

The Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications

Olivier Coibion
UT Austin
and NBER

Yuriy Gorodnichenko
U.C. Berkeley
and NBER

Gee Hee Hong
Bank of Canada

This Draft: June 8th, 2014

Abstract: We study the cyclical properties of sales, regular price changes and average prices paid by consumers (“effective” prices) using data on prices and quantities sold for numerous retailers across many U.S. metropolitan areas. Inflation in the effective prices paid by consumers declines significantly with higher unemployment while little change occurs in the inflation rate of prices posted by retailers. This difference reflects the reallocation of household expenditures across retailers, a feature of the data which we document and quantify, rather than sales. We propose a simple model with household store-switching and assess its implications for business cycles and policymakers.

Keywords: Sales, Price Changes, Store-Switching, Inflation Measurement.

JEL codes: E3, E4, E5.

We are grateful to Yury Yatsynovich for excellent research assistance, Martin Eichenbaum, Nir Jaimovich, Anil Kashyap, Joe Vavra and seminar participants at UC Berkeley, Booth, Cornell-Penn State Macro Workshop, Duke, Federal Reserve Board, John Hopkins, Maryland, Michigan, NBER Summer Institute, New Economic School, Ohio State, UT Austin, Zurich University, and Central European University for helpful comments. The views expressed in this Working Paper are those of the authors and do not necessarily represent those of the Bank of Canada, or any other institution with which the authors are affiliated.

I Introduction

Explaining the apparent non-neutrality of money has led macroeconomists to study a variety of frictions, such as nominal and real wage rigidities or information rigidities. But the most commonly emphasized potential source of monetary non-neutrality remains “sticky prices” as epitomized by Woodford (2003). In part, this likely reflects the ubiquitousness of sticky prices in daily life. For example, Starbucks has raised its brewed coffee prices only once per year since the 2007-2009 recession began, despite the fact that the spot price of (Robusta) coffee beans rose 50% between February of 2007 and February of 2008, then fell 33% over the next 24 months, before again rising 50% by February 2011.¹ This annual frequency of updating prices is common, and the infrequency of many price changes has been well-documented in the literature.²

Yet while firms may choose to change their prices infrequently, this need not imply that the “effective” prices actually paid by households are themselves sticky. Chevalier and Kashyap (2011), for example, argue that if households respond strongly to sales, then “effective” price flexibility due to consumers reallocating their expenditures across goods or time could potentially undo much of the macroeconomic effects of the underlying price rigidities commonly observed in regular prices. While a significant body of existing work quantifies how the treatment of sales affects the measured degree of price rigidity in posted prices (e.g. Bils and Klenow 2004, Nakamura and Steinsson 2008, Eichenbaum et al. 2011, and Kehoe and Midrigan 2012), evidence on the extent to which sales prices affect the effective prices paid by households remains limited. In large part, this shortcoming reflects data limitations: measuring effective prices paid by households requires data on both quantities and prices, whereas most data-sets include only the latter.

Using a panel dataset of both prices and quantities sold at the universal product code (UPC) level across different stores in 50 U.S. metropolitan areas from 2001 to 2011, we build on this literature by studying the *cyclical* in both the prices posted by retailers as well as the effective prices actually paid by consumers. Consistent with Gali and Gertler (1999), Williams (2006), Roberts (2006) and others documenting the lack of a strong relationship between inflation and economic activity in U.S. macroeconomic data, we find little cyclical sensitivity in the inflation rate of prices posted by retailers. In contrast, and consistent with the notion of significant consumer reallocation of expenditures in response to economic conditions, we document that effective price inflation is indeed more cyclically sensitive than inflation in posted prices. The difference is also quantitatively large: a 2% point rise in the unemployment rate lowers inflation in effective prices by 0.2-0.3% points relative to inflation in posted prices for a given UPC. To assess whether this sensitivity of effective prices paid by households is driven by sales, we also measure the cyclical of regular and sales price changes, as well as the proportion of a given good bought on sale. We find little evidence of cyclical in the properties of either regular price changes or in the frequency of sales, nor do we find that the share of goods bought on sale rises

¹ See “Starbucks to Raise Prices” in the January 4th, 2012 edition of the Wall Street Journal.

² See e.g. Carlton (1986), Cecchetti (1986), and Kashyap (1995).

with unemployment. Thus, the greater flexibility in prices paid by households relative to the prices charged by firms does not stem from the cyclical response of sales.

If sales are not the source of the observed flexibility in the prices actually paid by households for a given good in a given metropolitan area, what is? We argue that the discrepancy between the cyclical changes in posted and effective prices reflects consumers reallocating their expenditures across stores in response to economic conditions. Intuitively, with some dispersion of prices across stores in any given period, deteriorating local economic conditions should lead price-sensitive consumers to reallocate some of their consumption expenditures toward low-price retailers, thereby lowering the average price paid for any given good.

We document three pieces of evidence consistent with this mechanism. First, within each metropolitan area and each UPC, we find that the share of goods purchased below a threshold of the price distribution rises with the local unemployment rate, i.e. households purchase a larger share of each good at lower price locations. Second, because we have data for a variety of identical goods sold across different stores within a given geographic area, we can quantify the extent to which some stores are systematically more expensive than others across all goods. Combined with measures of quantities purchased, we find that the share of expenditures within a category and market spent at high-price stores relative to low-price stores declines when the local unemployment rises. In other words, higher-price retailers experience relatively larger declines in expenditures during downturns than lower-price retailers. Third, we exploit a detailed panel dataset tracking *individual* households' expenditures at the UPC level for each store. Given that we can characterize the degree to which some stores are more expensive than others, these data allow us to quantify the extent to which individual households reallocate their expenditures across different retailers in response to economic conditions. We find robust evidence that individual households do indeed reallocate their consumption expenditures toward low-price retailers when local economic conditions deteriorate.

In short, our empirical results point to significantly more flexibility in the prices paid by households than in the prices charged by firms, but this flexibility appears to be driven by the reallocation of household expenditures across retailers rather than by sales. While previous work has considered the macroeconomic consequences of effective price flexibility due to sales pricing by retailers, the potential implications of retail-switching remain unexplored. As a result, we build on the effective price literature by integrating store-switching by households into a basic New Keynesian model. In this model, two retailers each purchase intermediate goods from monopolistic competitors, who are subject to sticky prices, and convert these goods into an identical final good. One "local" retailer charges a premium over the "discount" retailer. However, household expenditures at the latter are subject to iceberg costs which can be reduced by time-intensive shopping on the part of households. When the return to shopping effort is diminishing, shopping effort on the part of households will be countercyclical. As a result, periods of low employment are also periods of high shopping effort activity. This effort leads to lower effective prices at the "discount" retailer and households therefore reallocate more of their expenditures toward this retailer

when employment is low. Consistent with the data, when the level of economic activity is high, time-intensive shopping effort and expenditure reallocation lead to effective prices being high relative to posted prices and high-price retailers of otherwise identical goods accounting for a larger share of total retail revenues when the level of economic activity is high.

In the model, the effective prices paid by households are more flexible than the underlying posted prices, as suggested by Chevalier and Kashyap (2011). However, this effective price flexibility stems not from the presence of sales but from the reallocation of household expenditures across retailers and the time-variation in shopping effort. Higher elasticities of store-switching reduce the degree of monetary non-neutralities in the model as effective prices become increasingly flexible. However, when the degree of store-switching is calibrated to levels consistent with micro-level data, the reduction in monetary non-neutrality is relatively small so that the response of output to a monetary policy shock looks much more like that of a typical New Keynesian model with high levels of price stickiness than one with low levels of price-stickiness. Despite significant flexibility in the effective prices faced by households, the dynamics of the model are dominated by the underlying level of rigidity in posted prices. In this sense, our results are similar in spirit to those of Eichenbaum et al. (2011), Guimaraes and Sheedy (2011) and Kehoe and Midrigan (2012) but we emphasize the role of the reallocation of expenditures across stores by households (rather than sales) as the source of effective price flexibility.

Despite the fact that the quantitative implications of store-switching are limited in terms of the degree of monetary non-neutrality, the presence of shopping effort and store-switching does have several novel business cycle and policy implications. For example, the presence of shopping effort and store-switching alters the *effective* Frisch labor supply elasticity and the *effective* elasticity of the output gap with respect to real interest rates, when the latter are expressed in terms of input prices rather than final consumption goods prices. For the labor supply elasticity, when the utility return to labor is high, then the return to shopping effort must also be high since the marginal disutilities of labor and shopping effort are identical. High returns to shopping effort require low shopping intensity under diminishing returns. But if less time is spent on shopping, then the marginal disutility of labor is low so households will be willing to supply more labor. For the sensitivity of the output gap to real interest rate changes expressed in terms of posted prices, if real interest rates rise, then output will fall but be expected to rise via the dynamic IS curve. With countercyclical shopping effort, the time spent on shopping will therefore be expected to rise thereby lowering expected inflation in final goods prices. The latter will lead to relatively more consumption and output, thereby mitigating the initial decline in output.

The presence of store-switching and shopping effort also leads to new insights into some traditional measurement issues regarding inflation. For example, the Boskin Commission Report (1996) and Shapiro and Wilcox (1996) discuss the potential for store-substitution to cause biased estimates of inflation. Reinsdorf (1993), Hausman and Leibtag (2007), Diewert et al. (2009), and Greenlees and McClelland (2011) try to assess the average inflation bias in the U.S. from store-switching, but focus primarily on the entry of new potentially low-

price outlets and the fact that these entrants will only gradually be included in the CPI sample of stores. Triplett (2003) also highlights that the final consumption price index should incorporate retail-switching behavior on the part of households as well as the related costs faced by households due to shopping effort. One contribution of the paper is to show that the “effective” price index in our model, which uses time-varying expenditure weights across retailers, closely tracks the final consumption price index despite not directly including the household costs stemming from shopping effort. This suggests that statistical agencies could more precisely measure the cost-of-living by tracking actual prices paid by consumers rather than those charged by retail outlets.

We also show that shopping effort and store-switching have policy implications above and beyond the measurement of inflation. First, the loss function derived from the household’s expected utility over consumption and leisure displays a larger relative weight on output gap stabilization. This primarily reflects a reduced welfare cost of inflation volatility in the presence of shopping effort. Inflation volatility implies higher price dispersion which requires more labor, and therefore less leisure, due to the convexity of labor supply. However, countercyclical shopping effort implies reduced shopping intensity when hours worked are high, thereby mitigating the decrease in leisure and the welfare costs of inflation volatility. Also, we consider whether welfare would be enhanced if central banks responded to effective price measures rather than posted prices. The results vary with the policy regime: an inflation-targeting central bank is generally well-served by responding to posted price inflation, but under price-level targeting regimes higher welfare can be obtained by targeting deviations of effective prices from a target level than from targeting traditional posted price gaps. This again suggests that devoting more resources to constructing aggregate measures of effective prices *paid by households* could lead to improved policy rules and therefore greater economic stability and welfare.

The structure of the paper is as follows. Section 2 discusses the data used for the empirical analysis. Section 3 presents baseline empirical results on the cyclicity of effective and posted prices, as well as supporting evidence on the cyclicity of consumer store-switching. Section 4 develops the model with store-switching and shopping effort and derives policy and business cycle implications. Section 5 concludes.

II Data Description

We use an extensive dataset from Symphony IRI, a marketing and market research agency,³ which is discussed in more detail in Bronnenberg et al. (2008). This dataset contains weekly scanner price and quantity data covering a panel of stores in 50 metropolitan areas and 31 product categories from January 2001 to December 2011, with multiple chains of retailers for each market. The dataset also includes a panel of households in two of the metropolitan areas with detailed information on household characteristics and their purchases in these retailers. This household dimension can be linked to the store-level information of the dataset.

³ Information Resources, Inc. (“IRI”) has changed its name to SymphonyIRI Group, Inc. All estimates and analyses in this paper based on SymphonyIRI Group, Inc. data are by the author(s) and not by SymphonyIRI Group, Inc.

The metropolitan areas are typically defined at the Metropolitan Statistical Area (MSA) level. Two of the metropolitan areas are smaller than usual (Eau Claire, WI and Pittsfield, MA) but the household panel data were constructed only for these two areas. Within each metropolitan area, the dataset includes price and quantity data from a panel of stores. For example, data from San Francisco in 2005 covers 58 different retailers. Each outlet has a time-invariant identifier so one can track their prices and revenues over time. Retailers are not identified by name, but a well-known feature of these data is that Wal-Mart is not included.

For each retail outlet, weekly data are available at the UPC level. Goods are classified into 31 general product categories (e.g., beer, coffee) as well as more refined categories. Brand information is included (e.g., Budweiser, Heineken) but all private-label UPCs have the same brand identification so that the identity of the retailer cannot be recovered from the labeling information. We exclude private-label UPCs from our analysis because IRI changed the coding of these goods in 2007 and 2008 and, as a result, we could not construct consistent pricing series over time. Detailed information about each good is included (e.g., low-fat) as is information about the volume of the product (e.g., 6-packs vs. 12-packs, volume per container). Unfortunately, no data on costs are available.

Retailers report the total dollar value of weekly sales (TR), inclusive of retail features, displays and retailer coupons but not manufacturer coupons, for each UPC code as well as total units sold (TQ). The combination of the two yields the average retail price during that week:

$$P_{msctj} = \frac{TR_{msctj}}{TQ_{msctj}} \quad (2.1)$$

where m , s , c , t , and j index markets, stores, product categories, time, and UPC respectively. We refer to this measure as “posted” prices. We also compute good-specific monthly “posted” price inflation rates as $\log(P_{msctj,t}/P_{msctj,t-1})$. For both posted prices and posted price inflation, we then aggregate across all goods and stores within each product category using one of three weighting schemes: *i*) equal weights; *ii*) expenditure shares for each market and year (“market specific”); *iii*) cross-market expenditure shares for each year (“common”).⁴ Finally, we cumulate monthly inflation rates into annual inflation rates $\bar{\pi}_{mct}^{post}$.

Because the dataset also includes quantities sold, we can construct the quantity-weighted average price paid by consumers (“effective” price) for each good across all stores in a given metropolitan area as

$$\bar{P}_{mcjt}^{eff} = \frac{\sum_{s \in m} TR_{msctj}}{\sum_{s \in m} TQ_{msctj}} \quad (2.2)$$

This measure can change because individual prices change or because consumers reallocate their purchases of the good across stores. As with posted prices, we construct good-specific monthly inflation rates, then aggregate

⁴ In constructing fixed expenditure-weights across stores for a given UPC, we use the average share of revenues for each store over the course of that calendar year. Thus, our measure of posted price inflation updates the weights applied to each retailer for each UPC at the annual frequency. In contrast, the Bureau of Labor Statistics (BLS) updates retailer weights once every five years, so our estimates likely understate the difference between posted and effective prices relative to BLS procedures.

across all UPCs within each product category using equal, market-specific, or common weights. Because these weights are held fixed within each calendar year, our measure of “effective” price inflation does not include substitution across goods. Finally, we cumulate monthly inflation rates into the annual inflation rate $\bar{\pi}_{mct}^{eff}$.

We can also decompose posted price inflation in terms of “regular” price changes and “sales” price changes, given that retailers in the dataset flag goods on sale. Specifically, we adopt the following conventions to distinguish between regular and sales price changes. First, any change in prices between two periods when neither period has a sales flag is defined as a “regular” price change if the price difference exceeds 1% in value. Second, when a sales flag is listed and the price after the sale expires is the same as prior to the sale, we assume no change in regular prices in intermittent periods. Third, in addition to the sales flag provided in the dataset, we apply a sales filter similar to that in Nakamura and Steinsson (2008). Specifically, we consider a good on sale if a price reduction is followed by a price increase of the same magnitude within four weeks. The size of a sale is measured as the change in the log price from the time a price quote switches from a “no sale” flag to an “on sale” flag. Since sales are temporary price cuts, the size of sales is measured in negative values. For example a sale of 20% is recorded as -0.2. Appendix E provides more details on how pricing moments are constructed. The frequency of sales and regular price changes are computed at the individual UPC level each month and averaged across all UPCs within a category in a given market and month. To aggregate across individual UPCs within a product category, we again use three weighting schemes: equal weights, market-specific, and common. The average size of sales is approximately 25% across all goods. While less than 20% of goods are on sale in any given week, almost 25% of all goods purchased are bought on sale.⁵ The average weekly frequency of regular price changes is 5%, so that the average duration between regular price changes is approximately 5 months. While higher than the frequency found by Nakamura and Steinsson (2008), our data is sampled at the weekly frequency which can lead to higher estimates of the frequency of price changes than when data is sampled at the monthly frequency, as documented in Kehoe and Midrigan (2012). We also construct a UPC-level measure of the share of goods bought below a price threshold (bottom 20%) at regular prices within each metropolitan area and month. Appendix Tables 1 and 2 present values of each statistic for different categories of goods in the data.

III The Cyclicity of Effective, Regular and Sales Price Changes and Consumer Store-Switching Behaviors

While a large literature exists on measuring and quantifying different forms of price changes, we focus on the *cyclicity* of price changes and quantities purchased as well as posted-price and effective-price inflation rates. In addition, we examine reallocation of consumer spending across retailers over the data period.

⁵ This share of goods bought on sales is calculated as follows. The denominator is the total number of units of a UPC sold across all stores in a given market and month. The numerator is the total number of units of this UPC sold on sale across all stores in a given market and month.

3.1 The Cyclicalities of Posted and Effective Prices

To assess the cyclicalities of price changes with respect to economic conditions, we adopt the following baseline empirical specification:

$$Y_{mct} = \beta UR_{mt} + \lambda_t + \theta_{mc} + error \quad (3.1)$$

where m , c , and t index markets (e.g., Atlanta, Detroit), the category of the good (e.g., beer, coffee), calendar time (i.e., month); Y_{mct} is a variable of interest (e.g., effective inflation rate); UR_{mt} is the local seasonally-adjusted unemployment rate;⁶ θ_{mc} denotes the fixed effect for each market and category of good while λ_t denotes time fixed effects. Because the unemployment rate at the metropolitan level is only available at the monthly frequency, we estimate (3.1) at the monthly frequency. Since the error term in (3.1) is likely to be serially and cross-sectionally correlated, we use Driscoll and Kraay (1998) standard errors.⁷

When time fixed effects are included, estimates of β in (3.1) assess the strength of correlations between *local* business conditions and various pricing moments and thus are informative about cyclical properties of these moments. One may also entertain a causal interpretation of the estimates. Because most goods sold in stores are not produced locally, specification (3.1) should not suffer from endogeneity issues typically associated with regressions of prices on real economic activity. For example, unobserved productivity innovations for a specific product, like razors, are unlikely to be correlated with local unemployment rates. While aggregate shocks could lead to simultaneous movements in prices of goods and local economic conditions, controlling for time fixed effects should eliminate this endogeneity issue. Hence, a causal interpretation of β in (3.1) can stem from the fact that almost all products sold by retailers will be produced in other geographic areas so that local variation in unemployment will serve as a proxy for shocks to the local demand for consumer goods. One can also estimate (3.1) using detrended series rather than time fixed effects. For example, linearly detrending each series can help address the possibility of spurious correlations due to long-run trends in pricing moments and economic conditions without purging the data of aggregate cyclical effects. This can be useful in verifying that any identified cyclical behavior with respect to local economic conditions is also found with respect to aggregate business cycles. However, causal interpretations are potentially more problematic in this setting, so we focus on specifications with time fixed effects but verify that our results are robust to using detrended variables.

