

Microeconomic evidence on demand-side real rigidity and implications for monetary non-neutrality*

Günter W. Beck[†] and Sarah M. Lein[‡]

7 July, 2016

Abstract

To model the observed slow response of aggregate real variables to nominal shocks, most macroeconomic models incorporate real rigidities in addition to nominal rigidities. One popular way of modelling such a real rigidity is to assume a non-constant demand elasticity. By using a homescan data set for three European countries, including prices and quantities bought for a large number of goods, in addition to consumer characteristics, we provide estimates of price elasticities of demand and on the degree of demand-side real rigidities. We find that price elasticities of demand are about 4 in the median. Furthermore, we find evidence for demand-side real rigidities. These are, however, much smaller than what is often assumed in macroeconomic models. The median estimate for demand-side real rigidity, the super-elasticity, is in a range between 1 and 2. To quantitatively assess the implications of our empirical estimates, we calibrate a menu-cost model with the estimated super-elasticity. We find that the degree of monetary non-neutrality doubles in the model including demand-side real rigidity, compared to the model with only nominal rigidity, suggesting a multiplier effect of around two. However, the model can explain only up to 6% of the monetary non-neutrality observed in the data, implying that additional multipliers are necessary to match the behavior of aggregate variables.

JEL classification: E30, E31, E50, D12, C3

Keywords: Demand curve, price elasticity, super-elasticity, price-setting, monetary non-neutrality.

*We are very grateful to AiMark and Europanel for providing us with the data used in this paper. We are in particular thankful to Alfred Dijs and Bernadette van Ewijk for their invaluable assistance in dealing with the data and their very useful suggestions. We thank Dave DeJong, Bo Honoré, Daniel Kaufmann, Emi Nakamura and the participants of the 5th Ifo Conference on Macroeconomics and Survey data, the SSES annual meeting 2015, the EEA Mannheim, the Verein für Socialpolitik Münster, and at seminars at the University of Basel and the ETH Zurich for helpful comments. Günter Beck furthermore would like to thank the Institute for Monetary and Financial Stability (IMFS) for its hospitality while part of this research was carried out.

[†]University of Siegen, Hölderlinstrasse 3, D-57076 Siegen, Germany, e-mail: guenter.beck@uni-siegen.de.

[‡]University of Basel, Department of Business and Economics, Peter Merian-Weg 6, CH - 4002 Basel, Switzerland, e-mail: sarah.lein@unibas.ch and KOF ETH Zürich, Leonhardstrasse 21, 8092 Zurich.

1 Introduction

Empirical studies employing macro data generally come to the conclusion that aggregate nominal shocks have relatively large and persistent real effects.¹ To explain this sizable degree of monetary non-neutrality most state-of-the-art macro models assume some form of nominal rigidity, most often in the form of stickiness in goods prices. To match the empirically observed patterns of the response of real variables to nominal shocks a considerable degree of nominal rigidity is required though (Maćkowiak and Smets, 2009).

However, there is now broad-based evidence that prices at the micro level change quite frequently (Bils and Klenow, 2004; Nakamura and Steinsson, 2008; Dhyne et al., 2005) and that nominal price rigidity alone thus cannot explain the large real effects of monetary policy. Nakamura and Steinsson (2010), for example, show that a menu-cost model with only nominal rigidity explains only about 2% of the U.S. business cycle when calibrated to the observed frequency of price adjustment.²

Taken together, the micro-macro disconnect suggests that multiplier effects, which amplify the effect of nominal rigidity at the micro level and increase persistence of aggregate variables, might be important. One way of introducing such a multiplier effect is the assumption of some form of real rigidity (Ball and Romer, 1990). Because real rigidities slow down the response of the desired price to a shock, the micro price flexibility and the macro price rigidity can be reconciled.

There are two classes of real rigidities (Klenow and Willis, 2006). One is the class of supply-side real rigidities, which include real-wage rigidities (as used, e.g., by Blanchard and Gali, 2007), firm-specific inputs (as used, e.g., in Gertler and Leahy, 2008, or Altig et al., 2011) or the use of intermediate inputs as proposed by Basu (1995) and employed by Nakamura and Steinsson (2010). The other class is made of demand-side real rigidities, which relax the assumption of constant-elasticity demand curves a la Dixit and Stiglitz (1977). This concept was introduced into the macro literature by Kimball (1995), examples for the application of such so-called Kimball-type demand curves include Eichenbaum and Fisher (2005), Dotsey

¹See Christiano et al. (1999) for an overview of this literature.

²See Table 10 in Nakamura and Steinsson (2010).

and King (2005), or Gopinath and Itskhoki (2010), amongst others. One reason for the popularity of the demand-side approach is that it generates an almost arbitrary amount of persistence, if the parameters governing real rigidities are chosen appropriately.

If firms face demand curves with a non-constant elasticity of substitution, the response of prices to shocks is slowed down. The basic mechanism is that consumers demand relatively few goods with high relative prices but do not excessively demand goods with low relative prices (Klenow and Willis, 2006). That is, the price elasticity is a positive function of the relative price and profits decline more steeply around the profit maximizing relative price of one, compared to the case of a constant-elasticity of demand. This makes firms' profits more sensitive to prices of competitors and thereby gives rise to strategic complementarities in price-setting. Furthermore, with non-constant demand elasticities, desired markups become variable. Suppose the desired markup decreases if the relative price of a firm increases: in this case, a positive cost shock leads to a smaller response of a firm's price if the firm has a high relative price, as compared to a firm, which has a low relative price. This leads to a sluggish response to shocks in the aggregate, too. In particular, this channel has been emphasized in the open economy literature, because it can generate incomplete exchange-rate pass-through (see Burstein and Gopinath, 2013, for example).

The extent of real rigidity is measured by the super-elasticity, which is the price elasticity of the price elasticity (Klenow and Willis, 2006). Macroeconomics models embracing Kimball-type preferences have used a broad range of values for the calibration of the super-elasticity parameter. Chari et al. (2000) point out that the preference specifications employed by most of these studies imply highly implausible behavior on the side of consumers. To illustrate their argument, Figure 1 plots the demand and profit functions which result from parameters for the price elasticities and super-elasticities assumed in recent contributions to this literature.³

The values for the super-elasticity of demand assumed in most studies imply very strong responses of consumers to even very small changes in relative prices. In all cases but the one of Bergin and Feenstra (2000), an increase in the relative price of only around 10% implies a

³The examples are taken from the survey presented in Table 1 of Dossche et al. (2010).

Figure 1: Demand and profit functions for various values of the price elasticity and the curvature of demand as employed in the literature




Figure 1 plots the demand functions and the resulting profit functions implied by the values of the elasticities and super-elasticities of demand as employed in the papers referred to in the legend. The representation and functional form of the Kimball aggregator is taken from Klenow and Willis (2006). These article include Kimball (1995), Chari et al. (2000), Bergin and Feenstra (2000), Eichenbaum and Fisher (2005), Coenen et al. (2007), Woodford (2003) and Klenow and Willis (2006). The selection of papers is taken from Dossche et al. (2010), who provide an overview of the implied parameters for price elasticities and super-elasticities of these studies. The numbers in the brackets after the author name(s) indicate the value of the elasticity and super-elasticity assumed.

complete extinction of a firm’s sales.

Since the estimation of the super-elasticity parameter imposes high data requirements, there exist only very few empirical studies on this topic so far. Some indirect evidence on the extent of demand-side real rigidity is provided in Klenow and Willis (2006). They show that demand-side real rigidities, which generate a reasonable amount of aggregate persistence, also generate frequent price adjustments, but are difficult to reconcile with the large absolute average size of price adjustments.⁴ As a consequence, these authors conclude that demand-side real rigidities are arguably not very large.⁵

⁴This is because strategic complementarities lead to a larger desire of firms to keep their price close to their competitors prices. If an aggregate shock hits the economy and firms adjust to the shock, they do not fully incorporate the shock in the first price adjustment round due to the joint presence of nominal and real rigidities, which lead to a coordination failure. Instead, firms that adjust take into account the aggregate price level and therefore largely the price of non-adjusters. Over time, more and more firms adjust to the shock and the aggregate price level moves slowly. This also implies that prices of a single firm respond only partially to the aggregate shock and therefore it takes a couple of price adjustments to fully incorporate the aggregate shock. Prices at the micro level then change frequently, but only by small amounts, conditional on adjustment. One should thus observe many small price changes in the data.

