

On the Importance of Sales for Aggregate Price Flexibility*

Oleksiy Kryvtsov[†] and Nicolas Vincent[‡]

September 2014

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 evidence from the U.K. and U.S. CPI micro data that the frequency of sales is strongly countercyclical. Second, we build a general equilibrium model in which sales arise endogenously. In response to a monetary contraction, firms facing rigid regular prices post more sales, and households search more intensively. 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.

Keywords: Price dynamics, Sales, Inflation, Monetary neutrality.

JEL classification: E31, E32, E52, L11, L25, L81, 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, 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 NBER Price Dynamics 2014 workshop, the Duke Macroeconomics seminar series, the Shanghai Macroeconomic Policies and Business Cycles conference, the Bank of Canada 2014 Fellowship Learning Exchange, the 2014 Meetings of the European Economics Association and the 2014 Dynare Conference 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 and Tony Chernis 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 implications of retail pricing for aggregate price flexibility. These studies find that although temporary price discounts imply more frequent price changes by retailers,¹ they do not, on the whole, play a significant role for inflation dynamics. 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.

A few recent studies have examined the potential role of sales in macroeconomic models. Kehoe and Midrigan (2008), 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 responds 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 the model predictions are borne out in the data. This is partly due to the fact that the empirical evidence on the cyclical properties of sales is very limited. A recent exception is Coibion, Gorodnichenko and Hong (2014): using a scanner data set from U.S. grocery stores, they document that sales are mostly acyclical and that most adjustment is done through consumers switching from high- to low-end retailers.

In contrast to these studies, we find that retailers’ adjustment along the sales margin is

¹For 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.

²Eichenbaum, 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.

economically relevant for the response of aggregate price and output to changes in macroeconomic conditions.

Our contributions are both empirical and theoretical. First, we provide empirical evidence on variations in the sales margin over time. To this end, we use the publicly available micro data underlying the consumer price index (CPI) composed by the U.K. Office for National Statistics (ONS). The data contain monthly prices collected from local retail outlets for a wide range of consumer goods and services over the period from 1996 to 2012. We find that the frequency of sales in the U.K. data is strongly countercyclical: a 1 percentage point (ppt) rise in the unemployment rate is associated with a roughly 0.4 ppt increase in the fraction of products on sale. For example, during the Great Recession the fraction of sales more than doubled from 1.8% to 3.8% of observations. 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 for most products; it is not a product of the exit of low-sale-frequency items; and it survives the use of multiple controls and alternative macroeconomic indicators.

Our finding is not specific to the United Kingdom. An aggregate time series on the incidence of sales obtained from the Bureau of Labor Statistics shows that temporary discounts are also strongly countercyclical in the United States. This evidence is corroborated by series derived from Vavra (2014) based on the U.S. CPI micro data.³

Unlike the fraction of sales, the average size and duration of sales in the United Kingdom are acyclical and much less volatile. Nonetheless, the observed fluctuations in the fraction of sales may well be an important source of aggregate price fluctuations due to two factors. First, sale-related price drops are large, on average between 20% and 25% in the U.K. data. Second, each additional sale generates almost 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 two factors.

To this end, in the second part of the paper, we develop a general-equilibrium business

³In a recent study, Sudo et al. (2013) 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 is dominated by strong trends in all three series during the “lost decades,” making it difficult to determine whether the behavior of sales is due to their countercyclical use by retailers or to structural changes in the Japanese retail industry that are not related to business cycles.

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 and size 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. When matched to key moments of sale-price behavior, the model predicts that in response to, say, 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 model without time-varying consumer search, or the standard sticky-price model with a single price per variety. Accordingly, the size of the real effect in response to monetary shocks is twice as small and 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 Taylor-type price adjustment constraints, 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 on searching for low prices, leading to a larger number of bargain hunters. The resulting reallocation of the consumption basket toward lower-priced products amplifies the fall in both store-average and aggregate price levels. We also show that the model's central feature — countercyclical shopping intensity — is very much in line with evidence from time-use surveys, household panel scanner data sets, and Google Trends search intensity for sales-related terms.

The paper is organized as follows. Section 2 presents empirical evidence on the frequency and size of sales from the end of the 1990s to 2012 using product-level price quotes from the U.K. CPI data set; the findings for the United Kingdom are corroborated by evidence for the

United States. In Sections 3 to 5 we describe the model with sales and analyze its response to shocks under plausible parameterizations. Section 6 provides evidence on shopping activity over the business cycle as well as the cyclical nature of other forms of price discrimination. Section 7 concludes.

2 Empirical evidence on sales

2.1 Data

To construct the consumer price index (CPI), the Office for National Statistics (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.⁴ 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. The survey excludes the housing portion of consumer prices, such as mortgage interest payments, house depreciation, insurance and other house purchase fees.⁵ Also, expenditures for purposes other than final consumption are excluded, e.g., those for capital and financial transactions, direct taxes, and cash gifts.⁶

The goods and services in the CPI are classified into 71 classes, according to the international (European) classification of household expenditure, Classification of Individual Consumption by Purpose (COICOP). A CPI class represents a basic group category, such as “Meat,” “Liquid Fuels” or “New Cars.” Each item in a given class is assigned an item weight that reflects its relative importance in households’ consumption expenditures.⁷ Changes in expenditure weights over time reflect changes in the expenditure composition of households’ consumption baskets.

Prices are collected across 13 geographical regions (e.g., London, Wales, East Midlands). There are four levels of sampling for local price collection: locations, outlets within location, items within section and product varieties. For each geographical region, locations and outlets

⁴Detailed descriptions of the data underlying the CPI, the statistical methodology used, collection and validation of prices, and calculation of weights, can be found in the “Consumer Price Indices Technical Manual” (2012). The price quote data are available via the ONS website: <http://www.ons.gov.uk/ons/datasets-and-tables/index.html>.

⁵These prices are used to construct the retail prices index (RPI).

⁶CPI inflation was used by the government for its inflation target starting in December 2003. Before then, the index was published in the United Kingdom as the Harmonised Index of Consumer Prices (HICP).

⁷Weights are calculated based on the Household Final Monetary Consumption Expenditure (HFMCCE) and ONS Living Cost and Food Survey (LCF).

are based on a probability-proportional-to-size systematic sampling with a size measure based on the number of employees in the retail sector (locations) and the net retail floor space (outlets). The data set contains around 150 locations with an average of more than 90 outlets per location.

Representative items are selected based on a number of factors, including expenditure size and product diversity, variability of price movements, and availability for purchase throughout the year (except for certain goods that are seasonal). There are currently over 510 items in the basket. Examples of representative items include: onions, men’s suit, single bed. Finally, for each item-outlet-location, individual products and varieties are chosen by price collectors based on their shelf size and regular stock replenishment.

Most prices are collected monthly, except for some services in household and leisure groups, and seasonal items. For the purpose of this paper, the sample period includes 212 months, from February 1996 till September 2013. The total raw number of observations is over 24.4 million, or about 115,000 per month, though we make some adjustments that are described later in order to make the database amenable to analysis.

In addition to posted prices, the data set also contains information about some characteristics of goods during price collection. Prices for goods that are on special offer (available to all consumers) or on temporary sale comprise 6.4% of observations in the data set (4.5% weighted). It is important to note that these “sales” 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). Forced substitutions happen 8.0% (5.5% weighted) of the time, with about a quarter (three-quarters) of them corresponding to substitutions for items that are non-comparable (comparable) to previously priced items. An item can be temporarily out of stock (2.2% or 1.5% weighted) or permanently missing (0.5% or 0.3% weighted). Finally, a small subset of goods has distinct seasonal patterns and is treated separately: they include some items of clothing, gardening products, holiday products and air fares. For those seasonal goods for which prices are not available, such as clothing, gardening and food, prices are imputed based on prices observed at the end of the previous season or based on prices observed for in-season goods in the same item category; in addition, weights are adjusted in accordance with the availability of such goods throughout the year.

We clean the raw database to make it suitable for analysis. First, we create identification numbers that uniquely define item/region/store combinations. Second, if there is more than one observation for the exact same item in a given month, we ignore them.⁸ We also delete observations that are not coded as valid by the ONS. Third, we need to deal with product substitutions. To do this, we split the price time series of a given item every time we encounter a substitution flag. The resulting benchmark data set contains a total of 20.7 million observations across about 2.3 million unique items. It should be noted that for our empirical analysis, we will mostly focus on items that have at least ten price quotes (17.1 million observations).

Value-added taxes (VAT) are included in the price quotes. The implication is that any change in the VAT rate will automatically lead to a price change. VAT changes, however, were very few over our sample period. The only exceptions are three major changes to the standard VAT rate in the later part of the sample. First, in response to the Great Recession, the Conservative government announced a widespread reduction in the VAT rate from 17.5% to 15% effective 1 December 2008. Then, the rate was brought back to 17.5% starting 1 January 2010. Finally, the standard VAT rate was raised from 17.5% to 20% on 4 January 2011. We control for these events in our analysis.⁹

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 data set.

First, we present results using the “sales flag” from the ONS. 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.” Once again, it is important to point out that this indicator does not reflect coupons, volume discounts or promotions linked to loyalty cards; it therefore arguably represents a lower bound on the incidence of sales. Despite the

⁸Discussions with the ONS indicate that most of these duplicates are due to multiple observations across different locations. Unfortunately, because the ONS did not include the location variable as part of its public database, it is not possible for us to link observations across time in the presence of duplicates, leaving us with no choice but to delete them. Using only items that have never registered duplicates over their lifetime has no noticeable impact on our results.

⁹Moreover, VAT changes are more akin to regular price changes, while our focus is on temporary sales. For this reason, they have little incidence on our results.

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 no recent change in the posted price, or even, in some very rare instances, 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 actual sales.¹⁰ That being said, the results using the raw sales indicator are very similar.

Second, we apply a V-shaped sales filter, similar to the one used by Nakamura and Steinson (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.¹¹ Under this definition, the price increase at the end of the sale need not be as large as the price drop at the beginning of the sale.

We also consider 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. For the sales flag and both V-shaped filters, the unobserved regular price during a sale is assumed to be equal to the last observed regular price.

Finally, we compute a reference price similar in spirit to Eichenbaum, Jaimovich and Rebelo (2011) using a seven-month window. More precisely, 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/lead in the adjustment of reference prices, we then apply a procedure similar to Kehoe and Midrigan (2008) 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.

¹⁰To 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. We should, however, stress that our results are only strengthened if we use the raw sales flag series from the ONS.

¹¹If 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 two consecutive drops in the regular price has no impact on our results.

2.3 Preliminary statistics

In Table 1 we report some basic statistics on price dynamics in our data set. Unless otherwise stated, all moments are weighted using the official CPI category weights. The fraction of price changes is 17.2% over the sample period and the average size of a price change is 11.1% in absolute terms. Price increases are more likely than price decreases (10.7% vs 6.5% 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 14.6%, 12.1% and 7.3%, 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.¹²

Turning our attention to the main object of study, Table 2 reports a number of basic statistics for sales, using CPI weights for all calculations. The first row shows that the frequency of sales varies significantly depending on the definition used, from 2.6% for the ONS sales flag to 6.2% for the sales based on the reference price filter.¹³ Differences are also visible for both the average and median sale sizes: they are much higher for the sales flag (between 20% and 23%) than the three other filters (between 6% and 12%).

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 data set (we drop the alternative V-shaped filter, since it shows no significant difference).¹⁴ We set the bound of the histogram at 60%, since larger sales are rare. Reassuringly, one can immediately notice 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 some 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

¹²See, for example, Klenow and Kryvtsov (2008) or Nakamura and Steinsson (2008) for the United States and Alvarez et al. (2006) for Europe.

¹³Though 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.

¹⁴The size of a sale corresponds to the difference between the sale and the regular price, and is therefore not limited to the onset of the sale (initial price drop).

according to the ONS indicator. This seems to indicate 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.

Finally, it is important to recall that because the data set does not take into account some popular price-promotion strategies (coupons, loyalty card discounts, “buy-two-get-one-free” deals, etc.), our sale frequencies are likely to be downwardly-biased estimates of the true prevalence of temporary discounts.

2.4 The time-series behavior of sales

We next look at the behavior of the frequency of sales during the past 15 years in the United Kingdom. For each category and month we compute the proportion of items on sale, 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 very strong seasonal patterns of sales, we also report on the right 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 3.4% at the beginning of the sample in 1997, then declines to a trough of about 1.8% in 2006 before rising back to more than 3.6% by 2011. Also, it is strikingly countercyclical: the fraction of sales moves very closely with the unemployment rate, rising as the economy is slowing down. While we formally control for this potential issue later on, neither series seems to exhibit any time trend that may bias our conclusions.

Next, we look into the robustness of our main result. First, we show that the countercyclicity of sales is not an artifact of using the sales flag from the ONS. Figure 3 compares the evolution of the fraction of items on sales based on the sales flag as well as the three other filters we described earlier: both the benchmark and more-restrictive V-shaped sales filters, as well as the one based on the 7-month reference price. For all series we focus on sales of at least 10%. While there are some differences at the beginning of the sample (a potential artifact of left-censoring issues with sales filters), the patterns tell a very similar story: there are large swings in the incidence of sales over the sample period, and the prevalence of sales

tends to be higher when the economy is weaker.

