Price dispersion and inflation: New facts and theoretical implications

Viacheslav Sheremirov[a]

[a]*Federal Reserve Bank of Boston, Research Department T-9, 600 Atlantic Ave, Boston, MA 02210, United States*

Abstract

In workhorse macroeconomic models, price dispersion is a central determinant of welfare, the cost of business cycles, optimal inflation, and the tradeoff between inflation and output stability. While price dispersion increases with inflation in the models, this relationship is negative in the data—due to sales prices. The comovement of price dispersion and inflation for *regular* prices is positive. A model with sales can quantitatively match the comovement in the data, whereas a range of similar models without sales cannot, even for regular prices. These findings have important implications for welfare calculations, optimal inflation, and the effects of monetary shocks.

Keywords: Inflation, Price dispersion, Sales, Sticky prices, Welfare
JEL: E31, E37, E40, E52, E58

⋆I would like to thank Urban Jermann (Editor), the associate editor, and the referee for detailed comments and suggestions, as well as Vladimir Asriyan, Olivier Coibion, Benjamin Faber, Pierre-Olivier Gourinchas, Erik Hurst, Amir Kermani, John Leahy, Maurice Obstfeld, Demian Pouzo, Andrés Rodríguez-Clare, David Romer, Yury Yatsynovich, and seminar participants at UNC Chapel Hill, the Bank of Canada, the Federal Reserve Bank of Boston, Simon Fraser University, the University of Illinois at Urbana-Champaign, the Federal Reserve Bank of Cleveland, the College of William and Mary, Vanderbilt University, UC Irvine, UH Manoa, and UC Berkeley. I am especially grateful to Yuriy Gorodnichenko for his invaluable advice, guidance, and support during this project. I acknowledge support from the UC Berkeley Graduate Divison Summer Grant, which I received while a graduate student there. I also thank Keith Barnatchez and Sandra Spirovska for research assistance and Suzanne Lorant for editorial assistance. The views expressed herein are those of the author and are not necessarily those of the Federal Reserve Bank of Boston or the Federal Reserve System.

Email address: `viacheslav.sheremirov@bos.frb.org` (Viacheslav Sheremirov)

March 2019

1. Introduction

The dispersion of prices for homogenous products is a salient feature of micro pricing data. What determines the level and time variation of price dispersion and what does it mean for aggregate analyses? In workhorse macroeconomic models (e.g., Christiano et al., 2005; Smets and Wouters, 2007), inflation is often perceived as an important source of price dispersion. The relation between inflation and price dispersion has significant implications for the dynamic properties of aggregate variables, welfare calculations, and the design of optimal policy.[1] However, different macroeconomic models make conflicting predictions about the level of price dispersion, as well as about its dynamic properties and sensitivity to inflation. In particular, a higher degree of price rigidity implies a stronger response of price dispersion to inflation. The nature of frictions is important too: models with time-dependent frictions produce stronger responses of price dispersion to inflation than models with state-dependent frictions. These contrasting predictions can help us to discriminate across alternative models, which, among other predictions, may differ substantially in their implications for the real effects of nominal shocks (e.g., employment and output responses to monetary policy).[2]

The relationship between price dispersion and inflation is especially important for determining the cost of inflation. Recently, there has been a renewed debate about raising the inflation target, due to secular stagnation and a decrease in the natural rate of interest. In the models employed to study the effects of such policies (e.g., Andrade et al., 2018), inflation reduces the probability of hitting the zero lower bound (ZLB) and allows the real rate to be more negative when the constraint is binding. Importantly, the cost of inflation in these models stems from price dispersion. If price dispersion responds strongly to inflation, raising the target is costly. In fact, many workhorse macro models produce a very large response of dispersion and, therefore, a relatively small optimal inflation rate (e.g., Coibion et al., 2012). However, if price dispersion does not react to inflation strongly, the cost of inflation is low, and the optimal inflation rate is relatively high (e.g., Burstein and Hellwig, 2008; Blanco, 2016). Hence, policymakers can increase welfare by raising the inflation target. Despite the strong link between inflation and price dispersion in these models, empirical evidence on this relationship is rather scarce.

In this paper, I examine the link between price dispersion and inflation in the data and in the models. First, I compute disaggregated inflation and price dispersion, using scanner data from U.S. grocery and drug stores during the period 2001–2011. These data allow me to measure price

[1] For example, Woodford (2003, Ch. 6) shows a negative effect of price dispersion on welfare in New Keynesian models with time-dependent price rigidity. The welfare loss is also present in models with state-dependent pricing and monetary search (e.g., Benabou, 1992; Diamond, 1993; Head and Kumar, 2005).

[2] In particular, in models with time-dependent pricing (e.g., Calvo, 1983), monetary shocks have a nontrivial effect on real variables, whereas in state-dependent models with fixed menu costs (e.g., Golosov and Lucas, 2007), they do not.

dispersion with the variation in prices for an identical product, instead of using indirect proxies such as relative price variability (RPV) or the dispersion of price changes. I then estimate the empirical comovement between price dispersion and inflation at the market–product category level. Disaggregated data enable me to pin down this relationship in the environment of low and stable aggregate inflation, as disaggregated inflation rates are more variable than the aggregate inflation rate. I then use my empirical estimates to discriminate between alternative models of price stickiness.

My first key finding is that predictions of many workhorse macroeconomic models about the relationship between inflation and price dispersion are at odds with the data. A major element that leads to such discrepancy is temporary price changes, or sales. The data feature a weak, positive correlation between price dispersion and inflation when sales are excluded. While workhorse sticky-price models can also produce positive correlation, this correlation in the Calvo model without sales is an order of magnitude larger than in the data, whereas, in a standard state-dependent model, it is significantly smaller than in the data. A hybrid model that combines various pricing frictions can match the positive comovement quantitatively, but only under counterfactual calibrations. None of these models, however, can match the negative correlation observed in the data with temporary sales. My second key finding is that the Calvo model with sales based on market segmentation can quantitatively match the positive correlation in the data without sales and the negative correlation in the data with sales. While this model produces a relatively low welfare cost of inflation, it also implies a relatively high degree of monetary non-neutrality. In light of the key role that price dispersion plays in welfare analysis, my results suggest that workhorse models should take sales—which are often ignored even in medium-scale models—seriously.

Specifically, my main empirical results are as follows: First, the dispersion of prices for exactly the same product sold across retailers within a narrow geographical area (MSA) is pervasive and cannot be fully explained by transitive price discounts or differences in store amenities. The average standard deviation of log prices is 9.5 log points. Even if temporary markdowns are excluded, price dispersion remains sizeable: 6.6 log points. Second, I document a negative correlation between price dispersion and inflation at the market–category level, an empirical regularity contrary to predictions of many macro models. Third, I show that this negative relationship is driven entirely by temporary sales: the correlation between inflation and the dispersion of *regular* prices is positive. Fourth, the negative comovement for posted prices and the positive comovement for regular prices hold not only within market–categories but also in the cross-section and in aggregate data—although they are somewhat more difficult to pin down statistically. Finally, the state of local labor markets, measured with changes in employment, exhibits only small comovement with price dispersion, and does not alter its relationship with inflation.