Given our measures of annual posted-price and effective-price inflation rates at the category/market level, we estimate their cyclical sensitivity to local economic conditions using equation (3.1). The results are presented in Table 1. Focusing on the results with both market-category and time fixed effects, we find little cyclical sensitivity of posted-price inflation measures to local economic conditions. This near acyclicity in posted price inflation to economic conditions is consistent with Gali and Gertler (1999), Williams (2006),

⁶ We find similar results when we use alternative measures of local business conditions, e.g., local employment statistics constructed by the BLS. We found that, wherever relevant, seasonal adjustment of the dependent variables in specification (3.1) makes no material difference for our results.

⁷ Driscoll-Kraay standard errors tend to be conservative. See Appendix Table 3 for other estimates of standard errors.

Roberts (2006) and others documenting the lack of a strong negative relationship between inflation and real economic activity in U.S. macroeconomic data, albeit that our results are at a microeconomic level. However, the results for inflation in effective prices are quite different: the inflation rate of prices actually paid by consumers drops more strongly when local economic conditions worsen. The difference in point estimates is statistically significant for all specifications which control for time fixed effects or time trends, and similar results are obtained when we weight categories and cities by their share in aggregate sales (columns 7-8). Furthermore, the difference between the two sets of coefficients is large: a 2% point increase in the local unemployment rate lowers the effective-price inflation rate by 0.2-0.3% at an annual rate below the inflation rate of posted prices. Hence, *the key finding is that effective prices paid by households are significantly more flexible and sensitive to economic conditions than the underlying prices charged by retailers.*

This differential cyclicity in effective-price and posted-price inflation could reflect a number of factors. For example, it could arise if firms rely more on sales during periods of low demand and households respond strongly to sales making effective-price inflation fall relative to posted-price inflation even if posted-price inflation varies little over the cycle. We assess this explanation in the next section by examining the cyclicity in frequency and size of sales versus the cyclicity in frequency and size of regular price changes. Another source could be that households use sales disproportionately more when economic conditions worsen, i.e. stocking up on goods when they are on sale. This could again generate greater cyclicity in effective prices than posted prices. We examine this explanation in the next section by examining the cyclicity of the share of goods bought on sale. A third channel is if households reallocate their expenditures across retailers, thereby generating variation in effective prices even if posted prices do not change. We examine this channel in detail in section 3.3.⁸

3.2 The Cyclicity of Sales and Regular Price Changes

The greater flexibility in the effective prices paid by households relative to those charged by firms is consistent with the logic of Chevalier and Kashyap (2011) in which the reallocation of consumer expenditures leads not just to greater effective price flexibility for households but also ultimately to diminished monetary non-neutralities. To investigate the source of this flexibility, we now consider whether sales are the key driver of this behavior. We estimate (3.1) using the frequency and size of sales for each UPC aggregated to the category/market level, as well as the share of monthly expenditures for each UPC done on sale, again aggregated to the category/market level. For comparison, we also assess the cyclicity of regular price changes, both in terms of their frequency and size. Since posted price inflation depends on both the frequency and size of price changes, this can be interpreted as decomposing the results for posted inflation rates.

⁸ Another channel could be the use of coupons, rebates and loyalty programs by consumers. Information about coupon and other rebate use is not available in our data, but some of these discounts are included in reported total revenues. As a result, measures of posted price inflation already incorporate some (but not all) of the usage of such rebates.

The first column of Table 2 documents results from estimating (3.1) at the category/market level using a simple average across all UPC products within a category, *excluding* all fixed effects. The results indicate that a higher local unemployment rate is associated with more frequent but smaller sales, and a larger share of goods bought on sale overall. However, it is unclear whether this larger role played by sales reflects the fact that sales become more prevalent when *i*) the unemployment rises (i.e. business cycle effects), *ii*) regions with higher unemployment rates on average also experience more frequent sales for other reasons (i.e. systematically more depressed areas may have on average more frequent sales), or *iii*) there is a comovement of trends in unemployment and properties of sales (e.g. the frequency of sales in our data was rising at a decelerating rate over 2001-2011). As a result, columns (2)-(8) present equivalent results controlling for geographic/category specific effects and/or time fixed effects as well as results for linearly-detrended series to address *ii*) and *iii*). Controlling for time fixed effects or linear time trends (columns 3-8) eliminates much of the cyclicity in sales: when one controls for time fixed effects, both the frequency of sales and the share of goods bought on sale see their coefficients fall and become insignificantly different from zero. Only the size of sales remains significant, with sales becoming smaller when the unemployment rate rises. Thus, sales pricing appears to be largely acyclical, or mildly procyclical if we focus on the size of sales. The share of goods bought on sale is similarly acyclical, suggesting that consumers do not disproportionately take advantage of sales when local unemployment rates are high.

Table 2 also presents results for regular price changes. After controlling for time fixed effects or time trends, there is only tentative evidence that either the frequency or size of regular price changes varies with local economic conditions. Weighted regressions point toward slightly reduced frequencies of regular price changes, coming largely from less frequent price increases, and slightly smaller average price changes. But these effects, even when statistically significant, are tiny in economic terms. In short, these results suggest that the properties of regular price changes are close to acyclical: neither the frequency nor size of these types of price changes varies in an economically important manner with local economic conditions.

3.3 Expenditure Reallocation across Retailers

If posted prices display little cyclical sensitivity and households do not stock up on sales to a greater extent in worse local economic conditions, what drives the cyclical sensitivity of effective prices at the level of the UPC? One potential channel is that consumers reallocate their expenditures toward lower-price retailers when economic conditions deteriorate, thereby driving down the average price they pay for any given good. Such a channel could reflect cyclical variation in the shopping effort of households (i.e. being more willing to drive out of their way when the opportunity cost of time is low) or non-homotheticities in preferences (i.e. more expensive stores provide a higher-quality shopping experience). We assess the evidence for store-switching in three ways. First, we check whether the share of units of a given UPC sold at low prices relative to those at high prices rises

with local unemployment rates. Second, we document that the share of total revenues across all goods going to high-price retailers declines when local unemployment rates rise. Third, we show using household-level information that individuals directly reallocate their expenditures toward lower-price retailers when local economic conditions worsen.

To document the reallocation of expenditures across retailers, we first construct the share of all units of a given UPC sold at regular prices in the bottom 20% of the cross-sectional distribution of all regular prices that month across all retailers in the metropolitan area.⁹ We then assess the cyclicity of this share using specification (3.1) and display results in the bottom row of Table 2. Across all specifications, we find that when local economic conditions deteriorate, a larger share of any given UPC is, on average, purchased at lower prices relative to the distribution of all regular prices in the metropolitan area. In other words, consumers seem to reallocate their expenditures toward lower-price retailers, at least at the level of an individual good. This reallocation can therefore potentially account for the cyclical sensitivity of effective prices even in the absence of cyclicity in posted prices.

To assess at the store-level the reallocation of purchases toward low-price stores when economic conditions are worse, we construct a time-varying measure of stores' relative prices as follows. First, for each UPC-level good j in category c and market m , we measure the log-difference (relative price) between the regular price of good j in store s and the median regular price for good j across all stores in market m that month. We then compute the average relative price for a store across the set of UPC products Ω . We consider several versions of Ω because different stores sell different goods: *i*) Ω includes all UPCs sold in every store in a given market (Ω_{max}); *ii*) Ω only includes UPCs sold in at least 90% of stores in a market (Ω_{90}); *iii*) Ω only includes UPCs sold in at least 75% of stores (Ω_{75}). The average relative price of a store for a given Ω is $\bar{R}_{mst,\Omega} = \sum_c \sum_{j \in \Omega} \omega_{mcstj} R_{mcstj}$ where ω is a weight (equal, market-specific, or common). The resulting measure captures how far a store's average price level is from the median price level in a given market and month.

Using $\bar{R}_{mst,\Omega}$, we run a series of store-level regressions:

$$Y_{mst} = q_{sm} + \alpha_1 UR_{mt} + \alpha_2 UR_{mt} \times \bar{R}_{mst,\Omega} + \alpha_3 \bar{R}_{mst,\Omega} + \lambda_t + error \quad (3.2)$$

where Y_{mst} is a moment considered for store s in market m at time t , q_{sm} is the store s fixed effect in market m , and λ_t is the time fixed effect. In Table 3, we first report results using the share of sales received by store s in market m in month t as the dependent variable. The aggregation of the price measures across goods is done via expenditure shares of each UPC product across all stores. We find significant interaction effects with a store's relative price: high-price stores experience losses in revenues relative to low-price stores when unemployment rates rise, consistent with the reallocation of expenditures toward low-price stores in worse economic conditions.

⁹ In Appendix Table 4, we document that our results are insensitive to the specific threshold used.

Table 3 also presents store-level results for pricing statistics, interacted with the store’s relative price. Most of the interaction effects are insignificant, indicating that there are few differences in the pricing behaviors of high-price and low-price retailers in response to economic conditions. This implies that the relative rise in revenues of low-price stores during downturns does not reflect very different pricing strategies across retailers but instead likely reflects reallocation of expenditures across retailers by consumers. One difference is that the decreased size of sales in the face of higher unemployment is driven by high-price stores, which likely reflects their decreased incentive to attract price-sensitive consumers if the latter are more likely to shop at low-price retailers when economic conditions deteriorate. High-price stores also have slightly more frequent regular price changes when unemployment is higher, although this reflects a higher frequency of both price increases and price decreases, and the economic magnitudes are very small.

Anecdotal evidence from stock prices is also consistent with significant store-switching across differently priced retailers and their performance during the recent recession. For example, Figure 1 shows that from December 1st, 2007 to June 29th, 2009 (the duration of the recent recession as dated by the NBER), the stock price of Whole Foods—a specialty and high-price retailer—declined by 87 (log) percent whereas that of Walmart fell only 2.5 percent. Similarly, Family Dollar—a prominent low-priced retailer—saw its stock price rise by about 40 percent over the course of the recession. For comparison, the stock price of Safeway, a typical grocery chain, declined 59 percent, an almost identical amount as the 52 percent decline in the S&P 500 index over this same period. Thus, financial market valuations of these retailers also reflected sharp cyclical differences, at least during the recent downturn, consistent with the reallocation of expenditures by households toward low-price retailers in times of economic duress.

We also consider an alternative and more direct approach to quantifying the extent of consumers’ expenditure reallocation across stores. The household panel data from Symphony IRI tracks between 5,000 and 10,000 households in Eau Claire, WI and Pittsfield, MA from 2001 to 2011. About 1,500 households are continuously present between 2001 and 2011. During this time, households’ expenditures on each UPC product were tracked, including the location of each purchase. These detailed data therefore allow us to directly measure the extent of the store-switching phenomenon *at the level of individual households*.

Specifically, we construct an average relative price of stores at which household h shopped in month t and market m (either Eau Claire, WI or Pittsfield, MA)

$$\tilde{R}_{hmt} = \sum_{s \in m} \psi_{mhst} \bar{R}_{mst, \Omega} \quad (3.3)$$

where $\bar{R}_{mst, \Omega}$ is the average relative price of store s in market m at time t across the set of goods Ω as defined before and ψ_{mhst} denotes the share of the household’s expenditures spent at store s in market m and month t . A low (high) value of \tilde{R}_{hmt} means that household h made purchases in low (high) price stores in month t . Using \tilde{R}_{hmt} therefore provides a way to quantify the extent to which each household is reallocating their expenditures across retailers of different average price levels.

To assess whether individual households reallocate their expenditures across stores as local economic conditions change, we estimate the following specification across households in the two markets

$$\tilde{R}_{hmt} = q_{hm} + a_1 UR_{mt} + \lambda_t + error \quad (3.4)$$

where q_{hm} is a household h in market m fixed effect and λ_t is a time fixed effect. The results are presented in Table 4 using both a simple average across all UPCs consumed by a household as well as using a household's average expenditure-weights across UPCs. Because a store's average price can vary with the set of UPCs used in their construction, we produce these results for each of the definitions of Ω considered before. The results are similar regardless of the choice for Ω or weights used to construct each household's average relative price: the coefficient on unemployment is negative and statistically significantly different from zero at standard levels. *Thus, higher local unemployment rates are associated with households substituting their expenditures towards low-price retailers.*

The household data also includes characteristics of each household for that year, such as age of the head of household, income, and the number of household members. As a result, we can also investigate which types of households are most likely to engage in store-switching behavior. Specifically, we focus on the relationship between household income and store-switching via the following empirical specification:

$$\tilde{R}_{hmt} = q_{hm} + a_1 UR_{mt} + \sum_g a_g UR_{mt} \times D_{hm}^g + \lambda_t + error \quad (3.5)$$

where q_{hm} and λ_t are defined as in (3.4), D_{hm}^g is a dummy variable equal to one if the average annual income of household h in market m falls into the g^{th} quintile of the income distribution in the IRI sample of households.¹⁰ Coefficient a_1 indicates the degree of store-switching for the lowest-income quintile, whereas a_g indicates whether this degree of store-switching varies at higher income levels. The results are presented in Table 5 for different measures of Ω and weighting across goods. In each case, the degree of store-switching is increasing in the household's income quintile. This could reflect the fact that it may be easier for high-income households to devote additional resources to shopping (such as driving to more distant stores) than lower-income households.¹¹

IV Aggregate Implications of Store-Switching

The key message from our empirical results is that while significant sensitivity to economic conditions is present in the prices paid by households relative to those charged by retailers, this flexibility is driven primarily by store-switching on the part of households rather than sales. Should macroeconomists care about store-switching behavior? In this section, we present a stylized New Keynesian model in which households reallocate their expenditures across stores in light of changing economic conditions. The model delivers testable predictions

¹⁰ The income ranges for each quintile are: 1st quintile ranges from \$5,000 to \$22,500; 2nd quintile ranges from \$22,600 to \$31,600; 3rd quintile ranges from \$31,700 to \$50,000; 4th quintile ranges from \$50,100 to \$70,000; 5th quintile covers incomes of \$70,100 and above.

¹¹ Another possible explanation could be that low-income households already do most of their purchases at low-price stores and therefore have little room left for store-switching during downturns. But we find that the average relative price-level at which households shop varies little across income quintiles in our sample.

which are confirmed in the data. We then present and discuss some business cycle and policy implications of the model.

4.1 New Keynesian Model with Consumer-Expenditure Reallocation across Retailers

To assess the implications of the reallocation of expenditures across stores, we incorporate a decision on the part of households over how much to purchase from different retailers into an otherwise standard New Keynesian model as in Clarida et al. (1999) and Woodford (2003). Retailers sell identical composite goods but at potentially different prices. The “local” retailer charges a higher price to consumers, but purchases at the “discount” retailer come with an iceberg cost. This cost can be reduced by households via time-intensive shopping effort. Thus, households can reduce the effective price of their aggregate consumption at the expense of leisure time. Intermediate goods are produced under monopolistic competition subject to infrequent price adjustment. Because our main interest is in quantifying the macroeconomic implications of store-switching as a form of effective price flexibility, we abstract from sales in the model.

4.1.1 Household Problem

The representative household maximizes lifetime utility over consumption and leisure

$$E_0 \sum_{t=0}^{\infty} \beta^t \{ \log C_t + \log(1 - L_t - S_t) \} \quad (4.1)$$

where C_t is the consumption bundle of goods, L_t is labor supply, S_t is the shopping effort or time spent searching for better prices, and β is the discount factor. Labor is freely mobile across employers. Households live in location A and can purchase consumption goods $C_{A,t}$ from the “local” retailer and purchase $C_{B,t}$ from the “discount” retailer in location B . We assume that the total consumption bundle C_t is given by

$$C_t = \left(C_{A,t}^{\frac{\gamma-1}{\gamma}} + C_{B,t}^{\frac{\gamma-1}{\gamma}} \right)^{\gamma/(\gamma-1)} \quad (4.2)$$

where γ measures the elasticity of substitution across stores. The imperfect substitution among retailers guarantees positive purchases at each retailer and will be used to calibrate the degree of store-switching to what we observe in the data. The budget constraint of the household is

$$P_{A,t} C_{A,t} + \tau_t P_{B,t} C_{B,t} + H_t = (1 + i_{t-1}) H_{t-1} + W_t L_t + \Pi_t \quad (4.3)$$

where τ_t is the iceberg cost of traveling to another location (or searching another location), H_t is bond holding, W_t is wages, and Π_t represents profits from ownership of all firms.

We will assume that the iceberg cost is a function of S_t : $\tau_t = \tau(S_t)$, $\tau \geq 1$, $\tau(0) > 1$, $\tau' < 0$. To simplify algebra, we will further assume that $\tau(S_t)$ has constant elasticity ϕ

$$\left(\frac{S_t}{\tau_t} \frac{\partial \tau_t}{\partial S_t} \right) = \phi \leq 0.$$

Intuitively, one can think of τ as an information cost associated with shopping at the “discount” retailer which is reduced as the household devotes more time to search at this retailer. For example, households typically do their shopping in one primary retailer at which they may accumulate substantial store-specific knowledge, such as the location of different goods (Rhee and Bell 2002). Shopping at other retailers will require more time and effort in the absence of this information, which is captured by τ . However, households can reduce this cost via time-intensive shopping through which information about other retailers is acquired.

Both retailers face the same marginal cost over consumption goods given by P_t . The “discount” retailer sells at cost $P_{B,t} = P_t$, whereas the price at the “local” retailer is assumed to be a constant premium over this input price $P_{A,t} = \mu P_t$, possibly reflecting higher local taxes, high-quality service, the use of an additional scarce input (such as more expensive local real estate), or a lower level of productivity. Thus, the relative posted prices of the two retailers are assumed to be constant $\left(\frac{P_{A,t}}{P_{B,t}} = \mu\right)$, but the relative expenditure costs faced by households will vary with τ_t : $\frac{P_{A,t}}{\tau_t P_{B,t}} = \mu/\tau_t$. The efficient allocation of consumption by the household across retailers implies

$$\frac{C_{A,t}}{C_{B,t}} = \left(\frac{\tau_t P_{B,t}}{P_{A,t}}\right)^\gamma = \left\{\frac{\tau_t}{\mu}\right\}^\gamma \quad (4.4)$$

such that the demand for goods from the “local” retailer at location A relative to the demand for goods from the “discount” retailer at location B will fall when the iceberg costs associated with the “discount” store are low. Since iceberg costs are a function of shopping effort, time variation in shopping intensity will affect the relative prices at the two stores faced by households and will therefore underlie the reallocation of household expenditures across retailers. The efficient allocation of consumption expenditures on the part of the household also implies that the price of the final consumption bundle P_t^C is

$$P_t^C = \left(\{P_{A,t}\}^{1-\gamma} + \{\tau_t P_{B,t}\}^{1-\gamma}\right)^{1/(1-\gamma)} = \mu P_t \left(1 + \left\{\frac{C_{A,t}}{C_{B,t}}\right\}^{\frac{1-\gamma}{\gamma}}\right)^{1/(1-\gamma)}. \quad (4.5)$$

Because the final goods sold at the two retailers are identical, log changes in P_t will be equivalent to those of a fixed-expenditure-weighted price index like the CPI. Changes in the price of the final consumption good (i.e. “cost-of-living” index), however, will systematically differ from changes in P_t for two reasons. First, as emphasized by Triplett (2003), the price of the final consumption good depends on the time-varying effort devoted to shopping, which is a relevant cost from the household’s point of view but is clearly not captured in standard price indices. Second, the price of the final consumption good will vary with the reallocation of expenditures across the two retailers on the part of the household, even when the household buys identical individual goods at the two retailers, as long as prices differ across retailers.