Second, the volatility and cyclicity of sales are not driven by just a few large categories. The top-left plot of Figure 4 shows the evolution of the fraction of price quotes with a sales flag aggregated using equal weights across categories, as opposed to the weighted statistic reported so far. Another concern could be that some changes in the collection of data or the definition of sales coincided with the onset of the recession and therefore are leading to spuriously identify a high degree of cyclicity. While the fact that the results hold across multiple sales filters should allow us to dismiss this concern, we further note that the cyclicity of sales is not a feature that is common to all categories, and therefore it is unlikely to be generated mechanically. Figure 4 shows a significant rise in the incidence of sales (here defined as the fraction of items with a sales flag) for categories such as “Clothing and footwear,” “Furnishings and other household equipment,” and “Personal care.” Yet, for other categories such as “Books” or “Housing rents and maintenance,” there is no significant rise in sales around 2007–09.

It should also be noted that the rise in the sale frequency is broad-based: the number of products that experienced a rise in their sale frequency between the periods 2005–06 and 2009–10 is almost three times as high as those that saw a drop.

Third, the cyclicity of the fraction of products on sale could either be driven by fluctuations in the incidence of new sales or by changes in the average length of a sale over time; in other words, it may be that retailers put items on sale more often, or that they keep them on sale for longer. We find no evidence of the former in the data: the average duration of a sales spell remained very stable around 1.6 months over our sample period, with no discernible cyclicity or trend.¹⁵

Finally, Figure 5 shows the evolution of the average size of sales over our sample period, under our three main definitions of sales. Under all definitions, there seems to be a rise in the (absolute) size of sales over the sample period. Still, while the size of discounts is not perfectly constant, the fluctuations are significantly more limited in relative terms than what we saw earlier for the frequency of sales. For example, using the V-shaped filter, we find that the average size of sales goes from about 21% in 1997 to 26% by the end of 2012. More importantly, there is no noticeable cyclicity for this margin of adjustment, unlike for the

¹⁵That being said, the fact that we only have monthly frequency data makes it difficult to reliably identify changes in the average length of sales.

sales frequency, which basically doubled during the Great Recession.

2.5 Regression analysis

Our findings from the previous section show that firms do, at least under some circumstances, alter the prevalence of temporary sales in reaction to macroeconomic fluctuations. While our results appear robust, the graphical analysis used limits our ability to control for various factors that may bias our conclusions. For this reason, we turn our attention to a panel regression approach. The basic specification is standard and given by

$$s_{it} = \alpha_i + \beta u_t + X'_{it}\Phi + e_{it}, \quad (1)$$

where s_{it} is the dependent variable measuring the incidence of sales, u_t is the unemployment rate and X_{it} is a matrix of controls such as calendar month dummies or a time trend. First, we run the regressions at the category (COICOP) level. In this case, s_{it} is the fraction of items in category i whose posted price is at least 10% below the regular price at time t . All regressions include dummies for the months in which VAT rate changes occurred, as described in Section 2.1. Table 3 summarizes our results. Robustness checks are shown only for the sales flag series for conciseness, but findings are similar for the other filters.

Our results in Table 3 indicate that a 5 ppt increase in the unemployment rate raises the likelihood of observing a sale by about 1.7 percentage points, which is in line with our earlier graphical findings. Adding a time trend or calendar month dummies does not much alter the economic or statistical significance of this relationship, though it does increase somewhat the R^2 of the regression. Even when we include the lagged value of the dependent variable, the coefficient on the unemployment rate remains positive and strongly significant, leading to a slight increase in the long-term effect of unemployment on sales at 0.403 ($= 0.257 / (1 - 0.362)$).¹⁶ Running instead a regression where observations are weighted by their CPI weights does not materially affect the findings (column (5) of Table 3). Finally, the last two columns show our benchmark specification using the two other filters. There again, the unemployment rate is strongly significant statistically, and the economic effects are of the same order of magnitude.

So far, our analysis has been carried out using the unemployment rate as a macroeconomic

¹⁶Interestingly, while we do not report the results, we note that the best fit tends to be obtained by using the unemployment rate lagged between 6 and 8 months (though the difference in fit is very much marginal), in contrast to Figure 2, where the sales frequency seems to be slightly *leading* the cycle.

indicator. While this seems to be a sensible indicator of the financial and economic situation of households, it is hardly the only one. In Table 4 we report results for our benchmark specification, but with alternative business cycle indicators on the right-hand side. In keeping with our focus on the consumer side, we 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 of a sample of households. We focus, in order, on (1) the aggregate consumer confidence index, (2) the personal financial situation over the next 12 months, (3) the general state of the economy next year, and (4) the intention of making major purchases. To facilitate comparisons, we divide all these variables, as well as the unemployment rate, by their respective time-series standard deviations over the sample. The results in Table 4 show very similar responses across three indicators out of four: a one-standard-deviation change in the explanatory variable leads to a statistically significant change of between 0.4 and 0.5 percentage points in the frequency of sales.¹⁷

The regressions at a category level suggest that countercyclicality of sales applies at a disaggregate level and does not stem from mere aggregation across categories. To show that this point applies at any level of disaggregation, we also conduct panel regressions at the individual product level. This will help us to alleviate any concern that our results are driven by some composition bias, for example.¹⁸

Under this specification, s_{it} in equation (1) is an indicator equal to 1 if the item i is on sale in month t , and 0 otherwise. All p -values shown are based on robust standard errors clustered at the category level. Results reported in Table 5 contain no surprises: the unemployment rate is a statistically significant predictor of the occurrence of a temporary sale; the size of the relationship is in line with the previous graphical and regression analyses; and the results hold for all three filters considered.¹⁹

The basic message of this section is clear: there is robust evidence that U.K. firms reacted

¹⁷As a point of reference, the unemployment rate during the 2007–09 recession rose by a magnitude of four standard deviations.

¹⁸For example, it could be that items less prone to temporary discounts were more likely to disappear as consumers increased their search intensity during the recession, and were being slowly substituted for comparable brands with higher sales frequencies as part of the CPI sampling methodology. While interesting by itself, this exit-and-entry story is different than one by which firms respond to macroeconomic shocks by altering their use of temporary sales.

¹⁹In non-reported regressions, we find that sales frequency is roughly similar across regions of the United Kingdom; sales are significantly more prevalent in stores belonging to chains than independent shops; and that sales are most frequent in January, followed by July. We also verified, using similar regressions, that the size of sales is not responsive to the macroeconomic cycle.

to macroeconomic conditions in the past 15 years by adjusting the frequency of temporary sales. In the next section, we build a model with sales to investigate the macroeconomic implications of this finding. Before doing that, we corroborate the U.K. facts with some evidence from the U.S. data.

2.6 Supporting evidence from the U.S. CPI data set

In the previous sections, we have carefully analyzed the publicly available U.K. CPI micro data and documented a very strong countercyclicality of sales. Ideally, we would like to confirm this characteristic from other data sets. This is unfortunately complicated by the fact that empirical evidence on the cyclical behavior of sales has received very little attention in the literature, and that access to CPI micro data in most countries is very limited.

Yet, for the United States we were able to obtain aggregate evidence derived from the U.S. Bureau of Labor Statistics (BLS) CPI micro data. First, the Bureau of Labor Statistics provided us with the monthly fraction of items with a sales flag since January 2001, aggregated using CPI expenditure weights.²⁰ The time series is represented in the left-hand plot of Figure 6, alongside the U.S. civilian unemployment rate over the same period.²¹ 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 0.86 and two clear spikes in the sales' fraction in the midst of the 2001 and 2008–2009 recessions.

Second, the right-hand plot depicts a series derived from Vavra (2014). For his analysis, Vavra filtered out both temporary sales and product substitutions to focus on regular prices; Vavra kindly provided us with the time series for the combined frequency of sales and substitutions. While the series is somewhat more volatile, possibly due to the behavior of substitutions, the results confirm our findings based on the data provided directly by the Bureau of Labor Statistics. They are extended by an additional 12 years of data, showing a clear rise in the fraction of sales during the 1990–91 recession, similar to those in the two subsequent downturns.

In summary, the evidence depicted in Figure 6 strongly supports the notion that the countercyclicality of sales is not specific to the United Kingdom.

²⁰We would like to warmly thank Brendan Williams at the BLS for producing these series for us.

²¹To smooth out the strong seasonality in sales, we present the 12-month centered moving average.

3 A general-equilibrium model with sales

To understand the variation of sales and to quantify their importance for aggregate fluctuations, we cannot rely on the general-equilibrium models by Guimaraes and Sheedy (2011) or Kehoe and Midrigan (2008), since they predict, using examples of monetary shocks, that the fraction of sales is very stable, which is at odds with our empirical findings. We therefore develop in this section a general-equilibrium model in which sales arise due to price discrimination by retailers and search for low prices by consumers.²²

A model economy is populated by a measure one of infinitely lived, ex-ante identical households, indexed by k , who are consuming a continuum of differentiated-good varieties, indexed by j . These varieties are sold by a continuum of retail industries in isolated “locations” indexed by i . There is one retail industry per location, and each industry consists of measure one of monopolistically competitive retail firms, each firm selling many highly substitutable brands of a single variety j . For example, a beer retailer sells Budweiser, Bud Light, Michelob Light, Coors Light, etc. For simplicity, we will assume brands of the same variety to be perfect substitutes. In all, a given industry sells all varieties, and a given variety j is sold in every location. This set-up differs from a conventional model with one industry of monopolistically competitive retailers, each producing a single brand of a given variety.

In our set-up, a retailer prefers to post different prices for different brands at any point in time. The existence of a non-degenerate price distribution in our model stems from the retailer’s desire to price discriminate in the face of heterogeneous household reservation prices at the point of purchase; without such a motive, retailers would prefer a single price for all brands of the same variety. In the model, the dispersion of reservation prices stems from a costly search for low prices. We assume that a household faces a random fixed cost of searching in a retail store selling variety j , and that this cost is in terms of forgone hours worked. The fixed cost is i.i.d. across time, stores and households, drawn from twice continuously differentiable distribution $G(z)$, with $0 \leq z \leq z_{\max}$. Depending on the realized cost of searching, households will either choose to actively search for a low price for variety j (and

²²There 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).

become “bargain hunters” for that variety) or not to search (“workers”).^{23,24}

In each period t , the economy experiences one of infinitely many events s_t . Denote by $s^t = (s_0; \dots; s_t)$ the history of events up to and including period t . The probability density, as of period 0, of any particular history s^t is $\pi(s^t)$, and the initial realization s_0 is given. The only shocks in this economy are aggregate shocks to the money supply. We assume that the supply of money follows a random-walk process of the form

$$\log M(s^t) = \log M(s^{t-1}) + \mu(s^t),$$

where $\mu(s^t)$ is money growth, a normally distributed i.i.d. random variable with mean 0 and standard deviation σ_μ .

3.1 Household’s problem

At the beginning of each period, state s_t is realized, and each household is randomly and uniformly assigned to a “designated” location, i^* . For each variety j , the household then observes the distribution of prices in each store $j(i)$, denoted by $F_{j(i)}(P, s^t)$, draws the fixed cost of searching z_j and decides whether to search, i.e., becomes either a bargain hunter or a worker.

Bargain hunters and workers differ in two respects: (1) they face different technology matching them to a store $j(i)$, and (2) for each store $j(i)$, they face different distributions of prices. Specifically, we assume that workers buy their consumption of variety j at a random location i and at a price randomly drawn from all prices in the store; i.e., they draw one price from distribution $F_{j(i)}(P, s^t)$. Conversely, while bargain hunters must purchase variety j at their designated location i^* , active shopping allows them to face, on average, lower prices. Let $F_{j(i)}^B(P, s^t)$ denote the price distribution faced by bargain hunters in store $j(i)$. It is assumed that $F_{j(i)}(P, s^t)$ first-order stochastically dominates $F_{j(i)}^B(P, s^t)$ for each store; i.e., for any price P a worker is more likely to draw a price above P than a bargain hunter. Without

²³Consumers therefore are price discriminated by their search effort (third-degree price discrimination). Another common type of price discrimination is by price elasticity of consumers (second-degree price discrimination), an approach undertaken in Guimaraes and Sheedy (2011).

²⁴Alternatively, one could think of an environment populated by families, each composed of an infinite number of family members. In this set-up, families are randomly allocated across locations and each family member is responsible for shopping for a specific variety. In a given period, a fraction of family members with a low cost of searching will become bargain hunters, while the others will be workers. Both environments give rise to exactly the same equilibrium conditions.

any loss of generality, suppose $F_{j(i)}^B(P, s^t)$ can be obtained from $F_{j(i)}(P, s^t)$ by applying a functional F :

$$F_{j(i)}^B(P, s^t) = F [F_{j(i)}(P, s^t)] , \quad \forall j(i), (P, s^t). \quad (2)$$

Condition (2) can be interpreted as a characterization of the advantage that bargain hunters have over workers when searching for low prices in a retail store with price distribution $F_{j(i)}(P, s^t)$: they face a lower expected price than the unconditional average price, which is the price the workers face.²⁵ We show below that the differences in price distributions faced by bargain hunters and workers, coupled with the randomization of workers across locations, lead to price discrimination by retailers.

