The Cyclicality of Sales and Aggregate Price Flexibility*

Oleksiy Kryvtsov† and Nicolas Vincent‡

February 2017

Abstract

Macroeconomists traditionally ignore temporary price mark-downs ("sales") under the assumption that they are unrelated to aggregate phenomena. We challenge this view. First, we provide robust evidence from the U.K. and U.S. CPI micro data that the frequency of sales is strongly countercyclical, as much as doubling during the Great Recession. Second, we build a general equilibrium model in which sales arise endogenously. In response to a monetary contraction, the resulting fall in the aggregate price level can be significantly larger than if sales were ignored, implying a much smaller response of real consumption to monetary shocks. Third, building on the insights of the model, we show that the contribution of sales to fluctuations in the effective price level is significant. This is largely due to the increase in unit consumption spending during sales, a margin not captured by conventional CPI measures.

Keywords: Price dynamics, Sales, Price index measurement.

JEL classification: E31, E32, E52, L11, M31.

*The views expressed here are ours, and they do not necessarily reflect the views of the Bank of Canada. We would like to thank Susanto Basu, Nicoletta Berardi, Marty Eichenbaum, Gee Hee Hong, Kim Huynh, David Jacho-Chávez, Nir Jaimovich, Alejandro Justiniano, Sergio Rebelo, Leonard Sabetti, Gregor Smith, Joe Vavra and participants at the 2015 Canadian Macroeconomic Study Group, the 2015 Banque de France Price Setting and Inflation Conference, NBER Price Dynamics 2014 Workshop, the 2014 Shanghai Macroeconomic Policies and Business Cycles Conference, the Bank of Canada 2014 Fellowship Learning Exchange, the 2014 Meetings of the European Economics Association, the 2014 Dynare Conference, and the Université de Montréal, University of Toronto, Queen's University, Duke University, Mannheim University, Riksbank, Uppsala University, and Banque de France seminar series for helpful comments and suggestions. We also thank Ainslie Restieaux and colleagues in the Prices Division at the Office for National Statistics for valuable feedback regarding the U.K. CPI data, and Brendan Williams at the U.S. Bureau of Labor Statistics for providing us the series based on the U.S. CPI data. Claudiu Motoc, Tony Chernis, and Anderson Nzabandorra provided superlative research assistance. We are grateful to Glen Keenleyside for extremely helpful editorial assistance.

†Bank of Canada, 234 Laurier Street Ottawa, Ontario K1A 0G9, Canada. Email: *okryvtsov@bankofcanada.ca*

‡HEC Montréal, Department of Applied Economics, 3000 Côte-Sainte-Catherine, Montréal, Quebec H3T 2A7, Canada. Email: *nicolas.vincent@hec.ca*

1 Introduction

Price discounts, or "sales," are an essential feature of retail price behavior and an important factor for households' consumption decisions. A typical sale is associated with a large but temporary price drop that returns close to its pre-sale level. In the past decade or so, macroeconomists extensively employed detailed weekly and monthly price data for a broad variety of retail goods to study the implications of retail pricing for aggregate price flexibility. These studies find that although temporary price discounts imply more frequent price changes by retailers, they do not, on the whole, play a significant role for inflation dynamics.[1] More generally, the prevalent view in macroeconomics has been that retail price discounts have little to do with aggregate phenomena and should be ignored by macroeconomists.[2] In this paper, we challenge this view.

In the first part of the paper, we provide empirical evidence on variations in the incidence of sales over time based on data for both the United Kingdom and the United States. For the United Kingdom, we use the publicly available micro data underlying the consumer price index (CPI) composed by the Office for National Statistics (ONS). The data contain monthly price quotes collected from local retail outlets for a wide range of consumer goods and services over the period from 1996 to 2013. We find that the frequency of sales in the United Kingdom data is strongly countercyclical: a 1 percentage point rise in the unemployment rate is associated with a roughly 0.4 percentage point increase in the fraction of products on sale. For example, during the Great Recession the fraction of sales more than doubled from 1.9% to 3.9% of observations for our preferred measure of sales. This strong correlation between the business cycle and the use of temporary discounts by firms is extremely robust: it does not depend on how sales are identified; it is observed across a wide range of categories and all regions; it applies to both clearance and non-clearance sales; and it survives the use of multiple controls and alternative macroeconomic indicators. Unlike the fraction of sales, the average size and duration of sales in the United Kingdom are mostly acyclical and much less volatile.

Our finding is not specific to the United Kingdom. An aggregate time series on the incidence of sales obtained from the U.S. Bureau of Labor Statistics shows that temporary discounts are also strongly countercyclical in the United States. This evidence is corroborated by series derived from

[1] For example, Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008) find, using U.S. CPI micro data, that, on average, prices adjust every four to seven months, and that excluding sale prices increases price durations by around three to five months. Klenow and Malin (2011) provide an excellent survey of microeconomic evidence on price setting.

[2] Eichenbaum, Jaimovich and Rebelo (2011) argue that most high-frequency movements in prices have little to do with monetary policy. An important exception is Klenow and Willis (2007), who find that, in the United States, CPI micro data sale-related price changes respond to macro information in a similar way as regular price changes.

Vavra (2014) based on the U.S. CPI micro data since 1988, a period that covers three distinct U.S. recessions.

There are at least two reasons why cyclical variations in the fraction of sales may have a significant impact on aggregate price fluctuations. First, sale-related price drops are large, on average between 20% and 25% in the U.K. Second, each additional sale generates on average close to three times more revenue than a regular-price transaction, and the revenue share of goods on sale may increase during economic slumps as households shift their consumption toward sale prices. Our theoretical contribution is to quantify the aggregate price flexibility due to these factors.

To this end, in the second part of the paper, we develop a general equilibrium business cycle model with consumer search and price discrimination by monopolistically competitive retailers. Households face an independent, identically distributed (i.i.d.) time cost of searching for low prices. Those households who draw sufficiently low cost realizations become bargain hunters and find, on average, lower prices. Retailers, in a desire to attract bargain hunters, keep a positive fraction of brands on sale in the store. The higher is the return to households from bargain hunting, the larger is the fraction of price discounts. Retailers' revenue from posting sales increases with households' willingness to substitute between market work and searching for sales. When the elasticity of substitution is sufficiently high, both the fraction of sales and the fraction of bargain hunters are strongly countercyclical, amplifying procyclical aggregate price dynamics.

Our calibrated model successfully matches salient features of retail discount and search behavior highlighted in the literature, including the prevalence of large but temporary discounts ("V-shapes"); significant fluctuations in the average fraction of discounts, but little variation in their average size; and the average time spent by households shopping for lower-price goods together with the lower price they get in return. At the aggregate level, the model predicts that in response to an unanticipated monetary contraction, the increases in the fractions of sales and bargain hunters lead to an aggregate price decrease that is twice as large as the one in the standard sticky-price model with a single regular price per variety. Accordingly, the size of the real effect in response to monetary shocks is twice as small, and it can be even smaller depending on the response of the fraction of bargain hunters.

We show that the main mechanism underlying the importance of sale prices for aggregate price flexibility is based on the interaction between the retailers' price discounting and the households' search for low prices. At the time of the monetary contraction, some prices fail to decrease due to constraints on price adjustment, leading to an increase in retail markup. High profit margins make it desirable for retailers to increase their market share. In our model, they do so through an increase in the fraction of brands on sale. In turn, more aggressive price discounting by retailers increases the return to consumers from searching for low prices, leading to a larger number of bargain hunters. The resulting reallocation of the consumption weight toward lower-priced products amplifies the fall in both store-average and aggregate price levels. The model's central feature—countercyclical shopping intensity—is corroborated by evidence from time-use surveys and household panel scanner data sets obtained by Aguiar, Hurst and Karabarbounis (2013), and Kaplan and Menzio (2015), among others.

Our theoretical findings have additional implications. In the model, fluctuations in sales frequency lead to an important disconnect between the "effective" consumption-based price level and the official CPI measure. This is because sales account for an outsized share of quantity sold—an effect that is not captured in traditional fixed-weight price aggregates such as the CPI. The result is a potential mis-measurement of the price of the consumption basket over the business cycle. Using a simple calculation on actual data, we show that price discounts can have a significant impact on the effective price level, with a one-percentage-point increase in the fraction of sales decreasing the effective price level of food by approximately 1.5 percent. The impact of discounts is especially dramatic from 2007 to 2010 when the fraction of sales experienced its largest increase in the United Kingdom. During this period, the sale-price component of the effective price decreased by 3 percentage points in the aggregate, or by almost 1 percentage point per year, and three times as much in some broad product categories. As we show, these deflationary pressures can be severely underestimated by official CPI inflation measures.

A few recent studies have examined the potential role of sales in macroeconomic models. Kehoe and Midrigan (2015), using modified versions of standard sticky-price models, argue that sales are mostly irrelevant for the transmission of monetary shocks, since, due to their temporary nature, sales cannot offset aggregate shocks well. Guimaraes and Sheedy (2011) reach similar conclusions using a sticky-price model with sales stemming from consumer heterogeneity and incomplete information. In their model, strong strategic substitutability of sales at the micro level implies that their frequency and size barely respond to monetary shocks. Both papers, therefore, predict that the sale margin is not useful for retailers' price adjustment in response to changes in macroeconomic conditions. Neither paper, however, tests whether these predictions are borne out in the data, mainly because the empirical evidence on the cyclical properties of sales is very limited. There are, however, a few recent exceptions.

Coibion, Gorodnichenko and Hong (2015) use a scanner dataset from U.S. grocery stores and conclude that sales for grocery products are acyclical and that most adjustment is done through consumers switching from high- to low-end retailers. We exploit additional time series for food products obtained from the BLS to show that less-than-cyclical behavior of sales for food found in Coibion, Gorodnichenko and Hong (2014) does not overturn our finding of cyclical sales at the aggregate level. Likewise, Anderson *et al.* (2016) do not find evidence that sales respond to wholesale and commodity cost shocks or to changes in local unemployment rates. Similar to the dataset in Coibion, Gorodnichenko and Hong (2015) covering mostly food products, the data in Anderson *et al.* (2016) is limited in scope, providing information for groceries and general merchandise products from a single U.S. retailer. In contrast, our study is based on Consumer Price Index (CPI) micro price data and, therefore, has much broader coverage.

At the aggregate level, Berardi, Gauthier and Le Bihan (2015) find little time variation in sales based on French CPI micro data, a result that may be a reflection of the fact that sales are heavily regulated in France. The closest paper to ours in scope and findings is Sudo *et al.* (2014) that looks at the behavior of sales across a wide range of product categories in Japan since 1988. The authors find a rise in the frequency of sales in Japan during the 1990s and 2000s at the same time as hours worked were declining and the unemployment rate was rising. This evidence, however, is dominated by strong trends in all three series during Japan's "lost decades," making it difficult to determine whether the behavior of sales is due to their countercyclical use by retailers or rather due to structural changes in the Japanese retail industry unrelated to business cycles. Also, while their scanner dataset is quite extensive, ultimately it excludes important categories and covers only 17% of households' consumption expenditures.

We further explore the broad coverage of our data and look for cross-sectional correlations for sales but find they tend to be weaker than correlations across time. There is little evidence of a relationship between sales frequency and unemployment rate across U.K. regions, which may be due to the use of uniform national pricing strategies by large retailers, who account for the bulk of the use of sales in our dataset. When looking across goods, we find that more durable goods and goods with more concentrated businesses tend to have more countercyclical sales, in favour of theories in which retailers compete for market share or in which intertemporal demand effects are present. Discounts for goods with smaller mean absolute sizes of price changes or those with more volatile frequencies of price changes tend to be more countercyclical, in accordance models with fixed costs of price adjustment.

We conclude that properly accounting for sales can have important ramifications for our assessment of the cyclical behavior of aggregate price and real variables. Macroeconomists who calibrate their models using micro- or macro-level price data without properly accounting for the dynamics of sales may be led to underestimate the degree of aggregate price flexibility and overestimate the fluctuations of real variables over the cycle.

The paper is organized as follows. Section 2 describes the data and basic statistics on sales. Sections 3 and 4 document aggregate and disaggregate evidence on the cyclicality of sales for both the United Kingdom and United States. We build and study the predictions of a general equilibrium model with sales in Sections 5 to 7. In Section 8, we discuss the role of time-varying sales for aggregate price indices, and Section 9 concludes.

2 Data

2.1 Data sources

Throughout the paper, we provide evidence from two sources. Most of it is from the consumer price index (CPI) micro dataset of the United Kingdom's Office for National Statistics (ONS), which is publicly available. Access to CPI micro data in other countries is very limited. That being said, we also provide aggregate evidence for the United States on the incidence of sales computed for us by the United States' Bureau of Labor Statistics from their CPI-RDB database.[3] We provide below a brief description of the U.K. dataset, while postponing the details to the Supplementary Material.

To construct the CPI, the ONS surveys the prices for goods and services that are included in the household final monetary consumption expenditure component of the U.K. National Accounts. The survey includes prices for more than 1,100 individual goods or services a month, collected locally from more than 14,000 retail stores across the United Kingdom. It excludes the housing portion of consumer prices, such as mortgage interest payments, house depreciation, insurance and other house purchase fees. Goods and services in the CPI are classified into classes that represent basic group categories, such as "Meat," "Garments" or "New Cars"; and each CPI class is divided into finer categories, "items." For each item and stratum (given by region and shop type pairing in the U.K. data), ONS dataset provides sampling weights that reflect products' relative importance in households' consumption expenditures. These weights exhibit small variation over time due to annual revisions to capture permanent changes in the expenditure composition of households' consumption baskets. Unless otherwise noted, all weighted statistics are constructed using the CPI consumption expenditure weights.

[3]U.S. CPI micro data have been extensively studied in the past. Descriptions and the key stylized facts for the U.S. CPI micro data can be found in Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008), among other sources. In many aspects of retail price behavior they obtain very similar results to those documented here and in Bunn and Ellis (2012) for the United Kingdom.

The sample period in the ONS dataset includes 212 months, from February 1996 till September 2013. We make some adjustments to make the dataset suitable for our analysis. First, we delete observations that are not used for CPI construction by the ONS. Second, we deal with product substitutions by splitting the price time series of a given item every time we encounter a substitution flag. The resulting benchmark dataset contains a total of 20.7 million observations across about 2.3 million unique products, covering about 57% of the U.K. CPI basket. For our empirical analysis we will mostly focus on items that have at least ten price quotes (17.1 million observations).

2.2 Sales filters

The first challenge when studying temporary sales in micro price data is to identify them. Ideally, we want to discriminate between price drops that are temporary and drops in regular prices. We use three main ways of identifying sales in our dataset.

First, for both the United Kingdom and the United States, we present results using the "sales flag" available in the respective datasets. The ONS indicates that "sale prices are recorded if they are temporary reductions on goods likely to be available again at normal prices or end-of-season reductions." Despite the advantage of being made directly available by the statistical agency, there are some issues with the sales flag that require us to make some adjustments. For example, there are a few instances in which the occurrence of a sales flag is accompanied by zero change in the posted price, or even, in some rare instances, by a price increase. One possible explanation is that the retailer uses some advertising features to gain or retain customers despite not actually changing the price; another is misreporting or a coding error. In what follows, we adopt a conservative approach and present results based only on sale flags that correspond to price decreases.[4]

Second, we apply a V-shaped sales filter, similar to the one used by Nakamura and Steinsson (2008), among others. In this instance, a "sale episode" begins with a price drop and ends as soon as a price increase is registered, as long as this price increase occurs within three months. Under this definition, we do not require the price at the end of the sale to be at or above the price at the beginning of the sale. For the sales-flag and V-shaped filters, the unobserved regular price during a sale is assumed to be equal to the last observed regular price.[5]

[4]To be precise, a sales flag is deemed valid only if it coincides with a price quote that is lower than the price that was posted right before the start of the spell of sales flags, which we define as the regular price for the duration of the spell.