I then focus on implications of these empirical regularities for aggregate models. I find that the best match of empirical findings, for both posted and regular prices, comes from a Calvo model with sales, as in Guimaraes and Sheedy (2011), calibrated to match the observed frequency of sales. In this model, sales are a source of the price flexibility that does not interfere with the frequency of regular price changes. Without sales, the Calvo model overstates the comovement of price dispersion with inflation by a factor of 15, while the fixed menu cost (FMC) model (Golosov and Lucas, 2007) understates it by a factor of 5. Intuitively, in time-dependent pricing models, most firms cannot adjust their prices in response to an inflationary shock, while those few that can adjust do so by a lot, thereby yielding a strong response of price dispersion and a small response of inflation. In contrast, in state-dependent pricing models, an inflationary shock moves firms outside the $S\,s$ bounds and forces them to reset their prices, thereby inducing a strong impact on inflation and a weak impact on price dispersion. In fact, if menu cost is small, price dispersion may even decrease. The finding that a model with sales matches the properties of regular prices better than a similar model without sales implies that sales have an important interaction with regular prices that is lost when sales are omitted.

To examine in more detail the relative role of different pricing frictions, I turn to the smoothly state-dependent pricing (SSDP) model of Costain and Nakov (2011a,b). This model combines state-dependent frictions with time-dependent frictions, and is especially useful for my exercise not only because it produces aggregate dynamics similar to other hybrid models (e.g., Dotsey et al., 1999; Woodford, 2009), but also because it obtains the purely time- and state-dependent models as limiting cases. The SSDP model naturally comes closer to matching the comovement for regular prices in the data. But to perform as well as the model with sales, it requires parameterization inconsistent with other empirical evidence on price setting. In particular, it requires that price-setting be more state-dependent than suggested by the distribution of price changes.

My results suggest that papers that compute welfare in New Keynesian models (e.g., Coibion et al., 2012) should be more careful about choosing the right measure of price dispersion, as the degree of price dispersion and its comovement with inflation observed in the data are inconsistent with the Calvo model without sales. Using price dispersion observed in the data not only creates a level effect on welfare and the cost of business cycles, but also changes the shape of their relationship with trend inflation and the estimates of the optimal inflation rate. In my simulations of their model, the cost of business cycles increases by 40 percent and the optimal inflation rate decreases by 0.3 percentage points when the sensitivity of price dispersion to inflation is taken from the data.

Finally, Guimaraes and Sheedy (2011) and Kehoe and Midrigan (2015) emphasize that the shape of the output responses to monetary shocks is similar in models with and without sales, even though models without sales cannot fully match the micro price dynamics observed in the data. My analysis is consistent with this implication. However, if the size of the sales sector is relatively large, there could be quantitative differences between the models with and without sales. When the sales model is calibrated to match the sales sector in the scanner data, the output responses in the sales model become 20–25 percent smaller than in the model without sales. Overall, my findings support the literature advocating for a more prominent role for sales in macro models, including papers focusing on understanding recent inflation dynamics (e.g., Stevens, 2019).

This paper is closely related to Nakamura et al. (2018), an influential recent study of the cost of inflation during the Great Inflation of the late 1970s and early '80s. This paper, instead, focuses on a recent period of low-to-moderate inflation, which is more similar to the economic environment of today. The cost of inflation could differ between such environments, because—as Gagnon (2009), Alvarez et al. (2019), and others find in international data—when inflation is high, firms update their prices more often. Thus, it is important to understand the aggregate properties of price dispersion both when inflation is high and when it is low. I extend the analysis in Nakamura et al. (2018) further by employing the richness of scanner data, which (1) enable me to measure price dispersion for the same products directly; and (2) provide a cross-sectional dimension that helps to overcome insufficient variability in aggregate inflation during this period. Overall, both papers find that the mechanism behind the cost of inflation in workhorse models is likely incomplete. In addition, I examine mechanisms that can bring these models closer to the data.

Early studies of the relation between price dispersion and inflation include Van Hoomissen (1988) and Lach and Tsiddon (1992). Due to data availability, these studies focus on price *change* dispersion. Similarly, many subsequent studies focus on RPV, measured as the cross-sector standard deviation of inflation rates (e.g., Grier and Perry, 1996; Debelle and Lamont, 1997; Silver and Ioannidis, 2001; Konieczny and Skrzypacz, 2005; Choi, 2010). Most of these studies find a mildly positive or no relationship between inflation and RPV. In contrast, Reinsdorf (1994) measures price dispersion directly, and finds a negative relationship with inflation. Yet, his data coverage is limited to nine large metropolitan areas and focuses on a relatively short period in the early '80s.

This paper contributes to several other strands of literature. On the empirical front, it is related to papers on price dispersion in micro pricing data (e.g., Pratt et al., 1979; Sorensen, 2000; Lach, 2002; Kaplan and Menzio, 2015). While many of those papers focus on quantifying price dispersion, this paper also studies its properties in the context of aggregate dynamics. Next, this paper contributes to the empirical literature on aggregate price flexibility (e.g., Vavra, 2014; Kryvtsov and Vincent, 2014; Coibion et al., 2015). It exploits cross-sectional variation, which has recently received much attention (e.g., Beraja et al., 2016). On the theory front, it contributes to the literature analyzing pricing frictions in macro models (e.g., Sheshinski and Weiss, 1977; Benabou, 1988; Caplin and Leahy, 1997; Midrigan, 2011; Head et al., 2012; Alvarez and Lippi, 2014), by introducing a new testable prediction and using it to discriminate between models.

The paper proceeds as follows: Section 2 describes the data and measurement. Section 3 quantifies price dispersion in the data. Section 4 presents the empirical strategy and results. Section 5 shows that a model with sales can match the empirical comovement of inflation and price dispersion, while Section 6 shows that models without sales are at odds with the data. Section 7 discusses implications of the results for welfare and monetary policy. Section 8 concludes.

2. Data and measurement

I use scanner data provided by IRI, a market research company.[3] The data contain units and total sales of consumer goods at the Universal Product Code (UPC) level and weekly frequency across U.S. grocery and drug stores during the period 2001–2011. The dataset covers 50 geographical markets, most of which correspond to a single Metropolitan Statistical Area (MSA), and 31 product categories, comprising mostly food and personal-care products. About three-quarters of sellers are grocery stores, and the rest are drug stores. This sector covers 10–15 percent of the U.S. economy. I compute the price as a unit value. Information whether a good was on sale is provided; however, no cost information is available. All private-label UPCs are masked and, therefore, excluded from the calculations.

I compute price dispersion for a given product as a standard deviation of log prices across stores in a given market. This measure is motivated by sticky price models, and is standard in the empirical literature (e.g., Gorodnichenko et al., 2018). Specifically, let P_{ist} be the price of product $i \in \mathcal{G}_c$ in store $s \in \mathcal{S}_m$ in month $t \in \mathcal{T}_y$, where $\mathcal{G}_c$ is the set of goods in product category c; $\mathcal{S}_m$ is the set of stores in geographical market m; and $\mathcal{T}_y$ are the months in calendar year y. Price dispersion $\tilde{\sigma}_{imt}$ is computed as the standard deviation of $\log P_{ist}$ across $s \in \mathcal{S}_m$.[4]

I aggregate this measure across products within a given category, using the annual shares of total sales within markets. That is, if S_{isy} denotes the total annual sales of a given product at the store level, and $\tilde{S}_{imy} \equiv \sum_{s \in \mathcal{S}_m} S_{isy}$ at the market level, the aggregation is as follows:

$$\sigma_{mct} = \frac{\sum_{i \in \mathcal{G}_c} \tilde{S}_{imy} \tilde{\sigma}_{imt}}{\sum_{i \in \mathcal{G}_c} \tilde{S}_{imy}}. \tag{1}$$

My empirical analysis is conducted at the market–category level, but one can use a similar strategy to aggregate price dispersion further to the market or national level.