In our set-up, a household's utility from consuming a given variety does not depend on the location where it was purchased, and so we can skip the location index hereafter. Since the fixed cost is i.i.d., the household's search decision can be characterized by a cut-off fixed cost, $z_j^*(s^t)$, for which it is indifferent between searching or not for variety j . If the realized fixed cost is low, $z \leq z_j^*(s^t)$, the household will be a bargain hunter and search more intensively, drawing prices from $F_j^B(P, s^t)$. Otherwise, it will draw prices from $F_j(P, s^t)$.²⁶ Given the price drawn, $P_{k,j}(s^t)$, household k chooses how much to consume, $c_{k,j}(s^t)$. As we show below, conditional on the same price, consumption demand does not depend on shopper type (bargain hunter or worker). Therefore, the incentive of retailers to lower some prices (i.e., put a fraction of items on sale) stems solely from their willingness to increase their market share by attracting more bargain hunters.

In addition to searching and consuming, households supply working hours in a competitive labor market, and trade money and state-contingent securities. Let $\mathbf{B}_k(s^{t+1})$ denote a vector of state-contingent bonds paying household k one dollar in state s^{t+1} , and let $\mathbf{Q}(s^{t+1}|s^t)$ denote the vector of corresponding prices in period t and state s^t . We can write each price as $Q(s^{t+1}|s^t) = \frac{Q(s^{t+1})}{Q(s^t)}$, where $Q(s^t)$ is the date 0 price of a bond that pays one dollar if state history s^t is realized.

Household k chooses sequences of consumption varieties $c_{k,j}(s^t)$, total consumption $c_k(s^t)$,

²⁵One can assume explicit searching technology and derive the corresponding F -mapping. For example, we study the case when bargain hunters have two random draws per store, choosing the lowest price from two draws, as opposed to a single draw for workers. A simple version of this technology is studied in Section 4.

²⁶We need to assume that, as with workers, bargain hunters also face random prices; otherwise, each store would post not more than measure zero of its prices on sale, and all bargain hunters coming to the store would buy only at those prices.

leisure $1 - l_k(s^t)$, cut-off search cost $z_j^*(s^t)$, cash holdings, $M_k(s^t)$, and bond holdings $\mathbf{B}_k(s^{t+1})$ to maximize lifetime expected utility:

$$\sum_t \int \beta^t \pi(s^t) \left[\frac{c_k(s^t)^{1-\sigma}}{1-\sigma} - \psi(1 - l_k(s^t)) \right] ds^t.$$

Total consumption is derived from consumption of good varieties according to a constant elasticity of substitution (CES) aggregator

$$c_k(s^t) = \left(\int_j \{E_j [c_{k,j}(s^t) \mid z, P, s^t]\}^{\frac{\theta-1}{\theta}} dj \right)^{\frac{\theta}{\theta-1}}, \quad (3)$$

where θ is the elasticity of substitution across varieties, and $E_j[\cdot \mid z, P, s^t]$ denotes the expected value conditional on distributions of cost z and price P for variety j in date t and state s^t .

Utility maximization is subject to the budget constraint

$$M_k(s^t) + \int_{s^{t+1}} \mathbf{Q}(s^{t+1} \mid s^t) \mathbf{B}_k(s^{t+1}) \leq M_k(s^{t-1}) - \int_j E_j [P c_{k,j}(s^{t-1}) \mid z, P, s^{t-1}] dj \\ + W(s^t) \left[l_k(s^t) - \int_j \left(\int_0^{z_j^*(s^t)} z dG(z) \right) dj \right] + B_k(s^t) + \Pi(s^t) + T(s^t) \quad , \quad (4)$$

and a cash-in-advance constraint

$$P_k(s^t) c_k(s^t) \leq M_k(s^t).$$

In the budget constraint, $W(s^t)$ is the nominal wage, $\Pi(s^t)$ are dividends paid, and $T(s^t)$ are lump-sum government transfers; finally, $\int_0^{z_j^*(s^t)} z dG(z)$ denotes the expected cost of searching for consumer k in store j , expressed in terms of forgone hours worked. The household combines unspent money holdings from the previous period with labor income, return on bonds, dividends, and transfers, and splits them into current-period cash and bond holdings to be carried over to the next period.

The first-order conditions for consumption varieties, conditional on the realization of fixed search cost z and price P , yield standard CES demand schedules:

$$c_{k,j}(s^t) = \left(\frac{P}{P_k(s^t)} \right)^{-\theta} c_k(s^t),$$

where for bargain hunters (who drew $z \leq z_j^*(s^t)$) price P is drawn from distribution $F_j^B(P, s^t)$, and for workers ($z > z_j^*(s^t)$) price P is drawn from store distribution $F_j(P, s^t)$. Hence, conditional on the price drawn, consumption demand is the same regardless of shopper type; the difference between bargain hunters and workers stems from the different price distributions they face, with bargain hunters drawing, on average, lower prices. The price of total consumption for household k , $P_k(s^t)$, is defined as

$$\begin{aligned} P_k(s^t) &= \int_j E_j \left[P \frac{c_{k,j}(s^t)}{c_k(s^t)} \mid z, P, s^t \right] dj \\ &= \left(\int_j E_j \left[P^{1-\theta} \mid z, P, s^t \right] dj \right)^{\frac{1}{1-\theta}}. \end{aligned} \quad (5)$$

Note that, due to i.i.d. search costs, households will end up with identical total consumption $c_k(s^t)$ and price $P_k(s^t)$, so we will drop the household index k , denoting them $c(s^t)$ and $P(s^t)$, respectively.

To derive the optimality conditions for the search decision, denote by $\alpha_j(s^t)$ the probability of drawing a search cost smaller than the threshold level $z_j^*(s^t)$, i.e., $\alpha_j(s^t) = G(z_j^*(s^t))$. In other words, $\alpha_j(s^t)$ represents the probability that a household will be a bargain hunter in store j . In the problem above we can rewrite the consumption aggregator (3) and the budget constraint (4) as

$$c(s^t) = \left\{ \int_j \left(\alpha_j(s^t) \left(\int c_{k,j}(s^t) dF_j^B(P, s^t) \right)^{1-\frac{1}{\theta}} + (1 - \alpha_j(s^t)) \left(\int c_{k,j}(s^t) dF_j(P, s^t) \right)^{1-\frac{1}{\theta}} \right) dj \right\}^{\frac{\theta}{\theta-1}}$$

and

$$\begin{aligned} M_k(s^t) + \int_{s^{t+1}} \mathbf{Q}(s^{t+1} | s^t) \mathbf{B}_k(s^{t+1}) &\leq M_k(s^{t-1}) \\ &- \int_j \left[\alpha_j(s^{t-1}) \left(\int P c_{k,j}(s^{t-1}) dF_j^B(P, s^{t-1}) \right) - (1 - \alpha_j(s^{t-1})) \left(\int P c_{k,j}(s^{t-1}) dF_j(P, s^{t-1}) \right) \right] dj \\ &+ W(s^t) \left[l_k(s^t) - \int_j \Xi(\alpha_j(s^t)) dj \right] + B_k(s^t) + \Pi(s^t) + T(s^t), \end{aligned}$$

where $\Xi(\alpha_j(s^t)) = \int_0^{G^{-1}(\alpha_j(s^t))} z dG(z)$ is the expected search cost for a household facing probability $\alpha_j(s^t)$ of becoming a bargain hunter in store j . Hence, the first-order condition for the cut-off fixed cost, $z_j^*(s^t)$, can implicitly be written instead as the first-order condition

for the probability of searching in store j , $\alpha_j(s^t)$:

$$\begin{aligned} W(s^t)\Xi'(\alpha_j(s^t)) &= \left(\frac{\theta}{\theta-1}D_j^B(s^t) - 1\right) \left(\int P c_{k,j}(s^t) dF_j^B(P, s^t)\right) \\ &\quad - \left(\frac{\theta}{\theta-1}D_j^W(s^t) - 1\right) \left(\int P c_{k,j}(s^t) dF_j^W(P, s^t)\right), \end{aligned} \quad (6)$$

where $D_j^B(s^t)$ and $D_j^W(s^t)$ are factors stemming from price dispersion faced by bargain hunters (B) and workers (W), respectively, in store j :

$$D_j^q(s^t) = \left[\int \left(\frac{P}{P_j^q(s^t)} \right)^{-\theta} dF_j^q(P, s^t) \right]^{-1/\theta}, \quad q = B, W.$$

$F_j^W(P, s^t) = F_j(P, s^t)$, and $P_j^q(s^t)$ is the quantity-weighted mean price faced by a type- q household:

$$P_j^q(s^t) = \left[\int P \frac{P^{-\theta}}{\int P^{-\theta} dF_j^q(P, s^t)} dF_j^q(P, s^t) \right]^{-1/\theta}.$$

Since $\theta > 1$, $D_j^q(s^t) > 1$, and if the mass of low prices is not too high, bargain hunters face a higher consumption risk for a given variety (i.e., $D_j^B(s^t) > D_j^W(s^t)$), which altogether implies that the right-hand side of (6) is positive for all varieties.

If we denote expected consumption by $c_j^q(s^t) = \int c_{k,j}(s^t) dF_j^q(P, s^t)$, then expected demand can be written as

$$c_j^q(s^t) = \left(\frac{P_j^q(s^t) D_j^q(s^t)}{P(s^t)} \right)^{-\theta} c(s^t).$$

It has the form of a standard CES demand, in which the expected price $P_j^q(s^t)$ is marked up by the dispersion factor $D_j^q(s^t)$: for a given average price in a store, expected consumption is lower than implied by a standard CES demand, because of the uncertainty of consumption of brands across varieties.

Condition (6) equates the marginal expected fixed cost of searching in store j with the marginal utility flow from bargain hunting in that store. Since the latter is positive for any j , and since the distribution of the fixed search cost is continuous over $0 \leq z \leq z_{\max}$, the measure of households that draw a sufficiently low fixed cost for a particular store is always a positive-valued function of the state history, i.e., $\alpha_j(s^t) > 0$. This time-varying mix of shopper types is the key difference from Guimaraes and Sheedy (2011), who assume a constant fraction of bargain hunters. In our model, the fraction of bargain hunters is determined endogenously by

the value of searching for lower prices through equation (6). 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.

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

$$\psi P(s^t) c(s^t)^\sigma = W(s^t),$$

while the first-order conditions for state-contingent bonds yield

$$Q(s^{t+1}|s^t) = \beta \pi(s^{t+1}|s^t) \frac{c(s^{t+1})^{-\sigma}}{c(s^t)^{-\sigma}} \frac{P(s^t)}{P(s^{t+1})},$$

where $\pi(s^{t+1}|s^t) = \frac{\pi(s^{t+1})}{\pi(s^t)}$ is the probability of state s^{t+1} conditional on state s^t . Similarly, period 0 prices satisfy

$$Q(s^t) = \beta^t \pi(s^t) \frac{c(s^t)^{-\sigma}}{P(s^t)}.$$

3.2 Retailer's problem

Retailers of type j are endowed with a linear production technology that converts hours worked $n_j(s^t)$ into output of brands of variety j , $y_j(s^t)$:

$$y_j(s^t) = n_j(s^t).$$

Let $\alpha_{j(i)}^B(s^t)$ denote the number of bargain hunters buying consumption from retailer j in location i , and $\alpha_{j(i)}^W(s^t)$ denote the corresponding number of workers. Due to the random allocation of workers across stores selling the same variety, a retailer of type j takes the number of workers as given, and can only increase its market share by retaining bargain hunters in its location. We can therefore focus on a representative retailer j , omitting location index i . This allows us to denote the number of bargain hunters by $\alpha_j^B(s^t)$ and the number of workers by $\overline{\alpha_j^W}(s^t)$, where a bar indicates that this margin is not impacted by the retailer's pricing decision. Since the search cost is i.i.d. among households, and households are uniformly assigned across locations, the number of households who choose to become a bargain hunter for variety j in their designated location is equal to the number of bargain hunters in that store, i.e.,

$$\alpha_j^B(s^t) = \alpha_j(s^t). \tag{7}$$

The problem of a retail store selling variety j is to choose the distribution of prices $F_j(P, s^t)$ to maximize its profits in period t :

$$\begin{aligned} & \alpha_j^B(s^t) \int (P - W(s^t)) \left(\frac{P}{P(s^t)} \right)^{-\theta} c(s^t) dF_j^B(P, s^t) \\ & + \overline{\alpha_j^W}(s^t) \int (P - W(s^t)) \left(\frac{P}{P(s^t)} \right)^{-\theta} c(s^t) dF_j(P, s^t) - \kappa W(s^t) F_j(\overline{P}, s^t), \end{aligned} \quad (8)$$

subject to the constraints for the number of bargain hunters in store j , (6) and (7), and the constraint on the price distribution faced by bargain hunters (2). The last term in (8) captures the total fixed cost of posting prices below a given price \overline{P} , expressed in units of labor. Finally, we assume that retailers face costs of adjusting their nominal prices, which we specify explicitly below.

In Section 4 we demonstrate, using a solution of the retailer's problem for the case of a two-point price distribution, how the differences in price distributions faced by bargain hunters and workers as well as the randomization of workers across locations lead to price discrimination by retailers. First, lower prices create an incentive for households who draw low search costs to become bargain hunters in their designated locations. By itself, this is not a sufficient condition for price discrimination. For example, if each household shops in the designated location regardless of the search cost, the measure of households shopping in a given store is fixed so that for every extra bargain hunter the store loses a worker. In this case, due to the concavity of the profit function, a smaller price discount yields higher profits, and hence, it is optimal for each retailer to post the same price for all its brands. This is where the second assumption — randomization of workers across locations — comes into play. It implies that a retailer j in a particular location has no control over the number of workers shopping in its store. In this case, retailers do not internalize the effect their pricing decisions have on the number of workers in their store and compete for market share by posting low prices to attract bargain hunters.