⁵Nakamura and Steinsson (2010) show that this criticism does not apply to supply-side real rigidities using roundabout production structure. See Nakamura and Steinsson (2010) for a discussion of how the type of real rigidity and the size of price changes interact.

Direct evidence for single products or retailers on non-constant elasticity demand curves can be obtained by estimating demand curves and the curvature of demand. The latter is closely related to the size of the super-elasticity. Using Dominick's Finer Foods data, Goldberg and Hellerstein (2013) estimate a super-elasticity for beer of 0.8. Nakamura and Zerom (2010) match data for retail and wholesale prices with commodity price data for coffee and estimate a median super-elasticity of demand of 4.6. Using scanner data from six stores of a European retailer, Dossche et al. (2010) find super-elasticities in the range of 4 for goods with an elasticity of 3 or larger. This is the case for five product categories: smoked salmon, wine, airing cupboards, tap, and toilet soap. Verhelst and Poel (2012) largely confirm their result.

In this paper, we add to this literature by providing some broader-based evidence on the amount of demand-side real rigidity in the data. To this end, we estimate the price elasticity of demand and the super-elasticity of demand from a rich homescanner data set for 58 goods categories in three European countries: Belgium, Germany and the Netherlands. The national data providers normally use the data to sell consultancy services to companies in the area of marketing. This implies that the analysis of this data set is basically intended to serve the same objectives as the ones of our study, i.e. to investigate consumer demand behavior and to inform firms about the demand for their products. One advantage of homescanner data is that we can directly track consumers' shopping behavior over time and we can control for important consumer characteristics, such as income or age, which arguably influence the shape and location of demand curves.

We find the median demand elasticity in the range of 3 to 4, in line with the IO literature, but somewhat lower than the demand elasticities that are usually implemented in macroeconomic models.⁶ We furthermore document that there is heterogeneity across goods and households. Nevertheless, 75% of our elasticity estimates, which are allowed to vary over goods categories and households, are below ten. Our estimates of the super-elasticity suggest that it is between 1 to 2, in line with Bergin and Feenstra (2000) and the findings of Goldberg

⁶There are exceptions, which use lower demand elasticities, more in line with our estimates, such as Nakamura and Zerom (2010), or Gopinath and Itskhoki (2010), for example. However, setting the demand elasticity to ten seems to be the prevalent parameter choice in macro models.

and Hellerstein (2013). These estimates imply that the desired markup decreases by 1% with a 1% increase in the relative price.

To quantify the degree of monetary non-neutrality implied by our estimates we calibrate a menu-cost model augmented with demand-side real rigidity . We first calibrate the model with nominal rigidity only and then add the real rigidity with the elasticity and super-elasticity parameters equal to our median estimates. In line with Nakamura and Steinsson (2010), we find that the model with nominal rigidities alone generates a very small degree of monetary non-neutrality, only around 3% of the variance of detrended real consumption can be generated by the model. This is largely because the price changes in our data are quite frequent and sizeable. Adding the real rigidity raises the variance of output that can be generated by the model by a factor of 2.

Our results support previous empirical evidence and the criticisms of Chari et al. (2000) or Klenow and Willis (2006): demand-side real rigidities in the data are rather small. As expected, demand side real-rigidities and nominal rigidities alone are not large enough to match both micro and macro facts on prices. Arguably, as also conjectured by Gopinath and Itskhoki (2011), additional sources of aggregate persistence are more important.

The rest of this paper is structured as follows: in section 2 we describe our data set and provide some descriptive statistics on the frequency and size of price changes across the countries and sectors considered in our sample. Section 3 outlines the approach which we use to estimate demand curves, whereas section 4 presents the empirical results. Section 5 sketches the theoretical model and discusses the obtained quantitative results. Section 6 summarizes and concludes.

2 Data and descriptive statistics

We employ a unique and very rich database on European scanner-price data which has not been used in the macro literature before. The data has been made available by AiMark (Advanced International Marketing Knowledge), which is a non-profit cooperation that promotes research in the area of retail markets and to this end provides data originally

compiled by Europanel and its partners Gesellschaft für Konsumforschung (GfK), Kantar Worldpanel and IRI. As will be outlined in more detail below, our data is best comparable to the IRI data used in Coibion et al. (2013) with the major difference that their data set covers the U.S. whereas our data comprises information from three European countries.

2.1 Description of the data set

In each country of our sample (Belgium, Germany, and the Netherlands), the data providers maintain a representative panel of households. Each household in the panel is endowed with a scanning technology which it uses to scan all the products belonging to the categories of *fast-moving costumer goods* it buys at retail outlets including, e.g., all major supermarket chains (such as Rewe or Aldi in Germany or Albert Heijn and C1000 in the Netherlands), drugstores, small corner shops and internet stores. This scanning technology is similar to the one underlying the Nielsen HomeScan database.

For each product bought, the household scans the bar-code, which uniquely identifies the product via the Global Trade Item Number (GTIN)⁷ and enters the number and associated price for this product into the homescanner⁸. The data set contains a description of each product and a classification system of the goods into more aggregate product categories.⁹ The household also provides the name of the retailer where it bought a product. The products belong to the categories of *fast-moving costumer goods*, such as grocery products, home and personal care products, and beverages. In addition to the detailed data on the individual transactions, we also have access to information on household characteristics, which comprises the location of the household, its income group and the age structure.

Table 1 reports some sample information. As mentioned above, we have available scanner price data for the countries Belgium, Germany, and the Netherlands. The sample period

⁷The GTIN-12 code corresponds to the Universal Product Code (UPC) which is used in the U.S. and Canada. In Europe, the GTIN was formerly known as European Article Number (EAN).

⁸In case the product does not have a bar-code, the household enters this information manually.

⁹We constructed a common classification scheme for the products in our data-set for all countries which rests on the one employed by the national data providers. However, the grouping systems of the individual providers can slightly differ across countries. We thus constructed comparable categories of goods by using the classification scheme of Germany as a basis and assigning the categories of the other countries to their German counterpart. The classification was done using both the assistance by country representatives of GfK and the extensive documentation of the different classification schemes, to which we had access at the data providers offices.

is from 2005 to 2008. The number of households is around 2,000 for Belgium, 12,000 for Germany, and 4,000 for the Netherlands, and we observe more than 2 million individual shopping trips in the data set. The number of available observations is quite large: it ranges from around 2 millions for Belgium to more than 9 millions for Germany. These observations include the purchases of more than 190,000 unique products. The products come from 58 different categories (such as beer or butter, see Table 6 in the Appendix for a list of all categories included). To evaluate the representativeness of the data, we calculate a simple measure of inflation as a weighted average of monthly price levels at the category level as in Coibion et al. (2013). We calculate the average month-on-month inflation rate in our sample and compare it to the inflation rate in the national CPIs (Consumer Price Indices). Even though we do not match inflation exactly, the numbers are not very different, and we observe that they comove with official inflation rates (see Figure C in the Appendix).¹⁰

Table 1: *Database - Overview*

Country	Househ.	Products	Categ.	Purch.	Observ.	Infl.	CPI
Belgium	1,746	49,808	30	0.7	2.1	2.1	2.3
Germany	11,631	72,617	19	4.1	9.4	3.1	1.6
Netherlands	4,030	68,909	14	2.3	10.4	0.7	1.5
Total	17,407	191,334	58	7.0	21.8	2.0	1.8


Note: Purchases and observations in 1,000,000. Products is the number of different GTINs, which are included in our estimation sample, categories are the number of categories included in the sample. A purchase occasion is one shopping trip by one consumer to one retailer (identified by the shopping data and the retailer). Inflation refers to a simple measure of inflation calculated in our estimation sample, and CPI is the average month-on-month growth rate of the Harmonized Index of Consumer Prices (HICP) for 2005-2008 (source: Eurostat).