[5]If two price drops occur in a row before the posted price settles to a new regular level, our filter will identify the first price drop as corresponding to a sale. To assume instead that this episode corresponds to either two distinct sales or two consecutive drops in the regular price has no impact on our results. We also considered a more-restrictive filter whereby a V-shaped sale is initiated by a price drop that is followed within three months by a return to a price equal to or higher than the initial price level. Results are not presented here because they are very similar, but are

Finally, we compute a reference price similar to one in Eichenbaum, Jaimovich and Rebelo (2011), using a seven-month window. For a given month t we set the reference price equal to the modal price observed between $t-3$ and $t+3$, as long as there are at least four price observations within that window. To avoid identifying spurious sales that arise from a lag or a lead in the adjustment of reference prices, we then apply a procedure similar to Kehoe and Midrigan (2015) to ensure that a change in the reference price coincides with an actual price change. A price observation corresponds to a sale price whenever the posted price is below the reference price.

For discounted prices to be recorded by ONS agents, they have to be available to everyone (i.e., coupons and discounts that require a loyalty card are not taken into account) and on a single purchase (e.g., discounts implied by "buy-two-get-one-free" promotions are not recorded). Hence, our estimates of the incidence of sales are most likely conservative.

2.3 Basic statistics

In Table 1 we report some basic statistics on price dynamics in our dataset. Unless otherwise stated, all moments are weighted using the official CPI weights. The fraction of price changes is 15.8% over the sample period, and the average size of a price change is 11.9% in absolute terms. Price increases are more likely than price decreases (9.8% vs 6.0% of observations, respectively). Not surprisingly, we find lower price change frequencies if we focus on regular prices, i.e., price series that were purged of observations for which the posted price differs from the regular price. The probabilities of observing a price change are 13.2%, 10.9% and 7.5%, based on the sale flag, V-shaped and reference price filters, respectively. Hence, the reference price filter generates significantly stickier price series, largely because it filters out both upward and downward temporary price deviations. Overall, our basic statistics show that prices are stickier in the United Kingdom than in the United States, but more flexible than in Europe.[6]

Table 2 reports several basic statistics for sales. The first row shows that the frequency of sales varies significantly depending on the definition used, from 2.9% for the ONS sales flag to 5.6% for the sales based on the reference price filter.[7] Differences are also visible for both the mean and median sale sizes: they are much higher for the sales flag (between 20% and 24%) than the three other filters (between 8% and 14%).

available upon request.

[6]See, for example, Klenow and Kryvtsov (2008) or Nakamura and Steinsson (2008) for the United States and Álvarez et al. (2006) for Europe.

[7]Though direct comparisons can be difficult, the prevalence of sales seems to be much lower than in the United States. For instance, Klenow and Kryvtsov (2008) find that about 11% of price observations have a sales flag, while Nakamura and Steinsson (2008) find that sales account for about 7.4% of observations using a V-shaped filter.

To understand why this is the case, we show in Figure 1 the truncated distribution of the size of sales across all observations in our dataset. The discount size in a given month is given by (the absolute value of) the difference between log sale price and log regular prices in that month. We set the bound of the histogram at 60%, since larger sales are rare. Reassuringly, there are spikes in the distribution at the familiar discount points: 10% off, 20% off, 25% off, 33% off and 50% off. Second, the three distributions exhibit striking differences between −10% and zero: the mass closer to zero is significant for the V-shaped and reference price filters, while there are very few small sales according to the ONS indicator. This suggests that the sales filters commonly used in the literature have a tendency to pick up small price drops that are not advertised as sales by retailers. On the other hand, the three distributions are much more similar for sales larger than 10%, a sensible threshold. For this reason, we focus on sales of at least 10% in our analysis. Under this condition, sale frequencies and sizes become very similar across filters, as can be seen from the bottom portion of Table 2.

3 Aggregate time-series behavior of sales

We now turn our attention to the main objective of our empirical work and study the aggregate behavior of temporary sales over time.

3.1 Evidence from the United Kingdom

First, we look at the behavior of the frequency of sales over the 1996:02–2013:09 period in the United Kingdom. For each category and month we compute the proportion of items with a sales flag, and then aggregate them using CPI weights. In the left plot of Figure 2 we show the raw constructed series as well as the U.K. unemployment rate. Because of the strong seasonal patterns of sales, we also report on the right-hand plot the 12-month moving average centered around each month. Clearly, the fraction of items on sale is far from being constant over time: it is around 3.2% at the beginning of the sample in 1997, then declines to a trough of about 1.9% in 2006 before rising back to almost 4% by 2011. Also, it is strongly countercyclical: the fraction of sales moves very closely with the unemployment rate, rising as the economy is slowing down. Neither series seems to exhibit any time trend that may bias our conclusions; we formally control for this potential issue later on.

Figure 3 compares the evidence from the sales flag filter to that of the V-shaped sales and reference-price filters described earlier. Unlike Figure 2, all series now focus on sales of at least 10%. First, the top-right panel demonstrates that aggregate sales frequency is not cyclical because of larger weights on sales-heavy items during recessions: using time-invariant CPI weights makes no noticeable difference. Second, overall patterns are very similar across sales filters: all give rise to a highly cyclical sales frequency.

Conceivably, retailers could use different sales-related margins in response to aggregate shocks. First, the cyclicality of the fraction of products on sale could be driven by fluctuations in the incidence of new sales and by changes in the average length of existing sales over time. We find no evidence of the latter in the data: the average duration of sale spells remained very stable around 1.6 months over our sample period, with no discernible cyclicality or trend.[8]

Second, retailers could vary the size of price discounts over the business cycle. Figure 4 shows the evolution over our sample period of the average absolute size of sales, conditioning on sales of at least 10%. Under the three definitions, there is a noticeable increase in the absolute size of sales over the sample period. Nonetheless, in relative terms, the variation in relative terms is largely contained, unlike fluctuations in the frequency of sales. For example, using the V-shaped filter, we find that the average size of sales fluctuates between roughly 22% and 25%. But most importantly, there is no clear cyclical pattern for this margin of adjustment: at the height of the Great Recession, sales tended to become a little smaller based on the sales flag filter, but slightly larger according to the other two definitions.[9]

3.2 Evidence from the United States

Before analyzing further the robustness of these results for the United Kingdom, we turn our attention to the United States to see whether the strong countercyclicality of sales is also present in the U.S. Bureau of Labor Statistics (BLS) CPI micro data. Despite not having direct access to the data, we were able to obtain aggregate evidence from two sources. First, the BLS provided us with the monthly fraction of items with a sales flag since January 2000, aggregated using CPI expenditure weights.[10] The time series, smoothed out using 12-month centered moving average, is presented in the left-hand plot of Figure 5, alongside the U.S. civilian unemployment rate over the same period. The similarities with our results for the United Kingdom are striking: there is a very clear positive co-movement between the incidence of sales and unemployment, with a correlation coefficient of 0.88. The turning points in the two series coincide very closely, with two clear spikes in the sales' fraction in the midst of the 2001 and 2008–09 recessions.

[8]The monthly frequency of our data may hinder identification of changes in the average duration of sales.

[9]Our findings are unchanged if we consider sales of all sizes instead of focusing on a 10% threshold for the size of sales.

[10]We are grateful to Brendan Williams and the BLS for providing these series to us.

Second, the right-hand plot in Figure 5 depicts a series derived from Vavra (2014). For his analysis, Vavra filtered out both temporary sales and product substitutions to focus on regular prices.[11] He kindly provided us with the time series for the combined frequency of sales and substitutions. Despite the fact that the series is somewhat more volatile, possibly due to the behavior of substitutions, the results visually appear to confirm our findings based on the data provided directly by the BLS. Moreover, Vavra's series spans an additional 12 years of data from 1988 to 1999, showing a clear rise in the fraction of sales and substitutions during the 1990–91 recession, similar to those in two subsequent downturns.

3.3 Evidence in Coibion, Gorodnichenko and Hong (2015)

Our finding that the incidence of temporary sales in the United States varies significantly with the business cycle over the past 15 years may seem at odds with the conclusion of Coibion, Gorodnichenko and Hong (2015), henceforth CGH. Using the Symphony IRI scanner dataset, which covers multiple grocery stores in 50 U.S. metropolitan areas over the 2001–11 sample period, they find that the relationship between unemployment and the frequency of sales becomes small or non-significant once they include a linear time trend or time fixed effects in their panel regressions.

A possible explanation for the dissimilar findings is that the U.S. grocery products studied by CGH display a different behavior than most consumption products. To explore this possibility, we obtained from the Bureau of Labor Statistics the frequency of sales for food from the CPI micro data, both weighted and unweighted. Figure 6 demonstrates that the behavior of sales frequency for food products is indeed markedly different from that at the economy-wide level that we documented in Figure 2. For both groups—all products in BLS sample and only food—the frequency of sales exhibits rapid increases in the wake of the 2001 and 2008–09 recessions. However, as the U.S. unemployment rate declines in the middle of the 2000s, sales become less prevalent at the aggregate level, while the weighted (unweighted) sales frequency for food products is instead rising (flat). Hence, the cyclical fluctuations in the time series for sales in the food category are difficult to distinguish from an upward trend. Indeed, we show in the Supplementary Material appendix that controlling for a linear time trend in the regression of sales frequency on unemployment, as it is done in CGH, yields only a weak relationship for food products, but does not alter the strong positive correlation at the aggregate level.[12]

[11] Vavra uses a joint sales-flag and 3-month-V-shaped filter to extract sales.

[12] In an earlier version their paper, CGH report category-level results and find a significant positive relationship between sales frequency and unemployment for personal care products in their dataset. We find that the sales flag time series for the CPI category "Personal care products"—also provided to us by the BLS—indeed exhibits a strong

3.4 Regression analysis

To support the graphical evidence, we run regressions for the U.K (all three sales filters) and the U.S. (BLS and Vavra) time series. In each case, we include calendar month dummies to control for the strong seasonality of sales. For every regression in Table 3, the coefficient on the unemployment rate is statistically significant at the 1% level, even if we include a time trend or one lag of the dependent variable. The elasticity of the frequency of sales to fluctuations in the unemployment rate is highest in the U.K. data and the lowest for the U.S. data provided by the BLS, with the estimates based on the Vavra time series in the middle. Nonetheless, alongside Figures 2 and 5, these results indicate that the countercyclicality of the aggregate sales frequency is robust for both countries.

One may conjecture that retailers use sales in conjunction with other margins of price adjustments, for example to offset customers' response to longer-lived regular price increases. The results in Table 4 indicate that this is not the case: for all datasets, the coefficient on the unemployment rate remains mostly unaffected by the inclusion of the frequency of regular price changes.[13] Only for Vavra's U.S. time series do we find a significant relationship between the frequencies of regular price changes and sales. This does not alter our finding about the cyclicality of sales however, and the coefficients on the unemployment rate are even larger in this case than the coefficients reported in Table 3.

4 Disaggregate evidence

Next, we investigate the robustness of our findings at a more disaggregated level. In the process, we document some other salient features of temporary sales in the data.

4.1 Category- and region-level results

As a first exercise, we verify that the cyclicality of aggregated sales is not merely driven by a few large categories and is in fact widespread across product types. We run separate regressions of sales frequency on the unemployment rate and a time trend for the 35 U.K. CPI categories with data for the whole sample and an average frequency of sales of at least 0.5%. The left panel of Figure 7 shows a histogram of the regression coefficients on the unemployment rate at the category level.

cyclical pattern. In a regression that includes a time trend, we find that the coefficient on the unemployment rate is strongly significant at 0.29, compared to 0.23 for the aggregate but only 0.10 for food products.

[13]The results shown Table 4 are based on the sales flag filter; they hold for the other two definitions of sales or if we use lagged values.

The histogram confirms that the countercyclicality of sales is broad-based: out of 35 categories, only two have a negative coefficient, and neither is statistically significant based on Newey-West standard errors. Of the 33 for which the coefficient on unemployment is positive, 24 are significant at the 1% level and 2 more at the 5% level.

We also use real consumption at the category level instead of aggregate unemployment. Because the data is only available at the quarterly frequency, we aggregate the fraction of items on sale accordingly. We then run regressions of category-specific sales frequency on the category-specific log of real consumption and a linear trend. The distribution of coefficients on consumption is reported in the right panel of Figure 7. Once again, sales are countercyclical for most sectors. The three most sensitive categories are *Household Appliances, Spirits* and *Audio-Video Equipment*.

In addition, we explore how the category-level cyclicality of sales changes with observed category characteristics and pricing behavior, including the data on life expectancy for durable goods, obtained from Bils and Klenow (1998) for the United States; 5-firm concentration ratios for the United Kingdom in 2004 from Mahajan (2006); and category-specific moments for price-adjustment variables (inflation, frequency of price changes, and absolute size of price changes; their means and standard deviations over time, using posted and regular prices), see Supplementary Material for detail.

We find that more durable goods and goods with more concentrated businesses tend to have *more countercyclical* sales, although these correlations are relatively weak due to the small number of observations for which we matched the data for durability (12) and concentration ratios (29). Theories in which retailers compete for market share or in which intertemporal demand effects are present would be favored by such evidence (e.g., Hendel and Nevo, 2006).

In regressions on price-adjustment moments, two of them result in significant correlations across categories. Discounts for goods with smaller mean absolute sizes of price changes or those with more volatile frequencies of price changes (for either posted or regular prices) tend to be *more* countercyclical. This evidence is consistent with predictions of models with fixed costs of price adjustment, in which goods with smaller fixed costs and more flexible price adjustment, and in particular, more frequent sales, would need smaller price changes on average.

Next, we verify whether the cyclicality of sales is present across U.K. regions. To do so, we regress the frequency of sales on the monthly regional unemployment rate and a time trend. The coefficients from these regressions are plotted in Figure 8. The frequency of temporary discounts is significantly countercyclical for all 11 regions. The strongest relationship is for Scotland and East Anglia: a five-percentage-point increase in the unemployment rate is associated with a 2 percentage point increase in the sales frequency, or close to double its unconditional mean. The weakest, on the other hand, is in London, the North as well as Yorks & Humber, at half the size. These differences in semi-elasticity estimates are directly related to the fact that while the sales frequency of sales is very similar across regions, unemployment rates vary much more. In Section 4.3, we discuss how uniform national pricing strategies by large retailers can rationalize this finding.

Finally, we investigate the role of clearance sales. To do so, we compute the probability of observing a sale within the last two months of a quote line.[14] Figure 9 plots the time series of both clearance and non-clearance sales for the aggregate as well as three broad categories. First, temporary sales clearly tend to be more prevalent toward the end of price quote lines. Aggregating across all products, this probability is about 8% in the last two months of a typical price quote line, compared to 3.5% otherwise. There are wide differences in this pattern across product categories, however. For Clothing and Footwear, 18% of the observations in the last two months are flagged as sales by the ONS, compared to a frequency of around 8% for other periods, while the same numbers are 25% and 13% respectively for Audio and Video. Food products, on the other hand, do not display more sales towards the price quote line.

Second, for the aggregate as well as the three broad categories, Figure 9 indicates that the frequency of both clearance and non-clearance sales follow the same countercyclical pattern that we documented earlier. Hence, the cyclicality of overall temporary discounts is not driven by more prevalent clearance sales during recessions than booms, possibly as firms are left with unsold inventories.

4.2 Product-level results

To control for additional factors that may affect our conclusions, we supplement our analysis with a panel regression analysis at the product level. The basic specification is standard and given by

$$s_{it} = \alpha_i + \beta u_t + X_t' \Phi + e_{it} \, , \tag{1}$$

where s_{it} is an indicator equal to 1 if product i is on sale at time t, and 0 otherwise; u_t is the unemployment rate, and X_t is a matrix of controls such as calendar month dummies or a time trend. All regressions include dummies for the months in which VAT rate changes occurred, as

[14] A price line can end because the product is replaced by another comparable item at a specific store, because the ONS stops collecting its price, or if the store permanently shuts down. See the Supplementary Material for additional evidence on the prevalence of sales over the product life cycle.

described in Section 2.1.[15]

Table 5 summarizes our results. The unemployment rate is a statistically significant predictor of the occurrence of a temporary sale: a 5-percentage-point increase in the unemployment rate raises the likelihood of observing a sale by about 2 percentage points. Adding a time trend or one lag of the dependent variable has little effect on the economic or statistical significance of this relationship. The results are also robust to the use of the two other definitions of sales, with the strongest effects found for the benchmark V-shaped filter.