3.3 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:

$$\int_0^1 \left[l_k(s^t) - \int_j \Xi(\alpha_j(s^t)) dj \right] dk = \int_0^1 \int_0^1 [n_j(s^t) + \kappa F_j(\bar{P}, s^t)] dj di.$$

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

$$\begin{aligned} \int_0^1 [\alpha_{k,j}(s^t) c_{k,j}^B(s^t) + (1 - \alpha_{k,j}(s^t)) c_{k,j}^W(s^t)] dk = \\ \int_0^1 \left[\alpha_j^B(s^t) \int \left(\frac{P}{P(s^t)} \right)^{-\theta} c(s^t) dF_j^B(P, s^t) + \overline{\alpha_j^W}(s^t) \int \left(\frac{P}{P(s^t)} \right)^{-\theta} c(s^t) dF_j(P, s^t) \right] di. \end{aligned}$$

The equilibrium search flow condition states that the total number of households who choose to become workers for variety j is equal to the total number of workers buying variety j across all locations:

$$\int_0^1 (1 - \alpha_{k,j}(s^t)) dk = \int_0^1 \overline{\alpha_j^W}(s^t) di.$$

A *symmetric equilibrium* for this economy is a collection of allocations for households: $c(s^t), l(s^t), c_j(s^t), M(s^t), B(s^t), \alpha_j(s^t)$; prices and allocations for firms: $y_j(s^t), n_j(s^t), F_j(P, s^t), F_j^B(P, s^t)$; and aggregate prices $P(s^t), W(s^t), Q(s^t)$ that satisfy household and firm maximization and market clearing conditions.

4 Model mechanics

In this section, we describe the mechanics of our model under the assumptions of both flexible and sticky prices. For simplicity, throughout the rest of the paper, we will consider equilibria with discrete two-price distributions. That is, stores post price $P_j^L(s^t)$ for a fraction $\gamma_j(s^t)$ of their brands and price $P_j^H(s^t)$ for the remaining fraction of brands, $1 - \gamma_j(s^t)$. Workers buy a brand that they randomly draw in this store. Bargain hunters, on the other hand, face lower prices, on average: we assume that bargain hunters are able to increase the probability of their lowest price being $P_j^L(s^t)$ by a factor $f > 1$. This implies that bargain hunters find

the low price with probability $f\gamma_j(s^t)$ and the high price with probability $1 - f\gamma_j(s^t)$.²⁷

4.1 Model with flexible prices

We begin by studying an environment with fully flexible prices: firms face no constraints to change either P^L or P^H .

Denote by $\Pi(P, s^t) = (P - W(s^t)) \left(\frac{P}{P(s^t)}\right)^{-\theta} c(s^t)$ the profit function at price P , and by $R(P, s^t) = P \left(\frac{P}{P(s^t)}\right)^{-\theta} c(s^t)$ the revenue function at P . Then the retailer's expected profit and revenue from selling to workers in period t and state s^t are

$$\Pi_j^W(s^t) = \gamma_j(s^t) \Pi(P_j^L(s^t), s^t) + (1 - \gamma_j(s^t)) \Pi(P_j^H(s^t), s^t),$$

$$R_j^W(s^t) = \gamma_j(s^t) R(P_j^L(s^t), s^t) + (1 - \gamma_j(s^t)) R(P_j^H(s^t), s^t).$$

Similarly, the retailer's expected profit and revenue from selling brands to bargain hunters, $\Pi_j^B(s^t)$ and $R_j^B(s^t)$, are the same as above, where $\gamma_j(s^t)$ is replaced by $f\gamma_j(s^t)$.

The retailer's problem is to choose prices $P_j^L(s^t)$, $P_j^H(s^t)$ and the fraction of brands on sale, $\gamma_j(s^t)$, to maximize

$$\sum_t \int Q(s^t) \left[\alpha_j(s^t) \Pi_j^B(s^t) + \overline{\alpha_j^W}(s^t) \Pi_j^W(s^t) - \kappa W(s^t) \gamma_j(s^t) \right] ds^t,$$

subject to the two-price version of the constraint (6) on the number of bargain hunters:

$$W(s^t) \Xi'(\alpha_j(s^t)) = \left(\frac{\theta}{\theta - 1} D_j^B(s^t) - 1 \right) R_j^B(s^t) - \left(\frac{\theta}{\theta - 1} D_j^W(s^t) - 1 \right) R_j^W(s^t). \quad (9)$$

The dispersion factor $D_j^W(s^t)$ is a function of the size of discount $P_j^L(s^t)/P_j^H(s^t)$ and the fraction of low prices $\gamma_j(s^t)$:

$$D_j^W(s^t) = \frac{\left(\gamma_j(s^t) \left(\frac{P_j^L(s^t)}{P_j^H(s^t)} \right)^{-\theta} + 1 - \gamma_j(s^t) \right)^{1-1/\theta}}{\gamma_j(s^t) \left(\frac{P_j^L(s^t)}{P_j^H(s^t)} \right)^{1-\theta} + 1 - \gamma_j(s^t)},$$

and $D_j^B(s^t)$ is given by the same expression, with $\gamma_j(s^t)$ replaced by $f\gamma_j(s^t)$.

²⁷For example, if bargain hunters have two price draws (with replacement) and buy consumption at the lowest of the two prices, the probability of finding $P_j^L(s^t)$ in this case is $2\gamma_j(s^t)$, i.e., $f = 2$.

Optimality of price discrimination

To gain intuition for why it is optimal for retailers to post different prices, we demonstrate the effect of the search constraint (9) on retailers' pricing behavior. First, suppose that retailers cannot affect the search decision of households via this constraint, treating the number of bargain hunters $\alpha_j(s^t)$ as given. It immediately follows from concavity of the retailer's profit function that it is optimal to charge the monopoly price for all its brands, $P_j^L(s^t) = P_j^H(s^t) = P_j^*(s^t) \equiv \frac{\theta}{\theta-1}W(s^t)$, and collect monopoly profits $\Pi_j^*(s^t) = \Pi(P_j^*(s^t), s^t)$.

In contrast, when the retailer can attract bargain hunters, it effectively faces constraint (9). In this case, posting a single price for all of its brands is no longer optimal. Specifically, consider a small perturbation by a retailer that, instead of posting a single monopoly price $P_j^*(s^t)$ for all its brands, posts an arbitrarily lower price $P_j^L(s^t) < P_j^*(s^t)$ for a small fraction of brands $\gamma_j(s^t) = d\gamma$, attracting a small number of bargain hunters, $\alpha_j(s^t) = d\alpha$, and taking the number of workers as given, $\overline{\alpha_j^W}(s^t) \approx 1$. The variation of the retail firm j 's profit relative to monopolistic profit $\delta\Pi_j^*(s^t) \equiv \Pi_j(s^t) - \Pi_j^*(s^t)$ is, up to first order, given by

$$\delta\Pi_j^*(s^t) \approx \Pi_j^*(s^t) d\alpha + d\gamma [H(P_j^L(s^t), s^t) - H(P_j^*(s^t), s^t) - \kappa W(s^t)] .$$

The first term on the right-hand side is the gain in profits due to sales to the extra fraction $d\alpha$ of bargain hunters; this term would have been fully offset by the loss in profits due to the smaller number of workers, if that number had directly depended on the retailer's pricing decision. The second term is the loss in profits due to sales of the fraction $d\gamma$ of brands at a lower price $P_j^L(s^t)$ and the total fixed cost of posting sale prices. From the search constraint (9), we can express $d\alpha$, up to first order, as

$$d\alpha \approx d\gamma \frac{1}{\theta-1} \frac{f-1}{W(s^t) \Xi''(0)} (R(P_j^L(s^t), s^t) - R(P_j^*(s^t), s^t)) . \quad (10)$$

The increase in the fraction of bargain hunters, $d\alpha$, is proportional to the increase in the number of discounted prices, $d\gamma$. This is true because while a higher fraction of discounted items raises the probability of both bargain hunters and workers finding an item on sale, that probability rises faster for bargain hunters: a 1 percentage point increase in the fraction of sales leads to an $f > 1$ percentage point rise in the likelihood of bargain hunters finding P^L , but only a 1 percentage point rise for workers. Expression (10) also shows that bargain

hunting increases if wages fall, or if the size of a price discount increases.

Plugging (10) into the expression for profit variation, we get

$$\delta \Pi_j^* (s^t) \approx d\gamma \cdot \Pi_j^* (s^t) \left[\begin{array}{c} -\frac{\Pi(P_j^*(s^t), s^t) - \Pi(P_j^L(s^t), s^t) + \kappa W(s^t)}{\Pi_j^*(s^t)} \\ + \frac{f-1}{\theta-1} \frac{R(P_j^L(s^t), s^t) - R(P_j^*(s^t), s^t)}{W(s^t) \Xi''(0)} \end{array} \right].$$

If the expression in brackets is positive, i.e., if

$$\frac{f-1}{\theta-1} \frac{R(P_j^L(s^t), s^t) - R(P_j^*(s^t), s^t)}{W(s^t) \Xi''(0)} > \frac{\Pi(P_j^*(s^t), s^t) - \Pi(P_j^L(s^t), s^t) + \kappa W(s^t)}{\Pi_j^*(s^t)},$$

then the benefits stemming from the increase in the market share of the retailer from selling to bargain hunters outweigh the cost of a smaller profit margin from selling at lower prices, on average, and so the retailer can improve upon single-price monopolistic profit by posting two prices.

Fraction of sale prices

What is the desirable number of low prices for a retailer? The first-order condition for the number of brands on sale $\gamma_j (s^t)$ in the retailer's problem is

$$\left(\overline{\alpha_j^W} (s^t) + f\alpha_j (s^t) \right) \left[\Pi(P_j^H(s^t), s^t) - \Pi(P_j^L(s^t), s^t) \right] + \kappa W(s^t) = \frac{\partial \alpha_j (s^t)}{\partial \gamma_j (s^t)} \Pi_j^B (s^t). \quad (11)$$

The left-hand side is the marginal loss from increasing the fraction of sales. It consists of two terms: the first term is the marginal loss from selling more goods at a lower price, equal to the product of the marginal increase in the measure of transactions at a low price, $\overline{\alpha_j^W} (s^t) + f\alpha_j (s^t)$, and the loss in profits per each brand on sale, $\Pi(P_j^H(s^t), s^t) - \Pi(P_j^L(s^t), s^t)$. The second term is the marginal increase in the total fixed cost of posting sales, $\kappa W(s^t)$. The right-hand side is the marginal profit from increasing the fraction of sales: it is the product of the marginal increase in the number of bargain hunters, $\frac{\partial \alpha_j (s^t)}{\partial \gamma_j (s^t)}$, and the average profits per bargain hunter, $\Pi_j^B (s^t)$. By differentiating (9) with respect to $\gamma_j (s^t)$, we can show that $\frac{\partial \alpha_j (s^t)}{\partial \gamma_j (s^t)}$ is decreasing in $\gamma_j (s^t)$, because the utility loss due to price dispersion, given by $D_j^B (s^t)$ and $D_j^W (s^t)$, is growing faster for bargain hunters than for workers; or because the marginal cost of searching, $\Xi''(\alpha_j (s^t))$, is increasing in the number of bargain hunters. In turn, the retailer's expected profit per bargain hunter, $\Pi_j^B (s^t)$, is also decreasing with $\gamma_j (s^t)$,

since more brands are sold at a lower price. Figure 7 plots the marginal gain and loss from increasing the fraction of brands on sale as functions of the number of brands on sale.

Size of price discount

How does the firm pin down the levels of its high and low prices? The corresponding first-order conditions are

$$N_j^i(s^t) (MR(P_j^i(s^t), s^t) - MC(P_j^i(s^t), s^t)) + \frac{\partial \alpha_j(s^t)}{\partial \ln P_j^i(s^t)} \Pi_j^B(s^t) = 0, \quad (12)$$

where $i = H, L$, and $MR(\cdot, s^t)$ and $MC(\cdot, s^t)$ are the marginal revenue and marginal cost functions of a single-price firm, respectively:

$$\begin{aligned} MR(P, s^t) &= -(\theta - 1) \left(\frac{P}{P(s^t)} \right)^{-\theta} c(s^t), \\ MC(P, s^t) &= -\theta \frac{W(s^t)}{P} \left(\frac{P}{P(s^t)} \right)^{-\theta} c(s^t). \end{aligned}$$

$N_j^H(s^t)$ and $N_j^L(s^t)$ are the number of transactions at high and low prices; $\frac{\partial \alpha_j(s^t)}{\partial \ln P_j^H(s^t)}$ and $\frac{\partial \alpha_j(s^t)}{\partial \ln P_j^L(s^t)}$ are the marginal changes in the number of bargain hunters due to the increase in the high and low log prices, respectively.