The data structure allows us to track prices and quantities-bought trajectories for all households and products in the sample. To give an example of the data structure, we extract the purchases of one type of product for one household. Here, we choose one household from Germany, which frequently buys strawberry yogurt of the brand Ehrmann at the same retailer. The consumer is 40-45 years old and has an income of 2,250-2,499 EUR. We show

¹⁰As mentioned below, we had to make some sampling decisions, which are documented in Appendix A. All statistics reported here are based on the sample, which is included in estimation.

the quantity bought and the associated price at a given shopping occasion in Figure 2. The price trajectory shown in Panel a) looks very much like the price trajectories we observe in CPI micro data sets: the reference price changes about once per year, while temporary price cuts, and even increases, are not uncommon. We observe that the household usually buys one to three pots of yogurt, while at sales prices, the household sometimes buys four.

Figure 2: *Example of the data structure*



Note: Example: 500 grams *Ehrmann Almighurt* yogurt (strawberry flavour) bought on 168 occasions by one consumer aged 40-45 years, income 2,250-2,499 EUR, Germany.

For the econometric analysis, we first split the data in categories (yogurt, ketchup, beer, etc.). We then group the data by brand and rank the brands by expenditure share. We use the four top brands and construct a fifth good, which is a composite of all other brands in the category. This fifth good is the outside good in the estimation. Because we observe a product's price only when it is bought, we do not have direct observations for the prices of alternative brands, which would be available in the same retailer at a given shopping trip of a household. We therefore construct the alternative prices as described in section A, largely by matching observations of other purchases by other households of the alternative brands from the same retailer.

2.2 Descriptive statistics

The frequencies of price changes are comparable to those obtained for CPI and PPI micro data for the euro area. We find a median frequency of price changes (fpc) of 15.24 on a

monthly basis. Álvarez et al. (2006) report a frequency of 28% for unprocessed food, of 14% for processed food in the CPI, and 12% for non-durable, non-food items in the PPI (Table 2. Distributions of the sizes of price changes and the frequencies of price changes are shown in Figure 6 in the Appendix). Most of the products in our data set fall in one of these non-durable goods categories, which suggests that the mean frequency of 19% and the median frequency of 15% seem reasonable. Excluding sales, we find a median frequency of 13%, suggesting that sales in our data set do not change the frequency of price changes substantially. This might of course be due to the fact that we can identify sales only via a v-shaped sales filter. We do not have a sales flag in the data set. The share of price increases in all price changes ($fraction^{up}$) reported in the evidence collected in Dhyne et al. (2005) is 54%, very close to our median.

Table 2: *Sample statistics on the monthly frequency and size of price changes*

	Mean	Median
Including sales		
fpc	19.06	15.24
$fraction^{up}$	50.93	54.38
$size$	6.38	3.73
$size^{abs}$	18.91	17.47
$size^{up}$	20.66	17.02
$size^{down}$	-15.48	-18.09
Excluding sales		
fpc_{exs}	16.54	13.10
$fraction_{exs}^{up}$	53.38	54.55
$size_{exs}$	6.12	4.27
$size_{exs}^{abs}$	20.98	17.25
$size_{exs}^{up}$	25.14	19.27
$size_{exs}^{down}$	-15.48	-14.56

Note: The figures for the frequency are computed as the percentage of prices which are changed within a given month at a given retailer. The figures for the size correspond to the percentage change of a price conditional on a price change. Sales are identified using a simple V-shaped filter. The superscripts *abs*, *up*, *down* indicate absolute values, price increases and price decreases, respectively.

The average absolute size of price changes conditional on a price change, $size^{abs}$, is 18% in our data, which is larger than the statistics reported in Dhyne et al. (2005). They report an average size of price changes of 15 to 16% for unprocessed food, and of 7 to 8% for processed food for the euro area. One explanation for this result might be that many small price changes are erratic, as shown in Eichenbaum et al. (2014): they find that many small price changes in CPI micro data are due to measurement issues and that the median price change of 10%

found in their data, corrected for measurement bias, is roughly 30%, which is more in line with our statistics. In particular, we can clearly identify a product replacement, because any small change in a good requires the use of a new GTIN number.

3 Empirical Method

In this section, we briefly describe the econometric model used. We estimate demand elasticities using a nested multinomial logit model with random coefficients. We use the control function approach developed by Petrin and Train (2010) and applied to estimate demand elasticities by, for example, Gordon et al. (2012), to control for price endogeneity.

3.1 Discrete choice specification

The model consists of two nests. In the upper nest, consumers choose whether or not to buy in a given category c . Conditional on choosing that category, they decide on the brand within the category. For example, a household chooses to buy in the category pasta, and within this category, it chooses to buy a certain brand.

The probability of household i choosing a brand j within category c is the product of the conditional probability of choosing a brand j within the category c $P_{ij|c_i}$, given the category c is chosen, and the marginal probability of choosing category c , P_{ic} ,

$$P_{ic} = P_{ij|c_i} P_{ic}.$$

It is convenient to decompose the probabilities, because the marginal and the conditional probabilities take the form of logits (Train, 2009),

$$P_{ic} = \frac{e^{W_{ic} + \lambda_c IncVal_{ic}}}{\sum_C e^{W_{ic} + \lambda_c IncVal_{ic}}}$$

$$P_{ij|c_i} = \frac{e^{Y_{ij}/\lambda_c}}{\sum_J e^{Y_{ij}/\lambda_c}},$$

where $IncVal_{ic} = \ln \sum_J e^{Y_{ij}/\lambda_c}$ is the inclusive value, which is the expected utility household

i receives from the $J = 5$ choice alternatives within category c . λ_c is the log-sum coefficient, and Y_{ij} are the explanatory variables, that vary over alternatives within the category c . W_{ic} is a vector of variables that describe category c , and does not vary over j .

The conditional utility of household i purchasing brand j in category c at shopping occasion t in the *lower* nest is

$$U_{ijt|c_{it}=1} = \beta_{ij} - \alpha_i p_{ijt} + \delta x_{ijt} + \varepsilon_{ijt}. \quad (1)$$

where β_j is a vector of brand-specific effects, p_{jt} is the price of brand j , and x_{ijt} includes other controls. We allow α and β to vary over households, and model heterogeneity by assuming normal mixing distributions for β_j and a lognormal distribution for α . This captures heterogeneity in price elasticities and constant preferences for certain brands. The error term ε_{ijt} is distributed iid extreme value and the option of buying the outside good is normalized to ε_{i5t} .¹¹

Utility from purchasing in category c in the upper nest is given by

$$u_{ict} = \rho w_{ict} + \Psi IncV_{ict} + \nu_{ict}. \quad (2)$$

where w_{ict} counts the number of weeks since the household last purchased in the category, as in Gordon et al. (2012), and $IncValue_{ict}$ is the inclusive value from (1), $\ln \sum_J [\hat{\beta}_{ij} - \hat{\alpha}_i p_{jt} + \hat{\gamma}_i loyal_{ijt} + \hat{\delta} x_{ijt}]$. The parameter ρ describes a need to buy some products regularly, while Ψ captures the effect of utility household i obtains from the choices in category c . The error term ν_{ict} is distributed iid logistic. We use a maximum likelihood estimator, which constrains Ψ to lie between zero and one.

We use the control function approach to control for endogeneity in prices (Petrin and Train, 2010).

$$p_{jt} = \theta'_j Z_{jt} + \mu_{jt}, \quad (3)$$

where Z_{jt} is the mean price of the same brand in other NUTS1 regions within the same country (Hausman, 1996; Nevo, 2001).

¹¹We index the shopping occasion by t , which does not refer to a well-defined constant frequency, but to the date at which household i goes shopping and buys in category c .

The control function is a variable that captures the conditional mean of the correlation between the observed variables and the error term. It allows us to estimate price elasticities of demand and super-elasticities of demand in a very flexible manner while controlling for price endogeneity.¹² Suppose the error term is decomposed into $\varepsilon_{ijt} = E(\varepsilon_{ijt}|\mu_{jt}) + \tilde{\varepsilon}_{ijt} \equiv \lambda\mu + \tilde{\varepsilon}_{ijt}$, where $\tilde{\varepsilon}_{ijt}$ is by construction not correlated with μ_{jt} . Suppose $\varepsilon_{ijt} = \varepsilon_{ijt}^1 + \varepsilon_{ijt}^2$, where ε_{ijt}^1 is correlated with price. ε_{ijt}^1 and μ_{jt} are jointly normal and ε_{ijt}^2 , which is uncorrelated with price, is iid extreme value. Then $\varepsilon_{ijt}^1 = E(\varepsilon_{ijt}^1|\mu_{jt}) + \tilde{\varepsilon}_{ijt}^1$. The conditional distribution of ε_{ijt}^2 is same as the unconditional distribution, because it is independent.