We also verify that our findings are robust to the use of alternative business cycle indicators in equation 1 instead of u_t. First, we replace the unemployment rate with monthly retail sales volume, linearly detrended. This is a measure that is arguably a particularly relevant indicator of aggregate economic conditions faced byf retailers. Second, to capture the economic outlook of households, we also use consumer confidence indicators for the United Kingdom as compiled by the company GfK on behalf of the European Commission. In these monthly surveys, various questions are asked to a sample of households. We focus on the aggregate consumer confidence index as well as the question about the personal financial situation over the next 12 months. To facilitate comparisons, we normalize all indicators by dividing them by their respective time-series standard deviations over the sample. The results in Table 5 show responses in the same order of magnitude across all four indicators: a one-standard-deviation decrease in either measure of consumer confidence leads to a statistically significant increase of around 0.3 percentage points in the frequency of sales, and that increase is around 0.5 ppt for a one-standard deviation increase (decrease) in the unemployment rate (retail sales volume).[16]

4.3 Sales at the retailer level

Finally, we investigate whether there are systematic differences in the use of sales across retailers. For each price quote the ONS provides both a store identifier and a flag indicating whether the store belongs to an independent retailer (less than 10 locations) or a chain (10 locations or more), a category labeled "multiples" by the ONS. Chains account for about 64% of observations and 55% of CPI weights.

Figure 10 plots the incidence of discounts of at least 10%, identified by the sale flag, against the unemployment rate for both types of stores, aggregated using CPI weights. The differences in scale

[15] To allow for the inclusion of item fixed effects, we run linear regressions despite the fact that we have a binary 0/1 dependent variable. The results are very similar, and in fact generally slightly stronger, if we use probits instead (see the Supplementary Material).

[16] As a point of reference, the unemployment rate during the 2007–09 recession rose by a magnitude of four standard deviations.

are striking: sales are much more prevalent for larger retailers (left column) than for smaller retailers (right column). For example, using the sales flag and CPI weights, the average frequencies are 3.5% for multiples and 1% for independents (4.5% and 1.5% unweighted), while the time series standard deviations are 1.1% and 0.4% respectively.[17] Moreover, sales appear to be much more cyclical for larger retailers. For example, during the Great Recession, the frequency of sales doubled for multiples, from 2.5 in 2007 to 5 percent by 2011. On the other hand, the incidence of sales rose by barely a quarter of a percentage point for independents.

These differences could reflect different sets of consumer goods sold by small and large retailers. For example, in the United Kingdom, large retailers are predominant sellers of grocery products, and they less frequently offer services such as audio-video repairs and dentist services. To minimize these differences, we focus on food categories and run a product-level regression of the sales flag on product, month and store fixed effects. In Figure 11, we plot the distribution of the store fixed effects obtained from this regression, and separate the densities for stores belonging to independents and multiples. The first striking finding is the large degree of heterogeneity across stores in their use of sales, even after controlling for variation across products and time.[18] Second, the evidence in Figure 11 indicates that even when we include product fixed effects, sales are much more frequent for larger retailers, by about 3.2 percentage points. These findings confirm the evidence in Nakamura *et al.* (2013) who note, based on the Symphony IRI scanner dataset, that "a substantial fraction of the cross-sectional variation in the frequency of price change arises from variation in the prevalence of temporary sales," and that "larger stores tend to have more frequent price changes and temporary sales".

Despite our very robust time-series evidence, we did not find any significant relationship between sales and aggregate economic conditions *across* U.K. regions. In fact, as we discussed in Section 4.1, there is very little regional variation in the frequency of sales. The predominant role of chains in the use of discounts may explain these findings. For the United Kingdom, Dobson and Waterson (2008) document that since the early 2000s, "the major retailers [...] have eschewed the opportunity to customize prices on a store-by-store basis in favor of national pricing." The prevalence of uniform national pricing is also reminiscent of the findings of Nakamura *et al.* (2013) for the U.S. grocery market, who conclude that a large portion of the variation in the use of sales across chains is due

[17]While for chains the sales flag and V-shaped filters yield very similar time-series patterns, they display less similarities for smaller retailers (right-hand column). For example, the rise in sales frequency during the Great Recession is much more noticeable for the V-shaped series. One possibility is that independent stores have more limited means to display promotions, or the flagging of sales by price collectors is less reliable than for large chains. The statistical filter, on the other hand, does not suffer from this issue.

[18]Lloyd et al. (2014) reach similar conclusions using a scanner dataset for U.K. supermarkets.

to chain-level heterogeneity. This may in turn explain why Anderson *et al.* (2016) and Coibion *et al.* (2014) also do not find evidence of economic variation in sales in the U.S. cross-section data.

5 A general equilibrium model with sales

To understand the dynamics of temporary sales and their importance for aggregate fluctuations, we build a general equilibrium model in which discounts arise endogenously. In our framework, sales represent the optimal strategy of retail firms that compete for market share in an environment populated by consumers searching for low prices.[19] We show that the model generates realistic behavior of regular and sale prices at both the micro and macro levels. In particular, firms respond to aggregate nominal disturbances by adjusting the frequency of sales, in line with our empirical evidence. As a result, our model can be used to gauge the role of sales in shaping macroeconomic fluctuations, unlike recent general equilibrium models where the size and frequency of sales are more or less time-invariant.[20]

At its core, the model is a standard macroeconomic model with constant elasticity of substitution (CES) preferences and monopolistically competitive producers facing price-adjustment constraints. We augment the standard setup with a static consumer search problem that gives rise to sales. We now turn our attention to describing the main components of the model.

5.1 Static consumer search problem

The model economy is populated by a unit measure of infinitely-lived ex-ante identical households deriving utility from consuming goods of countably many differentiated varieties, indexed by j. Production occurs in a large number of spatially distinct markets, or "locations." In each location an industry of monopolistically competitive retail firms produce goods of all varieties.

Each retail firm sells a continuum of perfectly substitutable "brands" of a given variety j. For simplicity, we assume that the retailer chooses two price points for each brand it sells: a high (regular) price and a low (sale) price.[21] A retailer j in a particular location posts a fraction γ_j of its prices at level P_j^L and the remaining fraction, $1-\gamma_j$, at level P_j^H, where $P_j^H \geqslant P_j^L$. In this

[19] There is a vast economic and marketing literature that studies the dispersion of prices in markets with monopolistic sellers and consumers who face costs of searching for low prices. See, for example, Butters (1977), Salop and Stiglitz (1977), and Varian (1980).

[20] The general equilibrium models with sales by Guimaraes and Sheedy (2011) or Kehoe and Midrigan (2015) predict that the fraction of sales is very stable following aggregate shocks, which is at odds with our empirical findings from the U.K. and U.S.. Sudo *et al.* (2014) extend Guimaraes and Sheedy (2011) model to include time varying consumption weights and use it for understanding trends in sales behaviour in Japan.

[21] See Burdett and Judd (1983) for analysis of equilibrium price dispersion with many price points.

setup, households potentially face a non-degenerate distribution of prices for each variety—across brands in a given location and across locations. They can expoit this dispersion by spending time searching for lower prices.

To model consumer search, we assume that each household is composed of shoppers, each of whom is designated to shop for a specific variety. Shoppers can be of two types: bargain hunters or regular shoppers. Each shopper can consume from only one location in a period. The mix of shopper types is optimally chosen by the household head at the beginning of each period.

Search by *regular shoppers* is undirected and they are distributed randomly and uniformly across locations. Upon arriving to the location, the regular shopper responsible for variety j buys consumption from the retailer selling that variety according to a demand schedule that we describe later. We assume that regular shoppers randomly pick a single brand from all the brands of the same variety. This implies that they will buy at low price P_j^L with probability γ_j , and at high price P_j^H with probability $1 - \gamma_j$.

Bargain hunters know the distribution of prices for each variety across all locations, and they can search across locations by paying a fixed cost z in terms of lost labor income. We abstract from explicitly modeling the directed search process and focus instead on modeling its outcome. This outcome is the distribution of bargain hunters across varieties and locations that equates the opportunity cost of searching to the expected return on searching for low prices, which we specify below.[22] Upon arriving to their locations, bargain hunters buy consumption from the retailer selling their designated variety. Like regular workers, bargain hunters draw one brand for consumption, but face a higher probability of drawing the brand with low price: $f\gamma_j$, where $f > 1$, and lower probability of facing the high price, $1 - f\gamma_j$.

The fixed search costs z are realized at the beginning of the period and are i.i.d. across shoppers and across time. Let the distribution of z be given by a twice continuously differentiable distribution $G(z)$, with $0 \leq z \leq z_{\max} < \infty$. Since the fixed search cost is i.i.d., the shopping type for variety j is determined by a cut-off fixed cost, z_j^*: if the realized fixed cost is low, $z \leq z_j^*$, the shopper responsible for purchasing variety j becomes a bargain hunter; otherwise, if $z > z_j^*$, she is a regular shopper. Denote by α_j the probability that a shopper for variety j is a bargain hunter, i.e., $\alpha_j = G(z_j^*)$. By choosing the cut-off cost z_j^*, or equivalently, by choosing the probability of sending the variety-j shopper as bargain hunter, the household faces the expected search cost $\Xi(\alpha_j) = \int_0^{G^{-1}(\alpha_j)} z dG(z)$, expressed in terms of forgone hours worked. We focus on symmetric search outcomes whereby the shares of bargain hunters and regular shoppers responsible for variety j are equal across locations.

[22] A complete treatment of the general search model is provided in Lucas and Prescott (1974).

Shoppers cannot communicate the prices they draw from the retailer back to the household head. As a result, the household head cannot arbitrage price differences ex-post nor condition demand on the shopper type. Hence, in this setup, there is no ex-ante heterogeneity of households or retail firms. As we show in Section 6, retailers find it optimal to set a lower price for a subset of the brands they sell in order to attract bargain hunters.[23]

Let us denote by c_j^L the consumption level if a shopper finds a brand on sale, P_j^L, and by c_j^H consumption at the regular price, P_j^H. The households' consumption problem in each period comes down to choosing probabilities and consumption bundles $\left\{\alpha_j,\, c_j^L,\, c_j^H\right\}$ to minimize expenditures on total consumption and total search cost:

$$\min \sum_j \left[\alpha_j \left(f\gamma_j P_j^L c_j^L + (1-f\gamma_j) P_j^H c_j^H\right) + (1-\alpha_j)\left(\gamma_j P_j^L c_j^L + (1-\gamma_j) P_j^H c_j^H\right) + W\Xi(\alpha_j)\right] ,$$

subject to the CES consumption aggregate

$$c = \left\{\sum_j \left[\alpha_j \left(f\gamma_j P_j^L c_j^L + (1-f\gamma_j) P_j^H c_j^H\right)^{1-\frac{1}{\theta}} + (1-\alpha_j)\left(\gamma_j P_j^L c_j^L + (1-\gamma_j) P_j^H c_j^H\right)^{1-\frac{1}{\theta}}\right]\right\}^{\frac{\theta}{\theta-1}} ,$$

where W denotes nominal wages, and θ is the elasticity of substitution across varieties.

The first-order conditions for consumption varieties, conditional on price P_j^k, $k = \{L,\, H\}$, yield standard CES demand schedules:

$$c_j^k = \left(\frac{P_j^k}{P}\right)^{-\theta} c\,, \tag{2}$$

and an aggregate price index P

$$P = \left\{\sum_j \left[\alpha_j (P_j^B)^{1-\theta} + (1-\alpha_j)\,(P_j^R)^{1-\theta}\right]\right\}^{\frac{1}{1-\theta}} , \tag{3}$$

which is expressed in terms of the expected prices faced by bargain hunters and regular shoppers,

$$P_j^B = \left(f\gamma_j (P_j^L)^{1-\theta} + (1-f\gamma_j)(P_j^H)^{1-\theta}\right)^{\frac{1}{1-\theta}}$$

and

[23] Equilibrium price dispersion of this type is studied in models with ex-post heterogeneity in consumer information, e.g., Burdett and Judd (1983). Our setup is akin to one studied by Lester (2011). In his consumer search model buyers are heterogeneous with respect to ex-ante information about sellers. A fraction of buyers have perfect information about both the prices and locations of all sellers. These buyers choose the seller (or mix between sellers) promising the maximum expected utility. The remaining fraction of buyers cannot observe the prices posted by any particular seller. Since all sellers appear ex-ante identical to an uninformed buyer, he picks a seller to visit at random.

$$P_j^R = \left(\gamma_j (P_j^L)^{1-\theta} + (1-\gamma_j)(P_j^H)^{1\text{-}\theta}\right)^{\frac{1}{1-\theta}}.$$

The difference between bargain hunters and regular shoppers stems entirely from the price distributions they face: bargain hunters encounter, on average, lower prices. If we denote the expected consumption spending of a bargain hunter by $c_j^B = f\gamma_j P_j^L c_j^L + (1 - f\gamma_j) P_j^H c_j^H$ and of a regular shopper by $c_j^R = \gamma_j P_j^L c_j^L + (1-\gamma_j) P_j^H c_j^H$, then the first-order condition for the probability of bargain hunting for variety j, α_j, is

$$W\Xi'(\alpha_j) = \left(\frac{\theta}{\theta - 1} D_j^B - 1\right) c_j^B - \left(\frac{\theta}{\theta - 1} D_j^R - 1\right) c_j^R \tag{4}$$

where D_j^B and D_j^R are price-dispersion factors that are functions of the size of the price discount. Condition (4) equates the opportunity cost of searching given by the marginal expected fixed search cost (left-hand side) with the expected return on searching given by the marginal utility flow from bargain hunting for variety j (on the right-hand side), which is positive as long as $P_j^L < P_j^H$ and $f > 1$.[24] Combined with the fact that the fixed search cost is continuous over $0 \leq z \leq z_{\max}$, this ensures that the measure of bargain hunters for a particular variety is strictly positive, i.e., $\alpha_j > 0$.

This endogenous split of household types is the key difference from Guimaraes and Sheedy (2011), who assume a constant fraction of bargain hunters. We demonstrate below that the importance of sale prices for aggregate price flexibility crucially depends on the interaction of households' search for lower prices and retailers' decisions for setting those prices.

Another feature of our setup is that the individual searching and consumption decisions are easily aggregated. The remaining dynamic optimization problem of households is compeletly standard.

5.2 Household's dynamic optimization problem

In addition to searching and consuming, households supply working hours in a competitive labor market, and trade money and nominal bonds. Let B_t denote holdings of nominal bonds paying the household a gross return R_t in period $t+1$. In each period the timing of events is as follows. First, the aggregate money shock and household-specific fixed costs are realized. Retail firms post their prices, and households earn their incomes. Households then combine unspent money holdings from the previous period with labor income, return on bonds, dividends, and transfers, and split them into current-period cash and bond holdings to be carried over to the next period. Households

[24] This follows from the fact that under these assumptions, $D_j^B > D_j^R > 1$, and $c_j^B > c_j^R$, where $D_j^B = \left(f\gamma_j (P_j^L)^{-\theta} + (1 - f\gamma_j)(P_j^H)^{\text{-}\theta}\right)^{\frac{1}{-\theta}} / P_j^B$ and $D_j^R = \left(\gamma_j (P_j^L)^{-\theta} + (1-\gamma_j)(P_j^H)^{\text{-}\theta}\right)^{\frac{1}{-\theta}} / P_j^R$.

choose their shopper types and pay the fixed search costs. Shoppers are then matched with brands and varieties, and use cash to make their consumption purchases.[25]

Households choose sequences of total consumption c_t, hours worked l_t, cash holdings, M_t, and bond holdings B_t to maximize lifetime expected utility:

$$E_0 \sum_t \beta^t \left[\frac{c_t^{\sigma-1}}{1-\sigma} - l_t \right] ,$$

subject to the budget constraint

$$M_t + B_t \quad \leq \quad M_{t-1} - P_t c_t + W_t \left[l_t - \Xi_t \right] + R_{t-1} B_{t-1} + \Pi_t + T_t \, ,$$

and a cash-in-advance constraint

$$P_t c_t \leq M_t \, ,$$

where Π_t are dividends paid, T_t are lump-sum government transfers, and Ξ_t denotes household's expected total cost of searching, expressed in terms of forgone hours worked.