Conditions (12) imply that if the retailer could not affect the number of bargain hunters in the store (i.e., if $\frac{\partial \alpha_j(s^t)}{\partial \ln P_j^H(s^t)} = \frac{\partial \alpha_j(s^t)}{\partial \ln P_j^L(s^t)} = 0$), then the optimal low and high prices would both equate the single-firm marginal revenue and marginal cost, and both would be equal to the optimal monopolistic price $P_j^*(s^t)$, since, by definition, $MR(P_t^*(j)) = MC(P_t^*(j))$. In contrast, if the retailer can indeed attract bargain hunters, search condition (9) implies that $\frac{\partial \alpha_t(j)}{\partial \ln P_t^H(j)} > 0$ and $\frac{\partial \alpha_t(j)}{\partial \ln P_t^L(j)} < 0$; i.e., a higher high price and a lower low price make bargain hunting more attractive. Then from (12) it follows that $MR(P_j^L(s^t), s^t) > MC(P_j^L(s^t), s^t)$ and $MR(P_j^H(s^t), s^t) < MC(P_j^H(s^t), s^t)$; i.e., the retailer's optimal low (high) price is below (above) the monopolistic price, $P_t^L(j) < P_t^* < P_t^H(j)$, see Figure 8. We surmise from the price-setting mechanism that a more elastic search implies larger absolute values of derivatives $\frac{\partial \alpha_t(j)}{\partial \ln P_t^L(j)}$, $\frac{\partial \alpha_t(j)}{\partial \ln P_t^H(j)}$, and, therefore, a larger price discount (lower $P_t^L(j) / P_t^H(j)$).

Finally, note that as $\gamma_j(s^t) \rightarrow 0$, $P_j^H(s^t) \rightarrow P_j^*(s^t)$ and $P_j^L(s^t) \rightarrow \frac{1}{1 + \frac{f-1}{\theta-1} \frac{\Pi_j^*(s^t)}{W(s^t)\Xi''(0)}} P_j^*(s^t) < P_j^*(s^t)$; i.e., price dispersion does not disappear as the fraction of sale prices goes to zero.

This means that there is no equilibrium with a single price. The reason why firms find it optimal to deviate from having a single price for all brands is that for any positive measure of brands on sale, the return for the household of being a bargain hunter is strictly positive, given by the right-hand side of (9). If the marginal expected fixed cost of searching is not growing too quickly with the probability of searching (e.g., if $\Xi''(0)$ is low enough), then there will always be enough bargain hunters to justify posting low prices by retailers.

4.2 Model with sticky prices

Next, we turn to our benchmark version with sticky prices: we assume that retailers face Taylor-type price adjustment constraints for high (regular) prices. Low (sale) prices are assumed to be flexible, although this has virtually no effect on our results, as we show below. Hence, retailer j sets a new high price in period t , and keeps that price fixed for N periods, i.e. it faces the following pricing constraint:

$$P_j^H(s^t) = P_j^H(s^{t+1}) = \dots = P_j^H(s^{t+N-1}) . \quad (13)$$

Following the standard Taylor (1980) pricing, such price contracts are evenly staggered, so that in every period a measure $1/N$ of retailers resets their prices.

Retailer j 's problem, then, is to choose prices $P_j^L(s^\tau)$, $P_j^H(s^\tau)$ and the number of prices on sale $\gamma_j(s^\tau)$ to solve

$$\max \sum_{\tau=t}^{t+N-1} \int Q(s^\tau) \left[\alpha_j(s^\tau) \Pi_j^B(s^\tau) + \overline{\alpha_j^W}(s^\tau) \Pi_j^W(s^\tau) - \kappa W(s^\tau) \gamma_j(s^\tau) \right] ds^\tau ,$$

subject to the constraint on bargain hunters (9) and pricing constraints (13).

The first-order conditions yield equations for new (reset) high prices:

$$P_t^H(j) = \frac{\theta}{\theta - 1} \frac{\sum_{\tau=t}^{t+N-1} \int Q(s^\tau) C(s^\tau) P(s^\tau)^\theta N_j^H(s^\tau) W(s^\tau) ds^\tau}{\sum_{\tau=t}^{t+N-1} \int Q(s^\tau) C(s^\tau) P(s^\tau)^\theta N_j^H(s^\tau) \left(1 - \frac{\partial \alpha_j(s^\tau)}{\partial \ln P_j^H(s^\tau)} \frac{\Pi_j^B(s^\tau)}{(\theta-1) N_j^H(s^\tau) R_j^H(s^\tau)} \right) ds^\tau} .$$

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

5 Results

5.1 Parameterization

Table 6 provides the parameter values that we use in our quantitative simulations. The period is a month, so the discount factor $\beta = 0.96^{1/12}$. We set $\sigma = 1$, a standard value in business cycle literature, and $\psi = 1$, which together imply that nominal wages follow 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 sensitive to the level of θ .

We assume that regular prices change once every 12 months ($N = 12$), consistent with our data and existing studies. Although this degree of price stickiness is somewhat higher than that 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 (2008). We assume that sale prices are fully flexible. As we show below, this assumption has little effect on overall price flexibility, since firms find it optimal to keep the size of the price discount close to the steady state. The latter result is consistent with our findings that the size of the average price discount is very acyclical. It is also 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 search efficiency parameter f , the fixed cost of posting price discounts κ , and parameters of the distribution of the fixed cost of searching $G(z)$. Parameters of $G(z)$ determine the steady-state level and the slope of the marginal fixed cost function $\Xi'(\alpha)$, and thereby pin down the share of consumption at low prices and the elasticity of consumer search over the cycle. Since consumer-search elasticity is directly linked to retailers' incentives to price discriminate, we choose the average price discount $P^L/P^H = 0.78$ as a calibration target. Regarding the share

of consumption purchased at low prices, Gandon (2011) reports that for groceries in the United States, sales account for 17% of price quotes but 40% of revenue, so that the ratio of revenue share to the fraction of sales is 2.7; this ratio is another calibration target. To match these two targets, we choose $G(z)$ to be uniform: it assigns equal probabilities of any fixed cost draw between 0 and z_{\max} . For $z_{\max} = 0.31$, this parameterization implies that in the steady state the fraction of bargain hunters is $\alpha = 0.10$ and the elasticity of the marginal expected fixed cost of searching, $\epsilon_{\Xi} = \frac{\alpha \Xi''(\alpha)}{\Xi'(\alpha)}$, is equal to 1.

The search efficiency f and the fixed cost of posting price discounts κ jointly determine the fraction of sales γ and have little effect on the other calibration targets. We therefore fix f at 2, which means that bargain hunting allows households to double the probability of finding the low price. We then set κ to 0.08 to match the fraction of sale prices $\gamma = 0.05$. This number is in the upper range of values in the U.K. data presented in this paper, and in other European CPI data studies;²⁸ it is somewhat lower than between the 7 and 11% found for the United States by Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008).

To highlight the main mechanism in the benchmark model, we consider an alternative parameterization of the fixed cost distribution that implies that the search for low prices is very inelastic. Such distribution has a step-like shape, splitting the unit probability weight between only two realizations of the fixed cost: zero (with weight α) and z_{\max} (with weight $1 - \alpha$), where z_{\max} is a large number. Households who face the step-like cost distribution will choose to search with probability α regardless of the state of the economy. Hence, the key difference between the models with a uniform and a step-like fixed cost of searching is that in the former model the fraction of bargain hunters will move elastically over the business cycle, while it is virtually fixed in the latter model.²⁹ Figure 9 plots both fixed cost cumulative distribution functions (CDFs). Appendix B provides specifications for both distributions.

²⁸The fraction of sales is between 0.03 and 0.05 in Austrian CPI data (Baumgartner et al. 2005); 0.03 in Norwegian CPI data (Wulfsberg 2009), 0.02 in French CPI data (Baudry et al. 2004; Berardi et al. 2013).

²⁹In our numeric simulations for the inelastic case, we utilize a flexible two-parameter fixed cost CDF specification, borrowed from Dotsey, King and Wolman (1999), who use it in their state-dependent pricing model. In the inelastic case, the steady-state elasticity of the marginal expected fixed cost of searching $\epsilon_{\Xi} = 17.3$, which is more than 17 times the elasticity in the elastic case. More-volatile marginal costs of searching in the inelastic case decrease the volatility of the fraction of bargain hunters.

5.2 Responses to monetary shocks

Our baseline experiment is a negative 1% innovation to the money growth. Figure 10 provides the responses of output, prices, the average fractions of price discounts and the average fractions of bargain hunters in the models with elastic and inelastic fractions of bargain hunters. For output and price responses, the figure also provides the responses in the standard Taylor model with a single retail price.

The figure demonstrates a stark difference in the output and price responses between the economies with elastic and inelastic fractions of bargain hunters. The aggregate price level (top-right panel) in the model with a fixed fraction of bargain hunters mimics that for the standard Taylor model, falling steadily to -1% within one year, as all 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, just as in the standard Taylor model.

In contrast, the real effect on output is about half as large in the model with a moving fraction of bargain hunters, 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 model with inelastic bargain hunters, immediately drops by 0.5% at the time of the shock.

To show that these results do not depend on the persistence of the monetary shock, we repeat the impulse-response function simulations assuming that the money-growth process is serially correlated with persistence parameter equal to 0.536, as in Guimaraes and Sheedy (2011). Figure 11 shows that indeed all the insights from Figure 10 carry over to the case with the persistent shock: the endogenous reaction of bargain hunters leads to a large impact of the sales margin on the response of aggregate prices and output to shocks.

5.3 Understanding impulse responses

The marked differences in the degree of the aggregate price flexibility between the two models are due to the mechanism by which households' search for lower prices amplifies the incentives for price discounting by retailers. To understand this mechanism, let us review the factors that determine the size of the decline in the aggregate price level immediately after the shock. First, we log-linearize the expression for the aggregate price level (5) around the steady state

assuming that both γ and α are, on average, close to zero, and P^H is close to P . Denoting the average price discount by $\delta = P^L/P^H$, we obtain that the log deviation of the aggregate price level from its steady-state value is, approximately,

$$\widehat{P}_t \approx -\frac{P}{\theta - 1} \left(\delta^{1-\theta} - 1 \right) \tilde{\gamma}_t, \quad (14)$$

where a hat (tilde) denotes the log-linearized (difference-) deviation from the steady state, and steady-state values have no subscripts. Expression (14) 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 price discount. In the elastic (inelastic) case, the discount is equal to $\delta = 0.78$ (0.89), implying a response of the fraction of sales of $\tilde{\gamma}_t = 1.1$ (0.5) percentage points (bottom-left panel in Figure 10). Plugging these values into (14) gives a price level response of -0.5% (-0.1%) at the time of the shock, as shown in the figure. Hence, both the average discount and relative responsiveness of the fraction of sales contribute to the fivefold larger response of the price level in the elastic case relative to the inelastic case.³⁰

We next describe how the differences in both the size and the number of discounts stem from the difference in the elasticity of search behavior by households. First, use pricing equations (12) to approximate the average size of price discount:

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

As we already established in Section 4.1, a larger search elasticity (lower $\Xi''(\alpha)$) and a larger search efficiency f imply a larger discount. Keeping $f = 2$, and using the steady-state value for $\Xi''(\alpha)$ in the elastic (inelastic) case, 0.3 (0.6), we obtain, approximately, that the average discount is 0.85 (0.92), close to the accurate steady-state value of 0.78 (0.89).

Second, to understand the response in the fraction of sales, $\tilde{\gamma}_t$, note that if that fraction

³⁰The average price discount in the inelastic case is robustly smaller than the one in the elastic case. In the inelastic case, the retailer's return to price discrimination is smaller, since households are less attracted by discounts. Accordingly, retailers scale down price discrimination by reducing both the size and the number of discounts. While the average number of discounts can be easily recalibrated to match the one in the elastic case (by changing the fixed cost of posting sales κ), the average size of discounts is only responsive to demand elasticity θ . To match the size of discount in the elastic case, bringing δ from 0.89 to 0.78, requires decreasing θ from 5 to 1.5. In this case, as in other recalibrations we conducted, the impulse-response results are virtually unchanged.

were fixed, sticky regular prices would imply a high average markup at the store, increasing $\Pi_j^B(s^t)$, the expected profit per bargain hunter. The first-order condition for the fraction of sales (11) states that the benefit flow for the retailer (right-hand side) depends on how many extra bargain hunters an additional brand on sale generates, i.e., on the derivative $\frac{\partial \alpha_j(s^t)}{\partial \gamma_j(s^t)}$. In the elastic case, the steady-state value of this derivative is almost six times larger than in the inelastic case, 1.75 vs 0.29. Since a monetary shock does not directly affect the loss from posting sales (left-hand side of equation (11)), retailers must bring down their marginal benefit by increasing the number of sales (since both $\frac{\partial \alpha_j(s^t)}{\partial \gamma_j(s^t)}$ and $\Pi_j^B(s^t)$ decrease with $\gamma_j(s^t)$). Altogether, the response of $\tilde{\gamma}_t$ is twice as large in the elastic case as in the inelastic case.

Finally, 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 yields 1.6 percentage points for the elastic case and 0.2 percentage points for the inelastic case (lower-right panel in Figure 10).

Where does the flexibility of the aggregate price response come from? There are two contributing factors: in addition to the increase in the fraction of sale prices, there is also a shift of consumption weight toward sale prices. Figure 12 decomposes the response of the true price index from our benchmark experiment, denoted by “ P ,” equal to -0.56% on impact, into counterfactual responses with one or both factors sequentially turned off. 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. Therefore, the unmeasured variation in the aggregate price accounts for around one-third of the total variance of the aggregate price level.³¹ Turning off fluctuations — in addition to varying consumption weights — in the number of brands on sale (given by “ $P_{CPI}(regular)$ ”) brings the impulse response to a mere -0.07% on impact, close to the response in the standard one-price Taylor model or the model with inelastic search.