Utility then becomes

$$U_{ijt|y_{it}=1} = \beta_{ij} - \alpha_{ij}p_{jt} + \gamma_i \text{loyal}_{ijt} + \delta x_{ijt} + \lambda\mu_{jt} + \tilde{\varepsilon}_{ijt}^1 + \varepsilon_{ijt}^2, \quad (4)$$

where $\tilde{\varepsilon}_{ijt}^1 = \sigma\eta_{ijt}$ is normal with zero mean and constant variance and η_{ijt} is standard normal.

To estimate the model, we apply a three-step approach. First, we estimate control function and retain residual $\hat{\mu}_{jt}$. Second, we estimate the mixed logit for the lower nest with $\hat{\mu}_{jt}$ and income and age (interacted with brand-specific dummies), brand-region fixed effects, and a proxy for loyalty as explanatory variables summarized in x_{ijt} in equation (1).¹³ Third, we estimate the logit for the *upper* nest using the inclusive value calculated from the results obtained in the second step. We then calculate the elasticities and super-elasticities for each household within a category.

¹²An alternative way to estimate the model would be to use the BLP (Berry et al., 1995) approach as in Nakamura and Zeron (2010), for example. However, the BLP approach is difficult to use in our context when we observe just a small number of purchases per product, because market shares should be observed with some precision, as these are necessary to estimate product-specific constants, which should remove the endogeneity from the error term. This makes the control function approach more reliable. See Petrin and Train (2010) for a discussion.

¹³Region is defined by the first level NUTS (Nomenclature of Territorial Units for Statistics) regions, which reference the administrative divisions of European countries for statistical purposes. For Belgium, these are the three regions Brussels Capital Region, Flemish Region and Walloon Region. For Germany, these are the 16 German Länder. For the Netherlands, there are four regions: north-, east-, west- and south Netherlands. Loyalty is proxied by a count variable, which gives the number of purchases of the same brand in the same category in the past.

3.2 Demand elasticities and super-elasticities

Demand elasticities can be derived from the nested mixed logit model. We derive the elasticity of the upper and the lower nest in the Appendix. The elasticity for the upper nest is given by

$$\theta_u = \alpha_i P_{ij|c_i} (1 - P_{ic}) p_{ijt}. \quad (5)$$

The elasticity for the lower nest is

$$\theta_l = \frac{\alpha_i}{\Psi} p_{ijt} (1 - P_{ij|c_i=1}). \quad (6)$$

Because the probability can be written as the product of the marginal and the conditional probability, the total elasticity is just the sum of the upper and the lower elasticities $\theta = \theta_u + \theta_l$.

One advantage of the mixed logit model is that the price elasticity depends on the price and therefore allows us to derive super-elasticities directly from the estimates, without having to assume a quadratic functional form (Hellerstein, 2008). This is because the model allows for heterogeneity in consumers price sensitivity, which contribute to the curvature of demand. It nests the CES with a super-elasticity of zero as a special case (Goldberg and Hellerstein, 2013).

The super-elasticity of the upper nest is

$$\epsilon_u = 1 + \theta_l - \Psi \alpha_i P_{ij|c_i} P_{ic} p_{ij} \quad (7)$$

and for the lower nest

$$\epsilon_l = 1 - \frac{\alpha_i}{\Psi} p_{ij} P_{ij|c_i=1} \quad (8)$$


The total super-elasticity is $\epsilon = \epsilon_u + \epsilon_l$.

4 Empirical results

To summarize estimates, we first calculate all elasticities for all categories and countries for each household. Before we calculate statistics of the mean, we trim these estimates, cutting off the upper and lower 10%. This is because we estimate a very large number of category- and household-specific coefficients, and there are a few outliers, which would otherwise contaminate estimates of the mean.

The distributions of all estimates of the elasticity and super-elasticity for each category and household are shown in the upper panel in Figure 3. This is the distribution of all estimated coefficients, for all households and goods categories. The distribution of estimated elasticities is highly skewed, suggesting that many households have low price elasticities. Furthermore, the distribution is wide, suggesting that there is a substantial heterogeneity in price elasticities across categories and households. This does not carry over to the estimates of the super-elasticity, which we do not restrict to be strictly positive or negative. Here, we find a tighter distribution, showing that most estimates are in the range between zero and five (Panel b)). There are only a few estimates with a value of ten or larger, suggesting that the assumptions often used in many macroeconomic models reviewed in the introduction seem rather unlikely.

Figure 3: *Distribution of elasticities and super-elasticities*



Notes: The distributions are shown for a trimmed sample, where we cut off the upper and lower 5% of the estimates, both for elasticities and super-elasticities. As noted in the text, there are a few very large estimates, and the graph would not be readable, if the range would be extended to those estimates.

How large should the parameters for price elasticities and super-elasticities be in

macroeconomic models? In the median, the price elasticity is rather small, at a value of 3.85 (Table 3). The mean and weighted mean are a bit lower. Many macroeconomic models assume a price elasticity of ten. This seems a relatively large value, given that more than 80% of our estimates are below ten.

Table 3: *Descriptive Statistics Estimates*

	Mean	Median	W. mean
<hr/>			
Demand elasticity			
Total elasticity	1.91	3.58	2.08
Upper nest	0.19	0.13	0.20
Lower nest	1.65	2.99	1.84
<hr/>			
Super-elasticity			
Total super-elasticity	1.65	1.39	1.47
Upper nest	-0.55	0.10	-0.71
Lower nest	2.06	1.31	2.08
<hr/>			
Implied markup elasticities	1.01	0.70	0.88
<hr/>			

Notes: The first two columns show the mean/median values of all estimated elasticities/super-elasticities. The third column shows weighted means (by expenditure share). Elasticities and markup elasticities are multiplied by -1 .

Super-elasticities are clearly below ten. The median estimate for the super-elasticities is 1.39. This is close to the estimates for the beer market found in Goldberg and Hellerstein (2013), and somewhat lower than the estimates for the coffee market in Nakamura and Zerom (2010) or for the European retailer in Dossche et al. (2010). 90% of our estimates are below a value of five. Taking these pieces of evidence together, the data does not support assumptions of a super-elasticity parameter of ten or even larger.

As documented by Gopinath and Itskhoki (2010), the combination of demand elasticity and super-elasticity matter for the elasticity of the markup. They show that the markup (μ) elasticity can be expressed as $\frac{\partial \mu}{\partial \ln P} |_{P=1} = \frac{\varepsilon}{\sigma-1}$, implying that the markup elasticity increases in the super-elasticity but decreases in the demand elasticity. Our estimates suggest that super-elasticities are small, but demand elasticities are smaller than what is often assumed in macroeconomic models, too. The implied markup elasticities are very close to one, suggesting that firms decrease their markups by one percent if a firm's relative price increases by one percent (Table 3, last row). This is close to the calibrations of the baseline good in Gopinath and Itskhoki (2010), who assume a super-elasticity of 4 and a demand elasticity of 5.

5 A menu-cost model with Kimball-type preferences

In this section, we describe the theoretical model, which we use for calibration in order to quantitatively assess the role of our estimated demand-side real rigidity for monetary non-neutrality. The baseline specification of our theoretical model closely follows Nakamura and Steinsson (2010).

5.1 Model setup

As in Nakamura and Steinsson (2010), our economy is inhabited by a representative household, a continuum of firms and a monetary authority that controls the evolution of nominal GDP. The household supplies labor to firms, decides how to allocate income between aggregate consumption and saving and determines the amount it wants to consume of each good available in the economy. Firms produce goods using labor (subject to idiosyncratic changes in labor productivity) and set prices to maximize profits. Changing prices is subject to a cost. The monetary authority determines the growth rate of nominal GDP by injecting money into the economy. Deviating from Nakamura and Steinsson (2010) we incorporate real rigidities into the model in the form of a Kimball-type preference structure (which embeds the standard CES case used by Nakamura and Steinsson (2010) as a special case) rather than a roundabout production setup.

