis, the estimated model with prior based on the empirical distribution derived from posted prices (under $\alpha_0/K = 1000$) now generates significantly sluggish responses of the aggregate price level, which brings the model closer to the data and the benchmark. However, the improved performance requires an unusually small value (0.018) of ζ. Most importantly, disproportionately larger weights on flexible sectors in the model with sales ($\omega_1 + \omega_2 = 63.5\%$) continue generating early switches in the sign of the inflation response, and hence early peaks in the responses of the aggregate price level. This feature is clearly at odds with the data and hampers the ability of the model to fit lower-frequency dynamics.

5.3 No real rigidities

Finally, we estimate the model imposing an essentially dogmatic prior on the degree of real rigidities, thereby ruling out meaningful pricing interactions among firms, both within and across sectors. The dogmatic prior for ζ is the uniform distribution on [0.99, 1.01]; those for the other parameters are the same as in the baseline estimation.

Table 7 presents the results. Compared to the benchmark: i) the model without pricing interactions fits the data worse, ii) the weights are redistributed from high-frequency to low-frequency sectors, which leads to the average duration of price spells ($\bar{k}$) of 6.2 quarters, which is significantly larger than those implied by micro data with or without sales (2.5 and 4.3 quarters, respectively,) and iii) the standard deviation of shocks to natural output (σ_n) is smaller.

Once again, we reports the impulse response functions in Figure 6. The absence of real rigidities is partially compensated by the (perhaps unrealistically) bigger degree of nominal rigidities. In addition, the smaller standard deviation of natural output shocks (σ_n) also helps the model to produce muted response of the price level to the shocks. However, the lack of pricing interaction continues causing the price level to respond more and to peak earlier, as in the previous counterfactual analysis. Such discrepancy between the model and the data is more pronounced when nominal output shocks hit the economy (the left panel in the figure). The reason is that the estimated the AR(1) processes for m_t are similar regardless of real rigidities as nominal output is observed.

sales (BK) and priors with a uniform distribution deteriorates. In contrast, the fit of models estimated under priors based on regular prices remains essentially unchanged. This is consistent with our previous finding that sectoral weights estimated under flat priors are somewhat similar to the empirical distribution without sales and product substitutions.

Table 7: Parameter estimates without real rigidities

	No pricing interactions	*Benchmark*	*Empirical distributions* *With sales*	*W/o sales*
ζ	0.999 (0.99;1.01)	0.042 (0.02;0.11)	–	–
ω_1	0.026 (0.00;0.07)	0.264 (0.099;0.49)	0.395	0.273
ω_2	0.044 (0.01;0.12)	0.072 (0.01;0.212)	0.240	0.071
ω_3	0.013 (0.00;0.05)	0.020 (0.00;0.08)	0.116	0.098
ω_4	0.021 (0.00;0.08)	0.027 (0.00;0.11)	0.118	0.110
ω_5	0.162 (0.02;0.38)	0.144 (0.01;0.34)	0.037	0.059
ω_6	0.157 (0.02;0.35)	0.123 (0.01;0.35)	0.033	0.129
ω_7	0.234 (0.03;0.52)	0.120 (0.01;0.35)	0.030	0.061
ω_8	0.277 (0.06;0.52)	0.112 (0.01;0.32)	0.032	0.198
$\bar{k}$	6.192 (5.62;6.66)	4.394 (3.21;5.46)	2.54	4.25
σ_k	1.790 (1.40;2.18)	2.523 (2.11;2.89)	1.86	2.66
ρ_0	0.000 ($-$0.00;0.00)	0.000 ($-$0.00;0.00)	–	–
ρ_1	1.307 (1.17;1.43)	1.426 (1.27;1.58)	–	–
ρ_2	$-$0.343 ($-$0.47;$-$0.21)	$-$0.446 ($-$0.59;$-$0.30)	–	–
σ_m	0.005 (0.00;0.01)	0.005 (0.00;0.01)	–	–
δ_0	0.001 ($-$0.00;0.00)	0.002 ($-$0.00;0.01)	–	–
δ_1	0.367 (0.04;0.70)	0.541 (0.27;0.76)	–	–
δ_2	0.342 (0.07;0.59)	0.146 ($-$0.03;0.33)	–	–
σ_n	0.020 (0.01;0.03)	0.069 (0.03;0.17)	–	–
log ML	798.55	808.03		

Note: The first column report the posterior medians under the tight prior on ζ around one, and the second column reproduces the posterior medians under a flat prior from Table 3; numbers in parentheses correspond to the 5^{th} and 95^{th} percentiles; the last two columns reproduce the empirical distributions from Table 2.

6 Robustness

This section discusses the robustness of our findings to different assumptions on the exogenous processes and the number of sectors. In addition, we consider two alternative models of price setting: the Calvo (1983) model and what we term the "Random Taylor" model. The latter produces the exact same results as our model, and yet can speak to a much larger set of empirical facts about price setting derived from micro data.

6.1 Exogenous processes and the number of sectors

Our findings are robust to different prior assumptions for the parameters ρ_i, δ_i, σ_m, σ_n and ζ, as well as different de-trending procedures and specifications for the exogenous time-series processes. In particular, they are robust to using a Baxter and King (1999) filter or first-differences instead of removing linear trends from the data, and to assuming correlated (i.e. $\sigma_{mn} \neq 0$) or AR(3) exogenous processes (i.e., $p_1 = p_2 = 3$).

Also, unreported results with different number of sectors suggest that one needs to allow for "enough" heterogeneity in order to avoid compromising the empirical performance of the model. In particular, the fit of models with $K = 4$ (as in Coenen et al. 2007) is much worse than models with $K \geq 6$. While the differences in empirical performance among models with $K = 6$, $K = 8$ and $K = 10$ are relatively small, the evidence against the specifications with $K = 4$ is quite strong: posterior odds ratios favor models with $K \geq 6$ by an order of $10^5 : 1$.[34] Importantly, while $K = 8$ in the benchmark specification, our main findings are robust to the case of $K = 10$ or $K = 6$: the model-implied (macro-based) ω is closer to the empirical distribution that excludes sales and product substitutions, and the estimate of ζ is small, indicating the existence of strong strategic complementarities in price setting.

In a related exercise, we also estimate two-sector models that feature a flexible-price sector and a sticky-price sector with price spells of length k. While these models are still at odds with the microeconomic evidence on the distribution of price stickiness, adding a flexible-price sector to a homogeneous-firms Taylor model increases model fit significantly – as long as the two sectors are sufficiently heterogeneous (i.e., as long as k is large enough). This reinforces the importance of accounting for heterogeneity in price rigidity. In particular, when $k = 8$, log marginal likelihood is 807.39, which is comparable to

[34] The log marginal likelihood of the data (log ML) increases with the number of sectors, but at a diminishing rate: it is approximately 794, 806, 808 and 809, when $K = 4$, 6, 8 and 10, respectively.

that under the benchmark specification. These two-sector models, however, are too simplistic to allow a meaningful comparisons between the model-implied distribution of price rigidity ω and its empirical counterpart. Hence they do not serve our purposes, which go beyond statistical fit.[35]

The finding that different number of sectors (and thus different ω) can produce similar model fit does not invalidate our identification result on ω and ζ. As is clear from our discussion in Section 2.4, identification of those key parameters is *conditional* on the structure of the economy – including the number of sectors and the maximum duration of price spell (*K*). While the identified ω and ζ are in general specific to the (pre-specified) *support* of the distribution of stickiness, our main findings are robust as long as the support covers sufficiently many different values, in light of the micro evidence. This finding once again underscores our case for incorporating some prior information from the microeconomic evidence on price-setting – especially if one is interested, as we are, in more than statistical model fit.