The remaining optimality conditions are standard: the first-order condition for hours worked implies

$$P_t c_t^{\sigma} = W_t \, ,$$

while the first-order conditions for nominal bonds yield

$$1 = \beta R_t E_t \left(\frac{c_{t+1}^{-\sigma}}{c_t^{-\sigma}} \frac{P_t}{P_{t+1}} \right) ,$$

where $E_t(\cdot)$ denotes the expected value with respect to information available through period t.

Aggregate fluctuations in the model stem from shocks to the money supply, M_t^s. We assume that the supply of money follows a random-walk process of the form

$$\ln M_t^s = \ln M_{t-1}^s + \mu_t \, ,$$

where μ_t is log money growth, a normally distributed i.i.d. random variable with mean 0 and standard deviation σ_μ.

[25] For tractability, we assume that shoppers can perfectly share within-period consumption risk, and all available cash is spent.

5.3 Retailer's problem

A retail firm selling brands of variety j in a given location is endowed with a linear production technology that converts hours worked n_{jt} into output of brands of variety j, y_{jt}:

$$y_{jt} = n_{jt} \,. \tag{5}$$

Let α_{jt}^{B} and α_{jt}^{R} denote the fractions of bargain hunters and regular shoppers buying variety j in a particular location.

Due to the random and uniform allocation of regular shoppers across locations, retailers take the number of regular shoppers as given. By contrast, the mass of bargain hunters visiting the retailer is linked to the optimal choice of households by

$$\alpha_{jt}^{B} = \int_{i=1}^{1} \alpha_{jt}\left(i\right) di \,, \tag{6}$$

where the integral is taken over all households. Since the distribution of bargain hunters across locations depends on the distribution of prices across locations, retailers can entice active shoppers by lowering some of their prices.

In each period, upon realization of the aggregate shock, the retailer posts price P_{jt}^{L} for a fraction γ_{jt} of its brands and price P_{jt}^{H} for the remaining fraction $1-\gamma_{jt}$. The brands with a price discount are chosen randomly (among brands that are not subject to price adjustment constraints below) and independently across periods. Posting a price discount requires a fixed cost κ, expressed in labor units; total fixed cost for a retailer, therefore, is $\kappa W_t \gamma_{jt}$. Households of either type who are matched with the retailer in a particular location draw their brands and purchase according to the demands (2). The retailer produces using technology (5) to fully satisfy the demand at posted prices.

Finally, retailers face constraints on price adjustment. We assume that retailers face Taylor (1980) price adjustment constraints for high prices. Low prices are assumed to be flexible, although this has virtually no effect on our results, as we show below. Retailer j sets a new high price in period t and keeps that price fixed for T periods, facing the following pricing constraints:

$$P_{jt}^{H} = P_{jt+1}^{H} = ... = P_{jt+T-1}^{H} \,. \tag{7}$$

Such price contracts are evenly staggered across varieties j, so that in every period a measure $1/T$ of retailers resets their prices. To keep the model tractable, we assume that all price changes are perfectly synchronized for all retailers selling the same variety j.

Denote by $\mathit{\Pi}_t(P) = (P - W_t)\left(\frac{P}{P_t}\right)^{-\theta} c_t$ the period-t retailer's profit function at price P. The retailer's expected profits from selling to bargain hunters and regular shoppers in period t are $\Pi^B_{jt} = f\gamma_{jt}\mathit{\Pi}_t(P^L_{jt}) + (1 - f\gamma_{jt})\,\mathit{\Pi}_t(P^H_{jt})$ and $\Pi^R_{jt} = \gamma_{jt}\mathit{\Pi}_t(P^L_{jt}) + (1 - \gamma_{jt})\,\mathit{\Pi}_t(P^H_{jt})$, respectively.

The problem of a retail store selling variety j is to choose the sequences of its two price points P^H_{jt} and P^L_{jt}, the fraction of price discounts γ_{jt}, the expected fraction of bargain hunters in retailer's location α^B_{jt}, total labor inputs n_{jt}, and total outputs y_{jt}, to maximize the present discounted value of its profits over the duration of price constraints:

$$E_t \sum_{\tau=t}^{t+T-1} \left\{ \beta^{t-\tau} \frac{c_\tau^{-\sigma}}{P_\tau} \cdot \left[\alpha^B_{j\tau}\Pi^B_{j\tau} + \alpha^R_{j\tau}\Pi^R_{j\tau} - \kappa W_\tau \gamma_{j\tau}\right] \right\}, \tag{8}$$

subject to the supply constraint

$$y_{jt} \le \alpha^B_{jt} c^B_{jt} + \alpha^R_{jt} c^R_{jt}\,,$$

and the demand constraints (2), the technology constraint (5), the constraint for the number of bargain hunters in retailer's location (4) and (6), and price adjustment constraints (7).

5.4 Equilibrium

The market clearing condition for labor implies that the hours supplied by households in the labor market are equal to the total hours demanded by the retail firms across all locations (we omit summations across locations):

$$l_t - \Xi_t = \sum_j \left[n_{jt} + \kappa\gamma_{jt}\right]\,.$$

The equilibrium search flow condition states that the total number of shoppers for each variety j is equal to 1, and is equal across all locations:

$$\alpha^B_{jt} + \alpha^R_{jt} = 1\,.$$

In a symmetric equilibrium, each location receives the same measure of bargain hunters for each variety, i.e.,

$$\alpha^B_{jt} = \alpha_{jt}\,,$$

where α_{jt} is the average measure of bargain hunters in a particular location given by (4).

The market clearing condition for each good variety j states that the total consumption of that variety is equal to the total amount of that variety produced:

$$\alpha_{jt} c^B_{jt} + (1 - \alpha_{jt})\, c^R_{jt} = \sum_j y_{jt}\,.$$

Finally, money and bond markets also clear.

A *symmetric equilibrium* for this economy is a collection of allocations for households: c_t, l_t, c^H_{jt}, c^L_{jt},M_t, B_t, α_{jt}; prices and allocations for firms: y_{jt}, n_{jt}, P^H_{jt}, P^L_{jt}, γ_{jt}, α^B_{jt}, α^R_{jt}; and aggregate prices P_t, W_t, R_t that satisfy household and retailer maximization, and market clearing conditions.

6 Model mechanics

In this section, we summarize the mechanics of the model, focusing on the case of fully flexible prices. Detailed explanations are provided in the Supplementary Material.

6.1 Equilibrium price dispersion

There exists a locally unique equilibrium with price dispersion in this model. In the Supplementary Material we derive a sufficient condition for the existence of two-price equilibria, expressed in terms of the expected fixed cost $\Xi(\alpha)$, probability parameter f, the cost of posting sales κ, and demand elasticity θ. The retailer prefers to post two different prices, instead of a single price for all its brands when the expected returns to retailer from posting an extra fraction of brands on sale outweighs the opportunity cost of maintaining high average profit margin. That is the case when the return on searching is high, the cost of posting sales is low, and demand is inelastic. When the sufficient condition is satisfied, there are always varieties for which the search cost realizations are low-enough so that households choose the shoppers responsible for that variety to be bargain hunters. In turn, the presence of bargain hunters validates retailers' incentives to post sales. The Supplementary Material also shows that there is no single-price equilibrium when the sufficient condition is satisfied.

6.2 Fraction of sale prices

What is the desirable number of discount prices for a retailer? The retailer's first-order condition with respect to the fraction of brands on sale γ_{jt} equates the opportunity cost of increasing the fraction of sales with the marginal return from extra sales:

$$\left(\alpha^R_{jt} + f\alpha^B_{jt}\right)\left[\mathit{\Pi}_t(P^H_{jt}) - \mathit{\Pi}_t(P^L_{jt})\right] + \kappa W_t = \frac{\partial \alpha^B_{jt}}{\partial \gamma_{jt}} \Pi^B_{jt}\,. \tag{9}$$

The opportunity cost of increasing the fraction of sales is the sum of two terms. The first term is the marginal loss from selling more goods at a lower price, which is the product of the measure of brands purchased on sale, $\alpha_{jt}^R + f\alpha_{jt}^B$, and the loss per brand on sale, $\mathit{\Pi}_t(P_{jt}^H) - \mathit{\Pi}_t(P_{jt}^L)$. The second term is the marginal increase in the total fixed cost of posting sales, κW_t. The retailer's opportunity cost is increasing with γ_{jt} as more brands are sold at a loss and the total fixed cost of posting sales increases.

The marginal return from extra sales is the product of the marginal increase in the measure of bargain hunters, $\frac{\partial \alpha_{jt}^B}{\partial \gamma_{jt}}$, and the average profits per bargain hunter, Π_{jt}^B. As γ_{jt} grows, the inflow of bargain hunters dissipates because the household's marginal cost of becoming a bargain hunter, $\Xi''(\alpha_{jt})$, is increasing with the number of bargain hunters. The retailer's average profit per bargain hunter is also decreasing with γ_{jt}, since more brands are sold at a lower price. Hence, the retailer's return from extra sales is decreasing in γ_{jt}.

We find that both the cost and the return on extra sales are very flat with respect to the level of γ_{jt}, which implies that their shifts due to aggregate shocks lead to large changes in γ_{jt}—the feature that helps the model match important aspects of the data, as we discuss in the next section.

6.3 Size of price discount

How does the firm pin down the levels of its high and low prices? Let $MR(P)$ and $MC(P)$ denote marginal revenue and marginal cost per brand sold at price P. If the retailer could not affect the number of bargain hunters in its store, there would be no price dispersion, since the concavity of the firm's problem would imply that the optimal low and high prices would equate the marginal revenue and marginal cost, $MR(P^*) = MC(P^*)$, where P^* is the optimal monopolistic price.

By contrast, if the retailer can attract bargain hunters, search condition (4) implies that a lower sale price or a higher regular price increase the return to bargain hunting. From the retail firm's first-order conditions with respect to its prices it follows that the retailer's optimal sale (regular) price is below (above) the monopolistic price, and the differences between the marginal revenue and marginal cost at these price points are offset by marginal effects of the inflow of bargain hunters, as illustrated in Figure 12.

The Supplementary Material contains the full system of equilibrium conditions in the model with sticky prices. The model is solved by applying the Blanchard-Khan (1980) local solution method to the log-linearized system of stochastic equilibrium equations around the deterministic steady state.

7 Model Simulations

In this Section we use model simulations to analyze the importance of countercyclical sales behavior. After properly calibrating the model's key parameters, we demonstrate that fluctuations in retailers' sales behavior can have sizeable implications for the economy's response to aggregate nominal shocks.

7.1 Parameterization

Table 6 summarizes the parameter values that we use in our quantitative simulations. The period is a month, so the discount factor is $\beta = 0.96^{1/12}$. We set $\sigma = 1$, a standard value in the business cycle literature, which implies that nominal wages track the money supply, $W_t = M_t$. Following Kryvtsov and Midrigan (2013), we assume that the growth rate of the money supply in the benchmark model is serially uncorrelated; we also report the results for serially-correlated money growth to facilitate comparisons with findings in Guimaraes and Sheedy (2011).

We set the elasticity of substitution across good varieties equal to $\theta = 5$, implying a 25% markup, which is between the high disaggregate markup estimates consistent with the industrial organization literature and the low macro aggregate markup estimates (e.g., see Basu and Fernald, 1997). Aggregate fluctuations in our model are not very sensitive to the level of θ.

We assume that regular prices change once every 12 months ($T = 12$), consistent with our data and existing studies. Although this degree of price stickiness is somewhat higher than reported in Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008), it is consistent with the findings of Eichenbaum et al. (2011) and Kehoe and Midrigan (2015). We assume that sale prices are fully flexible. As we show below, this assumption has little effect on the degree of aggregate price flexibility, since firms prefer to keep the size of the price discount roughly constant. This result is consistent with our findings that the size of the average price discount is acyclical, and it is consistent with Klenow and Kryvtsov (2008), who find that sale prices are roughly as sticky as regular prices.

Parameters that are specific to our price-discrimination model are the bargain-hunting technology parameter f, the fixed cost of posting price discounts κ, and parameters of the distribution of the fixed cost of searching $G(z)$. Since the parameter κ determines the average fraction of sales γ, we set κ to 0.08 to match an average fraction of sale prices of $\gamma = 0.05$. This number is midway between the unconditional sales frequencies we found for the United Kingdom and the United States (see Figures 2 and 5), and is in the range found in other studies of micro price data [26]

The parameters of $G(z)$ imply a specific level and slope of the marginal expected cost function $\Xi'(\alpha)$, and jointly determine the fraction of bargain hunters α and the size of the price discount P^L/P^H via the optimal condition for the number of bargain hunters (4) as well as retailer's pricing equations for P^L and P^H. To calibrate the average fraction of bargain hunters, we match the relative share of consumption at low prices. Glandon (2011) reports that for groceries in the United States, sales account for 17% of price quotes but 40% of revenue. In other words, one percentage point of price discounts accounts for 2.7 percentage points of revenue share; this ratio is our calibration target. The other calibration target is the average size of price discounts, which is $P^L/P^H = 0.78$ in our data. We assume that $G(z)$ is a uniform distribution that assigns equal probabilities of any fixed cost between 0 and $z_{\max}$. Then the remaining two parameters, $z_{\max} = 0.31$ and $f = 2$, are such that the model matches both of these calibration targets. The value $f = 2$ means that, in a given retail store, bargain hunters draw a sale twice more frequently than regular shoppers.

7.2 Model properties

Before using the model for quantifying the importance of sales for macroeconomic fluctuations, we demonstrate that it matches salient features of retail discount and search behavior highlighted in the literature.

Realistic discount strategy. The model rationalizes large price differences as *discounts*. Larger high-low price difference leads to more shoppers searching for bargains: the measure of bargain hunters α is higher when either P^H is higher or P^L is lower. The model predicts that the optimal regular price P^H is close to the optimal price of a single-price retailer, only 2% above P^*, however, the optimal sale price P^H is much lower, 20% below P^*. Indeed, since the marginal cost and revenue functions are increasing and concave in the posted price, the firm prefers to push down its sale price rather than mark up its regular price, as illustrated in Figure 12. Moreover, the average sale price is 0.3% below marginal cost, and therefore discounted brands are sold at a loss. Nonetheless, the total profit of the retailer is positive since sales generate an inflow of bargain hunters, most of whom (a fraction $1 - f\gamma = 0.9$) draw regularly-priced brands. In that dimension, our framework is akin to loss-leader models of imperfect competition, see Lal and Matutes (1994).

V-shaped sales. Price discounts in the model are short-lived, like they are in the data. For

[26] The fraction of sales reported in the European CPI data studies includes: between 0.03 and 0.05 in Austrian CPI data (Baumgartner et al. 2005); 0.03 in Norwegian CPI data (Wulfsberg 2016), 0.02 in French CPI data (Baudry et al. 2004; Berardi et al. 2015). For the United States, sales frequencies found by Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008) are between 0.07 and 0.11.

example, we can characterize the incidence of V-shaped sales by computing the measure of regularly priced brands that go on sale for n months before reverting to a regular price. The measure of prices that are not under pricing constraints is the sum of prices that were on sale last period, γ, and those regular prices whose price contract ended last period, $(1-\gamma)\frac{1}{T}$, where T is the duration of the price contract. Since new sale prices are drawn independently, the probability of a regular-to-sale price change is $\frac{\gamma}{(1-\gamma)\frac{1}{T}+\gamma}$; and the probability of a sale-to-regular price change is $\frac{(1-\gamma)\frac{1}{T}}{(1-\gamma)\frac{1}{T}+\gamma}$. Hence, the measure of regular prices that experience a V-shaped sale lasting n periods is $(1-\gamma)\frac{1}{T} \times \left(\frac{\gamma}{(1-\gamma)\frac{1}{T}+\gamma}\right)^n \frac{(1-\gamma)\frac{1}{T}}{(1-\gamma)\frac{1}{T}+\gamma}$. Under our benchmark calibration, this implies that the average fractions of one-month and two-month V-shapes in any period are 0.019 and 0.007, compared to 0.018 and 0.008 in the U.K. data. This corresponds to 61% and 24% of all V-shaped observations in the model, versus 51% and 23% empirically.