5.1.1 Households

The representative household maximizes expected discounted life-time utility, which depends positively upon aggregate consumption, C_t , and negatively upon labor supply, L_t , and is given by

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{1}{1-\gamma} C_t^{1-\gamma} - \frac{\omega}{1+\psi} L_t^{1+\psi} \right], \quad (9)$$

where E_0 denotes the rational expectations operator conditional on information available to the households at date 0. β (with $0 < \beta < 1$) represents the subjective discount factor. The period utility function is assumed to be additive separable in consumption and labor supply. The parameter γ governs the degree of relative risk aversion while ψ determines the

convexity of the dis-utility of labor. ω is a weighting term determining the relative extent of the dis-utility of labor. The composite consumption good, C_t , is generated via a Kimball aggregator as specified in equation (13) below.

Households choose composite consumption and labor to maximize (9) subject to the following budget constraint

$$P_t C_t + E_t [D_{t,t+1} B_{t+1}] \leq B_t + W_t L_t + \int_0^1 \Pi_t(z) dz \quad t = 0, 1 \dots \quad (10)$$

This equation requires that aggregate consumption expenditure C_t and investment in financial assets $D_{t,t+1} B_{t+1}$ cannot be larger than available resources consisting of the stock of financial assets carried over from the previous period B_t , wage income $W_t L_t$ and profits distributed by firms $\int_0^1 \Pi_t(z) dz$. $D_{t,t+1}$ is the period- t price of a financial assets that pays off one unit in period $t + 1$.

The first-order conditions for the household's optimization problem are given by:

$$D_{t,T} = \beta^{T-t} \left(\frac{C_T}{C_t} \right)^{-\gamma} \frac{P_t}{P_T} \quad (11)$$

$$\frac{W_t}{P_t} = \omega L_t^\psi C_t^\gamma. \quad (12)$$

Equation (11) represents the standard inter-temporal Euler equation linking consumption growth to the real interest rate while equation (12) states that labor supply adjusts as a function of the real wage, given the marginal utility of consumption.

The composite consumption good, C_t , is created by the costless aggregation of a continuum of differentiated goods, $c_t(z)$, which are supplied by monopolistic firms. Following Kimball (1995) and Klenow and Willis (2006), we implicitly define the composite consumption good, C_t , using an aggregator of the form

$$\int_0^1 \Upsilon \left(\frac{c(z)}{C} \right) dz = 1, \quad (13)$$

where the function $\Upsilon(\cdot)$ satisfies the conditions $\Upsilon(1) = 1$, $\Upsilon'(\cdot) > 0$ and $\Upsilon''(\cdot) < 0$ and where time indices are dropped for notational ease. In our simulations below, we report the outcomes

of two specifications for the aggregation function $\Upsilon(\cdot)$. In the baseline case, we specify the function $\Upsilon(\cdot)$ to be given by

$$\Upsilon\left(\frac{c(z)}{C}\right) = \left(\frac{c(z)}{C}\right)^{\frac{\bar{\theta}-1}{\bar{\theta}}} \quad (14)$$

which amounts to assume standard CES preferences as proposed by Dixit and Stiglitz (1977). We denote this as the CES case below. In this specification, $\bar{\theta}$ denotes the (constant) elasticity of substitution.

Alternatively, we employ the Kimball aggregator function proposed by Klenow and Willis (2006) and used, for example, by Gopinath and Itskhoki (2010). We denote this as the NON-CES case below. In this case, the aggregation function $\Upsilon(\cdot)$ is given by

$$\Upsilon(x) = 1 + (\bar{\theta} - 1) \exp\left(\frac{1}{\bar{\epsilon}}\right)^{\bar{\epsilon}(\frac{\bar{\theta}}{\bar{\epsilon}} - 1)} \left[\Gamma\left(\frac{\bar{\theta}}{\bar{\epsilon}}, \frac{1}{\bar{\epsilon}}\right) - \Gamma\left(\frac{\bar{\theta}}{\bar{\epsilon}}, \frac{x^{\frac{\bar{\epsilon}}{\bar{\theta}}}}{\bar{\epsilon}}\right) \right], \quad (15)$$

with $x = \frac{c(z)}{C}$ and $\Gamma(u, z)$ denoting the incomplete gamma function.

Given an optimal decision about overall consumption expenditure, C , household choose the optimal amount of each good $c(z)$ by minimizing the overall cost of purchasing C . For the CES case, the optimal demand for good z is given by

$$c(z) = \left(\frac{p(z)}{P}\right)^{-\bar{\theta}} C, \quad (16)$$

showing that the demand for good z depends positively upon overall consumption C and negatively upon the price of good z relative to the overall price level P . $\bar{\theta}$ is a parameter of the aggregation function Υ and can be interpreted as the elasticity of substitution between good z and some other good z' .

When employing the Kimball aggregator function, the demand function for good z is given by

$$c(z) = \left[1 - \bar{\epsilon} \ln\left(\frac{p(z)}{P}\right)\right]^{\frac{\bar{\theta}}{\bar{\epsilon}}} C, \quad (17)$$

where again the demand for good z depends positively upon overall consumption demand C and negatively upon the relative price of good z . $\bar{\theta}$ and $\bar{\epsilon}$ are parameters of the aggregation

function Υ determining the steady-state size and behavior of the elasticity of demand. For $\bar{\epsilon} = 0$ this demand function reduces to the CES case.¹⁴

The major difference between the setups of Dixit and Stiglitz (1977) and Klenow and Willis (2006) consists of their differing implications for the behavior of the price elasticity of demand. To illustrate these differences we first define the price elasticity of demand, $\theta(p)$, as

$$\theta(p) = \frac{\partial \ln c(z)}{\partial \ln p} \quad (18)$$

and the super-elasticity of demand, $\epsilon(p)$, as

$$\epsilon(p) = \frac{\partial \ln \theta(p)}{\partial \ln p}, \quad (19)$$

where we allow both the elasticity and the super-elasticity of demand to be non-constant and depending on the relative price of a good.¹⁵

In the CES case, the steady-state value of the elasticity of demand is $-\bar{\theta}$ and the super-elasticity is 0, i.e., the price elasticity is constant for all values of the relative price of good z . For the NON-CES case, the price elasticity of demand is given by

$$\theta^{NON-CES} = -\frac{\bar{\theta}}{1 - \bar{\epsilon} \ln \left(\frac{p(z)}{P} \right)}, \quad (20)$$

whereas the super-elasticity of demand is given by

$$\epsilon^{NON-CES} = \frac{\bar{\epsilon}}{1 - \bar{\epsilon} \ln \left(\frac{p(z)}{P} \right)}. \quad (21)$$

In the steady state where $\frac{p(z)}{P} = 1$ holds we get:

$$(\theta^*)^{NON-CES} = \bar{\theta} \text{ and } (\epsilon^*)^{NON-CES} = \bar{\epsilon}. \quad (22)$$

¹⁴In deriving equation (17), we made use of an approximation result obtained by Gopinath and Itskhoki (2010, Appendix).

¹⁵Dossche et al. (2010) note that different authors use slightly different measures for the super-elasticity of demand. Our notation follows that of Dossche et al., which is also employed by Gopinath and Itskhoki (2010).

These equations show that for $\bar{\epsilon} > 0$ the elasticity of demand in the NON-CES case increases in the relative price of good z and will be larger than for the CES case. This difference in the behavior of the elasticity of demand has profound implications for the response of price setters to macroeconomic shocks, such as a change in nominal aggregate demand induced by the monetary authority, as documented in the simulation exercises below.

5.1.2 Firms

Monopolistically competitive firms denoted by z produce differentiated products via the production function

$$y_t(z) = A_t(z)L_t(z), \quad (23)$$

where $A_t(z)$ denotes firm-specific productivity and $L_t(z)$ is the amount of labor employed by firm z . We assume that firm-specific productivity follows a first order auto-regressive process of the form

$$\ln A_t(z) = \rho \ln A_{t-1}(z) + \epsilon_t(z), \quad \epsilon_t(z) \sim N(0, \sigma_{\epsilon,t}^2). \quad (24)$$

Firms aim to maximize the discounted value of expected profits,

$$E_0 \sum_{t=0}^{\infty} D_{0,t} \Pi_t(z), \quad (25)$$

where profits $\Pi_t(z)$ are given by

$$\Pi_t(z) = p_t(z)y_t(z) - W_t L_t(z) - KW_t I_t(z) - P_t U. \quad (26)$$

$y_t(z)$ denotes firm's z sales which are equal to $c_t(z)$ as discussed above. $p_t(z)y_t(z)$ represent firm's z revenue, $W_t L_t(z)$ are labor cost, $KW_t I_t(z)$ are costs ("menu costs") of changing prices ($I_t(z)$ is an indicator function taking the value 1 if the firm changes its price in period t and 0 otherwise). $P_t U$ denote fixed nominal costs of production. Based on our empirical results for the estimated elasticities we follow Nakamura and Steinsson (2010) and introduce this term to justify the co-existence of empirically estimated large markups with observed relatively small firm profits in the national accounts.