6.2 Results under the Calvo (1983) model

We also considered versions of the model with Calvo (1983) pricing. Mimicking our baseline analysis, the first step was to show that the model allows for identification of the cross-sectional distribution of price rigidity from aggregate data, and, given that result, that it also allows for separate identification of nominal and real rigidities. Indeed, all identification results go through, and the intuition is very similar to the one in the Taylor model. In the Appendix we provide a thorough discussion of identification, including the case with strategic interactions in price-setting decisions (i.e., index of real rigidities $\zeta \neq 1$).[36]

However, under Calvo pricing, not all of our conclusions are equally robust when it comes to relatively small samples. The reason is that, in the context of our semi-structural framework, identification of heterogeneity in price stickiness under Calvo pricing is weaker than under Taylor pricing. Building on Monte Carlo analysis and analytical insights from simple versions of these two pricing models, we found that clear-cut identification of the distribution of price stickiness depends on whether the observable driving process has high variance relative to the unobservable process.

While this "restriction" applies to both price-setting specifications, the identification problem is more acute under Calvo pricing. Based on Monte Carlo analysis, we found that with our sample size and

[35] The posterior median of the sectoral weights is $\omega = (\omega_1 = 0.583, \omega_8 = 0.417)$, which implies the economy's average duration of price spell is 4.50 quarters. Compared to the benchmark, the main difference is that ζ is smaller (0.034) and the exogenous process y_t^n is more persistent ($\delta_1 = 0.792$, $\delta_2 = 0.044$).

[36] As mentioned previously, the one (well-known) exception if the case with homogeneous firms.

the relative variances for the two exogenous processes implied by our point estimates, the likelihood criterion fails to provide a sharp discrimination between alternative (non-degenerate) distributions of price stickiness under Calvo pricing. This mirrors what we find in the data: under Calvo pricing, they do not allow too sharp a discrimination between different models with heterogeneity in price stickiness. In contrast, given the same sample size and relative variances for those two processes, the version of the model with Taylor pricing provides more information about the underlying distribution of price stickiness – as seen in previous sections.

However, despite that difficulty, our main findings *do* hold under the Calvo pricing model – at least qualitatively. First, on the comparison between models with heterogeneity in price stickiness and models with homogeneous firms, the estimated models provide clear evidence in favor of the former. Specifically, we find that a likelihood-ratio test of the homogeneous Calvo model against multisector versions of the model leads to rejection of the former at significance levels of less than 1%.[37] Second, all estimated models feature $\zeta < 1$, implying strategic complementarities in price setting. Finally, estimations under informative priors derived from the empirical distributions of price stickiness (as described in Section 3) also provide (qualitative) evidence in favor of the distribution that excludes sales and product substitutions.[38]

6.3 An alternative model

As we mentioned in Section 2, the standard Taylor model is, strictly speaking, at odds with the microeconomic evidence on the duration of price spells (e.g., Klenow and Kryvtsov 2008). This inconsistency may be viewed as a weakness of the Taylor model relative to alternatives – in particular the Calvo model, which naturally produces a non-degenerate distribution of the duration of price spells at the firm level.

However, this evidence does not invalidate the use of that model for our purposes. To show that this is the case, here we provide an alternative model in which the duration of price spells varies at the firm level. The model can match the empirical distribution of the durations of price spells. Yet, the aggregate behavior of the model is identical to the one presented in Section 2. Furthermore, this alternative model

[37] Because real and nominal rigidities are not separately identified in our Calvo model with homogeneous price stickiness, comparisons based on the log marginal likelihood are sensitive to the prior on the index of real rigidities (even though we use a very loose prior). Hence, in this case we find it more appropriate to use a (frequentist) criterion based only on the likelihood.

[38] That is, the log marginal likelihood of the data given the model is always higher under informative priors based on the distribution that excludes sales. However, the difference is smaller than in the model with Taylor pricing – about 1.5 log points – and does not decay as noticeably when we increase the degree of prior informativeness within the same range as we did for the Taylor model.

can match additional micro facts documented in the literature – in a similar fashion as the Calvo (1983) model.

There is a continuum of monopolistically competitive firms divided into N *economic* sectors (i.e., not necessarily identified by price stickiness). Sectors are indexed by $n \in \{1, ..., N\}$. The distribution of firms across sectors is summarized by a vector $(\varphi_1, ..., \varphi_N)$ with $\varphi_n > 0$, $\sum_{n=1}^{N} \varphi_n = 1$, where φ_n gives the mass of firms in sector n. Each sector has a (sector-specific) *stationary* cross-sectional distribution of price stickiness. Before setting its price, a firm j in economic sector n makes a draw for the duration of its next price spell, and then sets its price optimally. Notice that the price will be chosen according to the same policy as in the Taylor model (i.e., the optimal price for a spell that will last for a known duration). This implies that, at a given time, firms within a given sector can be further divided into different "groups" depending on the duration of price spells that they draw.

The (also stationary) cross-sectional distribution of price stickiness for the entire economy can be constructed by aggregating across sectors. It can be summarized by a vector of weights over stickiness groups, $(\omega_1, ..., \omega_K)$, with $\omega_k \equiv \sum_{n=1}^{N} \varphi_n \omega_{n,k} \in (0, 1)$. It is easy to show that $\sum_{k=1}^{K} \omega_k = 1$:

$$\sum\nolimits_{k=1}^{K} \sum\nolimits_{n=1}^{N} \varphi_n \omega_{n,k} = \sum\nolimits_{n=1}^{N} \sum\nolimits_{k=1}^{K} \varphi_n \omega_{n,k} = \sum\nolimits_{n=1}^{N} \varphi_n \sum\nolimits_{k=1}^{K} \omega_{n,k} = 1.$$

The exact details of how each firm draws the duration for the new price spell – that is, how firms move around different "stickiness groups" within a sector – are inconsequential for the aggregate dynamics implied by this model. What matters is our assumption that the cross-sectional distribution of price stickiness of each sector is stationary (i.e. $\omega_{n,k}$ is time-invariant), which guarantees the stationarity of the economy-wide distribution of price stickiness. In the Appendix we provide an example with a flexible scheme for drawing durations within each sector, which allows for persistence in the duration of price spells at the firm level.

We can write the log-linear approximate model implied by this "Random Taylor" pricing scheme as:

$$x_t(k) = \frac{1-\beta}{1-\beta^k} E_t \sum_{i=0}^{k-1} \beta^i \left(p_{t+i} + \zeta \left(y_{t+i} - y_{t+i}^n \right) \right),$$

$$p_t = \sum_{n=1}^{N} \varphi_n p_t(n),$$

$$p_t(n) = \sum_{k=1}^{K} \omega_{n,k} p_t\left(n,k\right),$$

$$p_t\left(n,k\right) = \int_0^1 p_t\left(n,k,j\right) dj = \frac{1}{k} \sum_{j=0}^{k-1} x_{t-j}\left(k\right).$$

Note that $p_t\left(n,k\right)$ does not depend on n. Thus, we can rewrite the aggregate price index as:

$$\begin{aligned} p_t &= \sum_{n=1}^{N} \varphi_n p_t\left(n\right) = \sum_{n=1}^{N} \varphi_n \sum_{k=1}^{K} \omega_{n,k} p_t\left(n,k\right) = \sum_{n=1}^{N} \varphi_n \sum_{k=1}^{K} \omega_{n,k} \frac{1}{k} \sum_{j=0}^{k-1} x_{t-j}\left(k\right) \\ &= \sum_{k=1}^{K} \underbrace{\sum_{n=1}^{N} \varphi_n \omega_{n,k}}_{=\omega_k} \frac{1}{k} \sum_{j=0}^{k-1} x_{t-j}\left(k\right) = \sum_{k=1}^{K} \omega_k \tilde{p}_t(k), \end{aligned}$$

where

$$\tilde{p}_t(k) \equiv p_t\left(n,k\right) = \frac{1}{k} \sum_{j=0}^{k-1} x_{t-j}\left(k\right).$$

That is, despite time-variation in the duration of price spells at the firm level, the Random Taylor model implies the exact same aggregate dynamics as our multisector Taylor pricing model. Moreover, it is easy to augment the model with other features that leave aggregate dynamics intact, and yet allow it to match additional micro facts.[39]

Hence, this alternative model provides a cautionary note on attempts to test specific models of price setting by confronting them with descriptive price-setting statistics (e.g., Klenow and Kryvtsov 2008).