The household budget constraint can be rewritten in terms of the price of the final consumption good

$$P_t^C C_t + H_t = (1 + i_{t-1})H_{t-1} + W_t L_t + \Pi_t. \quad (4.6)$$

As a result, the first-order conditions with respect to aggregate consumption and bond holdings will be standard conditional on being expressed in terms of the final consumption good and its price index P_t^C :

$$\text{Consumption:} \quad C_t^{-1} = q_t P_t^C \quad \Rightarrow \check{q}_t = -\check{C}_t - \check{P}_t^C \quad (4.7)$$

$$\text{Bonds:} \quad q_t = E_t q_{t+1} \beta (1 + i_t) \quad \Rightarrow \check{q}_t = E_t \check{q}_{t+1} + i_t \quad (4.8)$$

where ‘‘checks’’ denote log-linearized deviations from steady-state values. The optimality condition for labor is

$$\text{Labor:} \quad (1 - L_t - S_t)^{-1} = W_t q_t \quad \Rightarrow \eta_L \check{L}_t + \eta_S \check{S}_t = \check{W}_t + \check{q}_t \quad (4.9)$$

where $\eta_L \equiv \frac{\bar{L}}{1-\bar{L}-\bar{S}}$ is the steady-state ratio of labor supply to leisure and equivalently for $\eta_S \equiv \frac{\bar{S}}{1-\bar{L}-\bar{S}}$, and bars indicate steady-state levels of variables. Note that η_L and η_S are the steady-state elasticities of the marginal disutility of reducing leisure hours with respect to labor and shopping hours respectively. The first-order condition for labor includes shopping effort S_t since the latter also affects the marginal disutility of labor.

The optimality condition with respect to shopping effort is

$$\text{Shopping Effort:} \quad (1 - L_t - S_t)^{-1} = -\phi \left(\frac{\tau_t^{1-\gamma}}{\mu^{1-\gamma} + \tau_t^{1-\gamma}} \right) \frac{1}{S_t} \quad \Rightarrow \{\eta_L \check{L}_t + \eta_S \check{S}_t\} = \omega_S \check{S}_t \quad (4.10)$$

where $\omega_S \equiv \phi(1 - \gamma) \left(\frac{\mu^{1-\gamma}}{\mu^{1-\gamma} + \bar{\tau}^{1-\gamma}} \right) - 1$ and $\bar{\tau}$ are steady-state iceberg costs. This optimality condition states that the marginal disutility of shopping effort (the LHS) must equal the marginal benefit of shopping time (the RHS), which is the utility flow from the reduction in expenditures associated with lower effective prices at the ‘‘discount’’ retailer. ω_S denotes the steady-state elasticity of this utility flow with respect to shopping. A steady state with an interior solution for shopping requires the marginal return to shopping effort to be diminishing in the hours spent shopping ($\omega_S < 0$) which, as we document later, is consistent with the data.

Because the marginal disutility of labor and shopping effort are equal, an optimizing household will therefore also equalize the marginal returns to labor and shopping hours:

$$\check{W}_t + \check{q}_t = \omega_S \check{S}_t$$

so that with diminishing returns to shopping effort, hours spent shopping will tend to be low when the utility return to labor is high. A direct relationship between hours worked and hours spent shopping also follows directly from the first-order condition with respect to shopping:

$$\eta_L \check{L}_t = \varepsilon_S \check{S}_t \quad (4.11)$$

where, given $\eta_S = -\phi \frac{\bar{\tau}^{1-\gamma}}{\mu^{1-\gamma} + \bar{\tau}^{1-\gamma}}$, $\varepsilon_S \equiv \omega_S - \eta_S = \phi - \gamma \phi \frac{\mu^{1-\gamma}}{\mu^{1-\gamma} + \bar{\tau}^{1-\gamma}} - 1$. The parameter ε_S represents the steady-state elasticity of the *net* utility cost from higher shopping effort, combining the fact that shopping effort above the steady-state increases both the marginal disutility of hours worked or spent shopping and leads to lower marginal expenditure reductions from shopping effort. Note that diminishing marginal return to shopping ($\omega_S < 0$) is sufficient (but not necessary) for this net utility elasticity to be negative, so that shopping intensity will be

countercyclical with respect to hours: $\varepsilon_S < 0$.¹² Intuitively, if labor hours are low, the marginal disutility to shopping is low, so the household will increase shopping effort which will both increase the disutility of shopping and lower its marginal return.

One can use this relationship between hours worked and shopping time to eliminate shopping effort from the first-order condition with respect to labor, yielding

$$\eta_L \left(1 + \frac{\eta_S}{\varepsilon_S}\right) \check{L}_t = \check{W}_t + \check{q}_t \Leftrightarrow \check{C}_t + \eta_L \left(1 + \frac{\eta_S}{\varepsilon_S}\right) \check{L}_t = \check{W}_t - \check{P}_t^C. \quad (4.12)$$

Shopping effort being countercyclical with respect to labor increases the *effective* Frisch labor supply elasticity.¹³ This reflects the fact that when the utility return to labor is high, shopping intensity will tend to be low so that the return to shopping is also high. But if the household spends fewer hours shopping, then the marginal disutility of labor will be lower, and the household will therefore be willing to supply more labor than it would in the absence of a time-varying shopping effort. In addition, equation (4.12) underscores that a key factor in the household's intra-temporal substitution between consumption and leisure is the price of the final consumption good P_t^C , which can vary because underlying prices change or because the household reallocates its expenditures across the locations.

By the same logic, the consumption Euler equation is given by

$$\check{C}_t = E_t \check{C}_{t+1} - [i_t - (E_t \check{P}_{t+1}^C - \check{P}_t^C)] \quad (4.13)$$

so that the relevant real interest rate for the household's intertemporal decision-making is expressed in terms of expected changes in the price of the final consumption good P_t^C . Thus, properly measuring P_t^C should be important for understanding household consumption and saving decisions.

4.1.2 Firms

Retailer $j \in \{A, B\}$ purchases intermediate goods along a continuum of mass one and assembles them into a consumption good:

$$C_{j,t} = \left(\int_0^1 y_{j,t}(i) \frac{\sigma-1}{\sigma} di \right)^{\sigma/(\sigma-1)}$$

where $y_{j,t}(i)$ is an individual good i (think of this as a UPC) bought by the retailer in location j , and σ is the elasticity of substitution across individual goods (varieties). Note that because the aggregator is the same across locations, the composition of goods in the basket is exactly the same across locations and buyers. Cost minimization implies that the marginal cost paid by both retailers for producing one consumption good is

¹² Diminishing marginal return to shopping also ensures that leisure is countercyclical with respect to labor since log-deviations of leisure are equal to $-\frac{\omega_S}{\varepsilon_S} \eta_L \check{L}_t$.

¹³ The steady-state labor-leisure ratio is given by $\eta_L = \frac{\sigma-1}{\sigma} \left(\frac{\alpha}{\mu}\right)$ which is independent of shopping effort in the model.

$$P_t = \left(\int_0^1 p_t(i)^{1-\sigma} di \right)^{1/(1-\sigma)}.$$

Intermediate goods are produced by a continuum of monopolistic competitors, each of which sells to both retailers. Producer i therefore faces the following total demand:

$$y_t(i) = y_{A,t}(i) + y_{B,t}(i) = \left(\frac{p_t(i)}{P_t} \right)^{-\sigma} C_{A,t} + \left(\frac{p_t(i)}{P_t} \right)^{-\sigma} C_{B,t} = \left(\frac{p_t(i)}{P_t} \right)^{-\sigma} Y_t$$

where $Y_t = C_{A,t} + C_{B,t}$. Production for each intermediate goods firm i is $y_t(i) = Z_t N_t(i)^\alpha$, $\alpha \in (0,1)$, where $y_t(i)$ is the output of variety i , Z_t is the aggregate level of technology, and $N_t(i)$ is employment of firm i . Technology follows a stationary AR(1) process in logs with i.i.d. innovations denoted by ϵ_t^Z and persistence ρ_Z : $\check{Z}_t = \rho_Z \check{Z}_{t-1} + \epsilon_t^Z$.

Workers are hired from a common labor market such that $L_t = \int_0^1 N_t(i) di$ at nominal wage W_t . Firms can reset prices as in Calvo (1983) with the probability of being able to reset their price denoted by $1 - \theta$. When able to reset prices, firm i therefore chooses a reset price $P_t^*(i)$ to maximize expected profits

$$\max_{P_t^*} E_t \sum_{k=0}^{\infty} \theta^k \left\{ \Xi_{t,t+k} \left(P_t^*(i) y_{t+k|k}(i) - W_{t+k} N_{t+k|k}(i) \right) \right\}$$

where $y_{t+k|k}(i)$ is the level of output in period $t+k$ of the firm that reset its price in period t , $\Xi_{t,t+k}$ is the stochastic discount factor, subject to their demand curve and the production function. Combined with the evolution of the intermediate goods price level $P_t = (\theta P_{t-1}^{1-\sigma} + (1-\theta) P_t^{*1-\sigma})^{\frac{1}{1-\sigma}}$, this yields the log-linearized New Keynesian Phillips Curve (NKPC) in terms of input prices P_t for retailers

$$\pi_t = \beta E_t \pi_{t+1} + \kappa \bar{M}C_t \tag{4.14}$$

where $\kappa = \frac{(1-\theta)(1-\beta\theta)}{\theta} \Theta$, $\Theta = \frac{\alpha}{\alpha+(1-\alpha)\sigma}$, $\pi_t = \check{P}_t - \check{P}_{t-1}$, and real marginal costs are given by $\bar{M}C_t = \check{W}_t - \check{P}_t - \check{Y}_t + \check{N}_t$. In the absence of sticky prices, output in the model would simply track technology, so the log-linearized output gap is defined as $\check{X}_t \equiv \check{Y}_t - \check{Y}_t^n = \check{Y}_t - \check{Z}_t$.

4.1.3 The Central Bank

Monetary policy follows an interest rate rule in which interest rates respond to inflation, output growth and the output gap such that, after-log-linearization,

$$i_t = \rho_r i_{t-1} + (1 - \rho_r) [\chi_\pi \pi_t + \chi_x \check{X}_t + \chi_{\Delta y} \Delta \check{Y}_t] + \epsilon_t^r \tag{4.15}$$

where ρ_r is the degree of interest-smoothing, χ_π is the long-run response to inflation, χ_x is the long-run response to the output gap, $\chi_{\Delta y}$ is the long-run response to output growth, and ϵ_t^r is an i.i.d. monetary policy shock.

4.2 Calibration and Model Evaluation

We calibrate most of the parameters to micro facts or standard values and then assess the extent to which these parameter choices are consistent with the empirical properties of effective prices and store-switching documented in section 3. For example, using quarterly frequency, we set $\beta = 0.99$ and θ , the degree of price rigidity, to 0.70 such that firms update their prices every 10 months on average, as in Nakamura and Steinsson (2008). We subsequently experiment with lower values of θ to explore the quantitative importance of store-switching relative to the effects of not treating sales as regular price changes. We set α to be 0.60 and σ , the elasticity of substitution across intermediate goods, to be 10. Coefficients for the interest rate rule follow Coibion and Gorodnichenko (2011), such that long-run responses to inflation (χ_π), output growth ($\chi_{\Delta y}$) and the output gap (χ_x) are equal to 1.5, 0.5, and 0.1 respectively while the degree of interest smoothing ρ_r is 0.9. The persistence of technology shocks is 0.95. Following Coibion et al. (2012), the standard deviation of technology shocks is set to 0.0090 while that of monetary policy shocks is 0.0024.

With respect to the parameters governing store-switching and shopping effort, we first normalize the iceberg cost such that $\bar{\tau} \equiv \tau(\bar{S}) = \mu$. As a result, steady-state consumption of “local” versus “discount” stores is equalized. This appears to be in line with data since the 1990s. For example, in the 2002 Economic Census, discount stores and supercenters accounted for 73% of sales at general merchandise stores and 41% of sales at both general merchandise and food stores (Annual Benchmark Report for Retail Trade and Food Services (2005)). Discount stores and supercenters are also significantly cheaper on average. Cleeren et al. (2010) report that discount stores often charge up to 60% less than regular grocery stores for leading brands. Hausman and Leibtag (2007) report that supercenters are on average 27% cheaper than other grocery stores for selected food products. We set a slightly higher value of $\mu = 1.5$, such that “discount” stores in the model are 33% cheaper on average, because the ability of superstores and discount stores to achieve price discounts is likely to be even larger for non-grocery products (which are more durable and hence more amenable to bulk purchases) than for the grocery goods which Hausman and Leibtag (2007) focus on. These parameters imply a steady-state labor-leisure ratio of 0.4. The 2005 American Time Use Survey reports that the average American age 15 and older worked 3.7 hours a day while spending 5.1 hours on leisure and 8.6 hours on sleep. These values imply average labor-leisure ratios of 0.72 (not counting sleep as leisure) or 0.27 (counting sleep as leisure).¹⁴

Evidence on the elasticity of substitution across stores is limited. Walters (1991) examined both within and across-store substitution and found relatively low rates for the latter relative to the former. Kumar and Leone (1988) found cross-store elasticities of substitution to be two to three times smaller than within-store elasticities for narrow product categories. Rhee and Bell (2002) similarly document limited switching by households in terms of their primary store of choice. As a result, we set $\gamma = 4$, less than half the elasticity of substitution across varieties. For the elasticity of iceberg costs to shopping effort, Aguiar and Hurst (2005) report

¹⁴ In the data, households also spent 4.7 hours per day on other activities such as eating and attending school. Since these are not in the model, we abstract from them here.

that retired households spend 42% more time shopping than pre-retired households, leading to a 17% decline in total food expenditures (with no change in food consumption). We use a slightly lower elasticity of $\phi = -0.35$ (rather than $\phi = -0.4 = -17\%/42\%$) since some of the reduction in expenditures of retirees reflects a rise in time spent preparing food. With $\gamma = 4$ and $\bar{\tau} = \mu$, diminishing marginal return to shopping effort obtains for $-\frac{2}{3} < \phi < 0$, so our baseline value of ϕ is in the middle of the admissible range.

While there is only limited evidence on empirical values for γ and ϕ , we can assess the extent to which our parameter values are consistent with the properties of effective prices and store-switching documented in section 3. For example, we aggregate the average relative price at which households in Eau Claire and Pittsfield shop each month and regress these relative prices on monthly unemployment in each market, a market-specific fixed effect, and a (market-specific) dummy variable for months after December 2008 (to capture the unusual effects of financial constraints and possible shifts in the natural rate of unemployment). The coefficient on the unemployment rate is -0.24. When converted to output gap effects via Okun's Law (a 1% point increase in the unemployment rate is associated with a 2% decline in the output gap as found in Knotek (2007)), this estimate yields an empirical semi-elasticity of the average relative price at which households shop with respect to the output gap of about -0.12 with a standard error of 0.03. In our model with two retailers, the average relative price at which the household shops is

$$R_t^C = \left(\frac{C_{A,t} P_{A,t}}{C_{A,t} P_{A,t} + C_{B,t} P_{B,t}} \right) \times \left(\frac{\mu-1}{\mu+1} \right) + \left(\frac{C_{B,t} P_{B,t}}{C_{A,t} P_{A,t} + C_{B,t} P_{B,t}} \right) \times \left(\frac{1-\mu}{\mu+1} \right)$$

where $\frac{\mu-1}{\mu+1}$ and $\frac{1-\mu}{\mu+1}$ are the relative prices of retailers A and B respectively. The steady-state semi-elasticity of the average relative price at which the household shops with respect to the output gap ε_R can be expressed as

$$dR^C = \left[\frac{2\gamma\mu(\mu-1)}{(1+\mu)^3} \right] \phi \zeta \check{X}_t = \varepsilon_R \check{X}_t$$

where $\zeta \equiv \left(\frac{\sigma-1}{\sigma\mu} \right) \frac{1}{\varepsilon_S} < 0$ is the steady-state elasticity of shopping effort with respect to the output gap. Our baseline parameter values then imply $\varepsilon_R \approx 0.12$, which is almost identical to the empirical estimate. Hence, our baseline parameter choices for γ and ϕ yield a rate of expenditure reallocation across retailers which is in line with our empirical estimates of actual household behavior.

Our calibration is also consistent with the dynamics of shopping effort during the Great Recession. Aguiar et al. (2013) document that approximately 7-8 percent of time lost in market production was absorbed into shopping time. Given that the average amounts of time devoted to market production (32.5 hours/week) and

shopping time (5 hours/week) and the elasticity of output gap to employment ($\check{X}_t = \alpha \check{N}_t$, $\alpha = 0.60$), the implied ζ is between -0.86 and -0.75 which is close to the value of $\zeta = -0.92$ we have from our calibration.¹⁵

We can also evaluate whether the model qualitatively and quantitatively captures the dynamics of effective price inflation. In the model, low economic activity is associated with more shopping activity and therefore lower relative consumption prices:

$$\check{P}_t^C - \check{P}_t = \frac{\phi}{2} \check{S}_t = \omega_\tau \check{X}_t. \quad (4.16)$$

where $\omega_\tau \equiv \frac{1}{2}\phi\zeta$ is one-half the elasticity of iceberg costs with respect to the output gap and $\omega_\tau > 0$ whenever shopping effort is countercyclical with respect to hours ($\varepsilon_s < 0$). When the output gap is negative, shopping effort is high, thereby reducing the price of the final consumption good relative to input prices. Because the price of the final consumption good, rather than intermediate prices, is what matters ultimately for household welfare, the novel channel in this model is how the ratio of the two evolves with economic conditions. \check{P}_t in the model is equivalent, when log-linearized, to the CPI price index in the data. Unfortunately, there is no direct empirical equivalent to the price of the final consumption good \check{P}_t^C since it incorporates both the iceberg cost, which is unobservable, as well as the reallocation of expenditures across retailers. However, we can construct an alternative effective price index P_t^E equivalent to that considered in section 3, which incorporates the reallocation of expenditures across stores $P_t^E \equiv (P_{A,t}C_{A,t} + P_{B,t}C_{B,t})/(C_{A,t} + C_{B,t})$. After log-linearization, this price index evolves according to:

$$\check{P}_t^E - \check{P}_t = \frac{(\mu-1)}{(\mu+1)} \gamma \omega_\tau \check{X}_t. \quad (4.17)$$

Thus, our alternative and observable measure of effective prices \check{P}_t^E should be high relative to the CPI price index \check{P}_t when the level of economic activity is high ($\check{X}_t > 0$), exactly as would be the case with the unobservable final consumption price index \check{P}_t^C .

We can assess the validity of this prediction in our dataset by constructing an aggregate measure of the gap between the average effective price level and the average posted price level. The average effective price level is based on the average price paid by households for a given UPC in a metropolitan area, then aggregated across UPCs and geographic areas using constant expenditure weights. Changes in this index therefore reflect both changes in the individual prices as well as consumer reallocation of expenditures across retailers. The posted price index is constructed using fixed expenditure weights over all individual UPC products and geographic areas. Panel A of Figure 2 plots the difference between these two series and the aggregate unemployment rate. As predicted by the theory, the average effective price index declines sharply relative to the posted price index when the U.S. unemployment rate rises, while the reverse happens when the unemployment

¹⁵ One could alternatively introduce store-switching motives into the model by assuming non-homothetic preferences on the part of households and that high-price stores provide a “higher-quality” shopping experience. However, such a channel would not imply time-variation in shopping effort of the type documented by Aguiar et al. (2013).

rate declines. Hence, this figure illustrates how, even after aggregating across all regions and product categories, the effects of consumer reallocation across retailers leads to a non-trivial mis-measurement of the household consumption price index.