Our findings therefore broaden the insights from Chevalier and Kashyap (2012), who emphasize the importance of accounting for the shift in consumption weights toward sale prices for the degree of aggregate price flexibility. First, we quantify that the shift in weights

³¹Most of the shift in consumption weights toward lower prices after the shock is due to the increase in consumption per shopper, while the increase in the number of shoppers buying at low prices has a smaller impact.

in response to a shock can add about half of the variance implied by the constant-weight price index. Second — and this is our main contribution — we show that aggregate price flexibility stemming from fluctuations in the fraction of sale prices can be even higher than the one implied by the shift in consumption weights.

5.4 What do sales tell us about business cycles?

Our impulse-response results demonstrate the importance of accounting for the interaction between retailers’ price-discounting behavior and households’ search for low prices. A countercyclical search for low prices is necessary in making sales important for aggregate price flexibility. This prediction of our model has important implications regarding the nature of the real effects of monetary policy shocks. Namely, in our model, after a monetary contraction, markups must rise, or else retailers would not benefit from attracting low-margin shoppers (bargain hunters). We demonstrate this by examining factors driving the increase of the fraction of bargain hunters after the shock.

According to equation (9), the fraction of bargain hunters would increase after the shock if (a) the fraction of sales increases, (b) the price discount increases, or (c) wages decrease (lowering the opportunity cost of searching). We repeat impulse-response simulations for variations of the baseline economy in which we shut down each of these three factors. First, we assume that low prices P^L are sticky and can be adjusted every 12 months together with high prices. In this variation of the model, the average price discount is virtually unchanged after the shock. In the second variation, we assume that the fraction of sales is fixed at its steady-state level, $\gamma(s^t) = \gamma$. In the third variation, we assume that nominal wages are sticky, given by equation

$$W(s^t) = W(s^{t-1})^{11/12} M(s^t)^{1/12}, \quad (15)$$

so that it takes around a year for nominal wages to converge to their pre-shock level. In this variation, we will assume that retailers’ marginal cost is still flexible and equal to the money supply, as in the baseline model. Finally, in the fourth variation of the model, we assume that household wages are flexible, and it is the nominal marginal cost that is sticky, obeying the law of motion (15).³²

³²It is straightforward to extend our benchmark model to include capital, sticky nominal wages and nominal cost rigidities to explicitly account for reduced forms for sticky wages and marginal cost. Such an extension would not change the point we make here.

Table 7 provides concurrent responses to the same shock (1% negative impulse to money growth) in the baseline model (column A) and its four variations (columns B through E). First, as we noted in Section 5, sticky sale prices (column B) make virtually no difference for our results, since retailers find it optimal to vary the fraction of sales, and not the discount, to attract bargain hunters. This arises because the retailer’s value is nearly linear in the fraction of sales: when that fraction is small, on average, both retailer’s marginal gains and losses are very inelastic with respect to changes in the fraction, see Figure 7. When retailers cannot adjust the fraction of sales (column C), they can attract bargain hunters only by increasing the price discount. After the shock, the discount increases by 1.2 ppt (-0.06 in the baseline). Since the discount size is a less-effective adjustment margin for retailers, the implication is a smaller change in the fraction of bargain hunters than in the baseline model, 1.17 vs 1.62 ppt, respectively. As a result, the shift of consumption weights toward low prices is more limited, resulting in a stickier response in the aggregate price and a larger response in consumption, -0.68% (vs -0.44% in the baseline). Even in this case, monetary non-neutrality is smaller by one-third relative to the case with an inelastic consumer search. These simulations underline the importance of studying the entire arsenal of tools that retailers use to attract customers: if retailers cannot adjust the fraction of sales in response to a monetary shock, they will adjust along other margins (size of discounts, coupons, etc.).

Our model predicts that the quick fall of household wages after the shock matters relatively little for incentivizing households to look harder for price discounts (column D in Table 7). When nominal wages are sticky, they fall by only about 0.1% after the shock, as opposed to 1% in the baseline model. Since the marginal time cost of searching, in dollars, is equal to $W(s^t)\alpha(s^t)$, a 0.9% smaller fall in wages implies a smaller rise in the fraction of bargain hunters by only 0.1 ppt ($= 0.9\alpha$). In addition, retailers post sales less aggressively, amplifying the effect on the fraction of bargain hunters by 0.26 ppt, (i.e., the increase in the fraction is 1.36 ppt vs 1.62 ppt in the baseline model) — still a relatively small number.

In the three model variations considered so far (columns B, C and D), markups rise by roughly half of a percentage point after the monetary contraction, making it a good time for retailers to increase their market share. In the last simulation, we limit the rise in the markup on impact by assuming sticky marginal cost (column E in Table 7). In this case, retailers barely change their prices after the shock. The fraction of bargain hunters increases

by only 0.1 ppt due to the 1% fall in household wages, implying only a small effect on the aggregate price and, therefore, a large consumption response at -0.94%, very close to the inelastic case. Hence, in our model, countercyclical markups represent the central impetus that leads retailers to make more-intensive use of sales during slumps.³³

In sum, our findings imply that changes in households' shopping activity are associated, to a large extent, with retailers' pricing over the cycle, and less with changes in the opportunity cost of time. Coibion, Gorodnichenko and Hong (2014) and Nevo and Wong (2014) emphasize the role of the decline in the opportunity cost of time in explaining the increased shopping activity during slumps. Our paper, therefore, shows that the rise in return to shopping due to more frequent sales in recessions can also be a powerful driving mechanism behind fluctuations in households' shopping time.

6 Discussion

In this section, we report some empirical evidence on price search behavior over the business cycle and discuss how a broader view of price discrimination may allow us to reconcile our findings with those of Coibion, Gorodnichenko and Hong (2014).

6.1 Price search behavior during the Great Recession

In the previous section, we presented a model in which price discrimination gave rise to sales and altered the cyclical properties of aggregate variables. As we made clear, the key mechanism behind this result is that households search for bargains more intensively in recessions. In what follows, we discuss both existing and novel evidence in support of this central prediction of our model.

Time-use surveys have been a useful source of evidence on search and shopping activity over the business cycle or across households. For example, Aguiar et al. (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

³³A growing consensus in macroeconomic literature is that nominal cost rigidities, rather than countercyclical markups, account for most of the monetary non-neutrality, see Christiano et al. (2005). In contrast, our theory, combined with countercyclical consumer search in the data, implies that countercyclical markups are indeed important for the real effects of monetary shocks. Kryvtsov and Midrigan (2013) arrive at a similar conclusion, based on the study of the observed behavior of inventories.

hours are allocated to increased shopping time. This is in line with evidence from Aguiar and Hurst (2007), who document important fluctuations in the shopping margin over the life cycle, and Krueger and Mueller (2010), who find that unemployed individuals devote 15% to 30% more time to shopping than the employed. Kaplan and Menzio (2014) use the Kilts-Nielsen Consumer Panel Dataset to show that “households with more non-employed members pay less for the same basket of goods than households whose members are all employed: the average price index for non-employed households is between 1% and 4.5% lower than the price index for employed households.” Using a similar data set, Nevo and Wong (2014) 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. All this evidence supports the assumption that, on average, consumers search for lower prices more intensively during economic downturns.

Next, we turn to a new source of information on search behavior: Google Trends. A few years ago, Google made it possible to extract time-series indexes of how intensively Internet users in a given location search online for specific search phrases relative to the total number of searches done through their engine. While these indexes are automatically normalized (Google does not make the raw data available), this is not problematic for our objective of determining whether search for low prices was significantly higher during the 2008–09 recession.

The first challenge in this type of analysis is to pick relevant Google search terms. For example, using the term “sale” by itself is problematic: it is too general and will include searches that are not relevant for our purposes (items that are “for sale”) in addition to those that are (products that are “on sale”). For this reason, we settled on two search terms that we considered closely related to bargain hunting behavior: “clearance” and “on sale.” That being said, our findings are similar across a wide range of other search terms, such as “outlet,” “deal of the day,” “shoes on sale,” “half-off deals,” “promotional code,” etc.

In Figure 13, we plot the unemployment rate alongside Google search activity between January 2004 and December 2013, which is normalized to be equal to 100 in December 2006, for both the United States and the United Kingdom. The evidence is very clear: for both countries, during the Great Recession there is a strong rise in online search using terms related to bargain hunting. For example, search activity in the United Kingdom for the term “clearance” tripled between early 2007 and mid-2009. In the United States, the indexes for

both terms have been trending downward alongside unemployment since 2010.³⁴

In summary, evidence from both time-use surveys and from online search activity appears to support our central modelling feature: search activity for low prices rises during recessions.

6.2 Other forms of price discrimination over the cycle: the case of coupons

In Section 2, we documented a very robust link between aggregate economic downturns and a more intensive use of sales by U.K. retailers. In the United States, preliminary evidence from the CPI micro data also seems to support the countercyclicality of sales. Yet, a recent study by Coibion, Gorodnichenko and Hong (2014) finds that sales are mostly acyclical. They reach this conclusion using the IRI Marketing Dataset, a large scanner database that covers about 30 product categories across grocery stores in 50 distinct U.S. markets, with up to 150 stores per market (see Bronnenberg et al. 2008 for a description of the database). Based on the same data set, we confirm their findings about temporary sales: while some markets/categories exhibit an increase in the incidence of sales during the Great Recession, these movements are in general not particularly significant, either statistically or economically.

We do, however, find another interesting feature: the use of coupons by retailers (generally through weekly flyers), information on which is available in the data set, almost doubled during the 2008–09 recession. This can be seen in the left-hand plot of Figure 14, where we show the 12-month moving average of the frequency of retailer coupons between 2004 and 2011 alongside the unemployment rate.³⁵ Notice, however, that the frequencies are very low and at first sight unlikely to be of much economic significance.³⁶ This should not be surprising, since most coupons are distributed by manufacturers and would therefore not be accounted for in the IRI database. Nonetheless, our finding is corroborated by evidence from industry sources that points to substantial increases in both the distribution and redemption of coupons during the last recession (see, for example, Inmar Inc. 2013, or NCH 2013). It also coincides with a dramatic rise in Google searches for the term “coupons” over the same time period (right-hand plot of Figure 14).

³⁴For conciseness, we report only graphical evidence for two search terms in two regions. We also performed a more detailed regression-based analysis using seven search terms and U.S.-state-level data, and reached very similar conclusions. Results are available from the authors upon request.

³⁵Because the unemployment rate in the figure is constructed as a weighted average of the market-specific regional jobless rates, it is not exactly similar to the U.S. unemployment rates shown elsewhere in the paper.

³⁶The use of coupons varies widely across products, being particularly prevalent for categories such as carbonated beverages and cold cereal.

Our findings for coupon usage are corroborated by Nevo and Wong (2014), who use the Nielsen Homescan data set to document an increase in shopping intensity in the United States during the Great Recession. In particular, Nevo and Wong find that shopping activities increased by 1.5 to 2 percentage points, on average, and that U.S. counties that experienced larger rises in unemployment also experienced, on average, higher increases across various shopping activities: sales, coupons, generic products or large-size items.

In sum, our evidence suggests that price discrimination can take many forms, and is likely to vary across products and markets. For example, coupons are a much less prevalent price-promotion tool in the United Kingdom. While sales have been the focus of this paper due to their prominence in the discussion surrounding price dynamics, it may be worthwhile for macroeconomists to examine whether other pricing tools, such as coupons, have a significant role to play in the transmission of shocks.

7 Conclusions

Temporary price discounts (“sales”) are an important feature of the pricing behavior of retailers. The recent literature argues that sales are mostly irrelevant for macroeconomists, under the assumption that they are not significantly affected by business cycles and represent high-frequency phenomena that have little impact on the predictions of macro models. We revisit this debate and argue that the dynamics of the observed sales behavior can indeed be important for aggregate price flexibility. First, using the U.K. CPI micro data covering the period from 1996 to 2012, we demonstrate that sales (which normally would be filtered out by macroeconomists) are correlated with the business cycle: in particular, the frequency of sales is strongly countercyclical and doubled during the last recession. Analyzing aggregate time series obtained from the Bureau of Labor Statistics, we find that countercyclical sales are also a feature of price dynamics in the United States.

Second, we study the propagation of monetary shocks in an environment where sales respond to macroeconomic conditions in line with the empirical evidence. We build a general-equilibrium business cycle model with consumer search and price discrimination by retailers. Our model, calibrated to price-discounting behavior in the data, predicts that sale prices represent a significant source of price flexibility. In response to an unanticipated monetary contraction, the increase in consumer search activity and more-aggressive discounts by re-

tailers lead to a much faster decrease in the aggregate price level and a substantially smaller response of real output.

Our conclusion is that focusing on posted regular or reference prices may lead macroeconomists to miss important aspects of pricing over the business cycle. More generally, we believe that future research should be aimed at investigating the cyclical properties and aggregate implications of the numerous price-discrimination strategies used by firms. Furthermore, evidence on retailers' price-discounting behavior, combined with evidence on households' search for low prices, can provide useful insights for the debate on the driving forces of business cycles.