Next, I construct disaggregated inflation rates, using an enhanced Törnqvist index with annual sales as weights:

$$\pi_{mct} = \frac{\sum_{(i,s) \in \mathcal{G}_c \times \mathcal{S}_m} S_{isy} \log\left(P_{ist}/P_{is,t-1}\right)}{\sum_{(i,s) \in \mathcal{G}_c \times \mathcal{S}_m} S_{isy}}. \tag{2}$$

---

[3] I would like to thank IRI for making the data available. All estimates and analysis in this paper, based on data provided by IRI, are by the author and not by IRI. A detailed data description is provided in Bronnenberg et al. (2008) and Kruger and Pagni (2008).

[4] In practice, I compute price dispersion at the weekly frequency first, and then obtain $\tilde{\sigma}_{imt}$ as the average across weeks. I also use standard deviations across prices weighted by the quantity sold and obtain similar results.

This method is based on the aggregation of individual price changes, with weights analogous to those used for price dispersion. Therefore, the effect of a change in P_{ist} on price dispersion and inflation is not affected by differences in the aggregation procedures. This is important for the analysis of price-setting mechanisms. The aggregate inflation rate computed using this method is highly correlated with the consumer price index (CPI) inflation for food obtained from the Bureau of Labor Statistics (BLS): the correlation coefficient is about 0.8. Online Appendix A provides further details on the disaggregated inflation measure and its robustness to alternative aggregation procedures. For ease of interpretation, I use annualized rates.

Finally, I use the sales flag provided with the data. This indicator is based on a proprietary algorithm that identifies a temporary price reduction of 5 percent or more, and is comparable to popular alternatives (e.g., Nakamura and Steinsson, 2008; Kehoe and Midrigan, 2015). About 20 percent of products were on sale in a given week, with an average discount of about 25 percent.

3. Price dispersion

The average dispersion of prices is reported in column (1) of Table 1. I compute it separately for posted and regular prices, as well as for the sample before and after the onset of the Great Recession. The weekly price dispersion during the entire period is 9.5 log points. Approximately one-third of this measure is due to temporary price reductions: the average dispersion of regular prices is 6.6 log points. Thus, price dispersion is not entirely due to transitory changes (sales). This conclusion holds also when price dispersion is measured across chains (column 2).[5] These estimates are smaller than in Kaplan and Menzio (2015), likely due to differences in sample composition and data collection (e.g., 34 percent of their data come from warehouse clubs, discount stores, and dollar stores).

Table 1: Standard deviation of log price

	Across stores (1)	Across chains (2)	Net of fixed effects (3)	Grocery stores (4)	Drug stores (5)	Food (6)	Bev. (7)	Beauty (8)	Stockpiling High (9)	Stockpiling Low (10)
Posted price	9.5	12.7	8.6	8.8	9.8	9.7	9.8	9.1	9.1	9.4
2001–2007	9.7	12.4	8.7	9.2	9.5	10.0	10.2	9.2	9.1	9.9
2008–2011	9.1	13.1	8.4	8.1	10.3	9.2	9.1	8.9	9.1	8.6
Regular price	6.6	9.5	4.9	6.0	6.0	6.4	6.6	6.9	6.9	6.2
2001–2007	6.7	9.1	4.8	6.1	5.9	6.4	6.8	7.0	6.9	6.4
2008–2011	6.5	10.3	5.0	5.6	6.3	6.3	6.3	6.8	6.9	5.9

Notes: The table reports the average standard deviation of log prices for a UPC-level product in a given market and week. In column (3), fixed effects are removed as in Eq. (3).

[5]The chain price is defined as the average log price across stores in a given chain, computed using stores' annual sales as weights.

I also document a decrease in price dispersion across stores between the 2001–2007 and 2008–2011 period: from 9.7 to 9.1 log points for posted prices and from 6.7 to 6.5 log points for regular prices. While cross-store price dispersion fell during the Great Recession, cross-chain price dispersion rose. A decrease in price dispersion across stores and an increase across chains are consistent with significant mergers and acquisitions (M&A) in the sector taking place since 2007. They may also result from market segmentation (see, e.g., Kaplan et al., 2016; Chevalier and Kashyap, 2019) and increasing search intensity during the Great Recession.

Price dispersion cannot be fully explained by good, store, or time fixed effects. I account for these effects, as well as time-varying store effects and good–store effects, by estimating the equation

$$\log P_{ist} = \alpha_i + \gamma_s + \delta_t + \zeta_{st} + \eta_{is} + \varepsilon_{ist}, \tag{3}$$

and then computing the dispersion of ε_{ist}. Hence, this procedure removes the variation due to some stores charging consistently higher prices for a given good than other stores, as well as some stores being overall more expensive than others in a given month. These effects might be due to, among other things, a chain brand premium, store location, size, amenities, and marginal cost. Yet, I estimate that these effects account for only 10 percent of the standard deviation of posted prices and 25 percent of the standard deviation of regular prices (column 3). Hence, a substantial portion of price dispersion remains unexplained, pointing to welfare loss due to misallocation.[6]

Microeconomic factors such as the elasticity of demand, retailers' market power, product characteristics, and store-specific costs are known to affect the degree of price dispersion (e.g., Gorodnichenko et al., 2018). I find, however, that such factors may have a limited effect on its aggregate properties. In particular, the degree of price dispersion across drug stores is similar to that across grocery stores (columns 4–5), despite the fact that drug stores charge a "convenience" premium, indicating greater market power. The level of price dispersion is also roughly similar across food, beverages, and personal-care items (columns 6–8)—all of which differ in the demand elasticity and the degree of storability—and across the categories of goods that differ in the perceived degree to which they can be stockpiled (columns 9–10), defined as in Bronnenberg et al. (2008). These findings suggest that inventory management, often emphasized in the literature (e.g., Kryvtsov and Midrigan, 2013; Anderson et al., 2017), plays only a limited role in the data.

4. Comovement with inflation

As price dispersion is unlikely to completely stem from micro factors examined in the previous section, can it be explained by aggregate variables? To answer this question, I focus on comovement with inflation, emphasized by aggregate models.

[6] This conclusion holds even when I partial out store-good-year effects.

4.1. *Econometric strategy*

I estimate the comovement of price dispersion and inflation at the market–product category level. This specification is supported by multisector models with sticky prices (e.g., Carvalho, 2006). Panel data also enable me to account for the correlation structure of residuals, trends in variables, and potential time breaks. In particular, I estimate the following specification:

$$\sigma_{mct} = \beta\,\pi_{mct} + \gamma_{mc} + \tau_t + \delta'\,z_{mct} + \varepsilon_{mct}, \tag{4}$$

where σ_{mct} is price dispersion across stores in market m, category c, and month t; π_{mct} is the corresponding disaggregated inflation rate; γ_{mc} and τ_t are market–category and time fixed effects, respectively; z_{mct} is a vector of control variables including local labor market characteristics or lags of π and σ; and ε_{mct} is the error term. All variables are seasonally adjusted using the U.S. Census X-12-ARIMA filter. Standard errors are computed as in Driscoll and Kraay (1998), to account for serial correlation and correlation across groups.

Macroeconomic models often give rise to a nonlinear relationship between inflation and price dispersion, but numerous tests suggest that a parsimonious linear specification provides a useful summary of the comovement (Online Appendix C). This functional form is also supported by the data (e.g., see Fig. E.1 in Online Appendix E). The linearity implies that the slope coefficient β is a natural measure of comovement, which should not be interpreted in a causal way, since both inflation and price dispersion are endogenous variables. The identification of exogenous changes in inflation is beyond the scope of this paper.