7 Conclusion

If prices change frequently and each and every price change contributes wholly to offset nominal disturbances, then nominal price rigidity cannot be the source of large and persistent monetary non-neutralities. Hence, bridging the micro-macro gap on the extent of price rigidity requires that price adjustments somehow fail to perfectly neutralize monetary innovations.

This paper provides evidence on three features – i) real rigidities, ii) heterogeneity in price stickiness and iii) temporary sales – that can reconcile frequent individual price changes with sluggish adjustments of the aggregate price level. To that end, we estimate a semi-structural model that allows us to extract

[39]For brevity we do not present details of the argument here, and refer the interested reader to an early working paper version (Carvalho and Dam 2010).

information about real rigidities and cross-sectional heterogeneity in price stickiness and to discriminate between different distributions of price stickiness from aggregate data.

Our estimation results point to the presence of large real rigidities and a significant degree of heterogeneity in price stickiness. Moreover, the cross-sectional distribution of price stickiness implied by aggregate data is in line with a micro-based empirical distribution that factors out temporary sales. All three mechanisms play an important role in accounting for aggregate dynamics. Real rigidities, which induce strong pricing interactions across heterogeneous firms, are particularly important. This finding warrants additional research on the nature and sources of real rigidities, which are unspecified in our semi-structural model.

As a by-product, we develop a price-setting model that produces the same aggregate dynamics as our multisector model with Taylor pricing and, yet, can match various empirical facts on price setting – including the evidence of variation in the duration of price spells at the quote-line level. Hence the model provides a cautionary note on attempts to test specific models of price setting by confronting them with descriptive price-setting statistics (e.g., Klenow and Kryvtsov 2008).

Finally, our findings reinforce the general point that heterogeneity can matter for aggregate dynamics. Sectors in our model, however, are assumed to differ only in the degree of price stickiness. Other sources of sectoral heterogeneity, which may be correlated systematically with sector-specific price stickiness, will certainly have interesting implications for aggregate dynamics.[40] Heterogeneity along other dimensions is likely to generate the need for additional observables – such as sectoral price and quantity data – in order to identify the underlying cross-sectional distribution. We leave this potentially interesting endeavour for future research.

[40]For example, we refer the reader to Barsky et al. (2007) who show aggregate implications of flexible-price durable sectors, Nakajima et al. (2010) who analyze a model with cross-sectional heterogeneity in real rigidities, and Eusepi et al. (2011) who introduce sector-specific labor shares.

References

[1] Ball, L. and D. Romer (1990), “Real Rigidities and the Non-Neutrality of Money,” *Review of Economic Studies* 57: 183-203.

[2] Barsky, R., C. House, and M. Kimball. (2007), “Sticky-Price Models and Durable Goods” *American Economic Review*, 97(3): 984-998.

[3] Baxter, M. and R. King (1999), “Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series,” *Review of Economics and Statistics*, 81: 575-593.

[4] Bils, M. and P. Klenow (2004), “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy* 112: 947-985.

[5] Bonomo, M., C. Carvalho, R. Garcia, and V. Malta (2014), “Persistent Monetary Non-neutrality in an Estimated Model with Menu Costs and Partially Costly Information,” mimeo available at https://sites.google.com/site/cvianac2/papers.

[6] Bouakez, H., E. Cardia and F. Ruge-Murcia (2009), “The Transmission of Monetary Policy in a multisector Economy,” *International Economic Review* 50: 1243-1266.

[7] __________ (2014), “Sectoral Price Rigidity and Aggregate Dynamics,” *European Economic Review* 65: 1-22

[8] Carvalho, C. (2006), “Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks,” *Frontiers of Macroeconomics* 2, Issue 1, Article 1.

[9] Carvalho, C. and N. Dam (2010), “The Cross-Sectional Distribution of Price Stickiness Implied by Aggregate Data,” FRB of New York Staff Report No. 419.

[10] Carvalho, C. and J. W. Lee (2011), “Sectoral price facts in a sticky-price model,” FRB of New York Staff Reports No. 495.

[11] Carvalho, C. and F. Nechio (2016), “Factor Specificity and Real Rigidities” *Review of Economic Dynamics* 22: 208-222.

[12] Carvalho, C. and F. Schwartzman (2008), “Heterogeneous Price Setting Behavior and Aggregate Dynamics: Some General Results,” mimeo available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2182343.

[13] __________ (2015), “Selection and Monetary Non-Neutrality in Time-Dependent Pricing Models,” *Journal of Monetary Economics* 76: 141-156.

[14] Coenen, G., A. Levin, and K. Christoffel (2007), “Identifying the Influences of Nominal and Real Rigidities in Aggregate Price-Setting Behavior,” *Journal of Monetary Economics* 54: 2439-2466.

[15] Coibion, O., Y. Gorodnichenko, and G. H. Hong (2015), “The Cyclicality of Sales, Regular, and Effective Prices: Business Cycle and Policy Implications,” *American Economic Review* 105: 993-1029.

[16] Dixon, H. and E. Kara (2011), “Contract Length Heterogeneity and the Persistence of Monetary Shocks in a Dynamic Generalized Taylor Economy,” *European Economic Review* 55: 280-292.

[17] Dotsey, M., R. King, and A. Wolman (1997), “State Dependent Pricing and Dynamics of Business Cycles,” Federal Reserve Bank of Richmond Working Paper Series 97-02.

[18] M. Eichenbaum, N. Jaimovich, S. Rebelo, and J. Smith (2014), “How Frequent Are Small Price Changes?” *American Economic Journal: Macroeconomics* 6: 137-55.

[19] Eusepi, S., B. Hobijn, and A. Tambalotti. (2011), “CONDI: A Cost-of-Nominal-Distortions Index” *American Economic Journal: Macroeconomics*, 3(3): 53-91.

[20] Gelman, A., J. Carlin, H. Stern and D. Rubin (2003), *Bayesian Data Analysis*, 2nd edition, Chapman & Hall/CRC.

[21] Geweke, J. (1999), “Using simulation methods for Bayesian econometric models: inference, development and communication”, *Econometric Review* 18: 1-126.

[22] Guerrieri, L. (2006), “The Inflation Persistence of Staggered Contracts,” *Journal of Money, Credit and Banking* 38: 483-494.

[23] Guimaraes, B. and K. Sheedy, (2011), “Sales and Monetary Policy,” *American Economic Review* 101: 844-76.

[24] Imbs, J., Jondeau, E., and F. Pelgrin (2011), "Sectoral Phillips Curves and the Aggregate Phillips Curve," *Journal of Monetary Economics* 58: 328–344.