References

- [1] Aguiar, M. and Hurst, E. 2007. “Lifecycle Prices and Production.” *American Economic Review* 97(5), 1533–1559.
- [2] Aguiar, M., Hurst, E. and Karabarbounis, L. 2013. “Lifecycle Prices and Production.” *American Economic Review* 103(5), 1664–96.
- [3] Álvarez, L. J., Dhyne, E., Hoeberichts, M., Kwapil, C., Le Bihan, H., Lünemann, 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.
- [4] Basu, S. and Fernald, J. 1997. “Returns to Scale in U.S. Production: Estimates and Implications.” *Journal of Political Economy* 105(2), 249–83.
- [5] 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.
- [6] 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.
- [7] Berardi, N., Gautier, E. and Le Bihan, H. 2013. “More Facts about Prices: France Before and During the Great Recession,” Working papers 425, Banque de France.
- [8] Blanchard, O. and Kahn, C. 1980. “The Solution of Linear Difference Models under Rational Expectations,” *Econometrica* 48(5): 1305–1311.
- [9] Bronnenberg, B., Kruger, M. and Mela, C. 2008. “Database paper: The IRI Marketing Data Set.” *Marketing Science* 27(4), 745–748.
- [10] Butters, G. 1977. “Equilibrium distribution of sales and advertising prices.” *Review of Economic Studies* 44 (October): 465–91.
- [11] Chevalier, J. and Kashyap, A. 2012. “Best Prices,” mimeo, University of Chicago.
- [12] Christiano, L., Eichenbaum, M. and Evans, M. 2005. “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy”, *Journal of Political Economy*, 113, 1–45.

- [13] Coibion, O., Gorodnichenko, Y. and Hong, G. H. 2014. “The Cyclicity of Sales, Regular and Effective Prices: Business Cycle and Policy Implications,” mimeo, University of California, Berkeley.
- [14] Dotsey, M., King, M. and Wolman, A. 1999. “State-Dependent Pricing and The General Equilibrium Dynamics Of Money and Output,” *Quarterly Journal of Economics* 114(2): 655–690.
- [15] Eichenbaum, M., Jaimovich, N. and Rebelo, S. 2011. “Reference Prices, Costs, and Nominal Rigidities.” *American Economic Review* 101(1): 234–62.
- [16] Glandon, P. 2011. “Sales and the (Mis)Measurement of Price Level Fluctuations.” mimeo, Vanderbilt University.
- [17] Guimaraes, B. and Sheedy, K. 2011. “Sales and Monetary Policy,” *American Economic Review* 101(2): 844–76.
- [18] Inmar, Inc., 2013. “2013 Coupon Trends: 2012 Year-End Report.” <https://www.inmar.com/>.
- [19] Kaplan, G. and Menzio, G. 2014. “The Morphology of Price Dispersion,” mimeo, University of Pennsylvania.
- [20] Kehoe, P. and Midrigan, V. 2008. “Temporary price changes and the real effects of monetary policy,” Working Papers 661, Federal Reserve Bank of Minneapolis.
- [21] 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.
- [22] 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.
- [23] Klenow, P. and Willis, J. 2007. “Sticky Information and Sticky Prices,” *Journal of Monetary Economics* 54 (September): 79–99.
- [24] Krueger, A. and Mueller, A. 2010. “Job Search and Unemployment Insurance: New Evidence from Time Use Data.” *Journal of Public Economics* 94: 298–307.

- [25] Kryvtsov, O. and Midrigan, V. 2013. “Inventories, Markups, and Real Rigidities in Menu Cost Models.” *Review of Economic Studies* 80(1), 249–276.
- [26] Nakamura, E. and Steinsson, J. 2008. “Five facts about prices: a reevaluation of menu cost models.” *Quarterly Journal of Economics* 123(4), 1415.
- [27] NCH Marketing Services, Inc., 2013. “CPG Coupons: U.S. Market Analysis,” <https://www2.nchmarketing.com/ResourceCenter/>.
- [28] Nevo, A. and Wong, A. 2014. “The Elasticity of Substitution Between Time and Market Goods: Evidence from the Great Recession,” mimeo, Northwestern University, June.
- [29] Salop, S., and Stiglitz, J. 1977. “Bargains and ripoffs: A model of monopolistically competitive price dispersion.” *Review of Economic Studies* 44 (October): 493–510.
- [30] Sudo, N., Ueda, K., Watanabe, K. and Watanabe, T. 2013. “Working Less and Bargain Hunting More: Macro Implications of Sales during Japan’s Lost Decades,” mimeo, University of Tokyo.
- [31] Taylor, J. 1980. “Aggregate Dynamics and Staggered Contracts,” *Journal of Political Economy* 88(1): 1–23.
- [32] Varian, H. 1980. “A model of sales.” *American Economic Review* 70 (September): 651–59.
- [33] Vavra, J. 2014. “Inflation dynamics and time-varying volatility: new evidence and a new Ss interpretation,” *Quarterly Journal of Economics* 129(1), 215–258.
- [34] Wulfsberg, F. 2009. “Price Adjustments and Inflation: Evidence from Consumer Price Data in Norway 1975–2004,” Norges Bank WP 2009/11.




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




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




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

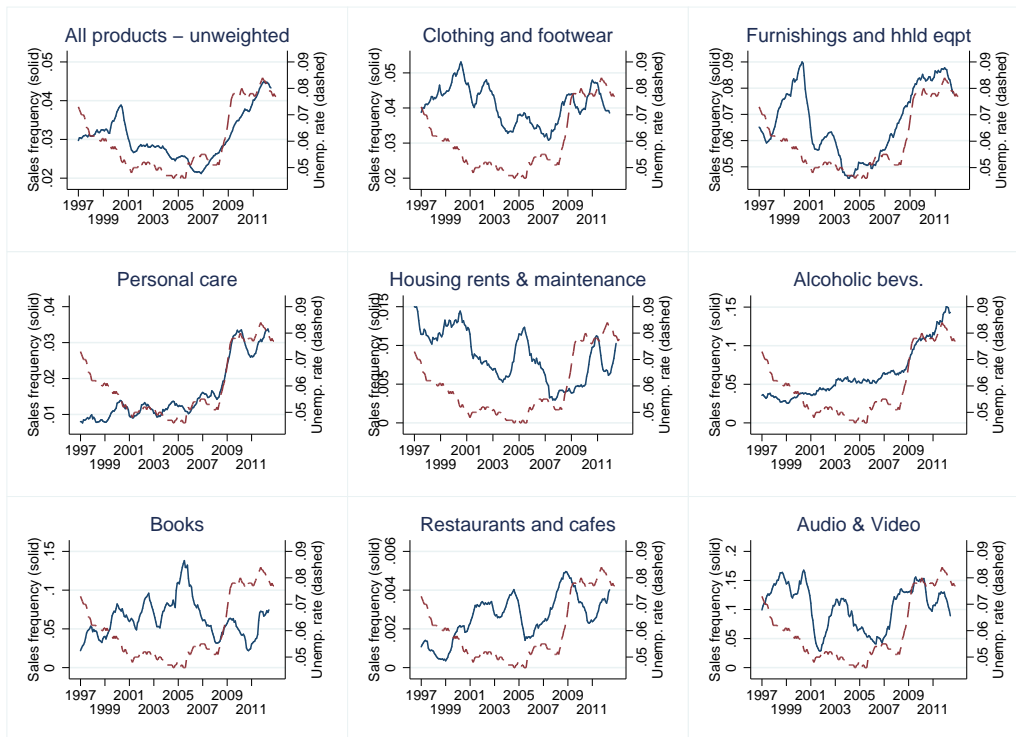


Figure 4: Frequency of sales (sales flag), selected categories




Figure 5: The evolution of the size of sales




Figure 6: Frequency of sales in the U.S. CPI data and unemployment rate




Figure 7: Profit loss and gain from 1 ppt increase in the fraction of low prices




Figure 8: Decision for high and low price levels




Figure 9: (Inverse) Distributions of the fixed cost of searching

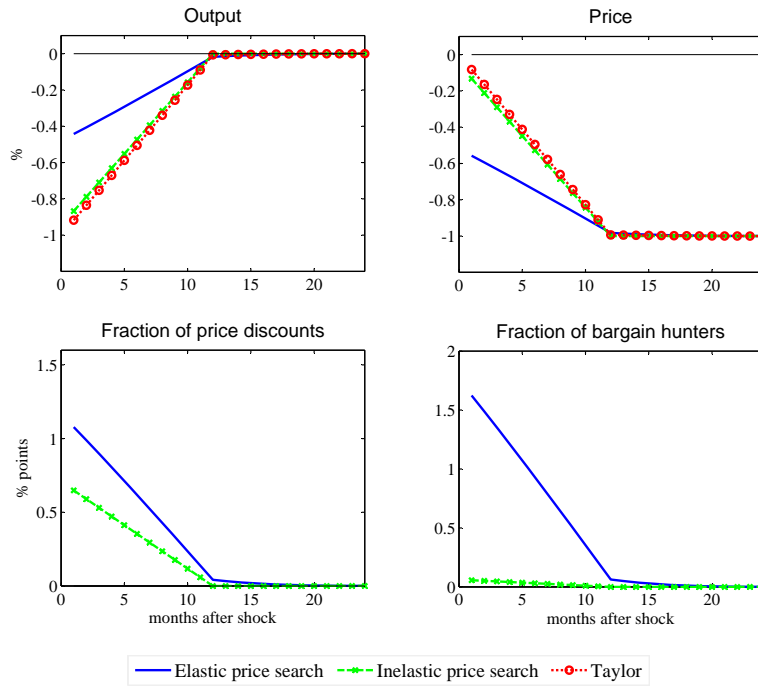


Figure 10: Responses to a negative 1% impulse to money growth




Figure 11: Responses to a negative 1% impulse to persistent money growth




Figure 12: Decomposition of aggregate price response to a negative 1% impulse to money growth

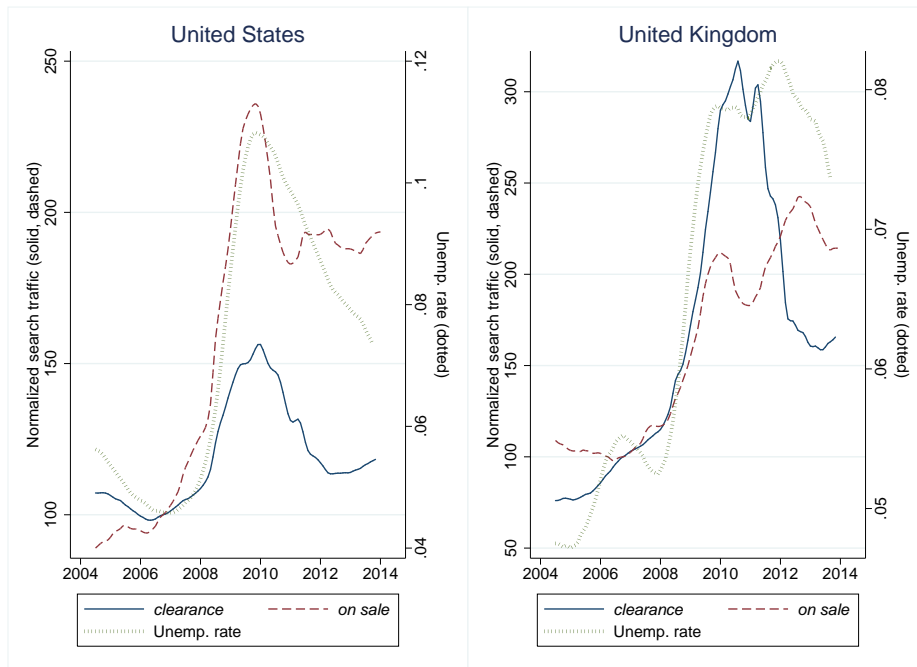


Figure 13: Google search activity for bargain hunting-related terms and unemployment rate (12-month centered moving averages)




Figure 14: Frequency of retailer coupons in the IRI Marketing scanner data set and Google search activity for term “coupons”

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
<i>All prices</i>				
	0.172	0.107	0.065	0.111
<i>Regular prices</i>				
Sales flag	0.146	0.095	0.050	0.080
V-shaped	0.121	0.084	0.037	0.087
Reference price	0.073	0.053	0.020	0.095

Table 2: Summary statistics for temporary sales

	Sales flag	V-shaped	V-shaped (alt.)	Reference price
<i>All sales</i>				
Frequency	2.6%	5.0%	3.9%	6.2%
Mean size	-22.4%	-11.9%	-11.2%	-10.4%
Median size	-20.0%	-7.0%	-6.1%	-5.6%
<i>Sales of at least 10%</i>				
Frequency	2.1%	2.0%	1.6%	2.5%
Mean size	-26.1%	-23.9%	-23.1%	-23.0%
Median size	-22.2%	-20.1%	-20.0%	-20.0%

Table 3: Panel regression results at the category level

	Sales flag					V-shaped	Ref. price
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
u_t	0.356***	0.377***	0.339***	0.257***	0.320***	0.466***	0.368***
$s_{i,t-1}$	-	-	-	0.362***	-	-	-
Month dum.	N	Y	Y	Y	Y	Y	Y
Time trend	N	N	Y	N	N	N	N
Fixed effects	Y	Y	Y	Y	Y	Y	Y
CPI weights	N	N	N	N	Y	N	N
R^2	0.022	0.049	0.052	0.175	0.041	0.041	0.034
F-stat	8.44	5.35	5.13	43.28	3.84	8.79	11.10