5.1.3 Monetary policy

Following Nakamura and Steinsson (2010) and Midrigan (2011) we assume that aggregate nominal GDP, $Y_t^N = P_t C_t$, grows at a constant long-run rate μ and is subject to temporary shocks. More specifically, we assume that the monetary authority controls the path of nominal GDP according to the process

$$\ln Y_t^N = \mu + \ln Y_{t-1}^N + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2). \quad (27)$$

5.2 Model solution

An equilibrium in this model is a set of policy rules for the endogenous variables that is consistent with the household's and firms' maximization, market clearing and the evolution of the exogenous processes for total factor productivity and nominal GDP. To solve for the equilibrium we first rewrite the firms' profit function employing both the labor demand and supply function and replacing firm's output by the corresponding demand function in real terms as follows:

$$\begin{aligned} \Pi_t^R(z) = & \left(\frac{p_t(z)}{P_t} \right) \mathcal{F} \left(\frac{p_t(z)}{P_t} \right) \frac{Y_t^N}{P_t} - \omega L_t^\psi C_t^\gamma \left(\frac{1}{A_t(z)} \right) \mathcal{F} \left(\frac{p_t(z)}{P_t} \right) \frac{Y_t^N}{P_t} - \\ & - K \omega L_t^\psi \left(\frac{Y_t^N}{P_t} \right)^\gamma I_t(z) - U, \end{aligned} \quad (28)$$

where $\mathcal{F}(\cdot)$ corresponds to equation (16) in the CES-Case and to equation (17) in the NON-CES-Case.

We solve the firm's optimization problem using dynamic programming.¹⁶ The state variables for the firm's optimization problem are given by the level of idiosyncratic productivity $A_t(z)$, aggregate real GDP as represented by the ratio of nominal GDP and the price level, $\frac{Y_t^N}{P_t}$, and the firm's relative price at the end of the previous period $\frac{p_{t-1}(z)}{P_t}$.

¹⁶To solve the model, we made intensive use of the Matlab programs developed in Nakamura and Steinsson (2010). We would like to thank these authors for making their code available.

Given these state variables, each firm maximizes the value function

$$V \left(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{Y_t^N}{P_t} \right) = \max \left\{ V^{NC} \left(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{Y_t^N}{P_t} \right), V^C \left(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{Y_t^N}{P_t} \right) \right\}, \quad (29)$$

where $V^{NC}(\cdot)$ denotes the value function when the firm does not change its price and $V^C(\cdot)$ denotes the value function when the firm changes its price. The expressions for these two functions are given by:

$$V^{NC} \left(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{Y_t^N}{P_t} \right) = \Pi \left(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{Y_t^N}{P_t} \right) + E_t \left[D_{t,t+1} V^{NC} \left(A_{t+1}(z), \frac{p_{t-1}(z)}{P_{t+1}}, \frac{Y_{t+1}^N}{P_{t+1}} \right) \right]$$

$$V^C \left(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{Y_t^N}{P_t} \right) = \max_{p_t} \left\{ \Pi \left(A_t(z), \frac{p_t(z)}{P_t}, \frac{Y_t^N}{P_t} \right) + E_t \left[D_{t,t+1} V^{NC} \left(A_{t+1}(z), \frac{p_t(z)}{P_{t+1}}, \frac{Y_{t+1}^N}{P_{t+1}} \right) \right] \right\}$$

To solve this optimization problem, the firm needs to form expectations about the future path of the state variables. This can be done in a straightforward manner for A_t and Y_t^N which both follow exogenous stochastic processes. In the case of the price index, P_t , however, one faces the following fixed point problem: the optimal decision of a firm depends on the path of the price level and this optimal decision in turn impacts the determination of the equilibrium path of the price level. To address this issue, we follow Nakamura and Steinsson (2010) and use the method by Krusell and Smith (1998) to approximate the distribution of relative prices by the first moments of the expected price distribution and postulate that firms use the formula

$$\frac{P_t}{P_{t-1}} = \Gamma \left(\frac{Y_t^N}{P_{t-1}} \right) \quad (30)$$

to form expectations of the change in the aggregate price level (i.e. the inflation rate).

Given this forecasting rule, our procedure to solve for the equilibrium proceeds as follows:

(1) We start by specifying a discrete grid vector for each of the three state variables, initialize

the stationary distribution and make a first guess of the forecasting rule $\Gamma\left(\frac{Y_t^N}{P_{t-1}}\right)$. (2) Given the forecasting rule, we then solve for the firms' policy function using value function iterations. (3) As a next step we update the stationary distribution using the policy function. (4) Finally, we update the forecasting rule and check whether it is consistent with the aggregate inflation rate implied by the firms' policy function. If this is the case we stop, otherwise we return to (2).

5.3 Calibration

To examine the quantitative implications of our empirical findings for the degree of monetary non-neutrality, we simulate the model for a variety of settings. The parameters to be calibrated are grouped into two classes. The first set of coefficients, reported in Table 4, include those which are common across all considered specifications. We choose standard values for these coefficients, identical to Nakamura and Steinsson (2010). We deviate from their specification only for the mean inflation rate and the standard deviation of money supply shocks. To calibrate these two values, we use CPI data over the period from 1999 to 2008 for Belgium, Germany and the Netherlands from Eurostat. Mean inflation, μ , and the standard deviation of money supply shocks, σ_η , are computed as the mean and standard deviation of the aggregate CPI inflation rate for these three countries, weighted by real GDP.

Table 4: *Parameters common across considered model specifications*

Parameter description	Value
Subjective discount factor, β	0.96 ^{1/12}
Relative risk aversion, γ	1
Elasticity of labor supply, ψ	0
Steady-state labor supply, L^*	1/3
Mean (monthly) inflation rate, μ	0.00154
Std. dev. of mon. pol. shocks, σ_η	0.00176

The second set of coefficients are not common among specifications (Table 5). Here, we modify the model along three dimensions.

First, we distinguish two alternative preference specifications: in the upper panel (labelled

CES), we report figures for an economy without demand-side real rigidity, i.e., we employ a standard Dixit-Stiglitz CES specification for the aggregation function $\Upsilon(\cdot)$. The value for the super-elasticity parameter, $\bar{\epsilon}$, is set equal to zero. In the lower panel of Table 5 (labelled NON-CES), we consider the specification including demand-side real rigidity. In this case, the parameter for the super-elasticity, $\bar{\epsilon}$, is set equal to the corresponding value found in the data.

Second, we compare the quantitative implications which result from employing either the estimated median (columns 1 and 2) or weighted mean (columns 3 and 4) elasticity and super-elasticity values. While the estimated super-elasticities of both specifications are comparable in size, there is a non-negligible difference in the elasticity values, implying that households will behave more responsive to deviations from the equilibrium price in the former case. This effect can be nicely seen in the corresponding demand curves (Figure 7 of the Appendix).

Third, we consider two different scenarios for the persistence of the idiosyncratic shocks. Gopinath and Itskhoki (2010) show that a higher shock persistence implies that firms will tend to adjust prices more frequently in response to shocks of similar sizes. As a result, smaller values for the menu costs and shock variances are needed in order to match observed price characteristics. In our simulations, we thus compare model outcomes for a persistence parameter of 0.7 (which corresponds to the setups of Klenow and Willis, 2006, and Nakamura and Steinsson, 2010) and 0.95 (which corresponds to the value taken by Gopinath and Itskhoki, 2010).