[25] Jadresic, E. (1999), "Sticky Prices: An Empirical Assessment of Alternative Models," IMF Working Paper # 99/72.

[26] Kehoe, P. and V. Midrigan (2014), "Prices Are Sticky After All," forthcoming in the *Journal of Monetary Economics*.

[27] Klenow, P. and B. Malin (2011), "Microeconomic Evidence on Price Setting," in: *Handbook of Monetary Economics 3A*, B. Friedman and M. Woodford (eds.), Elsevier.

[28] Klenow, P., and O. Kryvtsov (2008), "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123: 863-904

[29] Klenow, P. and J. Willis (2006), "Real Rigidities and Nominal Price Changes," Federal Reserve Bank of Kansas City Research Working Paper 06-03, available at http://www.kc.frb.org/Publicat/Reswkpap/rwpmain.htm.

[30] Kryvtsov, O., and N. Vincent (2014), "On the Importance of Sales for Aggregate Price Flexibility", mimeo available at https://sites.google.com/site/okryvtsov/.

[31] Lee, Jae Won (2009), "Heterogeneous Households, Real Rigidity, and Estimated Duration of Price Contracts in a Sticky Price DSGE Model," mimeo available at https://sites.google.com/site/jwlee7/research.

[32] Maćkowiak, B. and F. Smets (2008), "On Implications of Micro Price Data for Macro Models," ECB Working Paper Series 960.

[33] Nakajima, J., N. Sudo, and T. Tsuruga (2010), "How Well do the Sticky Price Models Explain the Disaggregated Price Responses to Aggregate Technology and Monetary Policy Shocks?" IMES Discussion Paper No. 2010-E-22.

[34] Nakamura, E. and J. Steinsson (2008), "Five Facts About Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics* 123: 1415-1464.

[35] __________ (2010), "Monetary Non-Neutrality in a multisector Menu Cost Model," *Quarterly Journal of Economics* 125: 961-1013.

[36] Sbordone, A. (2002), "Prices and Unit Labor Costs: A New Test of Price Stickiness," *Journal of Monetary Economics* 49: 265-292.

[37] Sims, C. (2002), "Solving Linear Rational Expectations Models," *Computational Economics* 20: 1-20.

[38] Taylor, J. (1979), "Staggered Wage Setting in a Macro Model," *American Economic Review* 69: 108-113.

[39] __________ (1980), "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy* 88: 1-23.

[40] __________ (1993), *Macroeconomic Policy in a World Economy: From Econometric Design to Practical Operation*, W. W. Norton.