4.2. *Empirical estimates*

Table 2 presents estimates of Eq. (4). The comovement of price dispersion and inflation is negative for posted prices (Panel A) and positive for regular prices (Panel B). Columns (1) to (3) show estimates for various combinations of fixed effects. Cross-sectional fixed effects control for heterogeneity across markets and product categories, and allow focusing on the comovement *within* a given market–category. Time fixed effects control for possible time trends. The baseline estimates used for comparison with models include both types (column 3). Quantitatively, a 1 percentage point increase in the annualized inflation rate is associated with a 0.022 log point *decrease* in price dispersion. Once sales are excluded, a 1 percentage point increase in the annualized inflation rate corresponds to a 0.026 log point *increase* in price dispersion.

I examine sensitivity of the baseline estimates to additional controls. First, I control for predetermined trends in the variables, using 12 lags of changes in inflation and price dispersion, and find qualitatively similar results (column 4). Second, I control for local business-cycle conditions, using the local unemployment rate and total employment. The former accounts for the flows between employment and unemployment, while the latter also includes migration and the flows out

Table 2: Comovement of price dispersion and inflation in the data

	Price dispersion					
	(1)	(2)	(3)[b]	(4)	(5)	(6)
Panel A: Posted prices						
Inflation	-0.057^{***}	-0.026^{***}	-0.022^{***}	-0.043^{**}	-0.022^{***}	-0.023^{***}
	(0.010)	(0.010)	(0.008)	(0.019)	(0.008)	(0.008)
Unemployment rate					−0.000	
					(0.000)	
Employment, log						0.025^{***}
						(0.006)
Market–category FE	N	Y	Y	Y	Y	Y
Time FE	N	N	Y	Y	Y	Y
Lags	N	N	N	Y	N	N
Observations	202, 788	202, 788	202, 788	182, 664	202, 788	202, 788
Panel B: Regular prices						
Inflation	0.050^{***}	0.029^{***}	0.026^{***}	0.052^{***}	0.026^{***}	0.026^{***}
	(0.009)	(0.004)	(0.004)	(0.008)	(0.004)	(0.004)
Unemployment rate					0.001^{*}	
					(0.000)	
Employment, log						−0.000
						(0.004)
Market–category FE	N	Y	Y	Y	Y	Y
Time FE	N	N	Y	Y	Y	Y
Lags	N	N	N	Y	N	N
Observations	202, 264	202, 264	202, 264	182, 192	202, 264	202, 264

Notes: This tables presents estimates of Eq. (4). The estimation sample is 2001–2011. The series are seasonally adjusted, using the X-12-ARIMA filter. The employment data are from the BLS. Driscoll and Kraay (1998) standard errors with serial correlation of up to 12 lags are in parentheses. Controls in column (4) include 12 lags of the change in inflation and price dispersion.

[b] denotes baseline specification.

*,**,*** denote the 10, 5, and 1 percent significance level, respectively.

of the labor force, which were particularly important during the Great Recession. I include the two measures separately to avoid collinearity. The slope of the local unemployment rate is insignificant for posted prices, and significant but quantitatively small for regular prices (column 5). The slope of log employment is positive for posted prices (column 6). Importantly, controlling for the state of local labor markets does not alter the estimates of the comovement of price dispersion and inflation.

Due to space constraints, many other robustness checks are relegated to Online Appendix E. For example, measurement choices related to data aggregation or the composition of stores do not have qualitatively relevant effects (Table E.2). Next, as Nakamura et al. (2018) propose to use the absolute size of price changes as a measure of price dispersion, I examine that measure's comovement with inflation. I confirm their finding that there is little response of the size of price changes to inflation: Once the baseline sets of fixed effects are included, the coefficients are small and insignificant (Table E.3, column 3). Hence, this exercise suggests that the price dispersion measured directly from the distribution of product prices has properties different from those of proxies motivated by the models.[7] Next, results with an alternative measure of inflation based on the Laspeyres index, as in Beraja et al. (2016), can be found in Table E.4. The estimates are quantitatively similar to the baseline for regular prices; for posted prices, the correlation is still negative but significantly less pronounced. This discrepancy highlights the importance of applying the same aggregation scheme to both inflation and price dispersion. As further checks, I report estimates of weighted regressions in Table E.7 and estimates of the comovement *between* markets and categories in Table E.8. Fig. E.2 shows that the comovement results hold both in cross-section and in time-series. In addition, I do not find that instances of negative inflation affect the relationship with price dispersion in a material way (Table E.10).

To summarize, the data reveal robust comovement of price dispersion and inflation, negative for posted prices and positive for regular prices. These results hold across a range of econometric specifications, sample periods, aggregation procedures, and dispersion measures. Aggregate variables other than inflation do not seem to have a strong association with price dispersion, or alter its relationship with inflation. Are these findings consistent with aggregate models?

5. A model with sticky prices and flexible sales

In this section, I show that a workhorse macroeconomic model with Calvo pricing that allows for sales can match the empirical comovement of price dispersion and inflation. Sales are modeled as in Guimaraes and Sheedy (2011, henceforth, GS), and following their approach, they are embedded in the general equilibrium model of Erceg et al. (2000). Since this model is well documented, I will summarize the key mechanisms that generates sales, and refer the reader to the original publication for full details.

In the GS model, a variety of product brands is embedded into a variety of product types, with each household having a preferred brand for some products but seeing brands of other products as close substitutes. Specifically, let $\mathscr{T}$ be a measure-one continuum of product types τ, and let $\mathscr{B}$ be a measure-one continuum of brands b for each product $\tau \in \mathscr{T}$. For a given household, there is a set of goods $\Lambda \subset \mathscr{T}$ of measure $\lambda \in (0,1)$ for which the household is *loyal* to a particular brand $\mathcal{B}(\tau) \in \mathscr{B}$, $\tau \in \Lambda$. For all other goods $\tau \in \mathscr{T} \backslash \Lambda$, the household is a *bargain hunter* (i.e., the brands are highly substitutable). The Dixit–Stiglitz consumption aggregator over products and brands is as follows:

$$C = \left(\int_{\Lambda} c\left(\tau, \mathcal{B}(\tau)\right)^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}\tau + \int_{\mathscr{T} \backslash \Lambda} \left(\int_{\mathscr{B}} c\left(\tau, b\right)^{\frac{\eta-1}{\eta}} \mathrm{d}b \right)^{\frac{\eta(\epsilon-1)}{\epsilon(\eta-1)}} \mathrm{d}\tau \right)^{\frac{\epsilon}{\epsilon-1}}, \tag{5}$$

[7] Nakamura et al. (2018) emphasize product heterogeneity observed within a narrow category (e.g., organic vs. regular milk) as a major obstacle to measuring price dispersion directly. In scanner data, however, we can measure price dispersion for *exactly the same* product (e.g., half-gallon of organic, low-fat milk of a given brand).

where $c(\tau, b)$ is consumption of brand b of product τ; ϵ is the elasticity of substitution between product types; and η is the elasticity of substitution between brands of the same product type for a bargain hunter, with $\eta > \epsilon$. (The relative size of the sales sector is $0 < \sigma \leq 1$.) Naturally, the elasticity of substitution across brands for a loyal customer is zero. The continuum of households with heterogenous preferences over brands they are loyal to ensures that households' idiosyncratic preferences do not matter for aggregate consumption of each product and brand.