Time-variation of discounts. The model predicts that over time, price discounts vary largely in their frequency and not their size, in line with the evidence presented in Section 3. This feature stems from the near-linearity of the retailer's profits in the fraction of sales when that fraction is sufficiently small. By contrast, the size of discounts is pinned down by the concavity of the profit function and is relatively costly for retailers to vary. Hence, in response to shocks, retailers prefer to adjust the fraction of price discounts instead of their size. We later demonstrate this feature using model responses to a monetary shock.

Returns to shopping. The calibrated model predicts that the average consumption-weighted (unweighted) price paid by bargain hunters is 2% (1.2%) lower than the price expected to be paid by regular shoppers. This prediction is in line with the household-level evidence. For example, Kaplan and Menzio (2015) use the Kilts-Nielsen Consumer Panel Dataset to show that households with more non-employed members face prices that are, on average, 1% to 4.5% lower than prices faced by households whose members are all employed. In addition, the model is qualitatively consistent with evidence on search and shopping activity across households derived from time-use surveys. In the model, the fixed costs of bargain hunting are uniformly distributed between 0 and 0.31 units of time. Assuming that households work on average 26 hours a week, as reported by Cociuba, Prescott and Prescott (2009), this amounts to a range between 0 and 10 hours a week. Given this cost, the calibrated model predicts that only one out of 10 households decides to become a bargain hunter, i.e., $\alpha = 0.10$. The break-even fixed cost that makes a marginal household indifferent between bargain hunting and working is 0.03 units of time, or 1 hour a week. Hence, the average cost per household is around 10 minutes per week. For example, Aguiar and Hurst (2007) document important fluctuations in the shopping margin over the household's life cycle; Krueger and Mueller (2010) find that unemployed individuals devote 15% to 30% more time to shopping than the employed.

7.3 Responses to monetary shocks

Our baseline experiment is a negative 1% innovation to the growth rate of the money supply. Figure 13 compares the responses of output and the aggregate price level in the model with sales to the model without sales, i.e., the standard Taylor model with a single retail price.

The figure demonstrates a stark difference in the output and price responses in the two models. The aggregate price level (top-right panel) in the Taylor model falls steadily to –1% within one year, as additional sticky-price cohorts of retailers get a chance to lower their prices after the shock. Accordingly, the fall in output (top-left panel) is immediate, down to around –0.9% at the time of the shock, dissipating to zero after one year. By contrast, in the model with sales the real effect on output is about half as large, with output falling by about 0.5% on impact and going back to its pre-shock level within a year or so. Such a response in output is due to a faster fall in the aggregate price level, which, unlike in the standard single-price Taylor model, immediately drops by 0.5% at the time of the shock.

The marked difference in the degree of the aggregate price flexibility between the two models is due to the interaction between discounting by retailers and bargain hunting by households. We explain the workings of this mechanism by reviewing the factors that determine the size of the decline in the aggregate price level immediately after the shock.

Let us first write a log-linear approximation for the aggregate price level P_t in (3) assuming that both γ and α are close to zero. Denoting the average price discount by $\delta = P^L/P^H$, the log deviation of the aggregate price level from its steady-state value P is, approximately,

$$\widehat{P}_t \approx -\frac{P}{\theta - 1}\left(\delta^{1-\theta} - 1\right)\tilde{\gamma}_t \,, \tag{10}$$

where a hat (tilde) denotes the log-linearized (difference-) deviation from the steady state. Expression (10) indicates that the response of the average fraction of sales γ_t affects the aggregate price response by shifting the weight from high to low prices. The impact of this shift depends on both the size of the response of the fraction of sales itself and on the size of the shift in consumption, given by the factor $\delta^{-\theta}$. The response of the fraction of sales is $\tilde{\gamma}_t = 1.1$ percentage points (bottom-left panel in Figure 13), while $\delta = 0.78$ and $\theta = 5$ by calibration. Plugging these values into (10) gives a price level response of –0.5% at the time of the shock, as shown in the figure. Hence, both the average discount and the size of the response of the fraction of sales contribute to the fivefold difference of price responses in the model with and without sales.

We next explain how the magnitudes of the average discount δ and the response in the fraction of discounts $\tilde{\gamma}_t$ are associated with bargain hunting by households. First, we use the pricing equations for P^L and P^H to approximate the average size of price discount:

$$\delta \approx \left(1 + \frac{f-1}{\theta(\theta-1)} \frac{1}{\Xi''(\alpha)}\right)^{-1} . \tag{11}$$

The firm's desired price discount is larger when search elasticity is higher (lower $\Xi''(\alpha)$), when firm has more monopoly power (lower θ), and when search efficiency is higher (larger f).

Second, we note that sticky regular prices and falling wages after the shock drive up the average markup at the store, increasing the expected profit per each shopper. This raises the retailer's return from each additional sale relative to the loss per sale and induces the retailer to increase the number and the size of price discounts until the return and the loss from an additional sale are equal. As we mentioned above, retailers mostly vary the fraction of price discounts, and not their size: the response of the fraction of sales is $\tilde{\gamma}_t = 1.1$ percentage points, and the average size of price discounts (not shown on the figure) is virtually unchanged, increasing by only 0.07 percentage points.[27]

The increase in sales generates the additional inflow of bargain hunters: from (4) the response in the fraction of bargain hunters is approximately

$$\tilde{\alpha}_t \approx \tilde{\gamma}_t \frac{1}{\theta-1} \frac{f-1}{W\Xi''(\alpha)} \left(\delta^{1-\theta} - 1\right) ,$$

which, given our calibrated parameter values, yields $\tilde{\alpha}_t = 1.6$ percentage points (lower-right panel in Figure 13). The prediction of the model that households increase their search for low prices during economic downturns is corroborated by ample evidence from time-use surveys on search and shopping activity over the business cycle. For example, Aguiar, Hurst and Karabarbounis (2013) find a significant increase in time spent shopping for a typical household during the 2008–09 recession based on the American Time Use Survey. They document that about 7% of forgone market hours are allocated to increased shopping time. Using a similar data set, Nevo and Wong (2015) estimate that household consumption declined by 60% less than market expenditures due to the reallocation of time from market work to home production and shopping.

[27] Since the optimal discount size is roughly time-invariant, imposing price-adjustment constraints on sale prices does not materially alter the results.

Coibion, Gorodnichenko and Hong (2014) and Nevo and Wong (2015) emphasize the role of the decline in the opportunity cost of time in explaining the increased shopping activity during slumps. By contrast, in our model, the increase in households' shopping activity is mostly associated with higher returns on shopping due to more prevalent price discounting by retailer, and less with changes in the opportunity cost of time. In the robustness analysis in the Supplementary Material, we compare two extensions of the baseline model: with sticky nominal wages and with sticky nominal marginal costs. We show that shutting down the quick fall of household wages after the shock matters relatively little for households search behaviour; but shutting down the fall in markups removes the incentives for retailers to post sales during slumps. Hence, our paper shows that the rise in return to shopping due to more frequent sales in recessions can be a powerful driving mechanism behind fluctuations in households' shopping time.[28] The Supplementary Material also demonstrates that the main insights from the model carry over to the case with the persistent shock.

8 Measuring the impact of discounts on the aggregate price level

In our model, the flexibility of the aggregate price response is due to the increase in the fraction of sale prices and a shift of consumption weights toward discounted items. Because the shift in consumption weights is generally not observed by statistical agencies, fluctuations in sale prices may have important implications for aggregate price measurement. In this section, we explore the quantitative importance of this margin both theoretically and empirically.

First, we decompose the contribution of price discounts to fluctuations in the aggregate price level in the model. Figure 14 shows the response of the effective price index specified in equation (3) to the monetary contraction from the previous section. This effective price index, denoted by "P," falls to –0.56% on impact. To gauge the impact of sales, we compare this response to two counterfactual responses. The price index that ignores the variations in consumption weights, "P_{CPI}," is the model's counterpart to a standard fixed-weight CPI index: it decreases by 0.37% after the shock. Ignoring sales altogether, as in the standard Taylor model, brings the impulse response to a mere –0.07% on impact, shown by the graph denoted by "$P_{CPI}\,(regular)$". Hence, fluctuations in price discounts (keeping consumption weights constant) account for about two-thirds of the variance of the effective aggregate price level relative to the price level ignoring sales; and the

[28] A growing consensus in the macroeconomic literature is that nominal cost rigidities, rather than countercyclical markups, account for most of the monetary non-neutrality, see Christiano, Eichenbaum and Evans (2005). In contrast, our theory explains how countercyclical fluctuations in the number of sales in response to monetary shocks are associated with countercyclical retail markups and their importance of the responses after monetary shocks. Kryvtsov and Midrigan (2013) arrive at a similar conclusion, based on the study of the observed behavior of inventories.

unmeasured variation in consumption weights accounts for the remaining one-third (or half of the constant-weight price index variance).[29]

The advantage of using a model-based approach to measure the importance of sales lies in the explicit specification of the mechanisms that can account for their cyclicality. We argued in the previous section that these mechanisms are realistic and can be corroborated by the evidence on retailers' discounting and households' search behaviour. Nonetheless, some of the simplifying assumptions in the model may be material for the quantitative importance of sales. For example, our analysis assumes symmetry across consumption sectors; homogeneity of price discounts in the data, both in terms of frequency and size, and that nominal demand shocks are the main drivers of business cycles. While the model can be extended along these dimensions, we undertake a complementary approach based on available micro price data. Specifically, we quantify empirically the wedge between the standard (CPI) aggregate price index and the effective price index, assuming that the shift in consumption expenditures toward discounted items is based on the constant-elasticity-of-substitution demand system. Unlike the model-based approach, this empirical exercise takes the observed dynamics of price discounts as given and measures the impact of shutting down those dynamics, keeping all else equal. We focus on the main idea and results here, leaving the details to the Supplementary Material.

First, express any posted (log) price of an item i in month t, $p_{i,t}$, as the sum of the regular (log) price, $p^R_{i,t}$, and a (log) discount in the event of a sale, $\Delta_{i,t}$:

$$p_{i,t} = p^R_{i,t} + \Delta_{i,t}\,, \tag{12}$$

By definition, $\Delta_{i,t} < 0$ if there is a sale in month t, and $\Delta_{i,t} \equiv 0$ if there is no sale. This decomposition allows us to express the aggregate price index as the sum of the regular price index, p^R_t, and an aggregate discount term, Δ_t:

$$\begin{aligned} p_t &= \sum_i \omega_{i,t} p^R_{i,t} + \sum_i \omega_{i,t} h^{\Delta}_{i,t} \Delta_{i,t} \\ &= p^R_t + \Delta_t\,, \end{aligned} \tag{13}$$

where $\omega_{i,t}$ are the usual CPI item weights while $h^{\Delta}_{i,t}$ are additional weights for items on sale. To construct the official CPI, statistical agencies implicitly assume that $h^{\Delta}_{i,t} = 1$ for all i and t: the

[29] Most of the shift in consumption weights toward lower prices after the shock is due to the increase in consumption per household, while the increase in the number of households buying at low prices has a smaller impact.

weight on each quote is the same, irrespective of whether the item is discounted or regularly-priced. In light of the cyclicality of sales in the data, we study the implications of relaxing this assumption and allowing $h^{\Delta}_{i,t}$ to be above 1 and vary over time.

Ideally, one would estimate the $h^{\Delta}_{i,t}$ factors using data on both prices *and* quantities. Such data, however, are not available for a broad spectrum of products. Therefore, we opt for a simple alternative: we use the weights implied by a constant-elasticity-of-substitution demand system, similar to the one in our model. With an elasticity of substitution θ, the additional weight $h^{\Delta}_{i,t}$ is given by:

$$h^{\Delta}_{i,t} = \left(\frac{\exp p_{i,t}}{\exp p^{R}_{i,t}} \right)^{-\theta} = \left(\exp \Delta_{i,t} \right)^{-\theta} \ . \tag{14}$$

For instance, a value of $\theta = 5$ implies that the weight on a 25% discount should be about 4.2 times higher than the weight on a regular price. In the Supplementary Material we utilize a large scanner dataset for food to show that CES demand with $\theta = 5$ is a good approximation of the average elasticity of quantities to discount prices across products within a category.[30] This estimate is also in line with other studies. For example, Glandon (2011) reports that for groceries in the United States, sales account for 17% of price quotes but 40% of revenue, i.e., one percentage point of price discounts accounts for 2.7 percentage points of revenue share. Assuming a 20% discount, Glandon's results imply a discount weight of 4.1. Using IRI data, Chevalier and Kashyap found the elasticities for peanut butter to be above 4.5 across metropolitan markets, and between 2 and 10 for coffee. In sum, the amount of extra consumption generated by sales is substantial.

The final step is to use weights $\omega_{i,t}$ and $h^{\Delta}_{i,t}$ to construct a measure of the aggregate discount term Δ_t for the United Kingdom. An implicit assumption in this exercise is that the price elasticity of quantities estimated across products with regular and sales prices is useful for inferring the elasticity of aggregate price in response to time variation in the number of sales.[31] For example, if Coke goes on sale without Pepsi going on sale, then, under this assumption, the increase in demand for Coke rises by approximately the same amount as in the case when both Coke and Pepsi go on sale simultaneously. This assumption is reasonable since in the data the fraction of discounts is not correlated with regular prices or other discounts within the same category, and the size of discounts is relatively stable. Both of these features are accounted for in our model, which explicitly expresses the elasticity of the aggregate price with respect to the fraction of discounts in equation (10) as a

[30] We also explored an alternative approach where, instead of assuming CES demand, we measure weight factors $h^{\Delta}_{i,t}$ directly from the scanner dataset. This approach yields quantitatively similar results.

[31] We are grateful to an anonymous referee for pointing this out to us.

function of additional expenditure weights $\delta^{-\theta}$(the counterpart of $h_{i,t}^{\Delta}$ in equation 13).

In Figure 15, we compare our constructed series for three values of θ: 0 (equivalent to the CPI case, with $\omega_{i,t}^{\Delta} = 1$), 3 and 5. By construction, each series represents a deviation from the regular price index (p_t^R). Time series are provided for the aggregate as well as three broad categories.

First, there are significant differences in the average *levels* of the discount term across values of θ and categories. In the aggregate, the effective price level is 0.9% below the regular price index for $\theta = 0$ (the official CPI case), but 7% lower when $\theta = 5$. With an elasticity of subsitution of $\theta = 3$, the average contribution of sales ranges from -3% for *Food and Beverages* to -6.2% for *Household Furniture and Equipment*, reflecting mostly differences in the prevalence of sales across categories.

Without any fluctuations in the frequency or the size of discounts, these contributions would be stable over time. Figure 15 makes it clear that they are not. In the aggregate, for $\theta = 5$, sales contributed to a 3-percentage-point drop in the effective price level in less than three years, relative to the regular-price index. For *Personal care services,* the contribution falls from -5% in 2005 to -17% by 2010. In other words, ignoring sales altogether leads us to overestimate the effective price inflation of this category by about 2.4 percentage points per year over this period.

Comparing those fluctuations to the ones for $\theta = 0$ (i.e. sale and regular quotes get equal weights) indicates that the official CPI can in some instances significantly underestimate the role of sales: in this scenario, the range of the discount term Δ_t over the sample period is only about one percentage point for *Food and Beverages* and *Clothing and footwear* and 1.8 percentage point for *Personal care services*. These findings are consistent with Anderson et al. (2016) who document a small contribution of sales to U.S. inflation measured using the BLS methodology.

An important observation from the empirical analysis (Figure 15) is that the importance of discounts is larger than in the model (Figure 14). For example, in the empirical analysis, a 1-percentage-point increase in the fraction of sales is associated with a 1.56-percentage-point decrease in the aggregate price level for $\theta = 5$, whereas in the model that decrease is 0.45 percentage points. This difference is reconciled by Jensen's inequality: since the weights in (14) are convex and increasing in the absolute size of the log discount, larger sales have a disproportionately larger contribution to the discount term, Δ_t. For instance, recomputing Δ_t using only discounts smaller than 30% reduces the elasticity by a six-fold, from 1.56 to 0.26. Because in our model all sales have the same size, the effect of the spread in the size of discounts is not captured. The upshot is that the model provides a lower bound for the macroeconomic impact of sales.