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

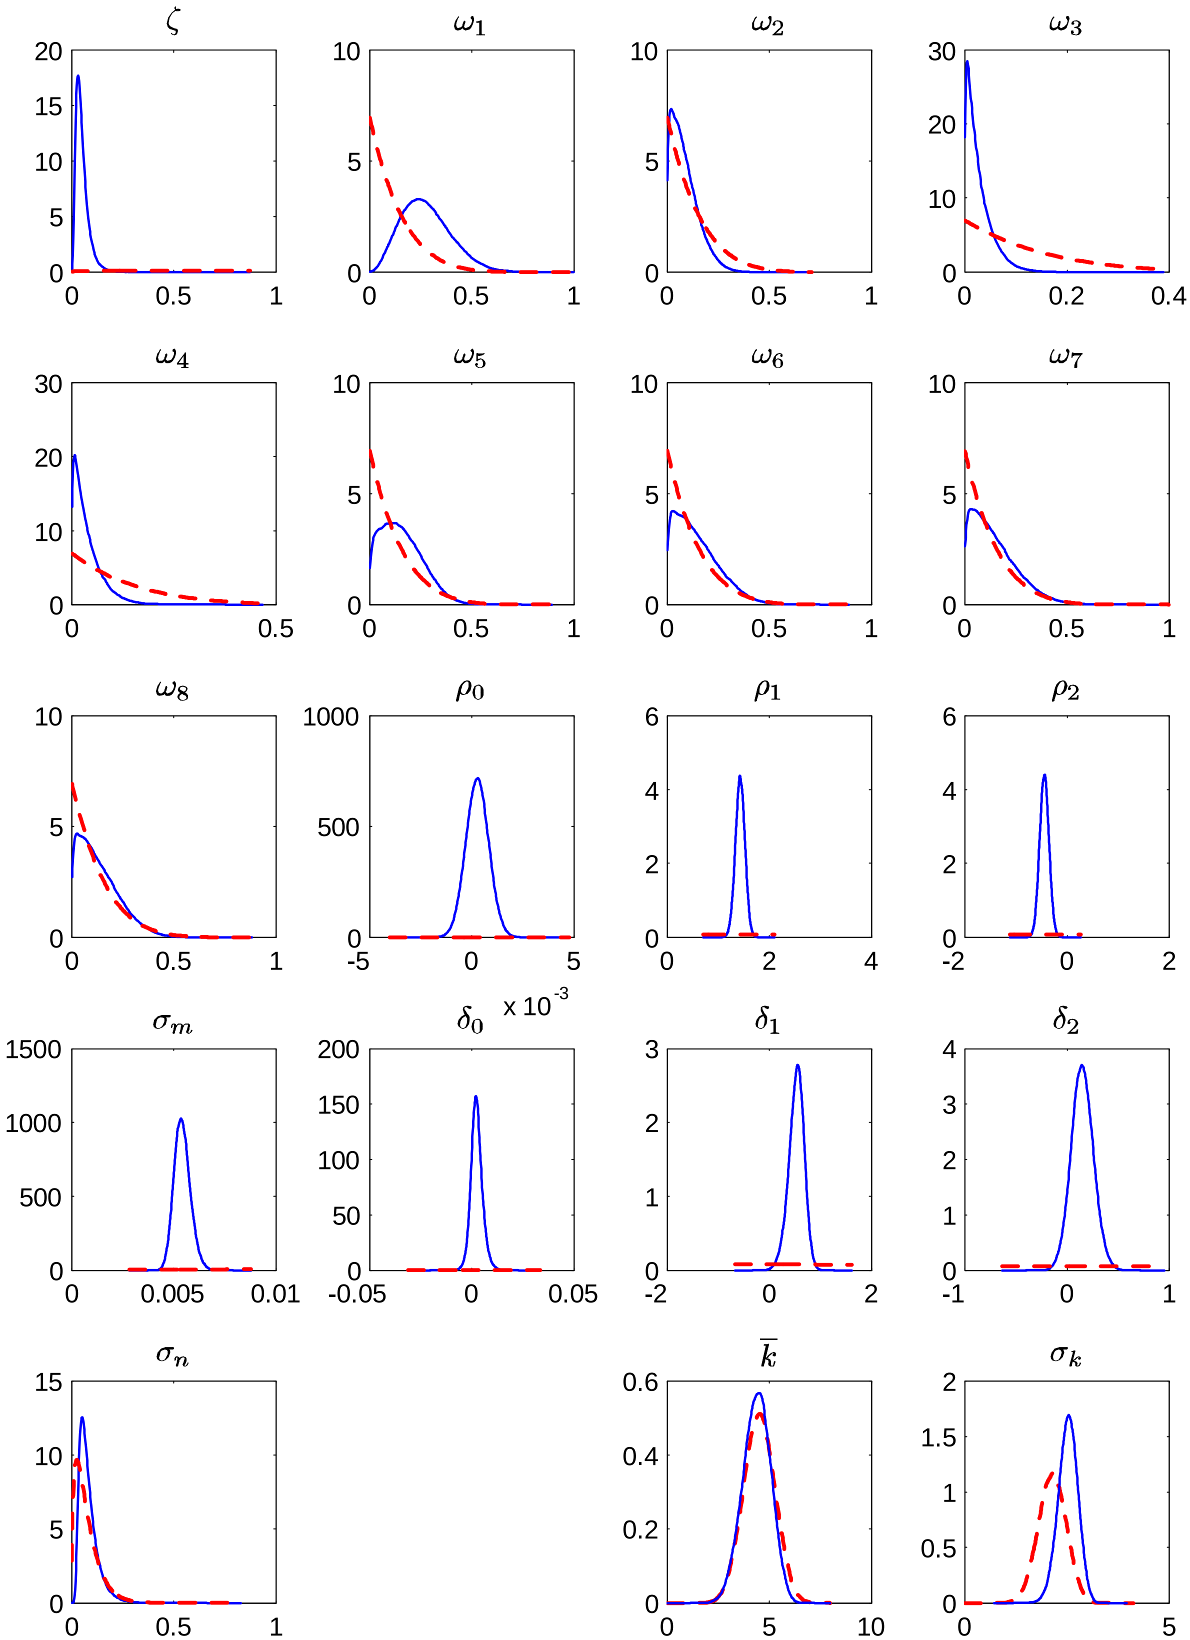

Figure 1: Marginal prior (dashed line) and posterior (solid line) distributions, flat prior