Furthermore, the quantitative magnitudes implied by the series shown in the figure are in line with the predictions of the model. At our baseline parameters, the elasticity of the difference between effective prices and posted prices ($\check{P}_t^E - \check{P}_t$) with respect to the output gap is approximately 0.13 in the model. In the data, regressing ($\check{P}_t^E - \check{P}_t$) on the unemployment rate and a dummy for post-2008 months yields a slope coefficient on the unemployment rate of -0.24 (s.e. 0.06), or equivalently (via Okun's law) an empirical sensitivity of ($\check{P}_t^E - \check{P}_t$) to the output gap of approximately 0.12 (s.e. 0.03), very close to the value implied by our calibrated model. Thus, both the qualitative and quantitative implications of this dimension of the model appear to be consistent with the data.

Another approach to comparing the store-switching channel in the model versus the data is to focus on quantities purchased rather than prices. In our model, the relative quantities purchased at "local" versus "discount" retailers follow

$$\check{C}_t^A - \check{C}_t^B = \gamma\phi\zeta\check{X}_t \quad (4.18)$$

such that when the output gap is low, households will tend to purchase relatively more goods from low-price stores than high-price stores. Again, this reallocation of expenditures will reflect the higher degree of shopping effort when economic activity is low. Given that the relative posted prices at the two stores are constant in the model, the cyclical reallocation in goods purchased will be mirrored in terms of relative dollar expenditures across the two retailers. The share of total revenues going to "local" retailers, defined as $s_{A,t} \equiv (C_{A,t}P_{A,t}) / (C_{A,t}P_{A,t} + C_{B,t}P_{B,t})$, will vary with economic conditions according to

$$ds_{A,t} = \frac{\mu}{(\mu+1)^2} \gamma\phi\zeta\check{X}_t \quad (4.19)$$

where $ds_{A,t}$ is the deviation of $s_{A,t}$ from its steady-state value. To evaluate this additional prediction, we construct the share of total revenues going to "high-price" retailers in our dataset, where "high-price" retailers are defined as in section 3.3 and have an average relative price above zero. Panel B of Figure 2 plots the time series for this share, aggregated across all metropolitan areas, as well as the aggregate unemployment rate. Consistent with the prediction of the model and the results of section 3, the revenue share of "high-price" retailers is strongly counter-cyclical. Thus, the cyclical behavior of relative revenues of "high-price" and "low-price" retailers in the data is also consistent with the predictions of the model with time-varying shopping effort and store-switching. One caveat is that the figure also illustrates how the revenue share of "high-price" retailers has a downward trend over the sample, which reflects the trend growth of discount retail stores in the U.S. during this time period.

Given our calibration of the model, the implied semi-elasticity of the share of revenues going to high-price retailers with respect to the output gap is 0.31. If we again construct equivalent empirical estimates by

regressing this share in the data on the aggregate unemployment rate and a dummy for post-2008 months, we find a slope coefficient on the unemployment rate of -0.86 (s.e. 0.30), or equivalently a semi-elasticity with respect to the output gap of 0.43 with a standard error of 0.15. The predicted value from our calibrated model is well within the confidence intervals of our empirical estimates. In short, the model can capture, both qualitatively and quantitatively, the dynamics of effective prices and store-switching observed in our data.

4.3 Business Cycle and Policy Implications of Store-Switching

In this section, we turn to the implications of shopping effort and store-switching for business cycle dynamics and policy in the model. We focus on four questions: 1) how does the presence of shopping effort and store-switching affect the degree of monetary neutrality, 2) how does it affect the relative importance of stabilizing inflation and output gap volatility for welfare, 3) how can we measure the price of the final consumption good which reflects time-variation in expenditures across retailers, and 4) does it matter whether monetary policy-makers target traditional inflation measures or are there welfare gains from targeting “effective” price measures which incorporate store-switching on the part of households? We address each question in turn.

4.3.1 Effective Price Flexibility and the Degree of Monetary Non-Neutrality

The first question is whether, with prices being effectively more flexible than suggested by posted prices, the degree of monetary neutrality will be lower than expected from the degree of price stickiness. This logic was suggested by Chevalier and Kashyap (2011) in the context of sales, but store-switching similarly generates flexibility in the effective prices paid by households which could potentially undo the stickiness in regular price changes. Our model displays precisely the property suggested by Chevalier and Kashyap (2011), namely that monetary non-neutrality is negatively related to effective price flexibility. We show this in two ways. First, in the special case when the interest rate rule includes only a response to inflation and no interest-smoothing (i.e. we replace (4.15) with $i_t = \chi_\pi \pi_t + \epsilon_t^r$), we can show analytically (see Appendix C) that the contemporaneous effects of monetary policy shocks on the output gap are strictly decreasing in the degree of store-switching. This follows from the fact the model can be reduced to the same three equations typical of New Keynesian models. The standard NKPC for input prices can be expressed in terms of the output gap

$$\pi_t = \beta E_t \pi_{t+1} + \tilde{\kappa} \check{X}_t \quad (4.20)$$

where $\tilde{\kappa} \equiv \kappa \left\{ \left(\frac{\sigma-1}{\sigma\mu} \right) + \frac{1}{\alpha} \right\}$. The Phillips curve is unaffected by the presence of store-switching in the model. In contrast, the Euler equation for the output gap is

$$\check{X}_t = E_t \check{X}_{t+1} - \frac{1}{1+\omega_\tau} (i_t - E_t \pi_{t+1} - r_t^n) \quad (4.21)$$

when expressed in terms of final goods prices. One can verify that ω_τ is increasing in γ and the absolute value of ϕ , so higher elasticities of substitution across retailers and higher elasticities of iceberg costs to shopping effort effectively decrease the elasticity of the output gap, holding future gaps constant, to contemporaneous real

interest rates. This effect reflects the countercyclicality of the shopping effort: if consumption prices are expected to rise, then the output gap should be high but falling and therefore the shopping intensity should be low but expected to rise. The expected increase in shopping time offsets part of the rise in consumption prices reducing the expected decline in gaps. As a result, when the Euler equation is in terms of input prices, the sensitivity of the output gap to real interest rates is diminished. Because this sensitivity plays a crucial role in determining the magnitude of the response of the output gap to monetary policy shocks, it is through this channel that increased store-switching in the model reduces the degree of monetary non-neutrality.

Second, we provide quantitative results for more general interest rate rules (using 4.15). Specifically, we consider the dynamic effects of a contractionary monetary policy shock in our model. Panel A in Figure 3 plots the impulse responses of inflation and the output gap from our model after a contractionary monetary policy shock. The responses include those from a standard New Keynesian model without store-switching ($\phi = 0$), as well as in the case of endogenous shopping effort ($\phi = -0.35$). These responses illustrate a key implication of endogenous shopping effort: a contractionary monetary policy shock has smaller real effects due to the reduced elasticity of output with respect to real interest rates in the IS equation.

However, while store-switching qualitatively reduces the degree of monetary non-neutrality, in our baseline calibration of the model the effects are small: the reduction in the real effects of monetary shocks is only 10% on impact. For comparison, Panel A also plots the impulse responses of inflation and the output gap in the model without store-switching and the levels of price flexibility from Bils and Klenow (2004): $\theta = 0.40$ such that firms update prices every five months on average. This is a useful benchmark because it illustrates dynamics when sales are treated as regular price changes, leading to a higher frequency of price changes. On balance, the results in Figure 3 suggest that store-switching has much smaller effects in terms of monetary non-neutrality than the degree of price rigidity, i.e. the model with store-switching more closely resembles (in terms of the real effects of monetary shocks) the New Keynesian model with very sticky prices than one with more flexible prices. Thus, *the effects of monetary policy shocks on real economic activity are much more affected by the increased price stickiness associated with distinguishing between sales and regular price changes than by the household's store-switching/shopping-effort mechanism.*

The fact that store-switching has relatively small real effects is not sensitive to our specific calibration.¹⁶ With our baseline parameters, $\omega_\tau \approx 0.16$ and thus the IES is effectively reduced by only 14%. While ω_τ is increasing in γ and ϕ , implying that the effective IES (and therefore the real effects of monetary shocks) could be further reduced by higher elasticities to shopping effort and store-switching, these two elasticities are jointly bounded from above by the need for diminishing marginal returns to shopping effort ($\omega_S < 0$). If we set these

¹⁶ Appendix Figure 2 illustrates how the contemporaneous and average responses of the output gap to monetary policy shocks are decreasing in both γ and the absolute value of ϕ . However, the limited effects of store-switching on the degree of monetary non-neutrality are robust to reasonable variation in these parameters.

elasticities at their upper bounds, then $\omega_\tau = 0.6$ and therefore the IES is effectively reduced by 38% at most. Panel A of Figure 3 shows that, even in this limiting case ($\phi = \text{"max"}$), the real effects of monetary shocks are only moderately dampened, with the impact effect being of the same order of magnitude as with $\theta = 0.4$ and no store-switching, but with the persistence of the response practically unchanged relative to our baseline scenario.

Despite the small effects of store-switching and shopping effort on aggregate output, their implications for inflation dynamics and relative consumption at the two retailers are substantial. Panel B of Figure 3 illustrates how store-switching affects relative consumption across retailers in the baseline case of $\phi = -0.35$. While consumption at “local” stores falls sharply, consumption at “discount” stores is much less affected. This pattern reflects the extent of consumer reallocation of household expenditures in response to economic conditions. It is also consistent with the empirical response of the growth rate of revenues in high- and low-price stores to the unemployment rate reported in Table 3. Panel B also illustrates the response of different price and inflation measures to the contractionary monetary policy shock. The price of the final consumption bundle \check{P}^C declines much more sharply on impact than CPI/input prices \check{P} as a result of the increased shopping effort. Similarly, shopping-effort/store-switching leads to a significant difference in the behavior of inflation rates: final good inflation drops much more on impact than CPI inflation. Despite this much greater flexibility in the effective prices paid by households due to store-switching than would be expected from the underlying degree of price stickiness, the implications for monetary neutrality are limited. In this sense, our results are in the same spirit as Eichenbaum et al. (2011), Kehoe and Midrigan (2012) and Guimaraes and Sheedy (2011), but whereas they focus on sales pricing as a source of effective price flexibility, our results stem from the shopping effort and store-switching margins. But the common finding is that even with high flexibility of effective prices, the degree of monetary neutrality hinges primarily on the underlying degree of rigidity in regular prices.

4.3.2 Welfare Objective Function with Endogenous Shopping Effort and Store-Switching

The difference in posted and effective price movements due to the shopping effort and store-substitution margins implies that approximations of household welfare in New Keynesian models will be affected. The most common approach to quantifying welfare is through 2nd order approximations of utility, as in Woodford (2003), which deliver a loss function in terms of inflation and output gap volatility. In Appendix D, we show that the 2nd order approximation to utility in the model with shopping effort and store-switching and $\bar{\tau} = \mu$ is:

$$E(U_t) \approx -\frac{1}{2} \left(\frac{\sigma-1}{\sigma\mu} \right) \{ c_\pi(\phi) E(\pi_t^2) + c_x(\phi) E(\check{X}_t^2) \}, \quad (4.22)$$

$$c_\pi(\phi) \equiv \frac{\sigma}{\theta} \frac{\theta}{(1-\theta)^2} \left(\frac{\omega_S}{\varepsilon_S} \right), \quad c_x(\phi) \equiv \frac{1}{\alpha} + \left(\frac{\sigma-1}{\sigma\mu} \right) \frac{\frac{\gamma}{4}\phi^2 - \frac{\phi}{2} + \omega_S^2}{\varepsilon_S^2}.$$

Panel C of Figure 3 illustrates how the coefficients on inflation and output gap volatility vary with the elasticity of costs to shopping time ϕ . Higher elasticities systematically *lower* the coefficient on inflation volatility, i.e., $\partial c_\pi / \partial \phi < 0$. This reflects the elasticity of leisure with respect to price dispersion: without shopping effort, high

price dispersion requires high levels of labor supply and therefore less leisure, which reduces welfare. However, when shopping effort is present and countercyclical, higher levels of labor supply associated with price dispersion are offset by fewer hours spent shopping, so the sensitivity of leisure to inflation volatility is reduced. The coefficient on output gap volatility reflects the loss from volatility in leisure, which depends on the variance in labor supply, shopping effort, the covariance between the two, and changes in the dispersion of consumption across the two retailers. The interaction of these effects is nonlinear, such that low elasticities of shopping reduce the cost of a given level of output gap volatility but this effect is reversed at high elasticities. The overall effect is to make welfare more sensitive to output gap volatility relative to inflation volatility as the steady-state level of shopping effort increases. Thus, the presence of store-switching and shopping effort suggests that stabilizing output gap volatility should play a more prominent role in the objective function of policymakers than implied by standard models omitting this channel.

4.3.3 How Can We Measure the Price of the Final Good?

Households in our model form their labor supply and saving decisions based in part on the price of the final consumption good \check{P}_t^C , rather than intermediate prices \check{P}_t , and the novel channel in this model is how the ratio of the two evolves with economic conditions. \check{P}_t in the model is equivalent, when log-linearized, to the CPI price index in the data. Although it is infeasible to construct \check{P}_t^C with available data (we do not observe shopping effort), we described in section 4.2 that one can construct an alternative effective price index P_t^E equivalent to the empirical measure used in section 3 and which inherits the same cyclicity as the final good price index. In fact, we can show that the relationship in the model between the unobservable final good price index and the observable effective price index is given by:

$$\check{P}_t^E - \check{P}_t = \frac{(\mu-1)}{(\mu+1)}\gamma(\check{P}_t^C - \check{P}_t). \quad (4.23)$$

Thus, our alternative measure of effective prices \check{P}_t^E should be high relative to the CPI price index precisely when the unobservable final goods price index is also high relative to the CPI price index. Only the magnitude of the elasticity of relative price movements with respect to economic conditions will differ. Given our calibration of the model, $\frac{(\mu-1)}{(\mu+1)}\gamma = 0.8$ so \check{P}_t^E will capture 80% of the movements in the final consumption price index, despite the fact that iceberg costs are not directly observable. Thus, our “effective” price measure \check{P}_t^E which reflects the prices consumers actually pay will, at least in our calibration, closely track the otherwise unobservable final consumption price index \check{P}_t^C .

Our “effective” price measure, \check{P}_t^E , and the final consumption price index, \check{P}_t^C , differ from the CPI price index, \check{P}_t , in that the first two measures incorporate the reallocation of expenditures across retailers. By doing so, they can better address the so-called “outlet-substitution bias”, a CPI bias that arises by not correctly reflecting the changes in the composition of outlets where households shop. The issue of the outlet or store-switching bias

has, of course, long been recognized in the inflation measurement literature, going back at least to Hoover and Stotz (1964) who point out the possible upward bias in the CPI measurement that fails to account for growing market shares of chain food distribution and declining independent stores. Since the CPI did not correctly reflect these changes in the weights and instead used fixed weights that continued to assign higher weights to independent stores which provided food products at higher prices, they found a non-negligible upward bias in the CPI with such fixed weights. The seminal paper by Reinsdorf (1993) and more recent work by Greenlees and McClelland (2007, 2011) address the related bias that emerges from the growing presence of supercenters, other discount stores and warehouse club stores in consumers shopping patterns and the CPI weights adjusting in a lagged manner to these changes. On average, according to Reinsdorf (1993) and subsequent estimates of overall CPI new outlet bias, including those by Lebow, Roberts, and Stockton (1994), the Boskin Commission (U.S. Senate, 1996), Shapiro and Wilcox (1996), Hausman and Leibtag (2007), and Lebow and Rudd (2003), the price index measurement bias lies somewhere between 0.1 to 0.4 percentage points per year.

A key difference between our results and this previous literature is that we focus on the cyclical variation in store-switching rather than secular-trend biases, and we propose an alternative measure based on expenditure shares which can capture much of the cyclical biases associated with store-switching. As emphasized in Triplett (2003), the difference in our model between the final consumption price index faced by households and a fixed-expenditure weight index such as the CPI will reflect two sources: the reallocation of expenditures by households across retailers as well as the time-varying intensity of shopping effort on the part of households. Using our detailed dataset, we show how one can calculate an empirical “effective” price index, which according to the model should closely track the final consumption price index despite not directly including the household costs stemming from shopping effort.

These results can be viewed as part of a much broader literature focused on addressing a variety of substitution biases which arise in measuring price indices. For example, the price of a basket of goods can vary, even when prices are constant, if consumers change the brands or quality of goods they purchase within a category. In fact, as mentioned in Greenlees and McClelland (2011) and in the *Consumer Price Index Manual* published by the International Labour Office (2004), the substitution bias and the outlet bias are conceptually similar. And while we focus primarily on measuring the effects of store-switching, one can also exploit IRI data to measure the more traditional substitution bias within categories of goods at a single retailer, as illustrated in Appendix B. We find the same order of cyclical variation in effective price indices that allow for substitution across goods within a retailer as we find in effective price indices that allow for substitution across retailers for a single good, so the well-known “substitution bias” will occur not just in the long-run, but also over the course of the business cycle.¹⁷ Another margin of adjustment that consumers can use is changing the share of expenditures

¹⁷ We constructed measures of effective price inflation within a category and store, i.e. allowing for substitution across goods within a category and single store but shutting down substitution across stores. This yielded similar elasticities of

going to different categories of goods (e.g. buying more pasta and less fruit), which we do not explore in this paper. A fourth margin is intertemporal substitution of purchases, where consumers can “stock-up” on products when they are on sale. We find little evidence for this margin in our data. In summary, the wide range of margins that can cause the price of a consumer’s basket of goods to vary suggests that statistical agencies could measure the cost-of-living more precisely by tracking actual prices paid by consumers rather than prices posted by a fixed sample of retailers.

4.3.4 What Price Measure Should Policymakers Target?

A fourth question of practical interest implied by the model is whether policymakers should focus on standard fixed-expenditure weight price measures or ones that reflect the shopping effort, and therefore expenditure-reallocation, activities of households. While household welfare is explicitly associated with the price of the final consumption good P_t^C , the fact that these prices are effectively more flexible than posted prices P_t might suggest that responding to the latter will be welfare-improving via the logic of Aoki (2001).

To address this question, we quantify household welfare (4.22) when the central bank responds either to posted price inflation π_t in the standard interest rate rule, final good price inflation π_t^C or the closely related expenditure-weighted final goods inflation rate π_t^E holding the other parameters in the model (including those in the policy rule) constant. The resulting inflation and output gap variances are plotted in Panel A of Figure 4 for different long-run inflation response coefficients χ_π . Responding to posted price inflation π_t , as assumed in our baseline, leads to lower volatility in posted price inflation but higher levels of output gap volatility, with the latter effect being stronger when the inflation response is high. Given that inflation volatility is more heavily weighted in the loss function, these results lead to utility being higher when the central bank targets posted price inflation, although the differences are quantitatively small.