The values for the menu-cost parameters (*Menu costs*) and the standard deviations of the idiosyncratic shocks (*Std.dev. of id. shocks*) are chosen to match the empirical observations on the mean frequency and size of price changes as reported in the upper panel of Table 2. Menu costs are reported as a fraction of steady-state revenue.¹⁷

Looking at the absolute values of the menu costs implied by our model, they are broadly in line with those provided by independent evidence on menu costs as reported for example by Levy et al. (1997) who find that costs of changing prices amount to around 0.7% of revenue.

¹⁷The corresponding values derived for the price statistics excluding sales are reported in Table 7 of Appendix B.

For the CES case and the median elasticity estimate, menu costs are of a relatively moderate size, amounting to a maximum of around 4% as a ratio of total revenue, or 0.6% when the frequency of price changes is taken into account (column 1). As one would expect menu costs are a bit smaller when the shock persistence is larger given that price setters in this case will tend to change their prices faster everything else equal (column 2). However, the differences in menu costs are very small. Including demand-side real rigidities raises the size of the menu costs by a factor of close to 1.5.¹⁸ When we calibrate to the weighted mean elasticity, menu costs are around 1.8% of total revenue, even closer to the independent evidence.

The values for the standard deviation of idiosyncratic shocks are broadly in line with independent evidence, too, at least for the case with higher idiosyncratic shock persistence. Bachmann and Bayer (2014) report values for the standard deviation of idiosyncratic productivity of around 9% annually based on a firm-level data set for Germany. Our calibrations of the model with demand-side real rigidity yields a standard deviation of idiosyncratic shocks of 11.4%, which is close to the Bachmann-Bayer estimate. This corroborates the results in Gopinath and Itskhoki (2010), who show that a calibration with higher shock persistence and a moderate super-elasticity would not require unreasonably large idiosyncratic shocks.

5.4 Simulation results

Employing the calibrated parameter values, we next investigate the degree of monetary non-neutrality associated with and without demand-side real rigidity. Following Nakamura and Steinsson (2010), we measure monetary non-neutrality using two different indicators: first, we compute the area under the impulse response function of aggregate real consumption C following a shock to nominal GDP. This statistic, denoted by CIR (Cumulative Impulse Response) in Table 5, captures the overall effect of a nominal shock on real consumption spending: the larger and longer-lasting the response of C to a nominal shock, the higher the value for CIR will be. Second, we compute the variance of aggregate consumption spending (*Var. of C*) that results from model simulations where only aggregate nominal shocks hit

¹⁸A similar finding is reported by Klenow and Willis (2006), who document an increase in menu costs by a factor of two when comparing the CES with the NON-CES case.

Table 5: *Sample-specific preference and production-process parameters and simulation results (including sales)*

	Median elasticity		Weighted mean elasticity	
	Low persistence	High persistence	Low persistence	High persistence
CES				
Elasticity	3.58	3.58	2.08	2.08
Superelasticity	0	0	0	0
Menu costs	4.35	4.14	1.85	1.82
Std.dev. of id. shocks	9.11	7.46	9.11	7.39
<i>fpc</i>	15.25	15.24	15.19	15.21
<i>fraction^{up}</i>	55.34	54.52	54.06	53.68
<i>size^{abs}</i>	17.46	17.48	17.45	17.5
<i>size^{up}</i>	16.69	16.95	17.08	17.24
<i>size^{down}</i>	18.42	18.1	17.89	17.79
Model outcomes:				
CIR	0.007	0.006	0.007	0.007
Var. of C: model	0.007	0.006	0.008	0.007
Var. of C: data	0.245	0.245	0.245	0.245
Var. of C explained	0.029	0.025	0.033	0.029
NON-CES				
Elasticity	3.58	3.58	2.08	2.08
Superelasticity	1.39	1.39	1.47	1.47
Menu costs	6.77	6.26	4.31	4.2
Std.dev. of id. shocks	14.09	11.42	21.57	17.61
<i>fpc</i>	15.2	15.22	15.2	15.25
<i>fraction^{up}</i>	53.82	53.66	52.07	52.43
<i>size^{abs}</i>	17.45	17.45	17.48	17.47
<i>size^{up}</i>	17.15	17.2	17.76	17.63
<i>size^{down}</i>	17.79	17.73	17.18	17.3
Model outcomes:				
CIR	0.013	0.008	0.016	0.016
Var. of C: model	0.013	0.013	0.017	0.015
Var. of C: data	0.245	0.245	0.245	0.245
Var. of C explained	0.053	0.053	0.069	0.061

Notes: 1) In the low-persistence case, the persistence of idiosyncratic shock, ρ , is assumed to be equal to 0.7, in the high-persistence case we have $\rho = 0.95$. 2) Menu costs are given by the percentage share of steady-state revenue $\frac{\theta-1}{\theta} \frac{K}{Y^*}$. 3) The values for the elasticity and super-elasticity are taken from the estimation results as reported in Table 3. 4) Menu cost and standard deviation parameters are chosen so as to match the median frequency and size of price changes reported in rows *fpc* and *size^{abs}* of the upper panel (descriptive statistics including sales) of Table 2. 5) CIR (cumulative impulse response) and Var. of C (variance of aggregate real consumption) represent measures of monetary non-neutrality as described in the main text. The numbers for the variance of C are multiplied by 10^4 . (6) The variance of real consumption is the HP filtered quarterly real consumption (summed over the three countries included in the data sample) from 1995 to 2008. The source of the data is Eurostat.

the economy. Again, higher values of this statistic indicate a larger degree of monetary non-neutrality. To obtain an intuition of how much of the cyclical movements present in the data our model can explain, we compare the variance in real consumption induced by nominal shocks in our model to the variance of detrended real consumption observed in our sample countries (*Var. of C explained*).¹⁹

The results for the model without demand-side real rigidity (Table 5, upper panel) clearly illustrate that nominal rigidity alone cannot generate sizeable degrees of monetary non-neutrality, confirming several results including Golosov and Lucas (2007) or Nakamura and Steinsson (2010), for example. Both indicators of monetary non-neutrality show that the degree of real volatility implied by nominal shocks is very small, independently of which setting is considered. The proportion of the variance in real consumption explained range from 2.5% to 3.3%. As one would expect, monetary non-neutrality is somewhat higher for the weighted mean statistics, because they imply a higher markup elasticity of 1.4 as compared to only 0.6 for the median statistics. Moreover, a higher persistence in idiosyncratic shocks tends to lower the real effects of monetary policy because the relative importance of aggregate shocks for output is lower, all other things equal.

The responses for the model with demand-side real rigidity are larger on impact and more pronounced. This is shown in the impulse responses of real consumption to a nominal demand shock in Figure 4. This is true for both cases of idiosyncratic shock persistence where the effect is somewhat larger in the low-persistence setup. The explicit consideration of the empirically documented real rigidity thus not only increases the effect of a monetary policy shock on impact, but also adds persistence to the model's response to this shock. However, the duration of the observed real effects last only about 5 to 6 months and is thus much smaller than what one normally finds in the data.

The model in this paper cannot match the aggregate business cycle. Only around 5% to 6% of empirically observed volatility in real consumption can be explained by monetary-policy shocks in the model (Table 5, last row). The aim of the quantitative analysis is not, however, to obtain a model, which matches both micro and macro facts on prices. Clearly, this model

¹⁹Unlike Nakamura and Steinsson (2010), we chose consumption rather than GDP as the reference variable, since the prices underlying our data sample are almost exclusively related to private consumer goods.