For each firm producing a given brand, the equilibrium price is a mixed strategy between a regular price and a sales price. The regular price allows the firm to extract consumer surplus from loyal customers, whereas sales are conducted to attract bargain hunters. Because firms compete for the same pool of bargain hunters, sales are strategic substitutes. That is, firms have more incentives to have sales when other firms' products are not on sale. Regular prices are sticky à la Calvo, while sales are flexible: given the sticky regular price, a firm chooses the frequency and size of sales.

In this model, the relation between inflation and regular-price dispersion is positive, since it is driven largely by the Calvo process. As in the standard model, an increase in inflation implies that the firms that do not change their regular price fall further behind from the firms that do and from the aggregate price level. However, the relation between inflation and posted-price dispersion is more nuanced. With sticky prices and flexible sales, the firms whose regular price falls behind the optimal reset price have an incentive to increase the frequency of sales. With more firms having sales, price dispersion across firms may decrease. Thus, the model of sales has a potential to explain the negative comovement between inflation and price dispersion of posted prices observed in the data. Whether the comovement is negative under a plausible parameterization is a quantitative question. Note that the sensitivity of observed regular-price dispersion to inflation is smaller than in the model without sales because firms lagging behind the optimal price are more likely to have sales.

Market segmentation lies at the heart of this selection mechanism: When the relative price of selling to loyal customers falls behind the optimum due to regular-price stickiness, the relative price of selling to bargain hunters remains to be chosen optimally, and therefore the market of bargain hunters becomes relatively more attractive than the market of loyals precisely for the firms that have not adjusted their regular price in a long time. This selection effect—of which firms choose to conduct sales for bargain hunters—gives the sales model a state-dependence flavor.

Following Guimaraes and Sheedy (2011), I calibrate the sales sector to match the average frequency and size of sales in my data. Because the frequency of sales is higher in the scanner data than in the BLS data (20 percent and 7 percent, respectively), the calibrated elasticity of substitution across product types ϵ is slightly higher than in that paper (3.15 vs. 3.01), and the elasticity of substitution across brands η is somewhat lower (16.45 vs. 19.70). Note that the sales sector does not affect the frequency of price changes for regular prices, which is exogenous. I

Table 3: Sales model calibration

	GS notation (3)	Value (4)
Sales sector		
Elasticity of substitution (ES) between product types	ϵ	3.15
ES between brands for a bargain hunter	η	16.45
Fraction of loyals	λ	0.735
Size of sales sector	σ	0.255
Nonsales parameters		
Discount factor	β	$1.03^{-\frac{1}{12}}$
Intertemporal ES	θ_c	0.333
Frisch elasticity of labor supply	θ_h	0.7
Elasticity of output w.r.t. hours	α	0.667
ES between differentiated labor units	ς	20
Price stickiness	ϕ_p	0.889
Wage stickiness	ϕ_w	0.889
Monetary policy shocks		
Persistence	$\mathfrak{p}$	0.536
Volatility	Ω_m	0.02

Notes: The model is calibrated as in Guimaraes and Sheedy (2011), except for two parameters that determine the frequency and size of sales. In the GS original calibration, $\epsilon = 3.01$ and $\eta = 19.70$, leading to a significantly lower frequency of sales than observed in the scanner data.

Table 4: Price dispersion–inflation comovement in the sales model

		Alternative value (1)	Baseline value (2)	Regular prices (3)	Posted prices (4)
Data				0.026	−0.022
Baseline calibration				0.033	−0.033
Sensitivity to parameter values:					
Price stickiness	ϕ_p	0.650	0.889	0.001	−0.008
Wage stickiness	ϕ_w	0.650	0.889	0.029	−0.037
Monetary persistence	$\mathfrak{p}$	0	0.536	0.033	−0.035
Monetary volatility	Ω_m	0.20	0.02	0.182	0.056
ES between products	ϵ	3.01	3.15	0.012	−0.127
ES between brands	η	19.70	16.45	0.002	−0.138
Fraction of loyals	λ	0.950	0.735	0.014	−0.110
Size of sales sector	σ	1	0.255	0.027	−0.019
Calibration to infrequent sales:					
$\epsilon = 3.01$, $\eta = 19.70$ (GS)				−0.000	−0.135
and $\sigma = 1$				−0.003	−0.095
and $\Omega_m = 0.20$				0.047	−0.022

Notes: This tables presents the comovement of price dispersion and inflation for regular prices (column 3) and posted prices (column 4) in the Guimaraes and Sheedy (2011) model. The alternative parameter values used in the sensitivity exercises are in column (1), along with the corresponding baseline values in column (2).

use the same macro parameters as in the GS paper, which employs calibration from the previous literature. Table 3 summarizes parameter values, using exactly the same notation as in the original paper.

I then simulate the path for price dispersion and inflation in this model, and estimate a time-series analog of Eq. (4).[8] The results are presented in Table 4. Under the baseline calibration, the model comes fairly close to matching the comovement for posted prices and regular prices, not only qualitatively but also quantitatively. Although quantitative success clearly depends on parameter values, the direction of the effect is stable for a wide range of parameters. There are two interesting observations. First, the comovement is sensitive to the choice of the price stickiness parameter, but not very sensitive to the choice of wage stickiness. Intuitively, the degree of price stickiness is directly related to the benefits of sales; if regular prices are flexible, there is little advantage from having sales. Second, if shocks are very volatile, the relationship between price dispersion and inflation can become positive for posted prices as well. This is because the absolute size of price changes by adjusters is increasing with inflation, leading to larger price dispersion of regular prices. With a lot of large shocks, this effect dominates the negative effect on dispersion stemming from sales.

6. Sales and state-dependent pricing

In this section, I examine models with pricing frictions that do not allow for sales. I focus on two questions: (1) Can a Calvo model without sales match the comovement, for regular prices, for a plausible degree of aggregate price rigidity? (2) Can a model with state-dependent pricing, or with both frictions, match the data under a plausible calibration? To address these questions in a unified framework, I present simulations of a hybrid model that allows for both time- and state-dependent frictions. This model, referred to as smoothly state-dependent pricing, has been studied extensively in Costain and Nakov (2011a,b), and therefore I relegate the details to Online Appendix G.

To provide intuition on the role of pricing frictions, I overview the price-adjustment mechanism of the model. Under purely time-dependent pricing such as Calvo (1983), the probability of price-adjustment $\lambda(\cdot)$ is constant and independent of the loss from inaction L: $\lambda(L) = \bar{\lambda} \in (0, 1)$. In the fixed menu cost model, this probability is characterized by the indicator function of the loss being above some fixed cost $\alpha > 0$: $\lambda(L) = \mathbb{1}\{L > \alpha\}$. Under SSDP, $\lambda(L)$ is a smooth, increasing function, such that $\lim_{L\to 0} \lambda(L) = 0$, $\lim_{L\to\infty} \lambda(L) = 1$, and the derivative $\lambda_L > 0$. The following

[8] To account for seasonal adjustment in the data, I run the MA(12) filter on both series. I also estimate the same regressions without the filter, and obtain similar results.

functional form satisfies these conditions:

$$\lambda(L) = \frac{\bar{\lambda}}{\bar{\lambda} + \left(1 - \bar{\lambda}\right)\left(\frac{\alpha}{L}\right)^{\xi}}, \tag{6}$$

where $\xi > 0$ controls for the degree of state-dependence, with larger ξ corresponding to more state-dependence. The usefulness of this functional form comes from the fact that the SSDP model converges to Calvo as $\xi \to 0$, and to FMC as $\xi \to \infty$: $\lim_{\xi\to 0} \lambda(L) = \bar{\lambda}$ and $\lim_{\xi\to\infty} \lambda(L) = \mathbb{1}\{L > \alpha\}$. This pricing assumption is imbedded into a general equilibrium model similar to the one studied before, the details of which, including calibration, can be found in Online Appendix G.[9]

Table 5 compares the comovement of price dispersion and inflation in the SSDP model and the two limiting cases with that in the data. The comovement in the Calvo model without sales is more than 15 times greater (for the baseline calibration) than in the data for regular prices, and the negative comovement for posted prices is hard to achieve even in alternative calibrations (see Online Appendix Table G.2). In the baseline FMC model without sales, the comovement coefficient is very small: more than 4 times smaller than in the data. The SSDP model produces the comovement that lies between those in the Calvo and FMC models, and is about 5 times greater than in the data. The Calvo model with sales clearly outperforms any of these models.