However, this result does not hold under price-level targeting (PLT) regimes, which—as we know from Woodford (2003), Gorodnichenko and Shapiro (2007), Coibion et al. (2012), and others—are closer to replicating the optimal policy under commitment than inflation targeting. Specifically, we replace the baseline interest rate rule with one that allows for responses to deviations of price-levels from their targets:

$$i_t = \rho_r i_{t-1} + (1 - \rho_r) [\chi_p \check{P}_t^m + \chi_x \check{X}_t + \chi_{\Delta y} \Delta \check{Y}_t] + \epsilon_t^r \quad (4.24)$$

where \check{P}_t^m denotes the deviation of posted prices \check{P}_t , final goods prices \check{P}_t^C , or expenditure-weighted prices \check{P}_t^E from a target path. Panel B of Figure 4 plots the implied inflation and output volatilities for different response coefficient χ_p to each price level, as well as expected utility losses for different values of χ_p . While output gap volatility is still higher when the central bank responds to posted prices, posted price inflation volatility is now also higher in this case. This reflects the fact that much of the price response is immediate for \check{P}_t^C and \check{P}_t^E and

effective price inflation as in Table 1, as documented in Appendix B. Hence, the cyclical effect of retail-switching on effective prices is about the same order of magnitude as the cyclical effect of substitution across goods.

more delayed for \check{P}_t . Under PLT, the large immediate decline in effective prices after a disinflationary shock must be reversed when the central bank targets the level of effective prices. This requires significant declines in nominal and real interest rates by monetary policy-makers, which immediately raise contemporaneous output gaps via the dynamic IS curve relative to how they otherwise would have responded. This, in turn, offsets some of the initial disinflation via the NKPC. As a result, both inflation and output gap volatility will be significantly reduced when the central bank targets the “effective” or final consumption price level. This effect does not occur under inflation targeting because the latter does not require the policy-maker to undo the immediate decline in effective prices and therefore does not require the significant declines in nominal and real interest rates which serve to stabilize both inflation and output volatility under PLT. Hence, expected utility is now strictly greater when the central bank targets deviations of effective final goods price deviations rather than posted price levels.

V Conclusion

A key objective for macroeconomists is quantifying both the degree of price rigidities and the implications of these rigidities for macroeconomic dynamics. We shed new light on this issue by documenting that the effective prices paid by households are significantly more flexible than those charged by retailers: during economic downturns, the prices that households pay for a given good decline more sharply than the prices posted by individual retailers. This effective price flexibility reflects a reallocation of household expenditures across retailers rather than more frequent sales or purchases of items on sales. When we integrate time-varying shopping effort and store-switching into a New Keynesian model, the ability of consumers to reallocate their expenditures across retailers charging different prices provides an additional force toward price flexibility which leads to smaller real effects of monetary policy shocks. However, given our calibration, the quantitative implications of store-switching for monetary neutrality are relatively small. Thus, our results are similar in spirit to those of Eichenbaum et al. (2011), Kehoe and Midrigan (2012), and Guimaraes and Sheedy (2011), but whereas they focus on the role of sales in generating increased flexibility in prices paid by households, we focus on the store-switching margin for effective price flexibility. Nonetheless, store-switching does have non-trivial implications for the relative importance of stabilizing output gap versus inflation variation in welfare calculations, the question of which price measures should be targeted by policymakers, and the interpretation of estimated relationships in terms of structural parameters.

Another implication of store-switching and time-varying shopping effort is that there is a *cyclical* mismeasurement of the price of households’ final consumption basket in standard aggregate inflation measures. This points toward a need for statistical agencies to devote more resources to measuring the reallocation of expenditures, not just across goods as commonly emphasized with respect to the substitution bias, but also across retailers. While the construction of time-varying expenditure weights in real-time is unlikely to be feasible for statistical agencies, one approach would be to track the prices paid by households in real-time rather than the prices

charged by a fixed set of retailers, as suggested by Triplett (2003). To this end, scanner price data, which provide prices and quantities, may be a useful tool for constructing cost-of-living indexes although a number of challenges remain to be addressed in this area (see Feenstra and Shapiro (2003)). A related implication is that the ideal price index in the presence of time-varying shopping effort should measure not only the reallocation of expenditures across stores but also the shopping effort of households, since this represents a cost of the consumption basket. Further work on how this could be measured in practice would be of immediate practical relevance. One implication of our results, however, is that shopping effort and store-switching are tightly linked, so that a price-measure which incorporates expenditure reallocation across retailers will capture most of the missing cyclical variation in the final consumption price index.

Another fruitful extension in thinking about consumers' expenditure reallocation behaviors would be to consider the impact of the rise in online retailing: the latter likely reduces shopping costs substantially and therefore facilitates the reallocation of expenditures across retailers by households. With online retailing growing rapidly, expenditure-reallocation across retailers is likely to become increasingly important. Just as Walmart reshaped the retailing industry in the U.S., the rise of Amazon.com and other online stores will similarly transform household expenditure decisions. Our results suggests that the ways in which households reallocate their expenditures, as well as the intensity of their shopping search, is not innocuous for macroeconomic dynamics and optimal policy.

References

- Aguiar, M., and E. Hurst, 2005. "Consumption versus Expenditure," *J.P.E.*, 113(5), 919-948.
- Aguiar, M., E. Hurst, and L. Karabarbounis, 2013. "Time Use During the Great Recession," *A.E.R.*, 103(5), 1664-1696.
- Aoki, K., 2001. "Optimal Monetary Policy Responses to Relative Price Changes," *J.M.E.* 48(1), 55-80.
- Bils, M. and P. J. Klenow, 2004. "Some Evidence on the Importance of Sticky Prices," *J.P.E.* 112(5), 947-985.
- Boskin, M. J., E. Dulberger, R. Gordon, Z. Griliches, and D. Jorgenson, 1996. "Toward a More Accurate Measure of the Cost of Living," Final Report to the Senate Finance Committee, Dec. 4.
- Bronnenberg, B. J., M. W. Kruger, and C. F. Mela, 2008. "Database Paper: The IRI Marketing Data Set," *Marketing Science* 27(4), 745-748.
- Calvo, G., A. 1983. "Staggered Prices in a Utility-Maximizing Framework," *J.M.E.* 12(3), 383-398.
- Carlton, D. W., 1986. "The Rigidity of Prices," *A.E.R.* 76(4), 637-658.
- Cecchetti, S. G., 1986. "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines," *Journal of Econometrics* 31(3), 255-274.
- Chevalier, J. A., and A. K. Kashyap, 2011. "Best Prices," NBER Working Paper 16680.
- Clarida, R., J. Gali, and M. Gertler, 1999. "The Science of Monetary Policy: A New Keynesian Perspective," *J.E.L.* 37, 1661-1707.
- Cleeren, K., F. Verboven, M. G. Dekimpe, and K. Gielens, 2010. "Intra- and Interformat Competition among Discounters and Supermarkets," *Marketing Science* 29(3) 456-473.
- Coibion, O., and Y. Gorodnichenko, 2011. "Monetary Policy, Trend Inflation, and the Great Moderation: An Alternative Interpretation," *A.E.R.* 101(1), 341-70.

- Coibion, O., Y. Gorodnichenko, and J. Wieland, 2012. "The optimal inflation rate in New Keynesian models," *R.E.Stud.* 79(4), 1371-1406.
- Driscoll, J. C., and A. C. Kraay, 1998. "Consistent Covariance Matrix Estimation With Spatially Dependent Panel Data," *R.E.Stat.* 80(4), 549-560.
- Eichenbaum, M., N. Jaimovich, and S. Rebelo, 2011. "Reference Prices and Nominal Rigidities," *A.E.R.* 101(1), 234-262.
- Feenstra, R., and M.D. Shapiro, 2003. *Scanner Data and Price Indexes*. U. of Chicago Press.
- Gali, J., and M. Gertler, 1999. "Inflation Dynamics: A Structural Econometric Analysis," *J.M.E.* 44(2), 195-222.
- Gorodnichenko, Y., and M.D. Shapiro, 2007. "Monetary Policy When Potential Output Is Uncertain: Understanding the Growth Gamble of the 1990s," *J.M.E.* 54, 1132-1162.
- Greenlees, J.S., and R. McClelland, 2007. "Price Differentials Across Outlets in Consumer Price Index Data, 2002-2006." *Tenth Meeting of the Ottawa Group, Ottawa, Canada, October*. Vol. 404. 2007.
- Greenlees, J.S., and R. McClelland, 2011. "New Evidence on Outlet Substitution Effects in Consumer Price Index Data," *R.E.Stat.* 93(2), 632-646.
- Guimaraes, B. and K. Sheedy, 2011. "Sales and Monetary Policy," *A.E.R.* 101(2), 844-876.
- Hausman, J. and E. Leibtag, 2007. "Consumer Benefits from Increased Competition in Shopping Outlets: Measuring the Effect of Wal-Mart," *Journal of Applied Econometrics* 22(7), 1157-77.
- Hoover, E.D and M. S. Stotz, 1964, "Food Distribution Changes and the CPI," *Monthly Labor Review* 87 (January), 58-64.
- International Labor Office, *Consumer Price Index Manual: Theory and Practice* (Geneva: International Labour Office, 2004). Available at: <http://www.ilo.org/public/english/bureau/stat/guides/cpi/index.htm>.
- Kashyap, A., 1995. "Sticky Prices: New Evidence from Retail Catalogs," *Q.J.E.* 110(1), 245-274.
- Kehoe, P. J. and V. Midrigan, 2012. "Price Are Sticky After All," Federal Reserve Bank of Minneapolis Research Department Staff Paper 413.
- Knotek, E.S. II, 2007. "How Useful Is Okun's Law?" Federal Reserve Bank of Kansas City *Economic Review*, Fourth Quarter, 73-103.
- Kumar, V., and R. P. Leone, 1988. "Measuring the Effect of Retail Store Promotions on Brand and Store Substitution," *J. of Marketing Research* 25(2) 178-185.
- Lebow, D, Roberts, J and D. Stockton, 1994, "Monetary Policy and 'The Price Level'", Board of Governors of the Federal Reserve System Manuscript.
- Lebow, D and J. Rudd. 2003. "Measurement Error in the Consumer Price Index: Where Do We Stand?" *Journal of Economic Literature*, 41(1): 159-201.
- Nakamura, E., and J. Steinsson, 2008. "Five Facts About Prices: A Reevaluation of Menu Cost Models," *Q.J.E.* 123(4), 1415-1464.
- Reinsdorf, M. 1993. "The effect of outlet price differentials on the U.S. Consumer Price Index" in Price Measurements and Their Uses. M. Foss, M. Manser, and A. Young (eds.). NBER Studies in Income and Wealth, Vol. 57. Chicago: University of Chicago Press.
- Rhee, H., and D. Bell, 2002. "The Inter-Store Mobility of Supermarket Shoppers," *J. of Retailing* 78(4), 225-237.
- Roberts, J.M., 2006. "Monetary Policy and Inflation Dynamics," *International J. of Central Banking* 2, 193-230.
- Shapiro, M.D., and D.W. Wilcox, 1996. "Mismeasurement in the Consumer Price Index: An Evaluation," in NBER Macroeconomics Annual 1996, eds. B.S. Bernanke, J.J. Rotemberg, MIT Press, Cambridge, MA.
- Triplett, J. E., 2003. "Using Scanner Data in Consumer Price Indexes: Some Neglected Conceptual Issues," in Scanner Data and Price Indexes, eds. R.C. Feenstra and M.D. Shapiro, University of Chicago Press.
- Walters, R. G., 1991. "Assessing the Impact of Retail Price Promotions on Product Substitution, Complementary Purchase, and Interstore Sales," *J. of Marketing* 55, 17-28.
- Williams, J. C., 2006, "The Phillips Curve in an Era of Well-Anchored Inflation Expectations," Working Paper, Federal Reserve Bank of San Francisco.
- Woodford, Michael, 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Table 1. Cyclical properties of selected moments of price changes.

Market×Category fixed effects	N	Y	Y	Y	Y	Y	Y	Y
Month fixed effects	N	N	Y	linear trend	Y	Y	Y	Y
Weighted regression	N	N	N	N	N	N	Y	Y
Dependent variable	Equal weights to aggregate UPCs				Expenditure shares as weights to aggregate UPCs			
	(1)	(2)	(3)	(4)	Market specific	Common	Market specific	Common
Inflation								
Posted prices	-0.078* (0.040)	-0.079 (0.052)	-0.050*** (0.017)	-0.165** (0.065)	-0.070*** (0.025)	-0.071*** (0.026)	-0.013 (0.047)	-0.006 (0.048)
Effective prices	-0.120* (0.067)	-0.126 (0.087)	-0.219*** (0.024)	-0.288*** (0.105)	-0.201*** (0.031)	-0.205*** (0.033)	-0.136*** (0.027)	-0.146*** (0.028)
<i>p</i> -value	0.202	0.267	<0.001	0.006	<0.001	<0.001	0.026	0.006

Notes: The table reports estimated coefficients on local seasonally-adjusted unemployment rate in specification (3.1). Column (4) shows results after controlling for linear time trends. Number of observations is 204,476. *P*-value row shows *p*-values for the test that the sensitivity of inflation to unemployment rate is the same for posted and effective prices. The last two columns report regression results where observations are weighted by the (time-series) average expenditure share of a given market/category cell in total spending across all cities and categories. Driscoll and Kraay (1998) standard errors are in parentheses. ***, **, * denote significance at 0.01, 0.05, and 0.10 levels. See section 3.1 for details.

Table 2. Cyclical properties of selected moments of price changes.

Market×Category fixed effects	N	Y	Y	Y	Y	Y	Y	Y
Month fixed effects	N	N	Y	linear trend	Y	Y	Y	Y
Weighted regression	N	N	N	N	N	N	Y	Y
Dependent variable	Equal weights to aggregate UPCs				Expenditure shares as weights to aggregate UPCs			
	(1)	(2)	(3)	(4)	Market specific	Common	Market specific	Common
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sales								
Frequency	0.843*** (0.113)	0.717*** (0.143)	0.224*** (0.092)	0.005 (0.102)	0.126 (0.114)	0.125 (0.111)	0.005 (0.112)	-0.039 (0.112)
Size	0.120*** (0.032)	0.148*** (0.035)	0.303*** (0.066)	0.147*** (0.046)	0.348*** (0.068)	0.356*** (0.063)	0.182*** (0.065)	0.185*** (0.059)
Share of goods bought on sale	0.966*** (0.121)	0.718*** (0.157)	0.161 (0.104)	0.022 (0.127)	0.018 (0.131)	0.033 (0.127)	0.044 (0.124)	-0.037 (0.127)
Regular prices								
Frequency of changes								
All	0.287*** (0.064)	0.317*** (0.078)	-0.018 (0.032)	-0.124 (0.080)	-0.068* (0.040)	-0.060 (0.039)	-0.091*** (0.038)	-0.075** (0.036)
Positive	0.140*** (0.046)	0.149*** (0.057)	-0.041* (0.024)	-0.156** (0.070)	-0.066*** (0.027)	-0.061** (0.027)	-0.067*** (0.025)	-0.058*** (0.024)
Negative	0.137*** (0.020)	0.157*** (0.025)	0.026** (0.011)	0.030** (0.015)	0.001 (0.015)	0.005 (0.014)	-0.021 (0.017)	-0.014 (0.014)
Size of changes								
All	-0.194*** (0.048)	-0.245*** (0.056)	-0.109*** (0.044)	-0.135*** (0.042)	-0.036 (0.042)	-0.032 (0.040)	-0.062** (0.028)	-0.064** (0.030)
Positive	0.128*** (0.023)	0.167*** (0.026)	0.070 (0.050)	0.099*** (0.035)	0.085* (0.046)	0.087* (0.048)	-0.117*** (0.037)	-0.099** (0.043)
Negative	-0.166*** (0.051)	-0.239*** (0.064)	-0.036 (0.060)	0.082 (0.064)	-0.015 (0.047)	-0.012 (0.049)	0.135*** (0.038)	0.103*** (0.042)
Share of goods bought at cheap regular prices	0.362*** (0.034)	0.200*** (0.022)	0.126*** (0.042)	0.050*** (0.018)	0.191*** (0.055)	0.199*** (0.054)	0.116** (0.050)	0.084* (0.047)

Notes: The table reports estimated coefficients on local seasonally-adjusted unemployment rate in specification (3.1). Column (4) shows results after controlling for linear time trends. The last two columns report regression results where observations are weighted by the (time-series) average expenditure share of a given market/category cell in total spending across all cities and categories. Share of goods bought at cheap regular prices is the share of UPC units sold at low regular prices (i.e., prices at the bottom 20% of the cross-sectional distribution of regular prices in a given market and month) in total number of units sold at regular prices. Number of observations is 204,476. “Size of sales” and “size of negative price changes” are negative values. For example, a temporary price cut (sale) of 20% is recorded as -0.2. Driscoll and Kraay (1998) standard errors are in parentheses. ***, **, * denote significance at 0.01, 0.05, and 0.10 levels. See sections 3.2-3.3 for details and the note to Table 1.

Table 3. Cyclicity of pricing moments by store relative prices.

Dependent variable	UPC sample: Ω_{\max}		UPC sample: Ω_{90}		UPC sample: Ω_{75}	
	UR (1)	$UR \times \bar{R}_{mst, \Omega}$ (2)	UR (3)	$UR \times \bar{R}_{mst, \Omega}$ (4)	UR (5)	$UR \times \bar{R}_{mst, \Omega}$ (6)
Share of revenues	-0.133*** (0.016)	-0.202*** (0.053)	-0.133*** (0.016)	-0.299*** (0.062)	-0.132*** (0.016)	-0.379*** (0.069)
Sales						
Frequency	-0.325*** (0.078)	0.084 (0.436)	-0.312*** (0.077)	0.059 (0.509)	-0.316*** (0.077)	-0.174 (0.587)
Size	0.213*** (0.068)	1.538*** (0.291)	0.234*** (0.066)	1.570*** (0.344)	0.224*** (0.066)	1.636*** (0.421)
Share of goods bought on sale	-0.448*** (0.076)	-0.031 (0.530)	-0.440*** (0.076)	-0.020 (0.642)	-0.442*** (0.076)	-0.305 (0.755)
Regular price changes						
Frequency						
All	-0.017 (0.031)	1.283*** (0.268)	-0.021 (0.031)	1.393*** (0.304)	-0.024 (0.031)	1.581*** (0.354)
Positive	-0.017 (0.020)	0.830*** (0.190)	-0.021 (0.020)	0.897*** (0.216)	-0.023 (0.020)	1.045*** (0.240)
Negative	0.003 (0.013)	0.426*** (0.154)	0.003 (0.013)	0.473*** (0.167)	0.002 (0.013)	0.522*** (0.195)
Size						
All	-0.059*** (0.019)	-0.242 (0.293)	-0.067*** (0.020)	-0.337 (0.318)	-0.067*** (0.020)	-0.261 (0.360)
Positive	-0.054 (0.044)	-0.423* (0.238)	-0.060 (0.044)	-0.431 (0.264)	-0.059 (0.043)	-0.452 (0.307)
Negative	0.071* (0.039)	0.712** (0.369)	0.079** (0.039)	0.691* (0.384)	0.076** (0.039)	0.674 (0.436)

Notes: The table reports estimates of specification (3.2). Number of observations is 235,420. UR is the local seasonally-adjusted unemployment rate. $\bar{R}_{mst, \Omega}$ is the relative price of the store. Ω indicates what universe of goods is used to calculate relative prices of stores. Expenditure shares are used to aggregate relative prices across UPCs. “Size of sales” and “size of negative price changes” are negative values. For example, a temporary price cut (sale) of 20% is recorded as -0.2. Driscoll and Kraay (1998) standard errors are in parentheses. ***, **, * denote significance at 0.01, 0.05, and 0.10 levels. See section 3.3 for further details.

Table 4. Relative price of the store where households shop as a function of local unemployment rate.