In summary, our findings in this section demonstrate that properly accounting for sales can have important ramifications for the assessment of the cyclical behavior of aggregate price and consumption. First, even small fluctuations in the frequency of sales can have economically-relevant effects on the effective price level, in line with what we found in our model. Second, the size of these effects depends heavily on the expenditure-switching behavior of consumers between discounted and regular-price items; this channel is mostly absent from official consumer price indices. Third, larger discounts provide a disproportionately large impact on consumption and price levels.

The takeaways of these findings are two-fold. First, statistical agencies that construct consumer price indices may miss a large portion of the effects of sales on the effective or average prices paid by households. Second, macroeconomists who calibrate their models to the price data obtained without properly accounting for the dynamics of sales may be led to underestimate the degree of aggregate price flexibility and overestimate the fluctuations of real variables over the cycle.

9 Discussion and conclusions

A large body of literature in industrial organization has analyzed the motives behind the use of temporary discounts in the retail sector, such as price discrimination (Varian, 1980); discounts as loss leaders (Lal and Matutes, 1994); consumer stockpiling behavior (Pesendorfer, 2002) or inventory management under uncertainty (Aguirregabiria, 1999). Yet, despite their ubiquity in many micro-price datasets, macroeconomists have traditionally paid little attention to the role of sales, instead focusing on regular price moments to calibrate their models. We provide evidence from CPI micro data that the frequency of sales is strongly countercyclical in the United Kingdom and United States, and that it rose significantly in both countries during the Great Recession as well as past downturns. Furthermore, we demonstrate that the contribution of sales to fluctuations in the aggregate price level can be significant. Most of this contribution would not be captured my traditional measurement methods because they do not take into account the disproportionate weight of discounted items in households' consumption baskets. Our conclusion is that focusing on regular or reference prices may lead macroeconomists to miss economically relevant aspects of pricing over the business cycle.

The evidence presented in this paper contributes to several ongoing strands of the literature.

Three recent papers have investigated whether price discounts provide an important margin for retailers who face price adjustment constraints. In an extended version of the general equilibrium menu-cost model, Kehoe and Midrigan (2015) show that temporary price changes indeed help retailers to respond to large and transient idiosyncratic shocks, but not to persistent monetary shocks. Guimaraes and Sheedy (2011) also argue that sales are mostly irrelevant for the transmission of monetary shocks. They build a sticky-price model with sales stemming from consumer heterogeneity and incomplete information. In their setup, strong strategic substitutability of sales at the micro level implies that their frequency and size barely responds to aggregate shocks. Hence, both papers predict that retailers do not vary price discounts over the business cycles driven by nominal aggregate demand disturbances, which is at odds with our finding that sales are strongly countercyclical. On the other hand, Alvarez and Lippi (2015) argue that when firms can choose a 2-point "price plan" allowing them to deviate temporarily from the sticky reference price, the flexibility of the aggregate price increases significantly. Future business cycle models with incomplete price adjustment should further explore mechanisms that are compatible with our evidence.

Our findings offer some guidance as to which modeling features may be most appropriate. For example, the much higher use of sales by larger retailers in our data may be evidence that posting temporary discounts comes with significant fixed marketing costs, such as in-store displays or the distribution of promotional flyers. It could also be a sign that wholesalers are more likely to put in place trade promotion agreements of the type studied by Anderson et al. (2016) with larger, more visible retailers. Paying particular attention to large retail chains also appears increasingly justified not only due the prevalence of national pricing strategies, but also because of a trend toward higher concentration in the retail industry, as evidenced, for example, in Hong and Li (2015).

Moreover, even if we observe that sales are more likely to occur toward the end of a product life, the similar cyclical behavior of clearance and non-clearance sales seems to indicate that inventory decisions are not likely to be the main cause behind the more prevalent use of discounts during recessions. Instead, our findings appear to favor mechanisms in which the time-varying use of sales by retailers is driven by variations in the price sensitivity of consumers over the cycle. For example, Coibion, Gorodnichenko and Hong (2015) and Nevo and Wong (2015) emphasize the role of the decline in the opportunity cost of time in explaining the increased shopping activity during slumps.

The extent to which sales matter quantitatively depends—in addition to volatility of sales occurence—on consumers' shopping activity over the business cycle. There is now ample evidence that households increase their shopping activity and search activity for low prices during economic downturns. For example, Aguiar, Hurst and Karabarbounis (2013) find a significant increase in time spent shopping for a typical household during the 2008–09 recession based on the American Time Use Survey. They document that about 7% of forgone market hours are allocated to increased shopping time. Using Nielsen Homescan database, Nevo and Wong (2015) estimate that household consumption declined by 60% less than market expenditures due to the reallocation of time from market work to home production and shopping. Gordon, Goldfarb and Li (2013) document that price elasticity on average rises when the economy weakens. Further studies linking cyclical shopping activity to retailers' price discrimination behavior appears promising.

Finally, our results show the importance of investigating further the factors that may be driving a wedge between the true, model-relevant price index and the official fixed-basket CPI. More detailed data and better methods should be the focus of this endeavor. For example, Chevalier and Kashyap (2015) propose a methodology based on the lowest available "best price" approach arguing that it would improve inflation measurement without much extra cost of data collection.

References

[1] Aguirregabiria, V. 1999. "The Dynamics of Markups and Inventories in Retailing Firms," *Review of Economic Studies*, 66, 275–308.

[2] Aguiar, M. and Hurst, E. 2007. "Lifecycle Prices and Production." *American Economic Review* 97(5), 1533–1559.

[3] Aguiar, M., Hurst, E. and Karabarbounis, L. 2013. "Lifecycle Prices and Production." *American Economic Review* 103(5), 1664–96.

[4] Álvarez, L. J., Dhyne, E., Hoeberichts, M., Kwapil, C., Le Bihan, H., Lünnemann, P., Martins, F., Sabbatini, R., Stahl, H., Vermeulen, P. and Vilmunen, J. 2006. "Sticky Prices in the Euro Area: A Summary of New Micro-Evidence," *Journal of the European Economic Association*, 4(2-3): 575–584.

[5] Alvarez, F. and F. Lippi. 2015. "Price plans and the real effects of monetary policy," mimeo, December 2015.

[6] Anderson, E., Malin, B., Nakamura, E., Simester, D. and Steinsson, J. 2016. "Informational Rigidities and the Stickiness of Temporary Sales," mimeo Columbia University.

[7] Baudry, L., Le Bihan, H., Sevestre, P. and Tarrieu, S. 2004. "Price rigidity. Evidence from the French CPI micro-data," Working Paper Series 0384, European Central Bank.

[8] Baumgartner, J., Glatzer, E., Rumler, F. and Stiglbauer, A. 2005. "How frequently do consumer prices change in Austria? Evidence from micro CPI data," Working Paper Series 0523, European Central Bank.

[9] Berardi, N., Gautier, E. and Le Bihan, H. 2015. "More Facts about Prices: France Before and During the Great Recession," mimeo, Banque de France.

[10] Bils, M. and P. Klenow. 1998. "Using Consumer Theory to Test Competing Business Cycle Models," *Journal of Political Economy*, University of Chicago Press, vol. 106(2), pages 233-261, April.

[11] Bunn, P. and Ellis, C. 2012. "Examining the Behaviour of Individual UK Consumer Prices," *The Economic Journal* 122 (February): F35–F55.

[12] Butters, G. 1977. "Equilibrium distribution of sales and advertising prices." *Review of Economic Studies* 44 (October): 465–91.

[13] Chevalier, J. and Kashyap, A. 2015. "Best Prices: Price Discrimination and Consumer Substitution," mimeo, University of Chicago.

[14] Coibion, O., Gorodnichenko, Y. and Hong, G. H. 2014. "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications," *American Economic Review* 105(2015), 993-1029.

[15] Dobson, P.W. and Waterson, M. 2008. "Chain-Store Competition: Customized vs. Uniform Pricing", Warwick Economic Research Paper no. 840.

[16] Eichenbaum, M., Jaimovich, N. and Rebelo, S. 2011. "Reference Prices, Costs, and Nominal Rigidities." *American Economic Review* 101(1): 234–62.

[17] Glandon, P. 2011. "Sales and the (Mis)Measurement of Price Level Fluctuations." mimeo, Vanderbilt University.

[18] Gordon, B. R., Goldfarb, A. and Li, Y. 2013. Does Price Elasticity Vary with Economic Growth? A Cross-Category Analysis," *Journal of Marketing Research* L(2013): 4–23.

[19] Guimaraes, B. and Sheedy, K. 2011. "Sales and Monetary Policy," *American Economic Review* 101(2): 844–76.

[20] Hendel, I. and A. Nevo. 2006. "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica* 74(6): 1637–1673.

[21] Hong, G.H. and Li, N. 2015. "Market Structure and Cost Pass-Through in Retail," forthcoming in *Review of Economics and Statistics*.

[22] Kaplan, G. and Menzio, G. 2015. "The Morphology of Price Dispersion," *International Economic Review* 56: 1165–1206, November.

[23] Kehoe, P. and Midrigan, V. 2015. "Prices are Sticky After All," *Journal of Monetary Economics* 75: 35–53.

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

[25] Klenow, P. and Malin, B. 2011. "Microeconomic Evidence on Price-Setting," in *Handbook of Monetary Economics*, ed. by B. Friedman, and M. Woodford, pp. 231–284, Amsterdam, Holland. Elsevier.

[26] Klenow, P. and Willis, J. 2007. "Sticky Information and Sticky Prices," *Journal of Monetary Economics* 54 (September): 79–99.

[27] Krueger, A. and Mueller, A. 2010. "Job Search and Unemployment Insurance: New Evidence from Time Use Data." *Journal of Public Economics* 94: 298–307.

[28] Lal, R. and Matutes, C. 1994. "Retail Pricing and Advertising Strategies," *Journal of Business*, 67(3), July 1994.

[29] Lloyd, T.A., McCorriston, S., Morgan, C.W., Poen, E. and Zgovu, E. 2014. "Retail Price Dynamics and Retailer Heterogeneity: UL evidence," *Economic Letters* 114(2014): 434–438.

[30] Mahajan, S. 2006. "Concentration ratios for businesses by industry in 2004," *Economic Trends,* Office for National Statistics, 635, October 2006.

[31] Nakamura, A., Nakamura, E. and Nakamura, L. I. 2013. "Price dynamics, retail chains and inflation measurement" *Journal of Econometrics* 162(2011), 47-55.

[32] Nakamura, E. and Steinsson, J. 2008. "Five facts about prices: a reevaluation of menu cost models." *Quarterly Journal of Economics* 123(4), 1415.

[33] Nevo, A. and Wong, A. 2015. "The Elasticity of Substitution Between Time and Market Goods: Evidence from the Great Recession," NBER Working Paper No. 21318, July.

[34] Pesendorfer, M. 2002. "Retail Sales. A Study of Pricing Behavior in Supermarkets," *Journal of Business*, 75(1), 33-66.

[35] Salop, S., and Stiglitz, J. 1977. "Bargains and ripoffs: A model of monopolistically competitive price dispersion." *Review of Economic Studies* 44 (October): 493–510.

[36] Sudo, N., Ueda, K., Watanabe, K. and Watanabe, T. 2014. "Working Less and Bargain Hunting More: Macro Implications of Sales during Japan's Lost Decades," mimeo, University of Tokyo.

[37] Varian, H. 1980. "A model of sales." *American Economic Review* 70 (September): 651–59.

[38] Vavra, J. 2014. "Inflation dynamics and time-varying volatility: new evidence and a new Ss interpretation," *Quarterly Journal of Economics* 129(1), 215–258.

[39] Wulfsberg, F. 2016. "Inflation and Price Adjustments: Micro Evidence from Norwegian Consumer Prices 1975–2004," *American Economic Journal: Macroeconomics* 8(3):175–194, July.

Figure 1: Distribution of the size of sales for the three main filters

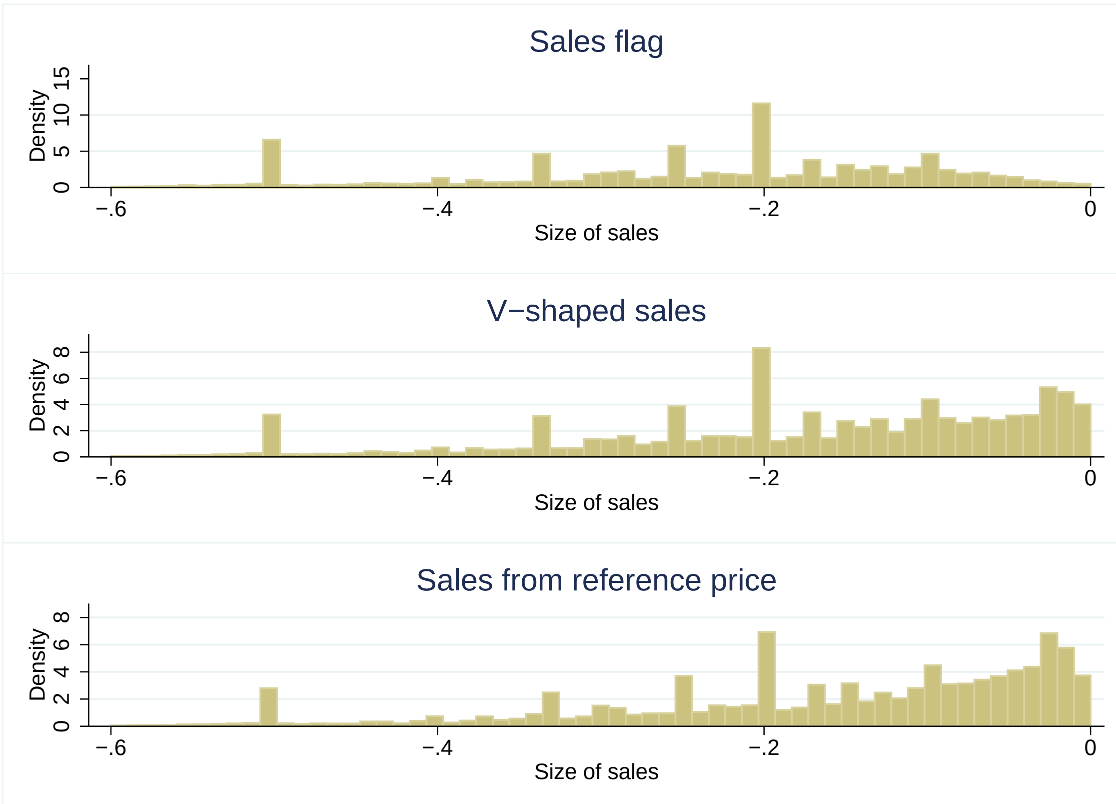

Notes: Histograms of the size of discounts for the sales flag, V-shaped sales and sales from reference price. Sample period is 1996:02–2013:09. U.K. Office for National Statistics data.

Figure 2: Frequency of sales flags (raw and 12-month moving average) and unemployment rate

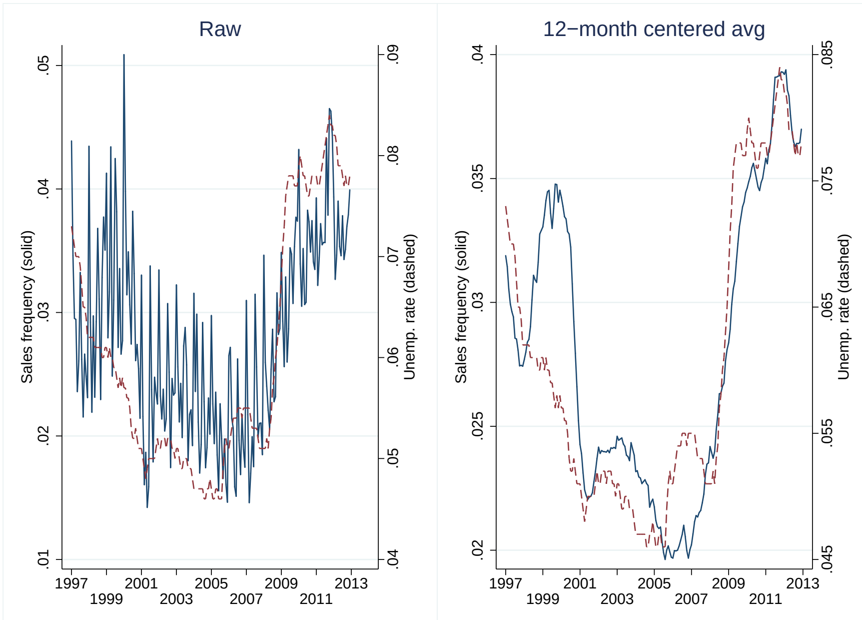

Notes: Monthly proportion of products with a sales flag. Aggregation is done using CPI quote weights. Sample period is 1996:02–2013:09. U.K. Office for National Statistics data.