Table 4: Panel regression results - alternative macroeconomic indicators

	Sales flag					
Unemployment rate	0.0050***					
Cons. confidence			-0.0041***			
Fin. situation next year				-0.0047***		
Economy next year					-0.0018*	
Major purchases						-0.0041***
R^2	0.051	0.042	0.048	0.029	0.044	
F-stat	5.36	4.44	4.75	4.21	4.49	

Table 5: Panel regression results at the product level

	V-shaped			Ref. price
	(1)	(2)	(3)	(4)
u_t	0.373***	0.309***	0.405***	0.272***
$s_{i,t-1}$	-	0.202***	-	-
Month dummies	Y	Y	Y	Y
Time trend	N	N	N	N
Fixed effects	Y	Y	Y	Y
F-stat	4.40	58.01	14.03	15.73

Table 6: Parameterization (benchmark model)

A. Calibrated Parameters			B. Targets (steady state)		
			Data	Model	
θ	elast. subst. goods	5	Markup	0.25	0.25
ε_{Ξ}	search cost elasticity	1	Price discount	0.78	0.78
κ	fixed cost of posting discounts *	0.08	Fraction of discounts	0.05	0.05
z_{max}	max fixed search cost *	0.31	Δ Revenue share of sales from $\Delta\gamma=1$ ppt	2.7	2.7

* - in units of time

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}	ser corr of money shock	0
σ_{μ}	stdev of money shock impulse	0.23

Table 7: Responses to a negative 1% impulse to money growth (on impact)

	A. Baseline	B. Sticky P^L	C. Fixed γ	D. Sticky wages (household)	E. Sticky marginal cost
Avg frac of sales, ppt	1.08	1.05	0.00	0.93	0.06
Avg discount, ppt	-0.06	-0.01	1.20	0.00	0.00
Wage (households), %	-1.00	-1.00	-1.00	-0.08	-1.00
Marginal cost, %	-1.00	-1.00	-1.00	-1.00	-0.08
Consumption, %	-0.44	-0.45	-0.68	-0.50	-0.94
Frac of bargain hunters, ppt	1.62	1.63	1.17	1.36	0.09
Avg markup, ppt	0.44	0.45	0.68	0.50	0.03

Note: Responses are given in the month of the impulse

Appendix A. Equilibrium conditions for sticky-price model

Denote the following auxilliary variables:

$$\begin{aligned}\Pi_{t,j}^L &= (P_{t,j}^L - W_t) \left(\frac{P_{t,j}^L}{P_t} \right)^{-\theta} c_t, & \Pi_{t,j}^H &= (P_{t,j}^H - W_t) \left(\frac{P_{t,j}^H}{P_t} \right)^{-\theta} c_t, \\ \Pi_{t,j}^W &= \gamma_{t,j} \Pi_{t,j}^L + (1 - \gamma_{t,j}) \Pi_{t,j}^H, & \Pi_{t,j}^B &= f \gamma_{t,j} \Pi_{t,j}^L + (1 - f \gamma_{t,j}) \Pi_{t,j}^H,\end{aligned}$$

$$\begin{aligned}R_{t,j}^L &= P_{t,j}^L \left(\frac{P_{t,j}^L}{P_t} \right)^{-\theta} c_t, & R_{t,j}^H &= P_{t,j}^H \left(\frac{P_{t,j}^H}{P_t} \right)^{-\theta} c_t, \\ R_{t,j}^W &= \gamma_{t,j} R_{t,j}^L + (1 - \gamma_{t,j}) R_{t,j}^H, & R_{t,j}^B &= f \gamma_{t,j} R_{t,j}^L + (1 - f \gamma_{t,j}) R_{t,j}^H,\end{aligned}$$

$$D_{t,j}^W = \frac{\left(\gamma_{t,j} \left(\frac{P_{t,j}^L}{P_{t,j}^H} \right)^{-\theta} + 1 - \gamma_{t,j} \right)^{1-1/\theta}}{\gamma_{t,j} \left(\frac{P_{t,j}^L}{P_{t,j}^H} \right)^{1-\theta} + 1 - \gamma_{t,j}},$$

$$D_{t,j}^B = \frac{\left(f \gamma_{t,j} \left(\frac{P_{t,j}^L}{P_{t,j}^H} \right)^{-\theta} + 1 - f \gamma_{t,j} \right)^{1-1/\theta}}{f \gamma_{t,j} \left(\frac{P_{t,j}^L}{P_{t,j}^H} \right)^{1-\theta} + 1 - f \gamma_{t,j}},$$

where $D_{1t,j}^B$ and $D_{1t,j}^W$ are derivatives of $D_{t,j}^B$ and $D_{t,j}^W$ with respect to $\frac{P_{t,j}^L}{P_{t,j}^H}$; and $D_{2t,j}^B$ and $D_{2t,j}^W$ are derivatives of $D_{t,j}^B$ and $D_{t,j}^W$ with respect to $\gamma_{t,j}$.

Equilibrium consists of $5N + 5$ equations for $5N + 5$ variables $P_{t,j}^L, P_{t,j}^H, \gamma_{t,j}, \overline{\alpha_{t,j}^W}, \alpha_{t,j}, P_t, c_t, W_t, R_t, M_t$ including:

- equation for $P_{t,0}^H$

$$P_{t,0}^H = \frac{\theta}{\theta - 1} \frac{E_t \sum_{\tau=t}^{t+N-1} \beta^\tau c_\tau^{1-\sigma} P_\tau^{\theta-1} \cdot N_{\tau,\tau-t}^H W_\tau}{E_t \sum_{\tau=t}^{t+N-1} \beta^\tau c_\tau^{1-\sigma} P_\tau^{\theta-1} \cdot N_{\tau,\tau-t}^H (1 - \Delta_{\tau,\tau-t}^H)},$$

- $N - 1$ equations for $P_{t,1}^H, \dots, P_{t,N-1}^H$

$$P_{t,0}^H = P_{t+1,1}^H = \dots = P_{t+N-1,N-1}^H,$$

- N equations for $P_{t,0}^L, \dots, P_{t,N-1}^H$

$$P_{t,j}^L = \frac{\theta}{\theta - 1} \frac{1}{1 + \Delta_{t,j}^L} W_t,$$

where

$$N_{t,j}^H = (1 - \gamma_{t,j}) \overline{\alpha_{t,j}^W} + \alpha_{t,j} (1 - f\gamma_{t,j}),$$

$$N_{t,j}^L = \gamma_{t,j} \left(\overline{\alpha_{t,j}^W} + \alpha_{t,j} f \right),$$

$$\Delta_{t,j}^H = \frac{\partial \alpha_{t,j}}{\partial \ln P_{t,j}^H} \frac{\Pi_{t,j}^B}{(\theta - 1) N_{t,j}^H R_{t,j}^H},$$

$$\Delta_{t,j}^L = - \frac{\partial \alpha_{t,j}}{\partial \ln P_{t,j}^L} \frac{\Pi_{t,j}^B}{(\theta - 1) N_{t,j}^L R_{t,j}^L},$$

$$\begin{aligned} \frac{\partial \alpha_{t,j}}{\partial \ln P_{t,j}^H} &= \frac{D_{t,j}^B R_{t,j}^B}{W_t \Xi''(\alpha_{t,j})} \frac{1}{\theta - 1} \left[- \frac{D_{1t,j}^B \frac{P_{t,j}^L}{P_{t,j}^H}}{D_{t,j}^B} - (\theta - 1) \frac{(1 - f\gamma_{t,j}) R_{t,j}^H}{R_{t,j}^B} \right] \\ &\quad - \frac{D_{t,j}^W R_{t,j}^W}{W_t \Xi''(\alpha_{t,j})} \frac{1}{\theta - 1} \left[- \frac{D_{1t,j}^W \frac{P_{t,j}^L}{P_{t,j}^H}}{D_{t,j}^W} - (\theta - 1) \frac{(1 - f\gamma_{t,j}) R_{t,j}^H}{R_{t,j}^W} \right], \end{aligned}$$

$$\begin{aligned} \frac{\partial \alpha_{t,j}}{\partial \ln P_{t,j}^L} &= \frac{D_{t,j}^B R_{t,j}^B}{W_t \Xi''(\alpha_{t,j})} \frac{1}{\theta - 1} \left[- \frac{D_{1t,j}^B \frac{P_{t,j}^L}{P_{t,j}^H}}{D_{t,j}^B} - (\theta - 1) \frac{f\gamma_{t,j} R_{t,j}^L}{R_{t,j}^B} \right] \\ &\quad - \frac{D_{t,j}^W R_{t,j}^W}{W_t \Xi''(\alpha_{t,j})} \frac{1}{\theta - 1} \left[- \frac{D_{1t,j}^W \frac{P_{t,j}^L}{P_{t,j}^H}}{D_{t,j}^W} - (\theta - 1) \frac{f\gamma_{t,j} R_{t,j}^L}{R_{t,j}^W} \right], \end{aligned}$$

- N equations for the fraction of bargain hunters $\alpha_{t,j}$

$$W_t \Xi'(\alpha_{t,j}) = \left(\frac{\theta}{\theta - 1} D_{t,j}^B - 1 \right) R_{t,j}^B - \left(\frac{\theta}{\theta - 1} D_{t,j}^W - 1 \right) R_{t,j}^W,$$

- N equations for the fraction of low prices $\gamma_{t,j}$

$$\left(\overline{\alpha_{t,j}^W} + f\alpha_{t,j} \right) [\Pi_{t,j}^H - \Pi_{t,j}^L] + \kappa W_t = \frac{\partial \alpha_{t,j}}{\partial \gamma_{t,j}} \Pi_{t,j}^B,$$

where

$$\frac{\partial \alpha_{t,j}}{\partial \gamma_{t,j}} = \frac{1}{W_t \Xi''(\alpha_{t,j})} \frac{f}{\theta - 1} [D_{2t,j}^B R_{t,j}^B + D_{t,j}^B (R_{t,j}^L - R_{t,j}^H)] - \frac{1}{W_t \Xi''(\alpha_{t,j})} \frac{1}{\theta - 1} [D_{2t,j}^W R_{t,j}^W + D_{t,j}^W (R_{t,j}^L - R_{t,j}^H)] ,$$

- N equations for the fraction of workers of variety j

$$\overline{\alpha_{t,j}^W} = 1 - \alpha_{t,j} ,$$

- equation for the price of aggregate consumption P_t

$$P_t = \left[\begin{array}{l} \frac{1}{N} \sum_{j=0}^{N-1} \alpha_{t,j} \left[f \gamma_{t,j} (P_{t,j}^L)^{1-\theta} + (1 - f \gamma_{t,j}) (P_{t,j}^H)^{1-\theta} \right] \\ + \frac{1}{N} \sum_{j=0}^{N-1} \overline{\alpha_{t,j}^W} \left[\gamma_{t,j} (P_{t,j}^L)^{1-\theta} + (1 - \gamma_{t,j}) (P_{t,j}^H)^{1-\theta} \right] \end{array} \right]^{\frac{1}{1-\theta}} ,$$

- the cash-in-advance constraint

$$M_t = P_t c_t ,$$

- a labor-market condition

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

- equation for the risk-free rate

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

- and an exogenous AR(1) process (in logs) for money growth $\mu_t = \frac{M_t}{M_{t-1}}$:

$$\ln \mu_{t+1} = (1 - \rho_\mu) \ln \mu + \rho_\mu \ln \mu_t + \varepsilon_{\mu t} ,$$

where i.i.d. innovations $\varepsilon_{\mu t}$ drawn from $N(0, \sigma_\mu^2)$.

Appendix B. Fixed search cost distributions

We utilize a flexible two-parameter fixed cost CDF specification, borrowed from Dotsey, King and Wolman (1999), who use it in their state-dependent pricing model:

$$\Xi'(\alpha) = z_{\max} \frac{\arctan\left(\frac{\alpha-c_1}{c_2}\right) - \arctan\left(\frac{-c_1}{c_2}\right)}{\arctan\left(\frac{1-c_1}{c_2}\right) - \arctan\left(\frac{-c_1}{c_2}\right)},$$

where $\Xi'(0) = 0$ and $\Xi'(1) = z_{\max}$. This implies

$$\begin{aligned}\Xi''(\alpha) &= \frac{1}{c_2 c_3} \frac{1}{1 + \left(\frac{\alpha-c_1}{c_2}\right)^2}, \\ [\Xi'(\alpha)]^{-1} &= G(z) = c_1 + c_2 \tan(c_3 \alpha - c_4), \\ \Xi'''(\alpha) &= -\frac{1}{c_2^2 c_3} \frac{2 \left(\frac{\alpha-c_1}{c_2}\right)}{\left[1 + \left(\frac{\alpha-c_1}{c_2}\right)^2\right]^2},\end{aligned}$$

where

$$\begin{aligned}c_3 &= \frac{1}{z_{\max}} \left(\arctan\left(\frac{1-c_1}{c_2}\right) - \arctan\left(\frac{-c_1}{c_2}\right) \right), \\ c_4 &= -\arctan\left(\frac{-c_1}{c_2}\right).\end{aligned}$$

In the model with inelastic search, we use the following parameter values: $c_1 = 0.107$, $c_2 = 10^{-4}$ and $z_{\max} = 1$.