Sample of UPCs used in calculating relative prices of stores	Equal weights to all goods	Goods are weighted by expenditure shares
	(1)	(2)
Ω_{max}	-0.520*** (0.206)	-0.898** (0.413)
Ω_{90}	-0.534*** (0.142)	-0.721*** (0.266)
Ω_{75}	-0.477*** (0.114)	-0.681*** (0.218)


Notes: The table reports estimates of specification (3.4). The dependent variable is the average (expenditure weighted) relative price of stores where a household shops in a given month. The table reports estimated coefficients on the local seasonally-adjusted unemployment rate. The number of observations is 742,879. Driscoll and Kraay (1998) standard errors are in parentheses. ***, **, * denote significance at 0.01, 0.05, and 0.10 levels. See section 3.3 for details.

Table 5. Relative price of the store where households shop as a function of local unemployment rate and household's income.

Sample of UPCs used in calculating relative prices of stores	Weights for aggregation of UPCs	UR_{ct}	$UR_{ct} \times D^2$	$UR_{ct} \times D^3$	$UR_{ct} \times D^4$	$UR_{ct} \times D^5$
		(1)	(2)	(3)	(4)	(5)
Ω_{max}	Equal	-0.474** (0.207)	-0.029*** (0.009)	-0.040*** (0.011)	-0.077*** (0.021)	-0.094*** (0.024)
	Expenditure shares	-0.821** (0.413)	-0.046*** (0.013)	-0.077*** (0.017)	-0.119*** (0.026)	-0.152*** (0.032)
Ω_{90}	Equal	-0.496*** (0.142)	-0.026*** (0.006)	-0.035*** (0.007)	-0.062*** (0.015)	-0.071*** (0.015)
	Expenditure shares	-0.681*** (0.265)	-0.034*** (0.009)	-0.039*** (0.009)	-0.061*** (0.015)	-0.073*** (0.019)
Ω_{75}	Equal	-0.438*** (0.114)	-0.025*** (0.006)	-0.035*** (0.006)	-0.063*** (0.013)	-0.072*** (0.014)
	Expenditure shares	-0.636*** (0.216)	-0.035*** (0.009)	-0.042*** (0.008)	-0.068*** (0.013)	-0.084*** (0.016)

Notes: The table reports estimates of specification (3.5). D^g is the dummy variable equal to one if a household is in g^{th} income quintile and zero otherwise. 5th quintile is the top income quintile. The number of observations is 742,879. Driscoll and Kraay (1998) standard errors are in parentheses. ***, **, * denote significance at 0.01, 0.05, and 0.10 levels. See section 3.3 for details.

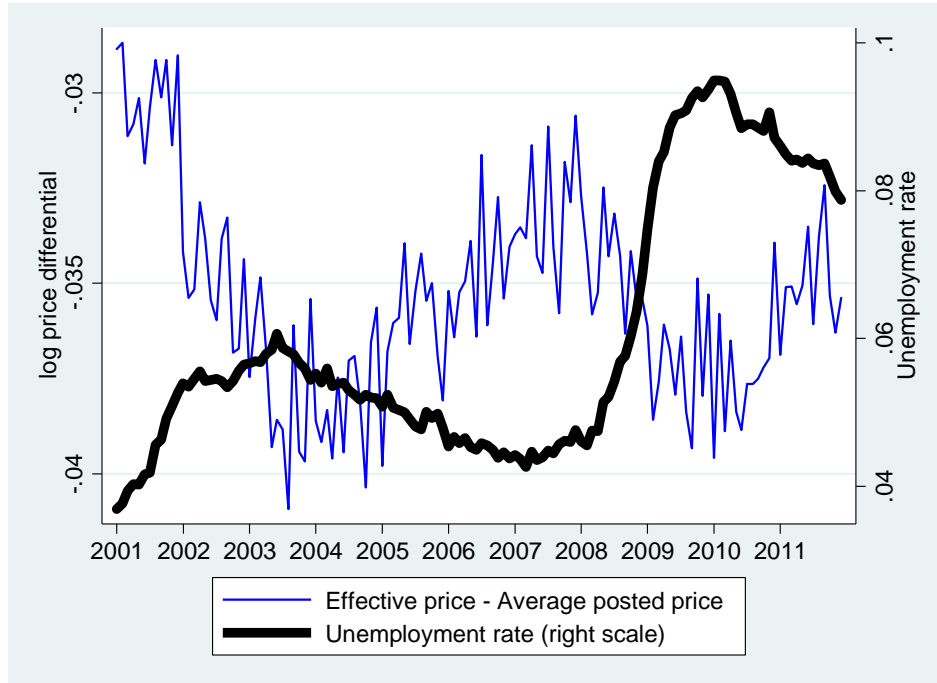
Figure 1: Retailer Stock Returns during the Great Recession



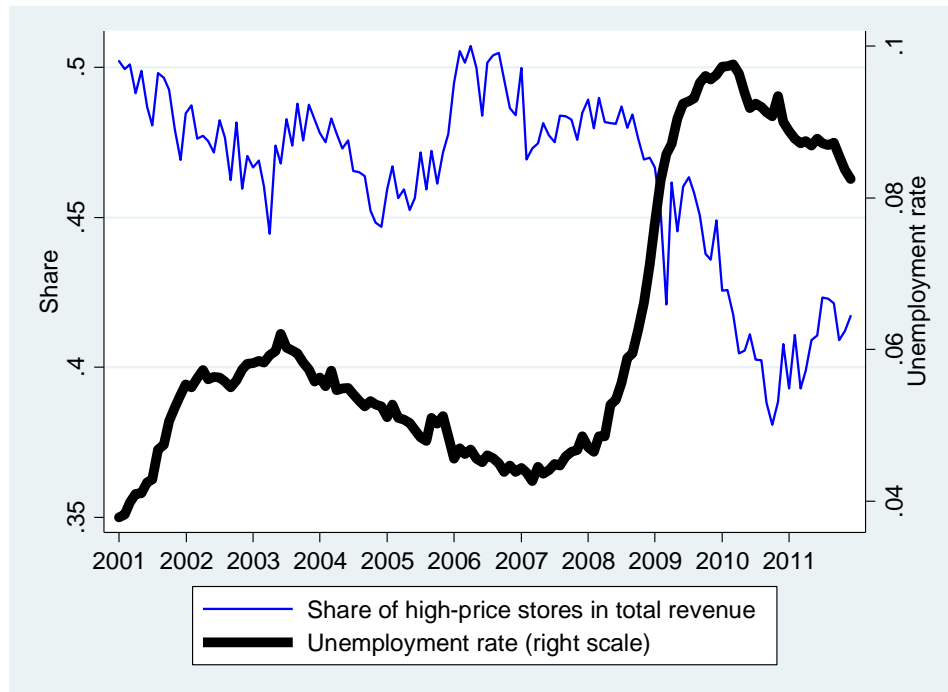
Notes: The figure plots the log-difference in stock prices relative to their December 2007 values for Wal-Mart, Family Dollar, Safeway, Whole Foods and the S&P 500 index. The shaded area indicates the duration of the Great Recession as defined by the National Bureau of Economic Research.

Figure 2. Economic Conditions and the Reallocation of Consumption Expenditures across Retailers.

Panel A: The Cyclicity of the Gap between the Effective Price Index and the Posted Price Index




Panel B: The Cyclicity of the Share of Total Retailer Revenues Coming from High-Price Retailers




Notes: Panel A plots the difference between the “effective” price index and the “posted” price index. The latter is a fixed-expenditure-weighted average of all UPC prices posted in each store and metropolitan areas in the data, where weights are average expenditure share of each UPC in each geographic area relative to total household expenditures. The former is the fixed-expenditure-weighted average over the average prices paid by households for each UPC across all retailers in a metropolitan area. Panel B plots the share of total retailer revenues arising from revenues at high-price retailers in the data. “High-price” retailers are defined as in the text for each metropolitan area. Total share is a fixed expenditure-weighted average across metropolitan areas of shares in each metropolitan area. See section 4.2 for details.

Figure 3. Business Cycle and Policy Implications of Endogenous Shopping Effort and Store-Switching.


Panel A: Dynamic Responses to Monetary Policy Shock with and without Endogenous Shopping Effort



Panel B: Substitution Effects after Monetary Policy Shock ($\phi = -0.35$)

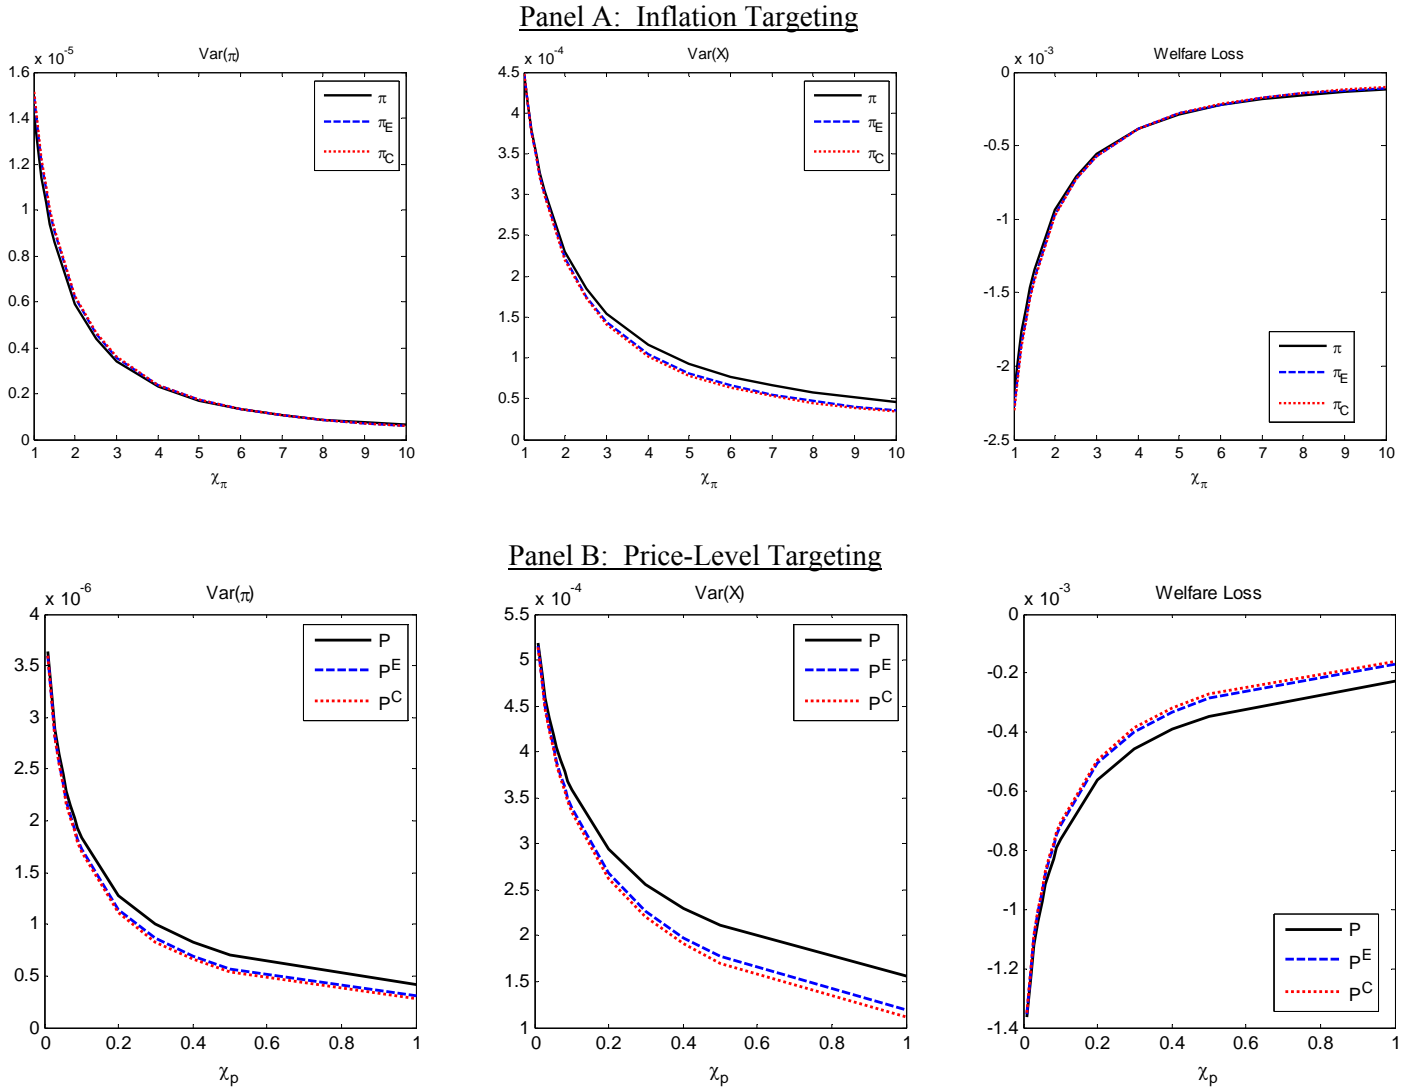


Panel C: Welfare Implications of Endogenous Shopping Effort



Notes: The top panel displays impulse responses of inflation, output gap and shopping effort to a contractionary monetary policy shock in the model described in section 4.1. θ is the degree of price stickiness and ϕ is the elasticity of iceberg costs to shopping effort. The middle panel plots the impulse responses under our baseline parameter values of inflation rates for final consumption goods (π^C), effective prices (π^E), and posted prices (π) in the left figure, price levels for consumption goods (P^C), effective prices (P^E) and posted prices (P) in the middle figure, and total consumption (C), consumption at the high-price retailer (C^A) and consumption at the low-price retailer (C^B) in the right figure. The bottom panel plots the coefficients on inflation volatility (left) and output gap volatility (middle) from the second-order approximation to utility in the model of section 4.1 and the right figure plots the ratio of the two. See section 4.3 for details.

Figure 4. Which Inflation Measure Should Policymakers Target?



Notes: The top panel displays the variance of inflation (left) and the variance of the output gap (middle) for different long-run responses to inflation by the central bank when the central bank responds to either posted price inflation (π), effective price inflation (π^E), or inflation of final goods prices (π^C). The right panel shows expected utility loss in each case. The bottom panel plots equivalent figures in the case of price-level targeting when the central bank responds to either posted prices (P), effective prices (P^E), or final consumption goods prices (P^C). See section 4.3 for details.

**NOT FOR PUBLICATION
APPENDIX MATERIAL**

APPENDIX A: ADDITIONAL TABLES AND FIGURES.

Appendix Table 1. Descriptive statistics by category.

Category	Sales			Regular prices						
	Frequency	Size	Share of goods bought on sale	Share of goods bought at cheap reg. prices	Frequency of changes			Size of changes		
					All	Positive	Negative	All	Positive	Negative
(1)	(2)	(3)	(5)	(5)	(6)	(7)	(8)	(9)	(10)	
All categories	0.200 (0.083)	-0.252 (0.081)	0.245 (0.105)	0.253 (0.048)	0.054 (0.038)	0.034 (0.031)	0.018 (0.014)	0.027 (0.054)	0.118 (0.055)	-0.126 (0.093)
Beer	0.154 (0.069)	-0.119 (0.035)	0.171 (0.077)	0.240 (0.032)	0.056 (0.033)	0.036 (0.024)	0.019 (0.012)	0.026 (0.030)	0.075 (0.028)	-0.061 (0.025)
Blades	0.158 (0.048)	-0.250 (0.060)	0.177 (0.053)	0.243 (0.047)	0.062 (0.032)	0.042 (0.026)	0.019 (0.012)	0.022 (0.045)	0.110 (0.039)	-0.166 (0.092)
Carbonated beverages	0.261 (0.061)	-0.246 (0.041)	0.317 (0.071)	0.249 (0.038)	0.061 (0.024)	0.038 (0.016)	0.022 (0.010)	0.038 (0.023)	0.132 (0.028)	-0.115 (0.031)
Cigarettes	0.059 (0.043)	-0.118 (0.062)	0.061 (0.045)	0.241 (0.045)	0.159 (0.085)	0.122 (0.082)	0.036 (0.022)	0.031 (0.039)	0.077 (0.036)	-0.083 (0.047)
Coffee	0.204 (0.070)	-0.239 (0.050)	0.238 (0.077)	0.252 (0.038)	0.060 (0.031)	0.039 (0.025)	0.019 (0.011)	0.031 (0.038)	0.123 (0.039)	-0.121 (0.052)
Cold cereals	0.222 (0.068)	-0.324 (0.060)	0.297 (0.083)	0.266 (0.045)	0.042 (0.019)	0.028 (0.015)	0.013 (0.007)	0.040 (0.030)	0.116 (0.044)	-0.106 (0.055)
Deodorants	0.211 (0.055)	-0.311 (0.054)	0.240 (0.062)	0.242 (0.038)	0.049 (0.021)	0.028 (0.014)	0.020 (0.011)	0.011 (0.060)	0.169 (0.047)	-0.225 (0.090)
Diapers	0.246 (0.070)	-0.168 (0.039)	0.270 (0.077)	0.232 (0.050)	0.073 (0.043)	0.037 (0.032)	0.033 (0.021)	-0.001 (0.041)	0.096 (0.035)	-0.107 (0.057)
Facial tissue	0.229 (0.076)	-0.294 (0.072)	0.298 (0.091)	0.264 (0.062)	0.041 (0.025)	0.025 (0.020)	0.015 (0.012)	0.035 (0.064)	0.120 (0.069)	-0.105 (0.085)
Frozen dinners	0.289 (0.077)	-0.294 (0.063)	0.377 (0.092)	0.262 (0.042)	0.043 (0.019)	0.028 (0.013)	0.014 (0.008)	0.042 (0.034)	0.129 (0.049)	-0.113 (0.054)
Frozen pizza	0.309 (0.085)	-0.278 (0.052)	0.400 (0.096)	0.264 (0.046)	0.050 (0.025)	0.031 (0.017)	0.017 (0.010)	0.042 (0.032)	0.129 (0.047)	-0.112 (0.051)
Household cleaning	0.169 (0.060)	-0.245 (0.058)	0.194 (0.068)	0.255 (0.048)	0.039 (0.019)	0.024 (0.014)	0.014 (0.010)	0.023 (0.053)	0.117 (0.048)	-0.138 (0.086)
Hot dogs	0.241	-0.340	0.333	0.269	0.038	0.027	0.010	0.058	0.128	-0.104