Figure 3: The evolution of the frequency of sales (alternative filters)

Notes: Monthly proportion of products on sale using the three types of filers. Aggregation is done using CPI quote weights. For the top right plots, the product-level weights are time invariant and correspond to their time-series average. Sample period is 1996:02–2013:09. U.K. Office for National Statistics data.

Figure 4: The evolution of the size of sales

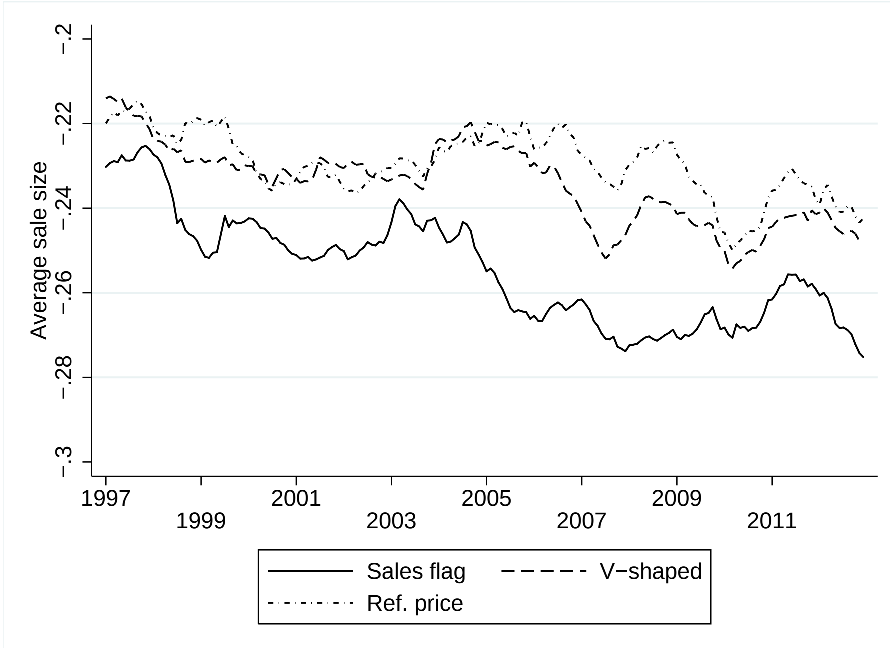

Notes: Weighted average of the size of discounts using the three types of filers. Aggregation is done using CPI quote weights. Data is monthly and the sample period is 1996:02–2013:09. U.K. Office for National Statistics data.

Figure 5: Frequency of sales in the U.S. CPI data and unemployment rate

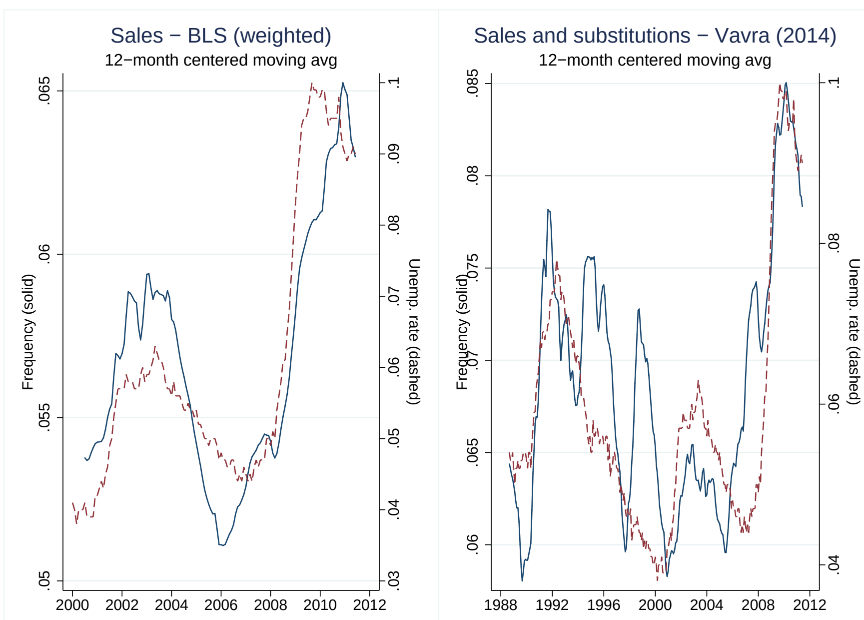

Notes: Monthly proportion of products with a sales flag (left) and sales flag or substitution (right). Aggregation is done using CPI weights. Sample periods are 2000-2012 (left) and 1988-2012 (right). U.S. Bureau of Labor Statistics data.

Figure 6: Frequency of sales for food products, U.S. CPI

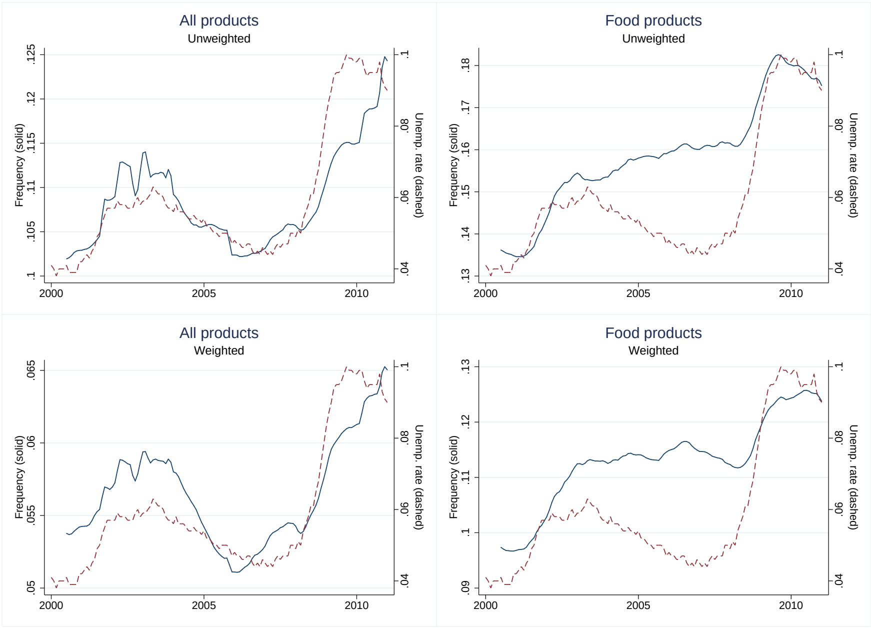

Notes: Monthly proportion of products with a sales flag. For the plots in the left column, aggregation is done using CPI quote weights. Sample periods is 2000-2012. U.S. Bureau of Labor Statistics data.

Figure 7: Cyclicality of sales at the category level, U.K. CPI data

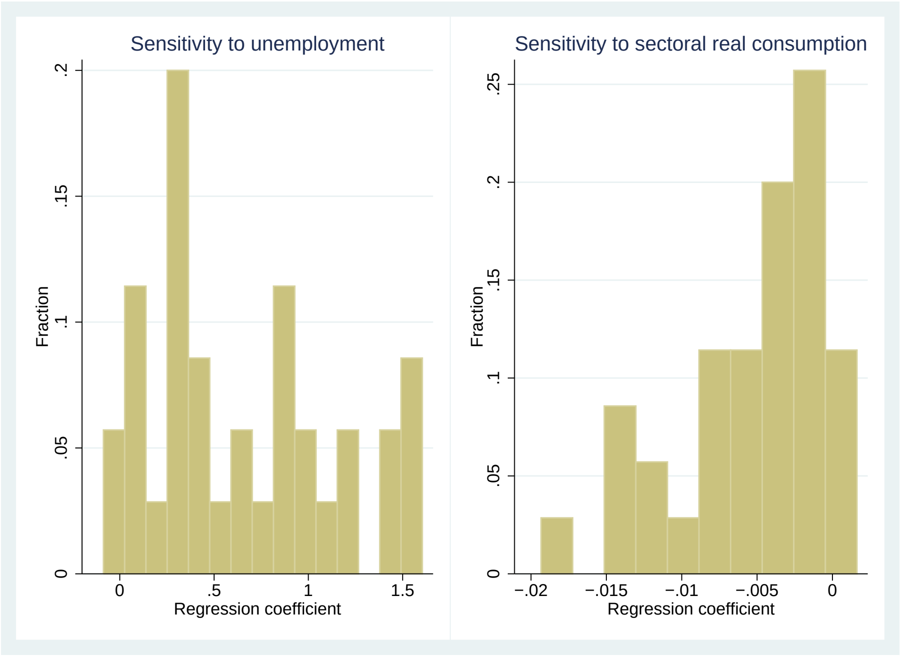

Notes: Distribution of coefficients from two regressions at the category level. Left: Sales frequency on monthly aggregate unemployment rate and a time trend. Right: Sales frequency on the log of quarterly category-level real consumption and a time trend. Sales are identified using the ONS sales flag and must be at least 10% in size. U.K. Office for National Statistics data.

Figure 8: Cyclicality of sales at the regional level, U.K. CPI data

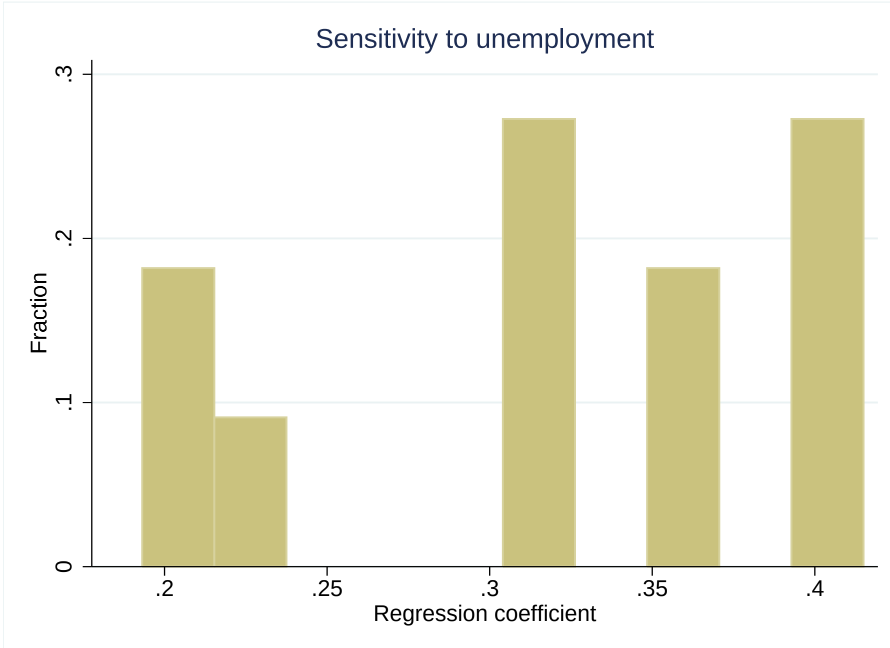

Notes: Distribution of coefficients from regional regressions sales frequency on monthly regional unemployment rate and a time trend. Sales are identified using the ONS sales flag and must be at least 10% in size. U.K. Office for National Statistics data.

Figure 9: Frequency of clearance and non-clearance sales

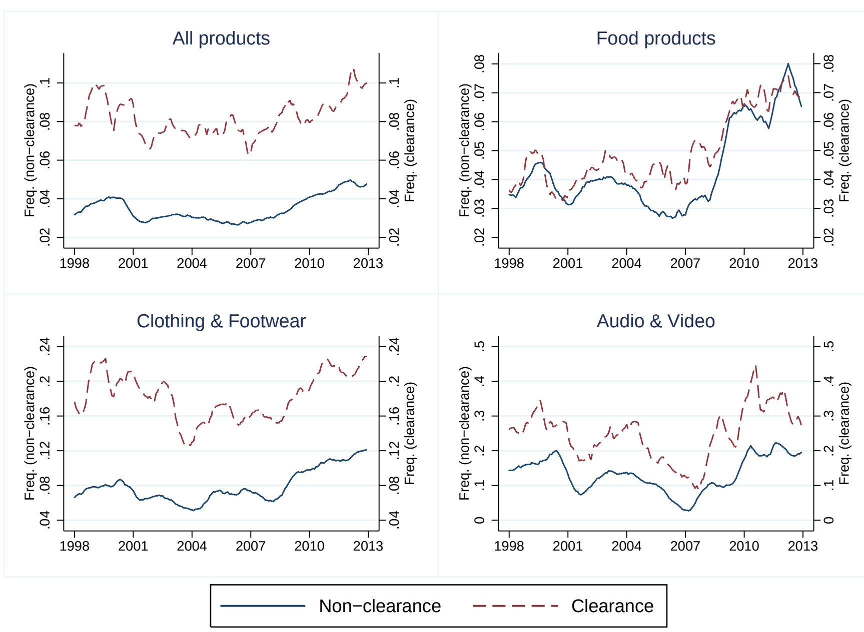

Notes: Clearance sales are defined as sales flag occuring within the last two months of a quote line. CPI quote weights are used for aggregation. U.K. Office for National Statistics data.

Figure 10: Frequency of sales by retailer size for food in the U.K. CPI data

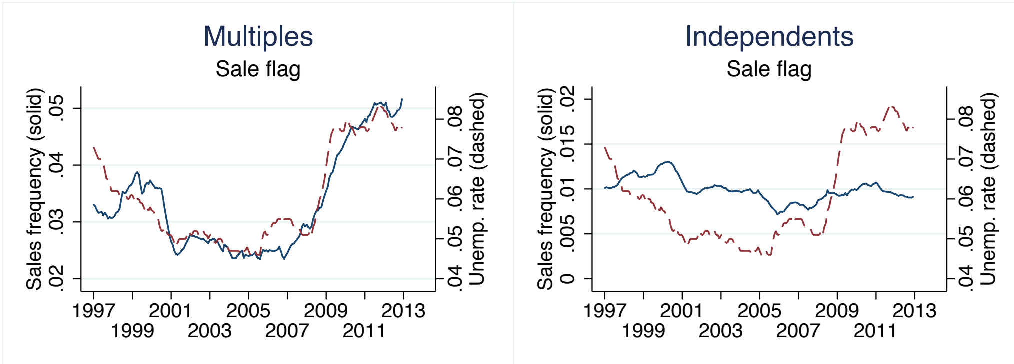

Notes: Monthly proportion of products with a sales flag, for "Multiples" (retailers with at least 10 outlets) and "Independents" (retailers with less than 10 outlets). Aggregation is done using CPI quote weights. Sample period is 1996:02–2013:09, U.K. Office of National Statistics data.