Table 5: Comovement coefficient across the models

	Regular prices (1)	Posted prices (2)
Data	0.026	−0.022
Calvo model with sales (GS)	0.033	−0.033
Calvo model without sales	0.385	—
Fixed menu cost model (GL)	0.006	—
Smoothly state-dependent pricing (CN)	0.137	—

Notes: This tables compares the comovement of price dispersion and inflation in the Calvo model with sales, as in Guimaraes and Sheedy (2011), with that in the models without sales. The Calvo model without sales overstates the comovement for regular prices, whereas the fixed menu cost model, as in Golosov and Lucas (2007), understates it. The hybrid Costain and Nakov (2011a) model is qualitatively closer to the Calvo model and also overstates the comovement.

What is the role of pricing frictions in these results? In the standard Calvo model, nominal shocks do not affect the number of firms that adjust their prices. If the frequency of price adjustment is small, very few firms change their prices, thereby having only a small effect on the aggregate price level. At the same time, firms that are able to reset their prices adjust them proportionally to the shock, thereby increasing price dispersion. Hence, nominal shocks have a relatively small effect on inflation and a relatively large effect on price dispersion. In terms of the estimated

[9]Note that, as wage stickiness plays only a quantitatively minor role for the dynamics of price dispersion in the GS model, and following the original paper, the SSDP model features flexible wages.

comovement, small changes in inflation are associated with large changes in price dispersion, implying a relatively large coefficient. To match the data, more firms have to adjust their prices, amplifying the response of inflation and dampening the response of price dispersion. Instead, in the FMC model the comovement is weak. Firms set their prices by the Ss rule, with a selection effect on which firms adjust: firms that are further away from the optimal price are more likely to adjust, thereby having a relatively large effect on inflation and a relatively small effect on price dispersion. If the menu cost is very small, most firms will adjust their prices to the same level, which may even lower price dispersion.

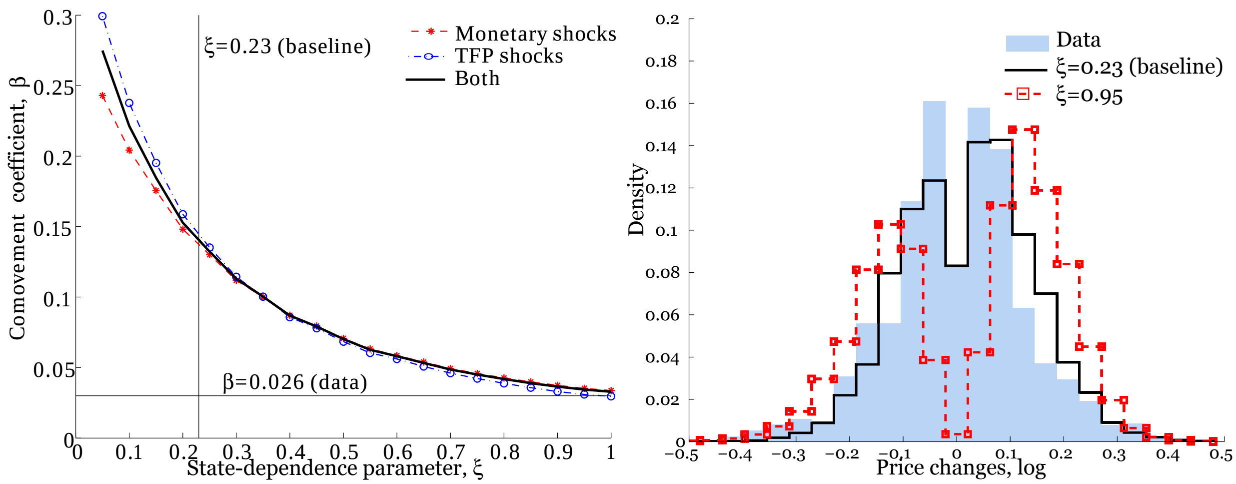

Panel A: Comovement coefficient

Panel B: Distribution of price changes

Figure 1: State-dependence, comovement, and price changes. Panel A shows that in order to match price dispersion dynamics in the data, the SSDP model needs to allow for more state-dependence than in the baseline calibration. Panel B shows that allowing for such degree of state-dependence leads to a counterfactual distribution of price changes (cf. Costain and Nakov, 2011a, Fig. 1).

As the SSDP model nests the Calvo and FMC models, the comovement is in between the two cases, and depends strongly on the smoothness parameter ξ. Under the baseline calibration, the comovement is larger than in the data, implying that the model should allow for more state-dependence to match the data (i.e., ξ should be higher). As the baseline value of this parameter is selected to match the empirical distribution of price changes (see Costain and Nakov, 2011a), increasing ξ will make the distribution look more like in state-dependent models, with fewer small price changes. A practical question is whether the increase in ξ needed to match the data is sufficiently small not to affect the distribution of price changes. Fig. 1 suggests that this is not the case. Panel A shows the comovement of price dispersion and inflation for different values of ξ. To match the comovement, ξ should be set to approximately 0.95, above its baseline value of 0.23. Note that the nature of shocks in this model is of the second order, because the comovement is determined largely by aggregate price stickiness. Panel B compares the distribution of price changes under the two values of ξ with the data, and shows that the distribution is bimodal under the larger ξ, with almost no price changes around zero. This discrepancy has been a major criticism of state-dependent models (e.g., see Midrigan, 2011).

To summarize, although the comovement of price dispersion and inflation clearly depends on the choice of parameter values, models without sales are unable to match the data without making counterfactual assumptions for a wide range of parameters. This result is not qualitatively sensitive to the type of shocks hitting the economy.

7. Implications for monetary policy and welfare

Why are price dispersion and its comovement with inflation important in practice? One reason is that this comovement is used for the calculations of welfare, the cost of business cycles, and the optimal inflation rate. The other reason is that the impulse-response functions of output to monetary shocks may potentially differ across models, and therefore one can use the comovement to discriminate between models with different degrees of monetary non-neutrality. In this section, I quantify these effects.

7.1. Welfare cost of inflation

The welfare analysis is based on Coibion et al. (2012), who derive a micro-founded welfare function in a Calvo model with trend inflation (and no sales). For zero trend inflation, the limiting case of this model is similar to Guimaraes and Sheedy (2011), with the size of the sales sector being zero, and to Costain and Nakov (2011a,b) with $\xi \to 0$. The advantage of looking at this model is that, due to the ZLB constraint, it gives rise to a non-zero optimal inflation target; thus, I can measure the effect of price dispersion on the optimal inflation rate.