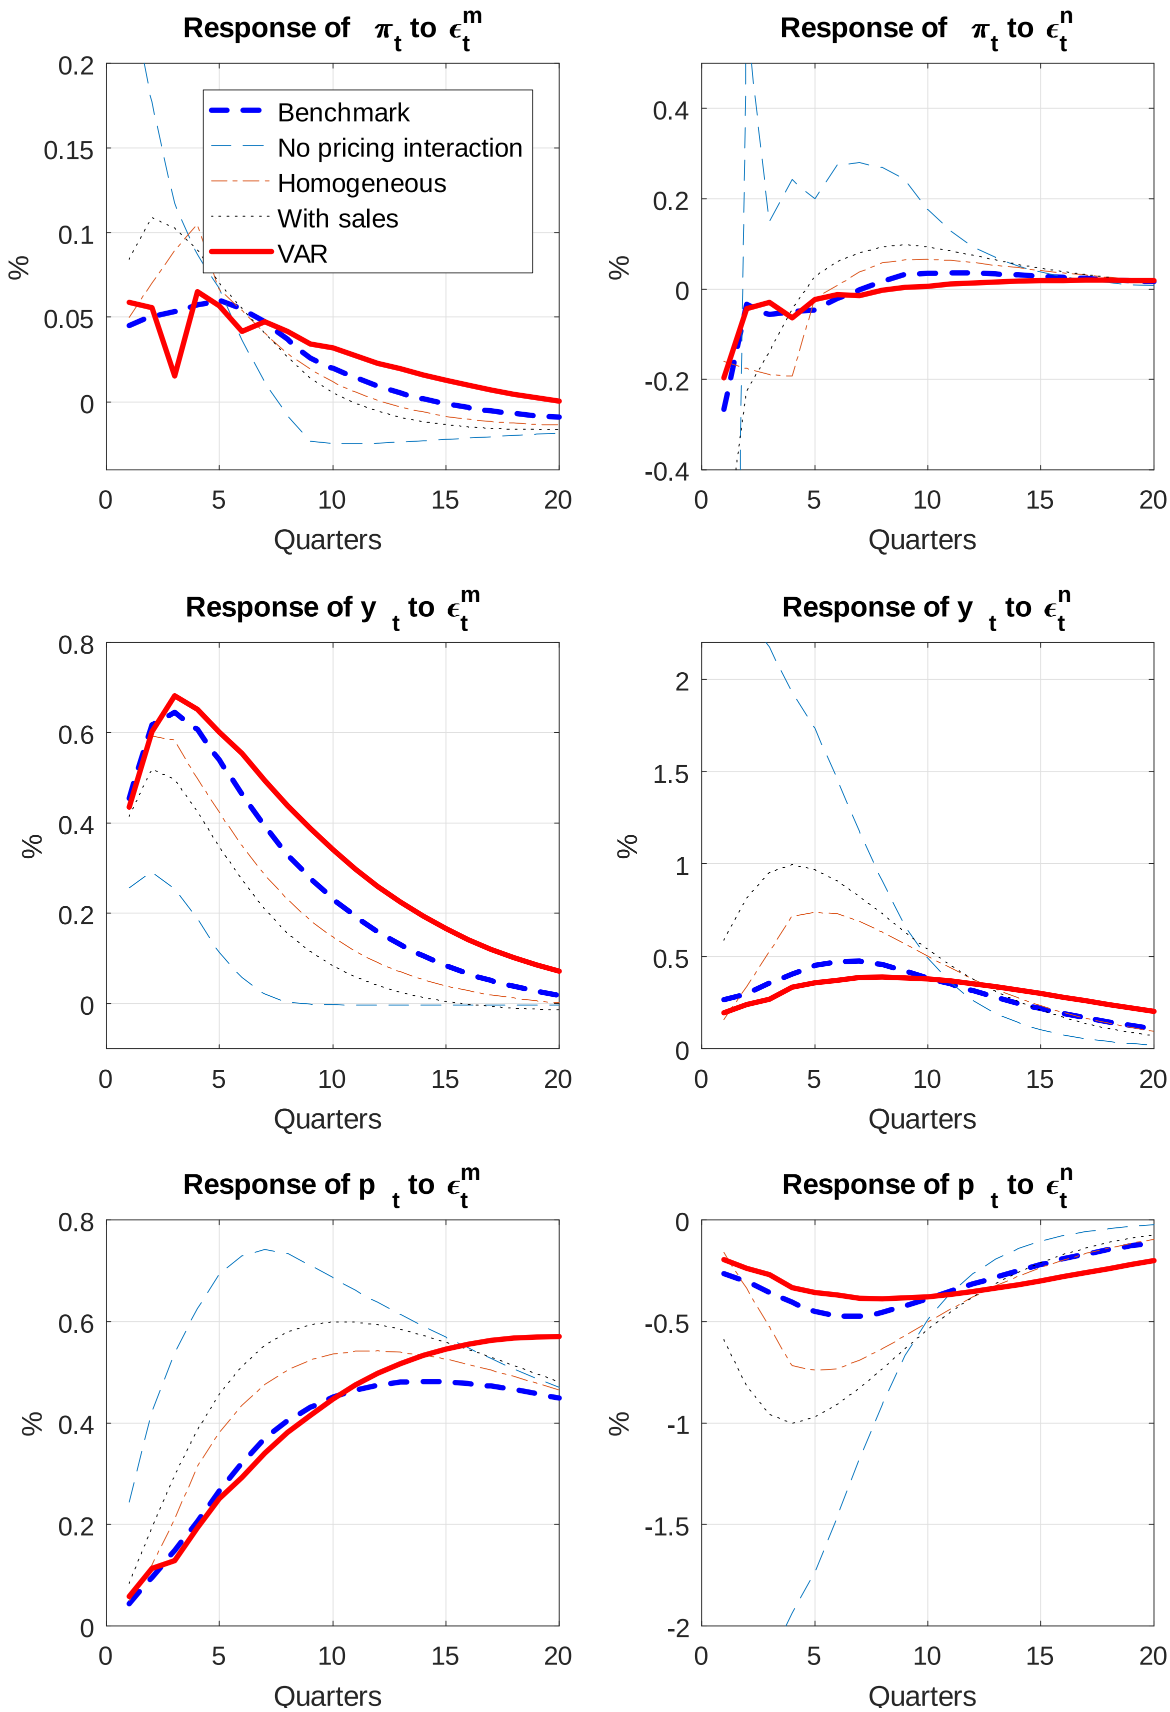

Figure 2: Impulse response functions of models and bivariate VAR

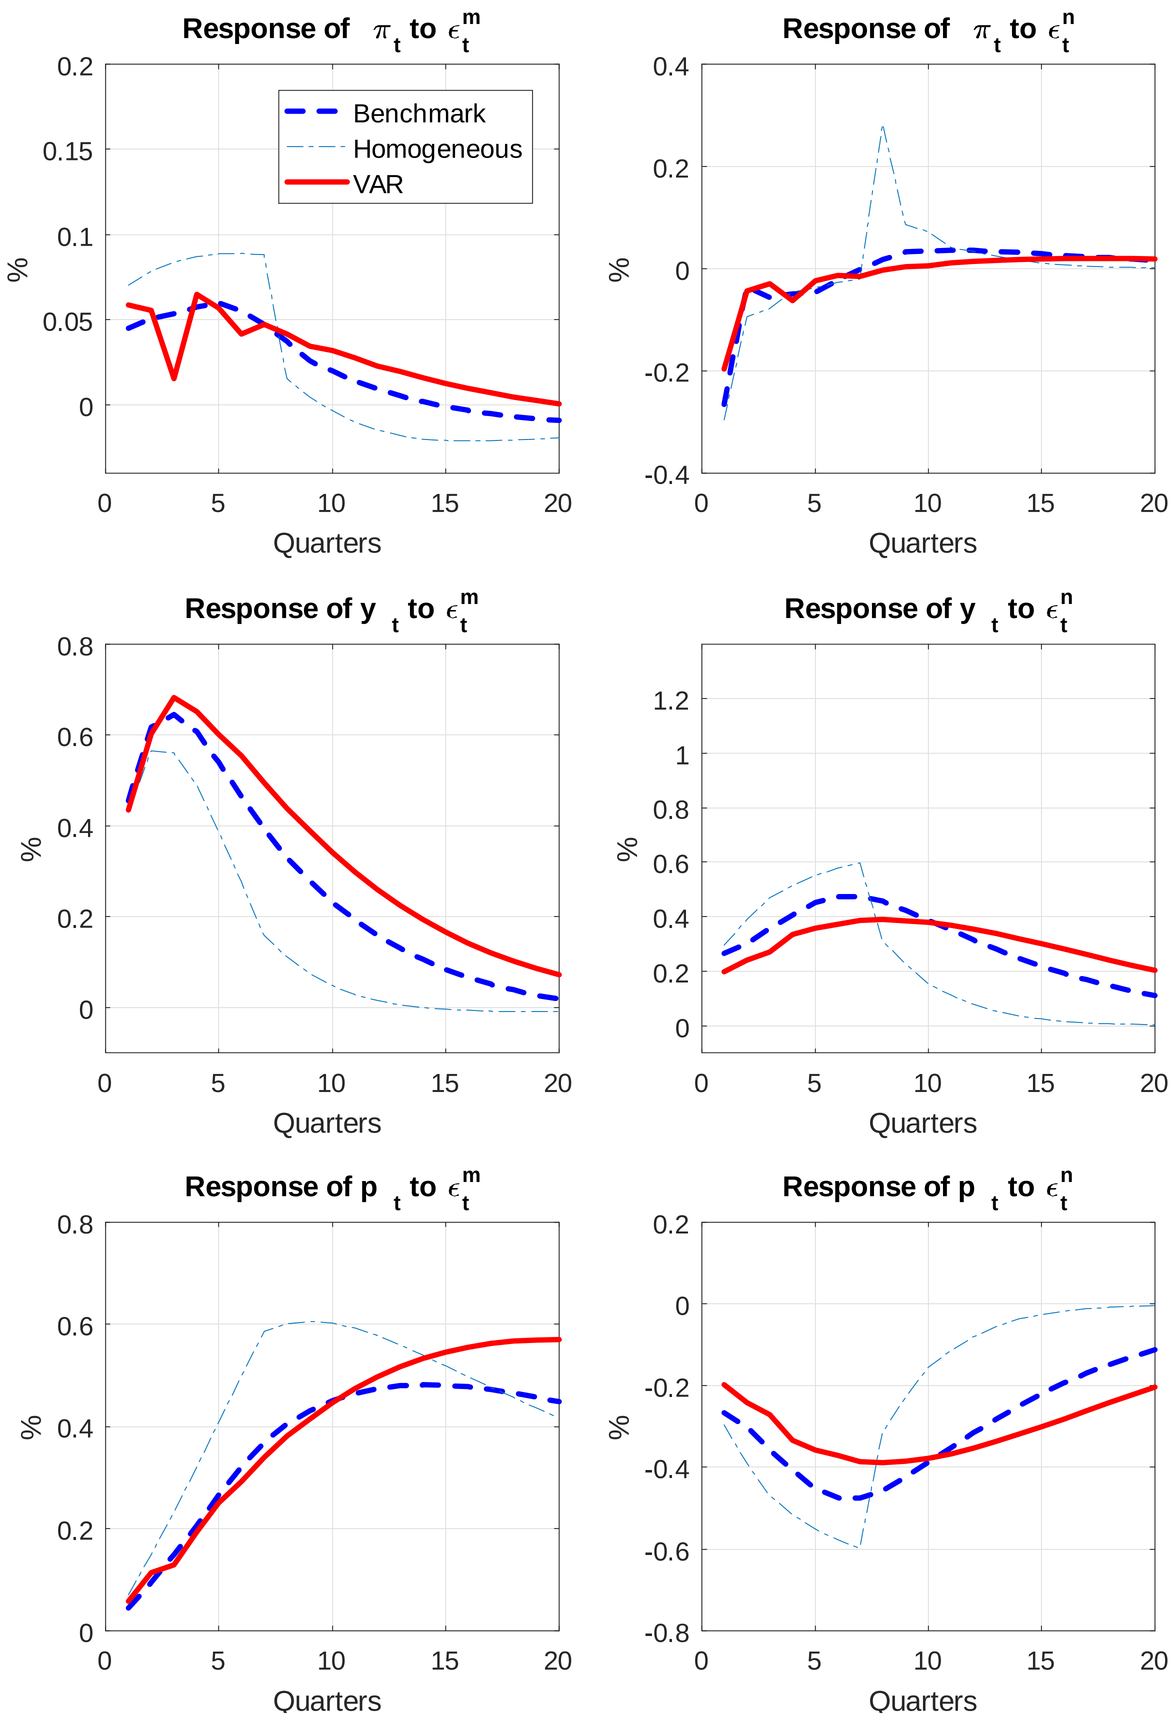

Figure 3: Impulse response functions of models and bivariate VAR: the role of heterogeneity in price stickiness