	(0.082)	(0.067)	(0.098)	(0.056)	(0.023)	(0.018)	(0.008)	(0.048)	(0.054)	(0.065)
Laundry and detergents	0.247	-0.282	0.323	0.258	0.049	0.030	0.018	0.032	0.125	-0.116
	(0.065)	(0.049)	(0.077)	(0.042)	(0.021)	(0.016)	(0.009)	(0.035)	(0.047)	(0.059)
Margarine and butter	0.183	-0.253	0.237	0.261	0.047	0.034	0.012	0.041	0.088	-0.062
	(0.071)	(0.074)	(0.085)	(0.049)	(0.032)	(0.028)	(0.009)	(0.037)	(0.041)	(0.046)
Mayonnaise	0.145	-0.267	0.183	0.257	0.047	0.035	0.012	0.038	0.095	-0.089
	(0.063)	(0.083)	(0.074)	(0.052)	(0.032)	(0.028)	(0.010)	(0.047)	(0.046)	(0.069)
Milk	0.137	-0.172	0.158	0.243	0.057	0.038	0.018	0.028	0.072	-0.056
	(0.058)	(0.042)	(0.065)	(0.038)	(0.031)	(0.026)	(0.013)	(0.022)	(0.026)	(0.028)
Mustard and ketchup	0.153	-0.206	0.171	0.261	0.051	0.035	0.014	0.036	0.112	-0.116
	(0.067)	(0.066)	(0.071)	(0.044)	(0.027)	(0.022)	(0.011)	(0.051)	(0.039)	(0.073)
Paper towels	0.214	-0.254	0.291	0.259	0.049	0.031	0.017	0.034	0.103	-0.086
	(0.065)	(0.054)	(0.075)	(0.047)	(0.028)	(0.023)	(0.012)	(0.044)	(0.046)	(0.059)
Peanut butter	0.154	-0.217	0.195	0.261	0.050	0.033	0.016	0.028	0.084	-0.071
	(0.065)	(0.077)	(0.076)	(0.056)	(0.033)	(0.028)	(0.013)	(0.043)	(0.044)	(0.057)
Photo	0.186	-0.302	0.206	0.232	0.062	0.034	0.027	-0.027	0.161	-0.266
	(0.065)	(0.100)	(0.069)	(0.061)	(0.047)	(0.033)	(0.027)	(0.139)	(0.093)	(0.181)
Razors	0.243	-0.227	0.260	0.224	0.069	0.039	0.028	-0.003	0.131	-0.202
	(0.078)	(0.074)	(0.083)	(0.052)	(0.038)	(0.027)	(0.023)	(0.089)	(0.067)	(0.135)
Salty snacks	0.223	-0.251	0.264	0.265	0.035	0.023	0.011	0.048	0.128	-0.106
	(0.057)	(0.038)	(0.065)	(0.031)	(0.015)	(0.010)	(0.006)	(0.025)	(0.031)	(0.036)
Shampoo	0.214	-0.275	0.236	0.239	0.064	0.037	0.025	0.004	0.158	-0.218
	(0.048)	(0.043)	(0.052)	(0.033)	(0.026)	(0.016)	(0.013)	(0.049)	(0.036)	(0.075)
Soup	0.178	-0.300	0.226	0.256	0.049	0.034	0.014	0.041	0.126	-0.131
	(0.079)	(0.077)	(0.101)	(0.043)	(0.023)	(0.018)	(0.009)	(0.040)	(0.041)	(0.068)
Spaghetti sauce	0.214	-0.265	0.268	0.268	0.056	0.037	0.018	0.035	0.119	-0.113
	(0.073)	(0.067)	(0.084)	(0.048)	(0.027)	(0.020)	(0.012)	(0.043)	(0.042)	(0.062)
Sugar and substitutes	0.117	-0.197	0.133	0.249	0.041	0.027	0.013	0.019	0.095	-0.110
	(0.066)	(0.078)	(0.074)	(0.057)	(0.026)	(0.020)	(0.013)	(0.066)	(0.059)	(0.103)
Toilet tissue	0.218	-0.250	0.313	0.273	0.052	0.033	0.018	0.029	0.090	-0.071
	(0.063)	(0.052)	(0.075)	(0.048)	(0.030)	(0.024)	(0.011)	(0.035)	(0.039)	(0.043)
Toothbrushes	0.194	-0.323	0.219	0.239	0.047	0.026	0.020	0.000	0.192	-0.276
	(0.056)	(0.062)	(0.065)	(0.039)	(0.021)	(0.013)	(0.012)	(0.075)	(0.062)	(0.106)
Toothpaste	0.215	-0.288	0.255	0.252	0.044	0.026	0.017	0.024	0.153	-0.187
	(0.058)	(0.046)	(0.069)	(0.040)	(0.020)	(0.015)	(0.009)	(0.046)	(0.043)	(0.077)
Yogurt	0.224	-0.249	0.290	0.260	0.039	0.025	0.013	0.034	0.095	-0.073
	(0.074)	(0.062)	(0.089)	(0.054)	(0.023)	(0.017)	(0.009)	(0.036)	(0.041)	(0.042)

Notes: UPCs have equal weights when aggregated to the category level.

Appendix Table 2. Moments for all (sales and regular) price changes.

Category	All price changes					
	Frequency			Size		
	All	Positive	Negative	All	Positive	Negative
	(1)	(2)	(3)	(4)	(5)	(6)
Total	0.238 (0.093)	0.124 (0.051)	0.114 (0.051)	0.001 (0.044)	0.206 (0.071)	-0.224 (0.079)
Beer	0.168 (0.082)	0.091 (0.046)	0.077 (0.040)	0.008 (0.026)	0.097 (0.028)	-0.102 (0.031)
Blades	0.205 (0.066)	0.112 (0.040)	0.093 (0.035)	0.001 (0.042)	0.197 (0.043)	-0.238 (0.059)
Carbonated beverages	0.315 (0.083)	0.163 (0.045)	0.152 (0.042)	0.006 (0.021)	0.201 (0.035)	-0.206 (0.037)
Cigarettes	0.199 (0.091)	0.141 (0.084)	0.058 (0.034)	0.022 (0.040)	0.083 (0.036)	-0.098 (0.050)
Coffee	0.231 (0.076)	0.122 (0.046)	0.109 (0.040)	0.005 (0.036)	0.198 (0.043)	-0.211 (0.047)
Cold cereals	0.241 (0.077)	0.125 (0.042)	0.116 (0.040)	0.002 (0.031)	0.254 (0.057)	-0.275 (0.060)
Deodorants	0.260 (0.069)	0.130 (0.036)	0.130 (0.039)	-0.002 (0.042)	0.278 (0.046)	-0.297 (0.053)
Diapers	0.313 (0.087)	0.152 (0.053)	0.161 (0.050)	-0.011 (0.032)	0.144 (0.032)	-0.159 (0.040)
Facial tissue	0.253 (0.082)	0.125 (0.044)	0.128 (0.048)	-0.009 (0.054)	0.247 (0.064)	-0.260 (0.068)
Frozen dinners	0.319 (0.078)	0.162 (0.042)	0.158 (0.042)	0.002 (0.032)	0.246 (0.059)	-0.254 (0.061)
Frozen pizza	0.354 (0.092)	0.181 (0.050)	0.173 (0.048)	0.004 (0.028)	0.232 (0.049)	-0.239 (0.052)
Household cleaning	0.179 (0.060)	0.091 (0.033)	0.088 (0.035)	-0.006 (0.045)	0.203 (0.045)	-0.227 (0.056)
Hot dogs	0.300 (0.089)	0.157 (0.051)	0.143 (0.047)	0.009 (0.048)	0.279 (0.063)	-0.293 (0.067)
Laundry and detergents	0.289 (0.078)	0.147 (0.043)	0.142 (0.041)	0.002 (0.028)	0.236 (0.044)	-0.246 (0.047)
Margarine and butter	0.220 (0.076)	0.116 (0.046)	0.104 (0.040)	0.007 (0.044)	0.200 (0.066)	-0.213 (0.069)
Mayonnaise	0.179 (0.074)	0.099 (0.047)	0.080 (0.039)	0.008 (0.057)	0.197 (0.065)	-0.226 (0.076)
Milk	0.179 (0.069)	0.096 (0.040)	0.084 (0.037)	0.005 (0.025)	0.131 (0.035)	-0.141 (0.039)
Mustard and ketchup	0.153 (0.061)	0.085 (0.038)	0.068 (0.037)	0.008 (0.050)	0.161 (0.043)	-0.186 (0.061)
Paper towels	0.253 (0.078)	0.130 (0.045)	0.123 (0.042)	-0.001 (0.034)	0.203 (0.047)	-0.216 (0.050)
Peanut butter	0.182 (0.069)	0.098 (0.043)	0.085 (0.038)	0.007 (0.048)	0.164 (0.058)	-0.178 (0.069)
Photo	0.232 (0.079)	0.116 (0.049)	0.116 (0.052)	-0.021 (0.084)	0.254 (0.075)	-0.297 (0.100)

Razors	0.269 (0.093)	0.137 (0.052)	0.132 (0.059)	-0.014 (0.063)	0.189 (0.057)	-0.230 (0.076)
Salty snacks	0.227 (0.068)	0.116 (0.036)	0.111 (0.034)	0.002 (0.021)	0.207 (0.036)	-0.217 (0.038)
Shampoo	0.264 (0.071)	0.133 (0.037)	0.130 (0.039)	-0.009 (0.034)	0.241 (0.035)	-0.268 (0.043)
Soup	0.207 (0.080)	0.111 (0.043)	0.096 (0.045)	0.009 (0.054)	0.237 (0.066)	-0.260 (0.071)
Spaghetti sauce	0.245 (0.074)	0.129 (0.041)	0.117 (0.042)	0.006 (0.042)	0.218 (0.056)	-0.231 (0.063)
Sugar and substitutes	0.136 (0.068)	0.071 (0.038)	0.065 (0.039)	-0.003 (0.056)	0.153 (0.057)	-0.180 (0.072)
Toilet tissue	0.270 (0.081)	0.139 (0.047)	0.131 (0.041)	0.000 (0.028)	0.198 (0.047)	-0.211 (0.048)
Toothbrushes	0.227 (0.065)	0.112 (0.034)	0.116 (0.037)	-0.011 (0.050)	0.287 (0.054)	-0.316 (0.061)
Toothpaste	0.246 (0.071)	0.124 (0.038)	0.122 (0.038)	-0.002 (0.033)	0.254 (0.039)	-0.271 (0.045)
Yogurt	0.271 (0.082)	0.137 (0.045)	0.134 (0.044)	0.001 (0.034)	0.204 (0.060)	-0.210 (0.062)

Notes: UPCs have equal weights when aggregated to the category level.

Appendix Table 3. Alternative standard errors.

Dependent variable	Equal weights to all UPCs	Standard errors		
	Point estimate	Driscoll- Kraay	Cluster by category/market	Cluster by month
	(1)	(2)	(3)	(4)
Sales				
Frequency	0.224	(0.092)***	(0.060)***	(0.047)***
Size	0.303	(0.066)***	(0.065)***	(0.037)***
Share of goods bought on sale	0.161	(0.104)	(0.070)**	(0.053)***
Regular price				
Share of goods bought at cheap regular prices	0.126	(0.042)***	(0.046)***	(0.023)***
Frequency of changes				
All	-0.018	(0.032)	(0.022)	(0.018)
Positive	-0.041	(0.024)*	(0.017)***	(0.014)***
Negative	0.026	(0.011)**	(0.010)***	(0.007)***
Size of changes				
All	-0.109	(0.044)***	(0.028)***	(0.030)***
Positive	0.070	(0.050)	(0.043)	(0.028)***
Negative	-0.036	(0.060)	(0.065)	(0.037)
Inflation				
posted prices	-0.050	(0.017)***	(0.014)***	(0.009)***
effective prices	-0.219	(0.024)***	(0.026)***	(0.013)***


Notes: The table shows the baseline point estimates (column 1) and standard errors (column 2) which correspond to results presented in column 3 of Tables 1 and 2. Columns 3 and 4 show alternative estimates of standard errors associated with column (1). Column (3) clusters standard errors by market and category (1550 clusters) which allows for arbitrary collation of errors across time. Column (4) clusters standard errors by month (132 clusters) which allow for arbitrary cross-sectional correlation. *, **, *** indicates statistical significance at 1, 5 and 10 percent.

Appendix Table 4. Share of Regular-Priced Goods Bought at Low Prices

Universe of UPCs	Cheap regular price			
	Bottom 10 percent (1)	Bottom 20 percent (2)	Bottom 25 percent (3)	Bottom 33 percent (4)
Panel A: equal weights to aggregate UPCs to category level				
Ω_{all}	0.145*** (0.041)	0.126*** (0.042)	0.199*** (0.055)	0.324*** (0.072)
Ω_{max}	0.320*** (0.072)	0.140 (0.089)	0.213** (0.105)	0.346** (0.141)
Ω_{90}	0.283*** (0.060)	0.280*** (0.066)	0.376*** (0.078)	0.427*** (0.104)
Ω_{75}	0.195*** (0.050)	0.180*** (0.060)	0.279*** (0.076)	0.380*** (0.095)
Panel B: expenditure shares (market specific) as weights to aggregate UPCs to category level				
Ω_{all}	0.170*** (0.046)	0.191*** (0.055)	0.285*** (0.070)	0.369*** (0.090)
Ω_{max}	0.297*** (0.077)	0.097 (0.099)	0.179 (0.115)	0.312** (0.147)
Ω_{90}	0.306*** (0.063)	0.319*** (0.076)	0.428*** (0.089)	0.434*** (0.112)
Ω_{75}	0.198*** (0.054)	0.194*** (0.068)	0.288*** (0.082)	0.331*** (0.105)


Notes: The table presents estimated coefficients for the cyclicalities of the share of goods bought at “cheap” regular prices, using alternative thresholds for identifying “cheap” prices (indicated in column headers), alternative groupings of UPCs used to calculate relative prices (indicated by rows) and different aggregation procedures to go from UPC level to category-market level (Panel headings). *, **, *** indicates statistical significance at 1, 5 and 10 percent.

Appendix Figure 1. Correlations between key moments



Notes: Figures report average (across time and goods) moments at the category level. Expenditure shares are used as weights to aggregate goods. Red line shows the best fit linear projection. See section 2 for details.

Appendix Figure 2. Sensitivity of Output Gap Response to Contractionary Monetary Policy Shock



Notes: This figure shows the sensitivity of the output gap response to a monetary policy shock as a function of two parameters ϕ (elasticity of iceberg cost with respect to shopping effort) and γ (elasticity of substitution across stores). The bottom panel is for the case with $\phi = -0.35$. The mean response is the average response of output gap over 20 periods.

APPENDIX B: WITHIN-CATEGORY SUBSTITUTION

While we focus on expenditure-switching across stores by households for a given UPC product, the literature on price measurement has long emphasized another margin of substitution, namely across goods. Our primary motivation for focusing on switching across stores for a given good is that, as in the construction of the CPI, it is helpful to consider the cost of a fixed basket of goods for welfare purposes. The substitution bias long emphasized in the literature, in which CPI inflation will be overstated because it ignores the possibility of consumers switching goods when relative prices change, instead involves a change in the composition of the basket which will have implications for welfare. Nonetheless, we also consider this additional margin here for two reasons. First, the substitution bias has primarily been considered as a source of long-run bias in inflation measurements, while the cyclical properties of this margin have not been considered. Second, comparing the degree of store-switching to the amount of cross-good substitution provides one metric to assess the relative importance of store-switching for the measurement of inflation.

To quantify the degree of substitution across goods, we first construct the quantity-weighted average “effective” price *across all goods j within category c* in store s and geographic area m as

$$\bar{P}_{mcts}^{eq} = \frac{\sum_{j \in c} TR_{msctj}}{\sum_{j \in c} TQ_{msctj} \times EQ_j} \quad (\text{B.1})$$

where EQ_j is the quantity equivalent of good j . Hence, in calculating \bar{P}_{mcts}^{eq} , all prices are converted into quantity-equivalent measures so that e.g. the price of a 6-pack of beer is comparable to a 12-pack and \bar{P}_{mcts}^{eq} measures the price of beer per liter. \bar{P}_{mcts}^{eq} can change because individual prices change or because consumers reallocate their consumption of goods within a given category. For category c , store s and market m , we compute the monthly inflation rate $\log(\bar{P}_{mcs,t}^{eq}/\bar{P}_{mcs,t-1}^{eq})$. Then, we aggregate across all stores in market m to get the average category-level inflation rate, using either equal or expenditure weights.¹⁸ Finally, we cumulate monthly inflation rates into annual inflation rate $\bar{\pi}_{mct}^{eq}$ which we refer to as the “within-category effective inflation rate”. While for $\bar{\pi}_{mct}^{eff}$ we fix the composition of the consumption basket but allow consumers to switch stores, for $\bar{\pi}_{mct}^{eq}$ we fix the store weights but allow consumers to substitute goods in the basket.

Because some categories include much more heterogeneity in goods than others, we consider two classification schemes for measuring the substitution of goods within categories. The first (and broadest) includes all UPCs within a category. The second allows substitution only within subcategories which approximately corresponds to adding another digit to the level of disaggregation. For example, we use all

¹⁸ Because quantity equivalents are not available or are not comparable for some categories, we exclude the following categories from this analysis: deodorants, frozen dinners, photos, and soups.

types of milk when we calculate \bar{P}_{mcts}^{eq} for the first classification. In contrast, the second classification considers separately such subcategories as whole milk, skimmed milk, 2% milk, etc.

The sensitivity of these inflation rates to economic conditions is then assessed using

$$\bar{\pi}_{mct}^{eq} = \beta UR_{mt} + \lambda_t + \theta_{m,c} + error \quad (5.2)$$

which is equivalent to the specification used to measure the sensitivity of effective across-store inflation rates to economic conditions. The results, presented in Appendix Table B.1 below, point to a statistically significantly negative relationship between unemployment rates and within-category effective inflation rates. Thus, as in our baseline results, this indicates significant substitution by households in response to changing local economic conditions but along a different margin, namely substituting across different goods within a category. Importantly, the quantitative magnitudes are of the same order as those identified for across-store substitution.

Appendix Table B.1. Within category substitution.

	Equal weights for all stores	Sales shares to aggregate stores
	(1)	(2)
Substitution within broad categories	-0.346*** (0.091)	-0.332*** (0.107)
Substitution within narrower categories	-0.309*** (0.064)	-0.397*** (0.101)

Notes: The table reports estimates of specification (5.2). The dependent variable is the within-category effective inflation rate. The table reports estimated coefficients on the local seasonally-adjusted unemployment rate. Number of observations is 94,851. Driscoll and Kraay (1998) standard errors are in parentheses. ***, **, * denote significance at 0.01, 0.05, and 0.10 levels.

APPENDIX C: STORE-SWITCHING AND MONETARY NON-NEUTRALITY

Our model can be reduced to

$$\text{Phillips curve: } \pi_t = \beta E_t \pi_{t+1} + \kappa x_t$$

$$\text{IS curve: } x_t = E_t x_{t+1} - \psi(i_t - E_t \pi_{t+1} - r_t^n)$$

where ψ is the elasticity of the output gap to real interest rates and depends on structural parameters, including store-switching parameters. Assume a simple Taylor rule

$$i_t = \chi \pi_t + \varepsilon_t$$

where ε is monetary policy shock which follows AR(1) $\varepsilon_t = \rho \varepsilon_{t-1} + v_t$, $\chi > 1$, and $0 \leq \rho < 1$.