Figure 11: Heterogeneity in the use of sales across retailers

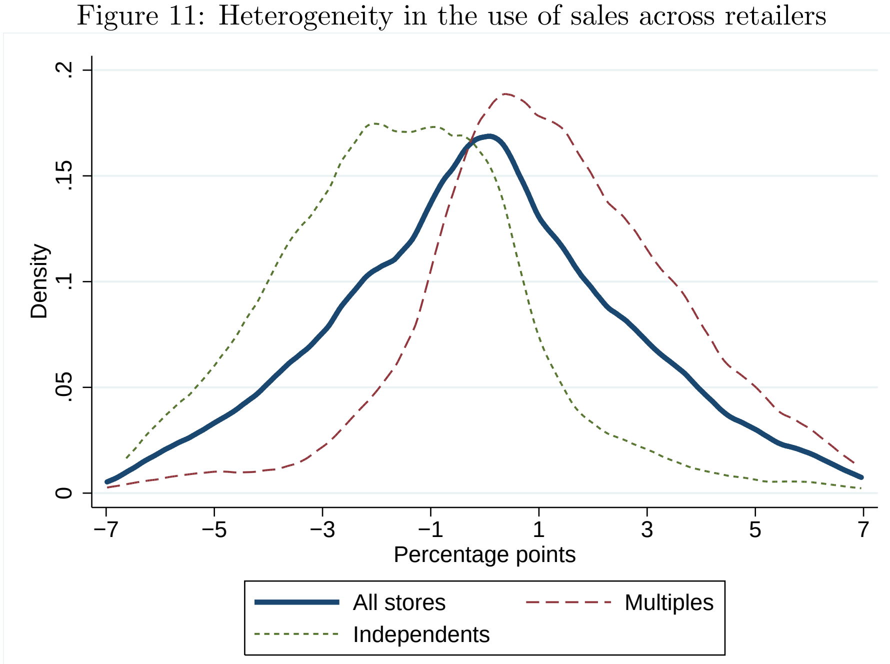

Notes: Kernel densities of the store fixed effects for all stores, "Multiples" (retailers with at least 10 outlets) and "Independents" (retailers with less than 10 outlets). The fixed effects are derived from a linear regression of the sale flag indicator on store, item and month dummies for the food category. The store fixed effects are centered at zero by construction. U.K. Office of National Statistics data.

Figure 12: Decision for high and low price levels

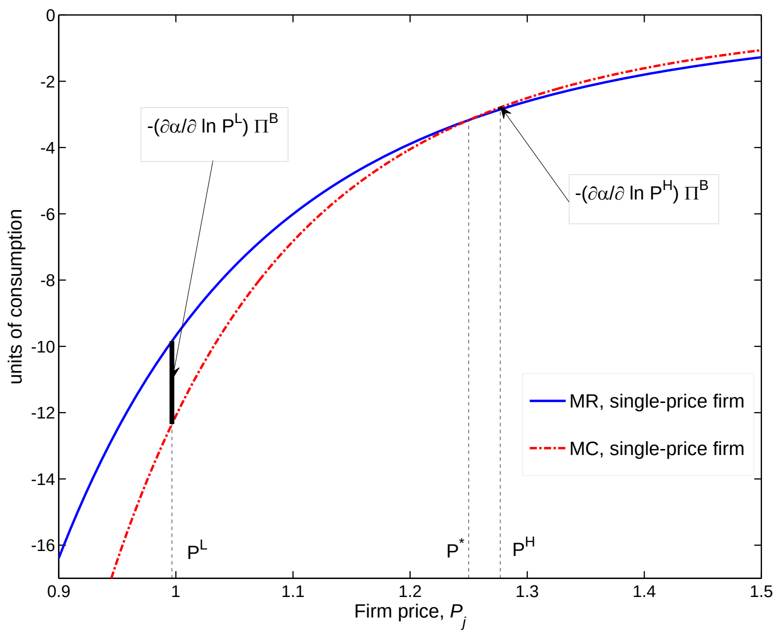

Figure 13: Responses to a negative 1% impulse to money growth

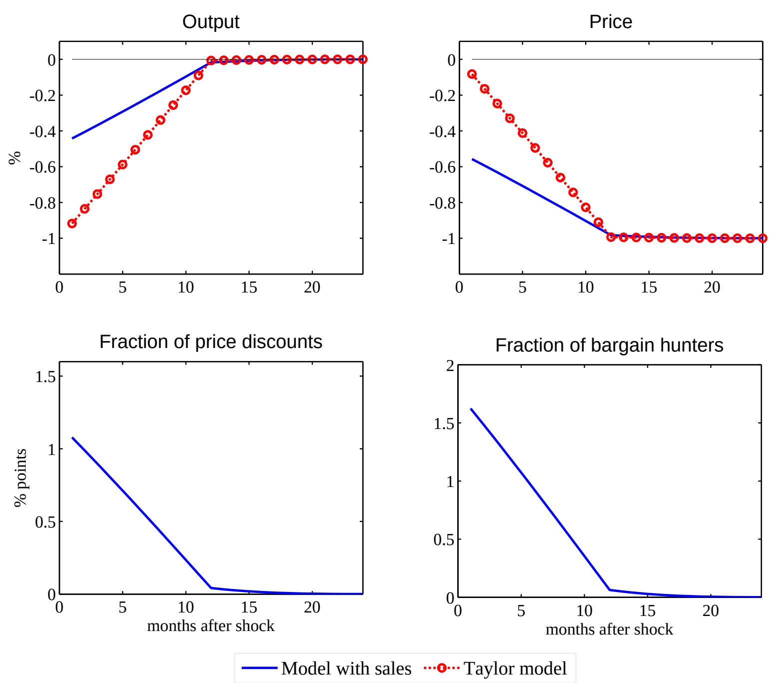

Figure 14: Decomposition of aggregate price response to a negative 1% impulse to money growth

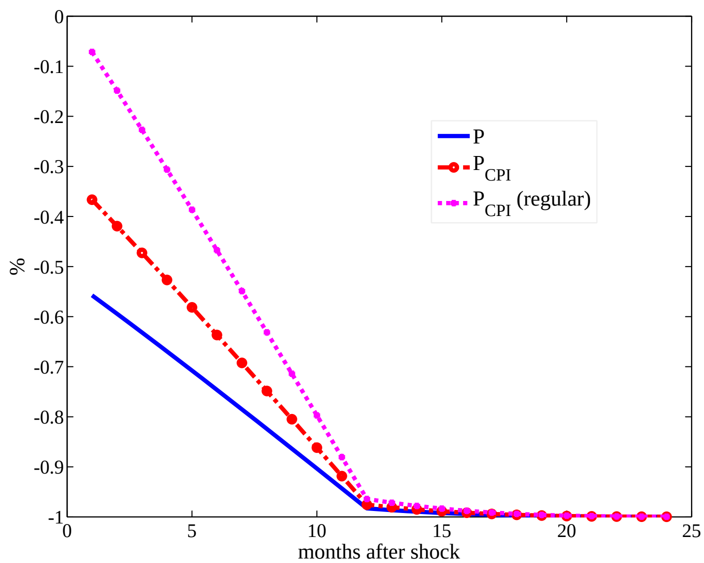

Figure 15: Contributions of sales to the effective price level in the United Kingdom.

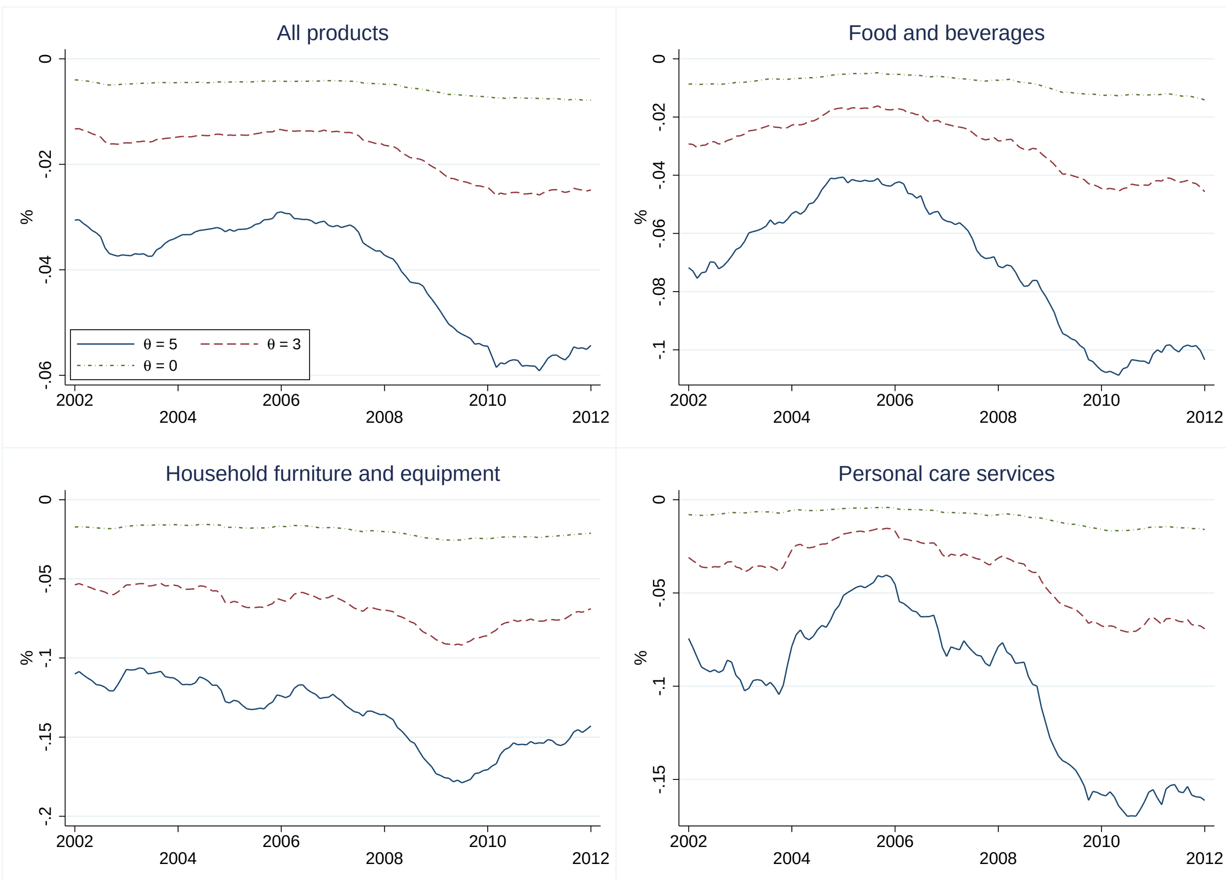

Notes: The time series are the 12-month centered moving averages of the aggregated discount term defined in (14) in Section (8), for various values of the elasticity of substitution θ. All series represent percentage deviations of the effective price index from its regular-price counterpart. Discount terms are based on V-shaped discounts and aggregated across products using official quote and stratum weights. " $\theta = 0$ " corresponds to the official CPI mesure in that it uses equal weights for regular-priced and discounted price quotes. The sample period is 1996:02–2012:12, U.K. Office of National Statistics data.

Table 1: Summary statistics for posted and regular price changes

	Frequency of price changes	Frequency of price increases	Frequency of price decreases	Abs. size of price changes
	All prices			
	0.158	0.098	0.060	0.119
	Regular prices			
Sales flag	0.132	0.086	0.046	0.085
V-shaped	0.109	0.076	0.033	0.086
Reference price	0.075	0.054	0.021	0.094

Notes: CPI data are from the U.K. Office for National Statistics, sample period is from February 1996 to September 2013. 'Sales Flag', 'V-shaped' and 'Ref.' refer to the sales filters described in the main text.

Table 2: Summary statistics for temporary sales

	Sales flag	V-shaped	Reference price
	All sales		
Frequency	2.9%	4.6%	5.6%
Mean size	-22.0%	-13.9%	-12.4%
Median size	-20.0%	-10.0%	-8.0%
	Sales of at least 10%		
Frequency	2.3%	2.3%	2.7%
Mean size	-25.6%	-23.6%	-23.1%
Median size	-21.9%	-20.1%	-20.0%

Notes: CPI data are from the U.K. Office for National Statistics, sample period is from February 1996 to September 2013. 'Sales Flag', 'V-shaped' and 'Ref.' refer to the sales filters described in the main text.

Table 3: Regressions at the aggregate level

	UK					
	Flag	*Flag*	*Flag*	*V-shaped*	*Ref.*	
u_t	0.393	0.143	0.358	0.464	0.328	
	(0.040)	(0.030)	(0.041)	(0.033)	(0.028)	
$sales_{t-1}$		0.651				
		(0.060)				
Linear time trend	N	N	Y	Y	Y	
R^2	0.63	0.79	0.65	0.80	0.72	
	US - BLS			**US - Vavra**		
	Mean	*Mean*	*Mean*	*Mean*	*Mean*	*Median*
u_t	0.177	0.153	0.233	0.338	0.320	0.440
	(0.021)	(0.024)	(0.030)	(0.043)	(0.046)	(0.060)
$sales_{t-1}$		0.148				
		(0.092)				
Linear time trend	N	N	Y	N	Y	Y
R^2	0.51	0.52	0.54	0.54	0.55	0.32

Notes: Linear regressions of the frequency of sales on the unemployment rate (u_t). 'Flag', 'V-shaped' and 'Ref.' refer to the sales filters described in the main text. 'Mean' and 'median' indicate mean and median frequencies respectively. All regression include calendar month dummies. Newey-West standard errors allowing for a maximum autocorrelation of 4 lags are indicated in parentheses. All coefficients are statistically significant at the 1% level except for the lag of the dependent variable in the 'US - BLS' panel which is not significant. The sample periods are 1996-2012 for the U.K., 2000-2011 for "U.S. BLS" and 1988-2011 for "U.S. - Vavra".

Table 4: Sales and regular price changes

	UK		**US - BLS**	**US - Vavra**	
	Flag	*Flag*	*Mean*	*Mean*	*Median*
u_t	0.362	0.362	0.208^{***}	0.357^{***}	0.483^{***}
	(0.025)	(0.041)	(0.033)	(0.057)	(0.058)
$freg_t$	-0.008			-0.11^{***}	-0.458^{***}
	(0.005)			(0.042)	(0.059)
$freg\_pos_t$		-0.011	-0.062^{*}		
		(0.007)	(0.032)		
$freg\_neg_t$		-0.002	-0.029		
		(0.008)	(0.037)		
Linear time trend	Y	Y	Y	Y	Y
R^2	0.67	0.67	0.63	0.57	0.46

Notes: Linear regressions of the frequency of the sales flag on the unemployment rate (u_t) as well as the frequency of regular price changes ($freg_t$), increases ($freg\_pos_t$) and decreases ($freg\_neg_t$). 'Mean' and 'median' indicate mean and median frequencies respectively. All regression include calendar month dummies. Newey-West standard errors allowing for a maximum autocorrelation of 4 lags are indicated in parentheses. *, ** and *** indicate coefficients that are statistically significant at the 10%, 5% and 1% levels respectively. The sample periods are 1996-2012 for the U.K., 2000-2009 for "U.S. BLS" and 1988-2011 for "U.S. - Vavra".

Table 5: Panel regression results at the product level

	Sales flag				V-shaped	Ref. price
	(1)	(2)	(3)		(4)	(5)
u_t	0.402	0.283	0.337		0.444	0.262
$s_{i,t-1}$			0.192			
Month dummies	Y	Y	Y		Y	Y
Time trend	N	Y	N		N	N
Fixed effects	Y	Y	Y		Y	Y
F-stat	3.77	3.64	42.39		11.48	16.26

	Alt. macro indicators (normalized)			
	(6)	(7)	(8)	(9)
u_t (normalized)	0.0050			
Retail sales vol.		-0.0054		
Consumer confidence			-0.0019	
Fin. situation next year				-0.0031
Month dummies	Y	Y	Y	Y
Time trend	N	N	N	N
Fixed effects	Y	Y	Y	Y
F-stat	3.77	4.10	3.90	4.04

Notes: Panel regressions at the item level: $s_{it} = \alpha_i + \beta u_t + X_t^{\prime}\Phi + e_{it}$, where s_{it} is a 0/1 sale indicator and u_t is an aggretate business cycle indicator, usually the unemployment rate. For the last panel of the table, the macroeconomic indicators are normalized by their standard deviation to facilitate comparisons. For all regressions, standard errors are clustered at the product category level. All coefficients are statistically significant at the 1% level.

Table 6: Parameterization (benchmark model)

A. Calibrated Parameters

θ	elast. subst. goods	5
ε_{Ξ}	search cost elasticity	1
κ	fixed cost of posting discounts *	0.08
z_{max}	max fixed search cost *	0.31

* - in units of time

B. Targets (steady state)

	Data	Model
Markup	0.25	0.25
Price discount	0.78	0.78
Fraction of discounts	0.05	0.05
Δ Revenue share of sales from $\Delta\gamma$=1ppt	2.7	2.7

C. Assigned Parameters

	period	1
β	discount factor	$0.96^{1/12}$
σ	risk aversion	1
f	search efficiency	2
$1/N$	frequency of price changes	1/12
ρ_{μ}	serr corr of money shock	0
σ_{μ}	stdev of money shock impulse	0.23