The second-order approximation of the utility function is

$$U_t = \Theta_0 + \Theta_1 \, \mathbb{V}\mathrm{ar}\,(\hat{y}_t) + \Theta_2 \, \mathbb{V}\mathrm{ar}\,(\hat{\pi}_t) + \text{h.o.t.}, \tag{7}$$

where y is the output gap; π is inflation; Θ_i, $i = 0, 1, 2$, are functions of the model parameters; and h.o.t. stands for higher-order terms. Notation $\hat{x}$ represents the log-deviation of variable x from its steady-state. Price dispersion affects welfare through the steady-state channel (Θ_0) and through the variability of inflation (Θ_2). This functional form holds across a range of models, while the exact definitions of the parameters Θ are specific to the model. The first-order approximation of steady-state price dispersion gives rise to a linear relationship with steady-state inflation: $\bar{\sigma} \simeq \sqrt{\alpha}/(1-\alpha)\,\bar{\pi}$.

As the relationship between price dispersion and inflation in the model differs from that in the data, I compare welfare under the model-based first-order approximation and under the relationship estimated in the data: $\bar{\sigma}^{\text{data}} = \hat{\gamma} + \hat{\beta}\,\bar{\pi}$, where $\hat{\gamma}$ and $\hat{\beta}$ are the estimates from Section 4. The constant γ is added in order to match the level of price dispersion in the data when steady-state inflation in the model equals trend inflation observed in the data.[10] In comparison to the approximation used in the model, $\bar{\sigma}^{\text{data}} > \bar{\sigma}^{\text{model}}$ and $\beta^{\text{data}} \ll \beta^{\text{model}}$. Table 6 compares welfare U_t and the cost of business cycles, $\Theta_1\,\mathbb{V}\text{ar}\,(\hat{y}_t) + \Theta_2\,\mathbb{V}\text{ar}\,(\hat{\pi}_t)$, obtained using $\bar{\sigma}^{\text{data}}$ and $\bar{\sigma}^{\text{model}}$. There are large differences in the welfare estimates, and the cost of business cycles is larger by 40 percent than the one based on the model-based approximation (−0.007 vs. −0.005). As $\bar{\sigma}^{\text{data}} > \bar{\sigma}^{\text{model}}$, inflation is relatively more costly, and the optimal rate of inflation goes down from 1.3 percent to 1 percent.

Table 6: Welfare cost of inflation

		Data (1)	Calvo model (2)
Output-gap variability	(a)	−0.000	−0.000
Inflation variability	(b)	−0.007	−0.005
Cost of business cycle	$(c) = (a) + (b)$	−0.007	−0.005
Steady-state loss	(d)	−0.238	−0.009
Total welfare loss	$(e) = (c) + (d)$	−0.245	−0.014
Optimal inflation rate, %		1.0	1.3

Notes: This table compares the welfare cost of inflation based on price dispersion–inflation comovement in the data (column 1) with that based on the Calvo model (column 2). The welfare function is based on the derivations in Coibion et al. (2012).

This example suggests that the mismatch in the key relationship between inflation and price dispersion used to calculate welfare, its components, and the optimal inflation rate can have a large quantitative effect on the calculations. Therefore, welfare calculations that rely on the relationship between inflation and price dispersion should be based on a model that matches it, or should use a direct measure of price dispersion obtained from the data.

7.2. Output response to monetary shocks

I compare output response to monetary shocks in the Calvo, FMC, SSDP, and GS models. As the SSDP model encompasses Calvo and FMC pricing, the comparison of the first three models was studied extensively in Costain and Nakov (2011a). Therefore, the value-added of this exercise is in the comparison of the three models that do not allow for sales—and cannot match the comovement of price dispersion and inflation—with the Calvo model *with* sales.

To make sure that sales are the only difference in the models' setup, I apply a scale adjustment to the GS model. First, I generate output responses to a 1 percentage point increase in money supply growth in the Guimaraes and Sheedy (2011) version of the Calvo model with and without sales. For each time horizon, I compute the ratio of the response in the two versions (i.e., by

[10]In the presence of time-invariant product–store effects, this measure of price dispersion may overstate the steady-steady welfare loss but not the cost of business cycles.

how much sales attenuate the output response). I then compute the response in the SSDP model by multiplying the response in the Calvo model without sales by this ratio. The resulting series represents the output response to a monetary shock in the Calvo model with sales, set up as in Costain and Nakov (2011a).

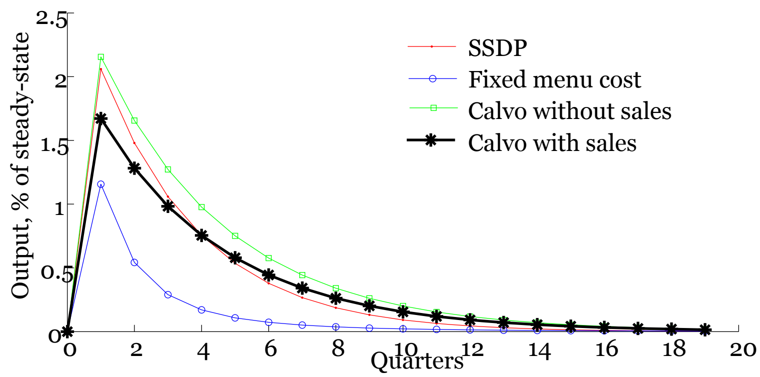

Figure 2: Output response to monetary shocks. The shock is a 1 percentage point increase in money growth. To allow for comparison across the setups, the response in the Calvo model with sales is obtained by multiplying, at each horizon, the corresponding response in the Calvo model without sales, as in Costain and Nakov (2011a), by the ratio of the responses in the Guimaraes and Sheedy (2011) models with and without sales.

As Fig. 2 demonstrates, the Calvo model with sales is characterized by a high degree of monetary non-neutrality, consistent with the results in the previous literature. Although sales add to aggregate price flexibility, resulting in a decrease of 20–25 percent in the output response, the persistence of the response is similar to that in the Calvo model without sales. The output response in the Calvo model with sales is smaller than in the SSDP model at short horizons, but larger at longer horizons. I conclude that (1) the sales model that matches the properties of price dispersion in the data is close to workhorse models without sales featuring a significant degree of time-dependence; and (2) when the sales sector is large, the quantitative differences between the output responses in the models with and without sales can become of practical importance.

8. Conclusion

Price dispersion is a central determinant of welfare, the cost of business cycles, the optimal rate of inflation, and the tradeoff between inflation and output stability. Many workhorse macroeconomic models give rise to dynamic properties of price dispersion that are inconsistent with the data. To get the models closer to the data, it is important to have a more realistic mechanism for the pricing decisions of firms. In particular, it is important to model a mechanism that gives rise to temporary price discounts. As this paper shows, a workhorse model that allows for sales does better than the alternatives in matching price dispersion dynamics. Although many dynamic properties of the sales model are qualitatively similar to those of the analogous model without sales, there are quantitative differences that can be important in practice. Welfare analyses are particularly sensitive to price dispersion properties. Overall, this paper makes the case for sales to be included in aggregate models, especially models used for quantitative predictions, such as those employed by central banks.

References

Alvarez, F., Beraja, M., Gonzalez-Rozada, M., Neumeyer, P.A., 2019. From hyperinflation to stable prices: Argentina's evidence on menu cost models. Quarterly Journal of Economics 134, 451–505.

Alvarez, F., Lippi, F., 2014. Price setting with menu cost for multiproduct firms. Econometrica 82, 89–135.

Anderson, E., Malin, B.A., Nakamura, E., Simester, D., Steinsson, J., 2017. Informational rigidities and the stickiness of temporary sales. Journal of Monetary Economics 90, 64–83.