We guess that the solution of the model to MP shocks (s.t. natural interest rate does not vary) is

$$\pi_t = a \varepsilon_t$$

$$x_t = b \varepsilon_t$$

Then the Phillips curve becomes

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t \Leftrightarrow a \varepsilon_t = \beta E_t a \varepsilon_{t+1} + \kappa b \varepsilon_t \Leftrightarrow b = a(1 - \beta\rho)/\kappa$$

The IS curve becomes

$$\begin{aligned} x_t = E_t x_{t+1} - \psi(i_t - E_t \pi_{t+1} - r_t^n) &\Leftrightarrow b \varepsilon_t = \rho b \varepsilon_t - s(\chi a \varepsilon_t + \varepsilon_t - a \rho \varepsilon_t) \\ &\Leftrightarrow (1 - \rho)b = -\psi a(\chi - \rho) - \psi \end{aligned}$$

Plugging in expression for b from above implies

$$\begin{aligned} (1 - \rho)b = -\psi a(\chi - \rho) - \psi &\Leftrightarrow \frac{(1 - \rho)a(1 - \beta\rho)}{\kappa} = -\psi a(\chi - \rho) - \psi \\ &\Leftrightarrow a \left[\frac{(1 - \rho)(1 - \beta\rho)}{\kappa} + \psi(\chi - \rho) \right] = -\psi \\ &\Leftrightarrow a = -\frac{\psi}{\left[\frac{(1 - \rho)(1 - \beta\rho)}{\kappa} + \psi(\chi - \rho) \right]} \end{aligned}$$

And therefore

$$b = \frac{a(1 - \beta\rho)}{\kappa} = -\frac{\psi(1 - \beta\rho)}{\kappa \left[\frac{(1 - \rho)(1 - \beta\rho)}{\kappa} + \psi(\chi - \rho) \right]} = \frac{-\psi(1 - \beta\rho)}{[(1 - \rho)(1 - \beta\rho) + \psi\kappa(\chi - \rho)]} < 0$$

Thus, $\frac{\partial b}{\partial \psi} < 0$, the degree of monetary non-neutrality (absolute value of b) is increasing in the elasticity of the output gap to the real interest rate.

In our model, this elasticity is given by (see equation 4.21):

$$\psi = \left[1 + \frac{1}{2} \left(\frac{\sigma - 1}{\sigma \mu} \right) \left(\frac{1}{1 - \frac{1}{\phi} - \frac{\gamma}{2}} \right) \right]^{-1}$$

such that a more negative value of ϕ (i.e. more elastic shopping effort) or a greater value of γ reduces ψ . Hence, increasing the degree of store-switching in the model reduces the elasticity of the output gap to the real interest rate and therefore lowers the degree of monetary non-neutrality in the model.

APPENDIX D: SECOND ORDER APPROXIMATION TO UTILITY WITH ENDOGENOUS SHOPPING EFFORT AND STORE-SWITCHING

The second-order approximation to utility is:

$$\begin{aligned}
U_t &= \log C_t + \log(1 - L_t - S_t) \\
&\approx \log \bar{C} + \frac{1}{\bar{C}}(C_t - \bar{C}) - \frac{1}{2} \left(\frac{1}{\bar{C}}\right)^2 (C_t - \bar{C})^2 + \log(1 - \bar{L} - \bar{S}) + \frac{-1}{1 - \bar{L} - \bar{S}}(L_t - \bar{L}) + \frac{-1}{1 - \bar{L} - \bar{S}}(S_t - \bar{S}) - \\
&\quad \frac{1}{2} \left(\frac{1}{1 - \bar{L} - \bar{S}}\right)^2 (L_t - \bar{L})^2 - \frac{1}{2} \left(\frac{1}{1 - \bar{L} - \bar{S}}\right)^2 (S_t - \bar{S})^2 - \left(\frac{1}{1 - \bar{L} - \bar{S}}\right)^2 (L_t - \bar{L})(S_t - \bar{S}) + h.o.t. \\
&= t.i.p. + (\check{C}_t + \frac{1}{2}\check{C}_t^2) - \frac{1}{2}\check{C}_t^2 - \eta_L(\check{L}_t + \frac{1}{2}\check{L}_t^2) - \eta_S(\check{S}_t + \frac{1}{2}\check{S}_t^2) - \frac{1}{2}\eta_L^2\check{L}_t^2 - \frac{1}{2}\eta_S^2\check{S}_t^2 - \eta_L\eta_S(\check{L}_t + \\
&\quad \frac{1}{2}\check{L}_t^2)(\check{S}_t + \frac{1}{2}\check{S}_t^2) + h.o.t. \\
&= t.i.p. + \check{C}_t - \eta_L\check{L}_t - \frac{1}{2}\eta_L(1 + \eta_L)\check{L}_t^2 - \eta_S\check{S}_t - \frac{1}{2}\eta_S(1 + \eta_S)\check{S}_t^2 - \bar{L}\eta_S\check{L}_t\check{S}_t + h.o.t.
\end{aligned}$$

where $\eta_L \equiv \frac{\bar{L}}{1 - \bar{L} - \bar{S}} = \left(\frac{\sigma-1}{\sigma}\right) \left(\frac{\alpha}{\mu}\right)$, $\eta_S \equiv \frac{\bar{S}}{1 - \bar{L} - \bar{S}} = -\frac{\phi}{2}$ when $\bar{r} = \mu$, and following Woodford (2003), $X_t - \bar{X} = \bar{X} \left(\check{X}_t + \frac{1}{2}\check{X}_t^2\right)$.

From the derivation of the model in the text with $\bar{r} = \mu$, we have $\check{Y}_t = \check{X}_t + \check{Z}_t$ and $\check{S}_t = \frac{\alpha}{\mu} \left(\frac{\sigma-1}{\sigma}\right) \frac{1}{\phi-1-\frac{1}{2}\gamma\phi} \check{L}_t$. Given the production function $Y_t(i) = Z_t N_t(i)^\alpha$, it follows that

$$\begin{aligned}
\Rightarrow N_t &= \int_0^1 N_t(i) di = \int_0^1 \left(\frac{Y_t(i)}{Z_t}\right)^{1/\alpha} di = \left(\frac{Y_t}{Z_t}\right)^{1/\alpha} \int_0^1 \left(\frac{Y_t(i)}{Y_t}\right)^{1/\alpha} di = \left(\frac{Y_t}{Z_t}\right)^{1/\alpha} \int_0^1 \left(\frac{P_t(i)}{P_t}\right)^{-\sigma/\alpha} di \\
&\Rightarrow Y_t = Z_t N_t^\alpha \left\{ \int_0^1 \left(\frac{P_t(i)}{P_t}\right)^{-\frac{\sigma}{\alpha}} di \right\}^{-\alpha} \\
&\Rightarrow \check{Y}_t = \check{Z}_t + \alpha \check{N}_t - d_t
\end{aligned}$$

where $d_t \equiv \alpha \log \left\{ \int_0^1 \left(\frac{P_t(i)}{P_t}\right)^{-\frac{\sigma}{\alpha}} di \right\}$. From the labor market clearing conditions, we have $N_t = L_t \Rightarrow \check{N}_t = \check{L}_t$. It follows that

$$\check{L}_t = \check{N}_t = \frac{1}{\alpha} (\check{Y}_t - \check{Z}_t + d_t) = \frac{1}{\alpha} \check{X}_t + \frac{1}{\alpha} d_t$$

Also, one can show (see Galí 2008) that $d_t = \alpha \log \left\{ \int_0^1 \left(\frac{P_t(i)}{P_t}\right)^{-\frac{\sigma}{\alpha}} di \right\} \approx \frac{1}{2} \frac{\sigma}{\theta} \text{var}_i(\log(P_t(i)))$. Define $\Delta_t \equiv \text{var}_i(\log(P_t(i)))$. Woodford (2003) shows that

$$\Delta_t = \theta \Delta_{t-1} + \frac{\theta}{1 - \theta} \pi_t^2 + hot$$

Hence, $E(\Delta_t) = \frac{\theta}{(1-\theta)^2} E(\pi_t^2)$. Thus, $E(d_t) \approx \frac{1}{2} \frac{\sigma}{\theta} \frac{\theta}{(1-\theta)^2} E(\pi_t^2)$.

Finally, for the link between consumption and output, given $C_t = \left(C_{A,t}^{\frac{\gamma-1}{\gamma}} + C_{B,t}^{\frac{\gamma-1}{\gamma}} \right)^{\gamma/(\gamma-1)}$, $(C_{A,t} + C_{B,t}) = Y_t$, and $\frac{C_{A,t}}{C_{B,t}} = \left\{ \frac{\tau_t}{\mu} \right\}^\gamma$, we have $C_t = Y_t \Phi_t$ where $\Phi_t \equiv \left(\left(\frac{\tau_t}{\mu} \right)^{\gamma-1} + 1 \right)^{\frac{\gamma}{\gamma-1}} / \left\{ \left(\frac{\tau_t}{\mu} \right)^\gamma + 1 \right\}$. So

$$C_t - \bar{C} = \bar{\Phi}(Y_t - \bar{Y}) + \bar{Y} \left. \frac{\partial \Phi_t}{\partial \tau_t} \right|_{\tau_t = \bar{\tau}} (\tau_t - \bar{\tau}) + 0(Y_t - \bar{Y})^2 + \frac{1}{2} \bar{Y} \left. \frac{\partial^2 \Phi_t}{\partial \tau_t^2} \right|_{\tau_t = \bar{\tau}} (\tau_t - \bar{\tau})^2 + \left. \frac{\partial \Phi_t}{\partial \tau_t} \right|_{\tau_t = \bar{\tau}} (\tau_t - \bar{\tau})(Y_t - \bar{Y})$$

Note that steady-state levels with $\bar{\tau} = \mu$ are given by $\bar{\Phi} = 2^{\frac{1}{\gamma-1}}$, $\left. \frac{\partial \Phi_t}{\partial \tau_t} \right|_{\tau_t = \bar{\tau}} = 0$, and $\left. \frac{\partial^2 \Phi_t}{\partial \tau_t^2} \right|_{\tau_t = \bar{\tau}} = -\frac{\gamma}{\mu^2} \frac{\bar{\Phi}}{4}$, so

$$C_t - \bar{C} = \bar{\Phi} \bar{Y} \check{Y}_t + 0 + 0 - \frac{1}{8} \gamma \bar{\Phi} \bar{Y} \check{\tau}_t^2 \\ \Rightarrow \check{C}_t = \check{Y}_t - \frac{1}{8} \gamma \check{\tau}_t^2 = \check{Y}_t - \frac{1}{8} \gamma \phi^2 \check{S}_t^2 = \check{X}_t + \check{Z}_t - \frac{1}{8} \gamma \phi^2 \zeta^2 \check{X}_t^2$$

Now substituting into our approximation to utility yields

$$U_t = \log C_t + \log(1 - L_t - S_t) \\ \approx t.i.p. + \check{C}_t - \eta_L \check{L}_t - \frac{1}{2} \eta_L (1 + \eta_L) \check{L}_t^2 - \eta_S \check{S}_t - \frac{1}{2} \eta_S (1 + \eta_S) \check{S}_t^2 - \eta_L \eta_S \check{L}_t \check{S}_t + h.o.t. \\ = t.i.p. + \left(\check{X}_t - \check{Z}_t + -\frac{1}{8} \gamma \phi^2 \zeta^2 \check{X}_t^2 \right) - \left(\frac{\sigma-1}{\sigma} \right) \left(\frac{\alpha}{\mu} \right) \left(\frac{\check{X}_t}{\alpha} + \frac{d_t}{\alpha} + \frac{1}{2} \left\{ \frac{\check{X}_t}{\alpha} + \frac{d_t}{\alpha} \right\}^2 \right) + \frac{\phi}{2} (\zeta \check{X}_t + \zeta d_t + \frac{1}{2} \{ \zeta \check{X}_t + \zeta d_t \}^2) - \frac{1}{2} \left(\left(\frac{\sigma-1}{\sigma} \right) \left(\frac{\alpha}{\mu} \right) \right)^2 \left\{ \frac{\check{X}_t}{\alpha} + \frac{d_t}{\alpha} \right\}^2 - \left(\frac{\sigma-1}{\sigma} \right) \left(\frac{\alpha}{\mu} \right) \frac{\phi}{2} \left\{ \frac{\check{X}_t}{\alpha} + \frac{d_t}{\alpha} \right\} \{ \zeta \check{X}_t + \zeta d_t \} + h.o.t. \\ = t.i.p. + \left[1 - \left(\frac{\sigma-1}{\sigma} \right) \left(\frac{1}{\mu} \right)^{\frac{1}{2} \frac{\phi(1-\gamma)-1}{\phi - \frac{1}{2} \phi \gamma - 1}} \right] \check{X}_t + \left[- \left(\frac{\sigma-1}{\sigma} \right) \left(\frac{1}{\mu} \right)^{\frac{1}{2} \frac{\phi(1-\gamma)-1}{\phi - \frac{1}{2} \phi \gamma - 1}} \right] d_t - \frac{1}{2} \left(\frac{\sigma-1}{\sigma} \right) \left(\frac{1}{\mu} \right) \left[\frac{1}{\alpha} + \frac{1}{4} \gamma \phi^2 \left(\frac{\sigma-1}{\sigma} \right) \left(\frac{1}{\mu} \right) \frac{1}{(\phi - \frac{1}{2} \phi \gamma - 1)^2} - \frac{\phi}{2} \left(\frac{\sigma-1}{\sigma} \right) \left(\frac{1}{\mu} \right) \frac{1}{(\phi - \frac{1}{2} \phi \gamma - 1)^2} \right] \check{X}_t^2 - \frac{1}{2} \left\{ \left(\frac{\sigma-1}{\sigma} \right) \left(\frac{1}{\mu} \right)^{\frac{1}{2} \frac{\phi(1-\gamma)-1}{\phi - \frac{1}{2} \phi \gamma - 1}} \right\}^2 \check{X}_t^2 + h.o.t.$$

Therefore expected utility is

$$EU_t \approx -\frac{1}{2} \left\{ \left(\frac{\sigma-1}{\sigma \mu} \right)^{\frac{1}{2} \frac{\phi(1-\gamma)-1}{\phi - \frac{1}{2} \phi \gamma - 1}} \frac{\sigma}{\theta(1-\theta)^2} \right\} E(\pi_t^2) - \frac{1}{2} \left\{ \left(\frac{\sigma-1}{\sigma \mu} \right) \left[\frac{1}{\alpha} + \frac{1}{4} \gamma \phi^2 \left(\frac{\sigma-1}{\sigma \mu} \right) \frac{1}{(\phi - \frac{1}{2} \phi \gamma - 1)^2} - \frac{\phi}{2} \left(\frac{\sigma-1}{\sigma \mu} \right) \frac{1}{(\phi - \frac{1}{2} \phi \gamma - 1)^2} \right] + \left\{ \left(\frac{\sigma-1}{\sigma \mu} \right)^{\frac{1}{2} \frac{\phi(1-\gamma)-1}{\phi - \frac{1}{2} \phi \gamma - 1}} \right\}^2 \right\} E(\check{X}_t^2) + t.i.p. + h.o.t.$$

which after rearranging yields the expression in the text.

APPENDIX E: CONSTRUCTION OF DATA MOMENTS

This appendix describes how we constructed moments for the empirical analysis in sections II through IV.

Frequency of sales. The IRI dataset provides a flag to indicate whether a given good was on sale in a given store in a given week. In addition to this flag, we use filters as in Nakamura and Steinsson (2008). Specifically, if a price is reduced temporarily (up to three weeks) and then returns to the level observed before the price cut, we identify this episode as a sale. When we apply this filter, we use two approaches to identify a price spell. In the first approach (approach “A”), we treat missing values as interrupting price spells. In other words, if a price was \$4 for two weeks, then the price was missing for a week, and then was again observed at \$4 for another three weeks, we treat the data as reporting two price spells with durations of two and three weeks. In the second approach (approach “B”), missing values do not interrupt price spells if the price is the same before and after periods of missing values. For example, in the previous example, approach “B” yields one price spell with a duration of five weeks. To identify the incidence of sales, we use the union of sales flags that we obtain from the IRI dataset directly and from applying approaches “A” and “B”. In the end, using approaches “A” and “B” does not materially change the incidence of sales identified by the sales flag provided in the IRI dataset.

The frequency of sales is computed at the monthly frequency as the fraction of weeks in a month when a good is identified by the sales flag as being on sale. For example, suppose that the time series for a price is observed for eight weeks $\{4,2,4,4,4,2,2,4\}$ with sales flag series $\{0,1,0,0,0,1,1,0\}$, that is, sales occur in weeks #2, #6, and #7 when the price is cut from \$4 to \$2. Then, the frequencies of sales in the first and second months are $\frac{1}{4}$ and $\frac{1}{2}$ respectively.

Size of sales. The size of the sale is computed as the (log) difference between the sales price (the incidence of a sale is identified by the sales flag) and the price preceding the sale. Since one may have more than one sale in a month, we take the average size of sales in a month. For example, suppose that the time series for a price is observed for eight weeks $\{4,2,4,4,3,4,2,4\}$ with sales flag series $\{0,1,0,0,1,0,1,0\}$, that is, sales occur in weeks #2 (the price is cut from \$4 to \$2), #5 (the price is cut from \$4 to \$3), and #7 (the price is cut from \$4 to \$2). Then, the size of the sale in the first instance of a sale is $\log(\$2/\$4)$, which is a negative magnitude. Since this is the only sale in the first month, the average size of sales in this month is $\log(\$2/\$4)$. In the second month, there are two sales with sizes $\log(\$3/\$4)$ and $\log(\$2/\$4)$. The average size of sales in the second month is taken as the arithmetic average $0.5 \cdot \{\log(\$3/\$4) + \log(\$2/\$4)\}$. In some cases, we can observe no price preceding a sale but we observe a price immediately after the sale. In this situation, we calculate the size of a sale as the log price during a sale minus the log price immediately after the sale.

Share of goods sold on sale. The share of goods sold on sale for a given good in a given market (or store) in a given month is calculated as the following ratio. The numerator is the number of units sold during episodes identified as sales by the sales flag. The denominator is the total number of sold units.

Frequency of regular price changes. A price change in a given week is identified as regular if the following criteria are satisfied: i) the sales flag does not identify this week as a period when a good is on sale; ii) the sales flag does not identify the preceding week as a period when a good is on sale; iii) the price change is larger than one cent or one percent in absolute value (or more than 0.5 percent for prices larger than \$5). The last criterion removes small price changes which could arise from rounding errors

and the like. Again, we use approaches “A” and “B” to identify the price in the preceding period. The incidence of regular price changes is the union of incidents identified by “A” and “B”. The frequency of regular price changes is computed at the monthly frequency as the fraction of weeks in a month when a good is identified as having a regular price change.

Size of regular price changes. The size of a regular price change is computed as the (log) difference between the price in the period identified as having a regular price change and the price in the preceding period (using approach “A” and “B” to identify the preceding period). Since one may have more than one regular price change in a month, we take the average size of regular price changes in a month.

Share of goods bought at cheap regular prices. For each month, market and UPC, we construct the cross-sectional distribution of regular prices. We define a regular price as cheap if the price falls in the bottom X% of the cross-sectional distribution (e.g., bottom 20%). The share of goods bought at cheap regular prices is the share of units sold at cheap regular prices in the total number of units sold at regular prices.

Weighting. To aggregate across goods to the category level, we employ three weighting schemes: *i*) equal weights; *ii*) expenditure shares for a given market and year (“market specific”); *iii*) cross-market expenditure shares for a given year (“common”). The market-specific expenditure share weights are calculated as

$$\omega_{mctj} = \frac{\sum_s TR_{msctj}}{\sum_k \sum_s TR_{msctk}}$$

and the common expenditure share weights are calculated as

$$\omega_{ctj} = \frac{\sum_m \sum_s TR_{msctj}}{\sum_m \sum_k \sum_s TR_{msctk}}$$

where m , s , c , t , and j index markets, stores, product categories, time, and UPC, and TR_{msctj} is the revenue from selling good j in the year covering month t .

To aggregate across goods and categories to the store level, we employ two weighting schemes: *i*) equal weights; *ii*) expenditure shares for a given store, good and year. The expenditure share weights are calculated as

$$\omega_{msctj} = \frac{TR_{msctj}}{\sum_k \sum_c TR_{msctk}}$$

where TR_{msctj} is the revenue from selling good j in the year covering month t .