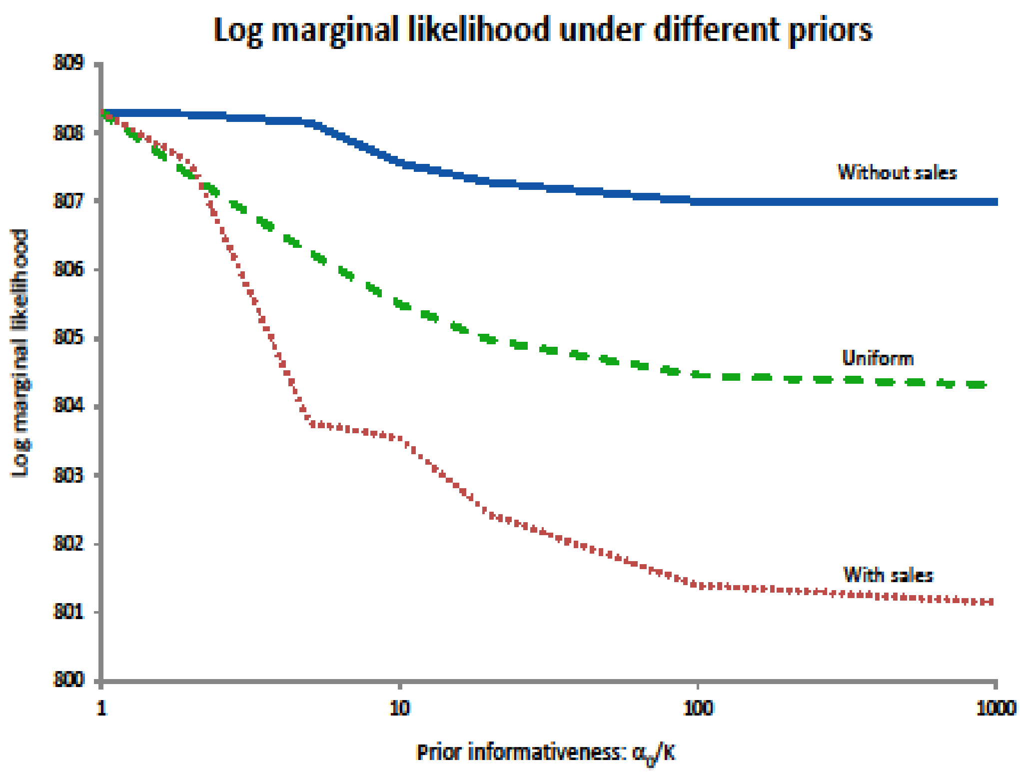

Figure 4: Log marginal likelihood of various models as a function of prior informativeness

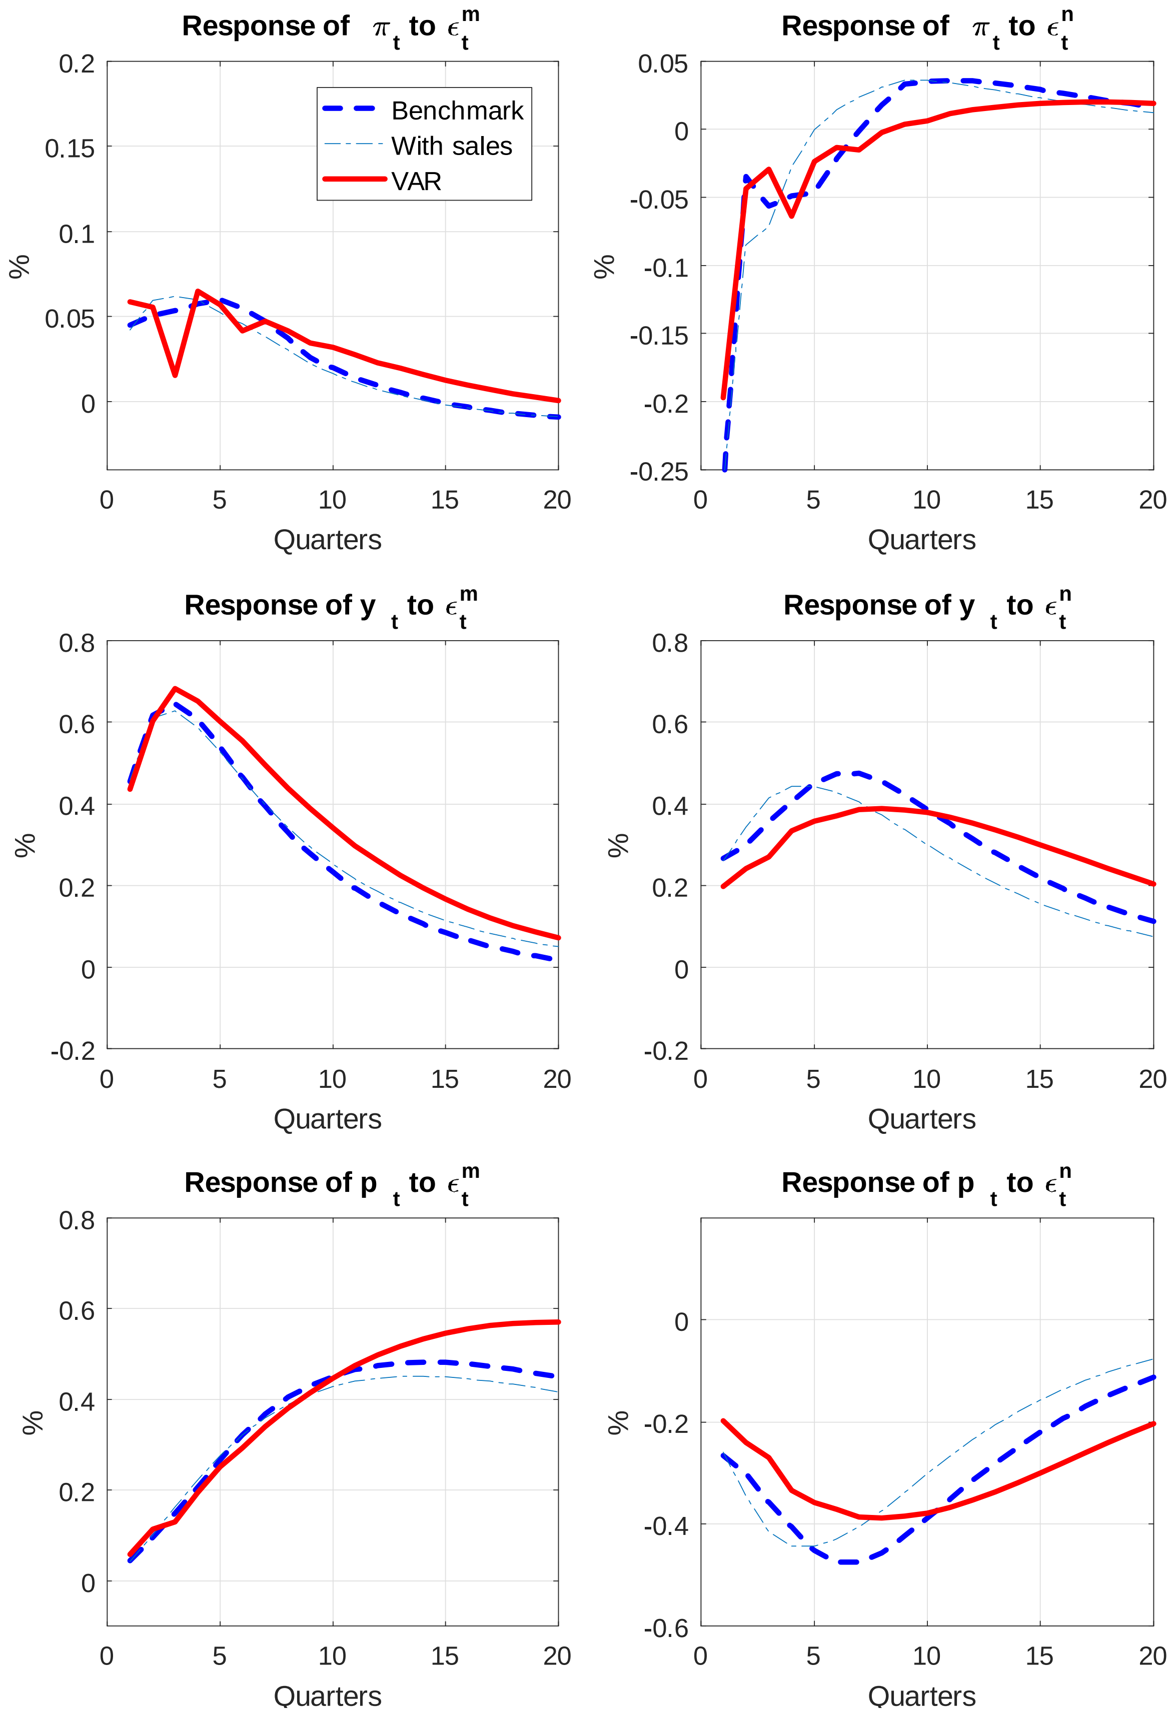

Figure 5: Impulse response functions of models and bivariate VAR: the role of sales