Andrade, P., Galí, J., Le Bihan, H., Matheron, J., 2018. The optimal inflation target and the natural rate of interest. Working paper 24328. National Bureau of Economic Research.

Benabou, R., 1988. Search, price setting and inflation. Review of Economic Studies 55, 353–376.

Benabou, R., 1992. Inflation and efficiency in search markets. Review of Economic Studies 59, 299–329.

Beraja, M., Hurst, E., Ospina, J., 2016. The aggregate implications of regional business cycles. Working paper 21956. National Bureau of Economic Research.

Blanco, J.A., 2016. Optimal inflation target in an economy with menu costs and zero lower bound. Working paper. University of Michigan.

Bronnenberg, B.J., Kruger, M.W., Mela, C.F., 2008. Database paper—the IRI marketing data set. Marketing Science 27, 745–748.

Burstein, A., Hellwig, C., 2008. Welfare costs of inflation in a menu cost model. American Economic Review 98, 438–443.

Calvo, G.A., 1983. Staggered prices in a utility-maximizing framework. Journal of Monetary Economics 12, 383–398.

Caplin, A., Leahy, J., 1997. Aggregation and optimization with state-dependent pricing. Econometrica 65, 601–626.

Carvalho, C., 2006. Heterogeneity in price stickiness and the real effects of monetary shocks. B.E. Journal of Macroeconomics 6, 1–58.

Chevalier, J.A., Kashyap, A.K., 2019. Best prices: Price discrimination and consumer substitution. American Economic Journal: Economic Policy 11, 126–159.

Choi, C.Y., 2010. Reconsidering the relationship between inflation and relative price variability. Journal of Money, Credit and Banking 42, 769–798.

Christiano, L.J., Eichenbaum, M., Evans, C.L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy 113, 1–45.

Coibion, O., Gorodnichenko, Y., Hong, G.H., 2015. The cyclicality of sales, regular and effective prices: Business cycle and policy implications. American Economic Review 105, 993–1029.

Coibion, O., Gorodnichenko, Y., Wieland, J., 2012. The optimal inflation rate in New Keynesian models: Should central banks raise their inflation targets in light of the zero lower bound? Review of Economic Studies 79, 1371–1406.

Costain, J., Nakov, A., 2011a. Distributional dynamics under smoothly state-dependent pricing. Journal of Monetary Economics 58, 646–665.

Costain, J., Nakov, A., 2011b. Price adjustments in a general model of state-dependent pricing. Journal of Money, Credit and Banking 43, 385–406.

Debelle, G., Lamont, O., 1997. Relative price variability and inflation: Evidence from U.S. cities. Journal of Political Economy 105, 132–152.

Diamond, P.A., 1993. Search, sticky prices, and inflation. Review of Economic Studies 60, 53–68.

Dotsey, M., King, R.G., Wolman, A.L., 1999. State-dependent pricing and the general equilibrium dynamics of money and output. Quarterly Journal of Economics 114, 655–690.

Driscoll, J.C., Kraay, A.C., 1998. Consistent covariance matrix estimation with spatially dependent panel data. Review of Economics and Statistics 80, 549–560.

Erceg, C.J., Henderson, D.W., Levin, A.T., 2000. Optimal monetary policy with staggered wage and price contracts. Journal of Monetary Economics 46, 281–313.

Gagnon, E., 2009. Price setting during low and high inflation: Evidence from Mexico. Quarterly Journal of Economics 124, 1221–1263.

Golosov, M., Lucas, Jr., R.E., 2007. Menu costs and Phillips curves. Journal of Political Economy 115, 171–199.

Gorodnichenko, Y., Sheremirov, V., Talavera, O., 2018. Price setting in online markets: Does IT click? Journal of the European Economic Association 16, 1764–1811.

Grier, K.B., Perry, M.J., 1996. Inflation, inflation uncertainty, and relative price dispersion: Evidence from bivariate GARCH-M models. Journal of Monetary Economics 38, 391–405.

Guimaraes, B., Sheedy, K.D., 2011. Sales and monetary policy. American Economic Review 101, 844–876.

Head, A., Kumar, A., 2005. Price dispersion, inflation, and welfare. International Economic Review 46, 533–572.

Head, A., Liu, L.Q., Menzio, G., Wright, R., 2012. Sticky prices: A new monetarist approach. Journal of the European Economic Association 10, 939–973.

Kaplan, G., Menzio, G., 2015. The morphology of price dispersion. International Economic Review 56, 1165–1206.

Kaplan, G., Menzio, G., Rudanko, L., Trachter, N., 2016. Relative price dispersion: Evidence and theory. American Economic Journal: Microeconomics, in print.

Kehoe, P., Midrigan, V., 2015. Prices are sticky after all. Journal of Monetary Economics 75, 35–53.

Konieczny, J.D., Skrzypacz, A., 2005. Inflation and price setting in a natural experiment. Journal of Monetary Economics 52, 621–632.

Kruger, M.W., Pagni, D., 2008. IRI academic data set description. Version 2.1. Information Resources Incorporated.

Kryvtsov, O., Midrigan, V., 2013. Inventories, markups, and real rigidities in menu cost models. Review of Economic Studies 80, 249–276.

Kryvtsov, O., Vincent, N., 2014. On the importance of sales for aggregate price flexibility. Working paper 14-45. Bank of Canada.

Lach, S., 2002. Existence and persistence of price dispersion: An empirical analysis. Review of Economics and Statistics 84, 433–444.

Lach, S., Tsiddon, D., 1992. The behavior of prices and inflation: An empirical analysis of disaggregated price data. Journal of Political Economy 100, 349–389.

Midrigan, V., 2011. Menu costs, multiproduct firms, and aggregate fluctuations. Econometrica 79, 1139–1180.

Nakamura, E., Steinsson, J., 2008. Five facts about prices: A reevaluation of menu cost models. Quarterly Journal of Economics 123, 1415–1464.

Nakamura, E., Steinsson, J., Sun, P., Villar, D., 2018. The elusive costs of inflation: Price dispersion during the U.S. Great Inflation. Quarterly Journal of Economics 133, 1933–1980.

Pratt, J.W., Wise, D.A., Zeckhauser, R., 1979. Price differences in almost competitive markets. Quarterly Journal of Economics 93, 189–211.

Reinsdorf, M., 1994. New evidence on the relation between inflation and price dispersion. American Economic Review 84, 720–731.

Sheshinski, E., Weiss, Y., 1977. Inflation and costs of price adjustment. Review of Economic Studies 44, 287–303.

Silver, M., Ioannidis, C., 2001. Intercountry differences in the relationship between relative price variability and average prices. Journal of Political Economy 109, 355–374.

Smets, F., Wouters, R., 2007. Shocks and frictions in US business cycles: A Bayesian DSGE approach. American Economic Review 97, 586–606.

Sorensen, A.T., 2000. Equilibrium price dispersion in retail markets for prescription drugs. Journal of Political Economy 108, 833–862.

Stevens, L., 2019. Coarse pricing policies. Working paper. University of Maryland.

Van Hoomissen, T., 1988. Price dispersion and inflation: Evidence from Israel. Journal of Political Economy 96, 1303–1314.

Vavra, J., 2014. Inflation dynamics and time-varying volatility: New evidence and an Ss interpretation. Quarterly Journal of Economics 129, 215–258.

Woodford, M., 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press, Princeton, N.J.

Woodford, M., 2009. Information-constrained state-dependent pricing. Journal of Monetary Economics 56, S100–S124.