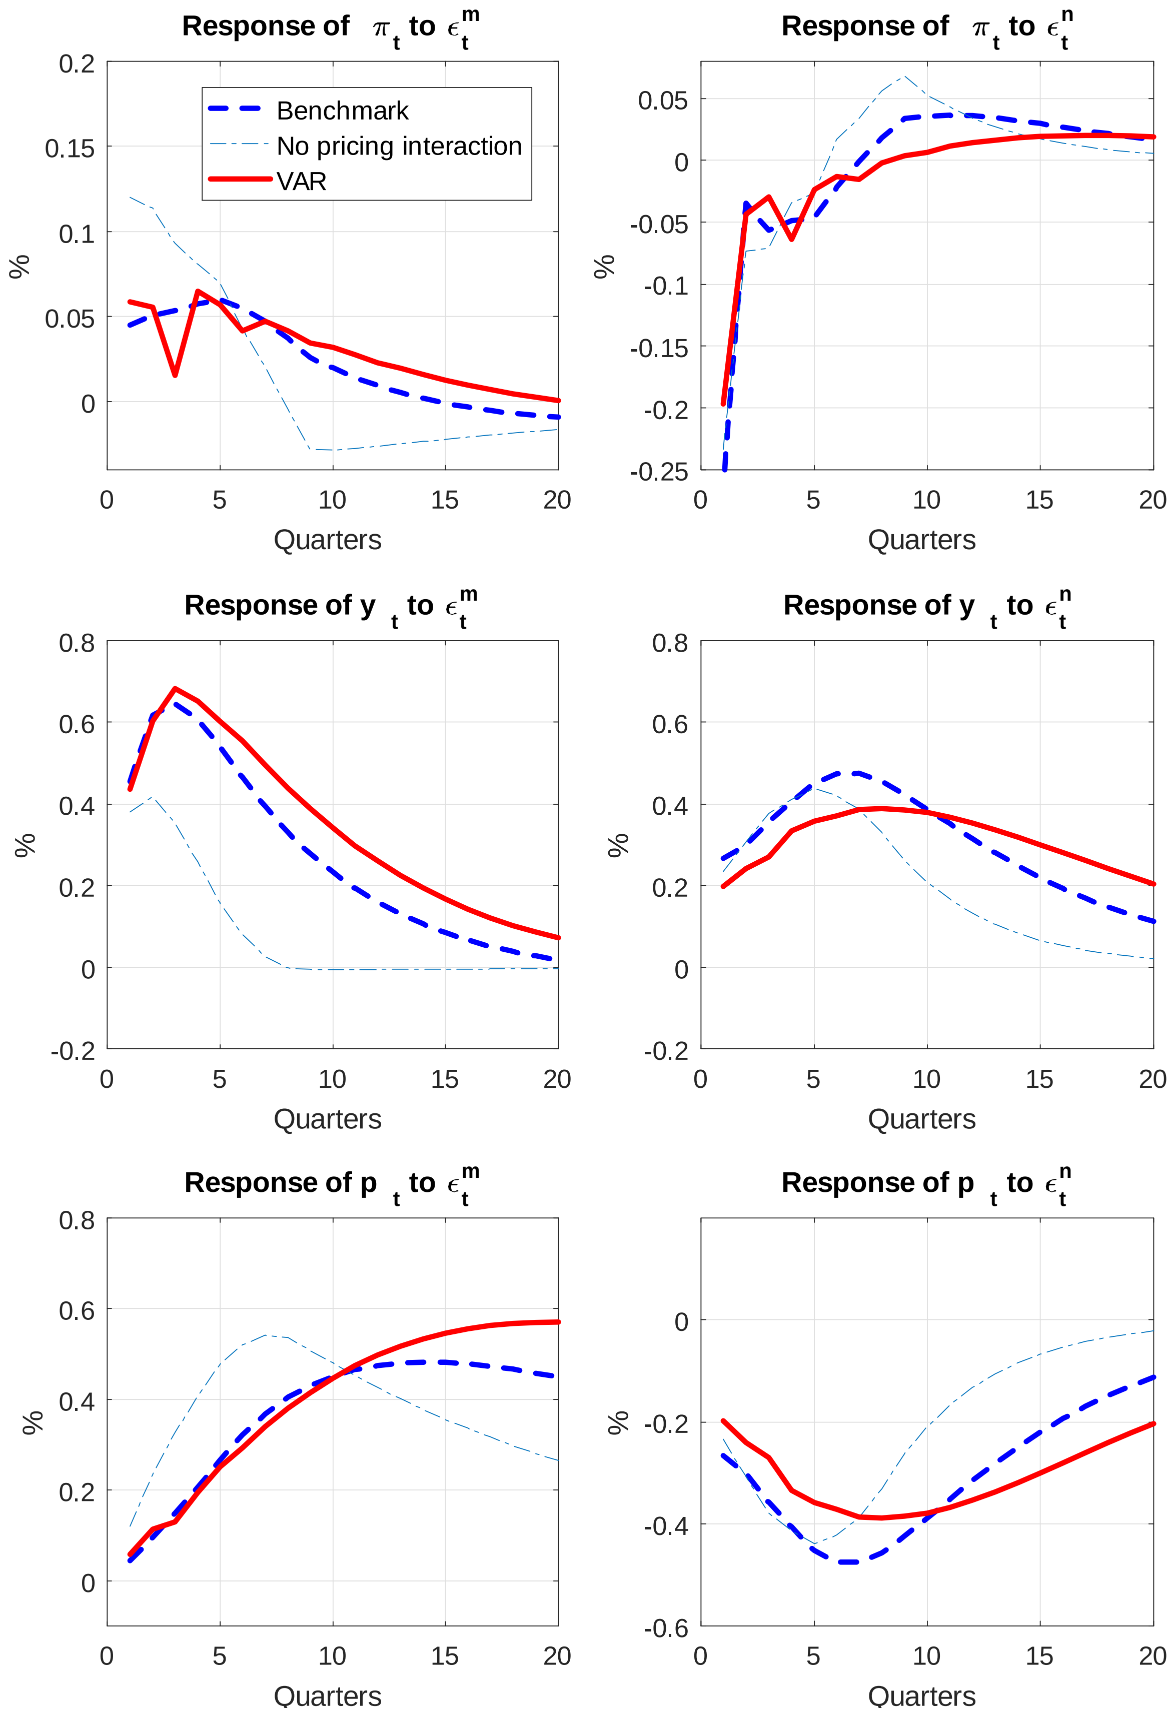

Figure 6: Impulse response functions of models and bivariate VAR: the role of real rigidities