Figure 4: *Response to monetary policy shock*

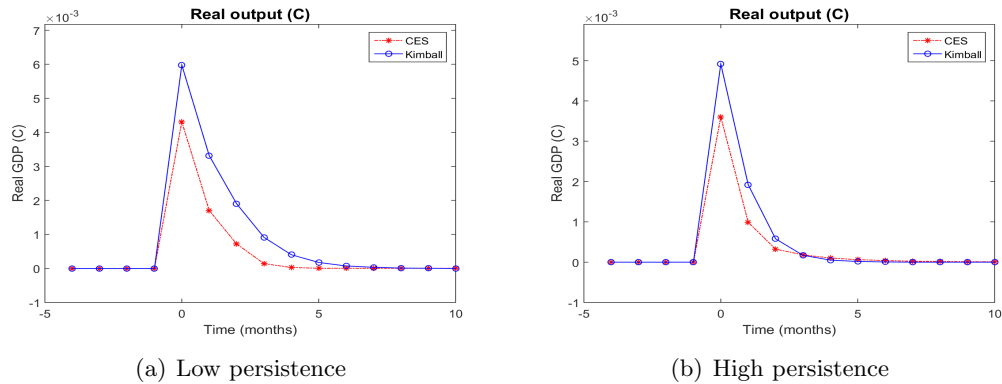


Figure 4 plots the impulse response of real consumption expenditures to a (positive) shock to nominal aggregate demand both for the CES and NON-CES (labeled "Kimball" in the graph) demand functions. The left panel considers the case where the persistence of idiosyncratic shocks is equal to 0.7, whereas the right panel considers the case of 0.95. Elasticity values are equal to their medians.

matches only the micro facts. The aim of the quantitative analysis is to show the marginal effect of adding demand-side real rigidity to a model with only nominal rigidity, and no additional features. This analysis suggests, that this effect is to roughly double monetary non-neutrality, raising the share of consumption explained from around 2.5 – 3% to 5 – 6%. Interestingly, the results are not only qualitatively but also quantitatively similar across all four considered specifications and amount to an increase of the real effects of monetary policy by a factor of around two. In other words, even the relatively moderate values empirically found for the super-elasticity imply a doubling of monetary non-neutrality.²⁰

6 Conclusions

The findings by Bils and Klenow (2004) that micro prices are changed relatively often and that both large price increases and decreases frequently occur has challenged the by then dominating view amongst most monetary economists that considerable nominal frictions exist in the economy. Since then, several attempts have been made to reconcile the micro evidence on relatively flexible prices and observed relatively large responses of real variables to nominal shocks. One of these attempts has been to introduce real rigidities resulting from non-constant

²⁰ Analogously to our results, Bergin and Feenstra (2000) only find a relatively moderate difference in the persistency effects of monetary policy shocks between the CES and their Non-CES specification when no further real rigidity is considered.

elasticity demand curves into macro models. While very flexible in its implementation, this approach has so far lacked broad-based microeconomic evidence. The aim of this paper is to provide this evidence and evaluate its quantitative implications.

To this end, we have employed a new data set on consumer retail transactions which contains detailed information on prices and quantities for three European countries and estimate a discrete choice model of demand to obtain estimates on the size and distribution of the elasticity and super-elasticity of demand. Our findings suggest values for the price elasticity parameters which are in the range between 2 and 4. While these numbers are well below the values most often used in the macro literature they tend to be in line with the ones found in the IO and marketing literature. Similarly, our findings for the super-elasticity parameter are much smaller than the values used in macro models with values being in the range between 1 and 2. Together with the demand elasticity estimates, these imply a markup elasticity with regards to the relative price of around 1.

To quantitatively assess the importance of demand-side real rigidity, we augment a model with only nominal rigidity and augment it with an empirically plausible demand-side real rigidity. Calibrating the model with and without the demand-side real rigidity and comparing the monetary non-neutrality generated by both versions of the model, allows us to obtain an estimate of the multiplier effect of demand-side real rigidity. Our results suggest that this multiplier effect is about two, because the model including demand-side real rigidity shows a two-times larger degree of monetary non-neutrality, compared to the pure nominal rigidity model. These calibrations imply plausible values for menu costs and idiosyncratic shock variances and can still match the observed frequencies and sizes of price changes. This is in particular true when a high persistence in idiosyncratic shock processes is assumed. However, demand-side real rigidities are not very large. Our results suggest that only up to 6% of the actual variance in real consumption can be explained by the menu cost model, featuring only nominal price rigidity and demand-side real rigidities. This suggests that other forms of multipliers, such as sectoral heterogeneity or supply-side real rigidities, are needed to match both the micro and the macro facts on prices.

References

- ALTIG, D., L. J. CHRISTIANO, E. MARTIN, AND J. LINDE (2011): “Firm-specific Capital, Nominal Rigidities and the Business Cycle,” *Review of Economic Dynamics*, 14, 225–247.
- ÁLVAREZ, L., E. DHYNE, M. HOEBERICHTS, C. KWAPIL, H. L. BIHAN, P. LÜNNEMANN, F. MARTINS, R. SABBATINI, H. STAHL, P. VERMEULEN, AND J. VILMUNEN (2006): “Sticky Prices in the Euro Area: A Summary of New Micro Evidence,” Diskussion Papers 02/2006, Deutsche Bundesbank.
- BACHMANN, R. AND C. BAYER (2014): “Investment Dispersion and the Business Cycle,” *American Economic Review*, 104, 1392–1416.
- BALL, L. M. AND D. ROMER (1990): “Real Rigidities and Non-Neutrality of Money,” *Review of Economic Studies*, 57, 183–204.
- BASU, S. (1995): “Intermediate Goods and Business Cycles: Implications for Productivity and Welfare,” *The American Economic Review*, 85, 512–531.
- BERGIN, P. R. AND R. C. FEENSTRA (2000): “Staggered Price Setting, Translog Preferences and Endogenous Persistence,” *Journal of Monetary Economics*, 45, 657–680.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.
- BILS, M. AND P. J. KLENOW (2004): “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112, 947–985.
- BLANCHARD, O. J. AND J. GALI (2007): “Real Wage Rigidities and the New Keynesian Model,” *Journal of Money, Credit and Banking*, 39, 35–65.
- BURSTEIN, A. AND G. GOPINATH (2013): “International Prices and Exchange Rates,” Working Paper 18829, National Bureau of Economic Research.
- CHARI, V. V., P. J. KEHOE, AND E. R. MCGRATTAN (2000): “Sticky Price Models of the Business Cycle,” *Econometrica*, 68, 1151–1180.

- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): “Monetary Policy Shocks: What Have We Learned and to What End?” in *Handbook of Macroeconomics*, Elsevier, 65–148.
- COENEN, G., A. T. LEVIN, AND K. CHRISTOFFEL (2007): “Identifying the Influences of Nominal and Real Rigidities in Aggregate Price Setting Behavior,” *Journal of Monetary Economics*, 54, 2439–2466.
- COIBION, O., Y. GORODNICHENKO, AND G. H. HONG (2013): “The Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications,” *American Economic Review*, 105, 993–1029, manuscript, University of California at Berkeley.
- DHYNE, E., L. ALVAREZ, H. L. BIHAN, G. VERONESE, D. DIAS, J. HOFFMAN, N. JONKER, P. LÜNNEMANN, F. RUMLER, AND J. VILMUNEN (2005): “Price Setting in the Euro Area: Some Stylised Facts from Individual Consumer Price Data,” Working Paper 524, European Central Bank.
- DIXIT, A. AND J. E. STIGLITZ (1977): “Monopolistic Competition and Optimum Product Diversity,” *American Economic Review*, 67, 297–308.
- DOSSCHE, M., F. HEYLEN, AND D. V. DEN POEL (2010): “The Kinked Demand Curve and Price Rigidity: Evidence from Scanner Data,” *Scandinavian Journal of Economics*, 112, 723–752.
- DOTSEY, M. AND R. G. KING (2005): “Implications of State-Dependent Pricing for Dynamic Macroeconomic Models,” *Journal of Monetary Economics*, 52, 213–242.
- EICHENBAUM, M. AND J. D. M. FISHER (2005): “Evaluating Models of Sticky Prices,” *Computing in Economics and Finance 2005* 175, Society for Computational Economics.
- EICHENBAUM, M., N. JAIMOVICH, S. REBELO, AND J. SMITH (2014): “How Frequent Are Small Price Changes?” *American Economic Journal: Macroeconomics*, 6, 137–55.
- GERTLER, M. AND J. LEAHY (2008): “A Phillips Curve with an S,s Foundation,” *Journal of Political Economy*, 116, 533–572.

- GOLDBERG, P. AND R. HELLERSTEIN (2013): “A Structural Approach to Identifying the Sources of Local Currency Price Stability,” *The Review of Economic Studies*, 80, 175–210.
- GOLOSOV, M. AND R. E. LUCAS (2007): “Menu Costs and Phillips Curves,” *Journal of Political Economy*, 115, 171–199.
- GOPINATH, G. AND O. ITSKHOKI (2010): “Frequency of Price Adjustment and Pass-through,” *The Quarterly Journal of Economics*, 125, 675–727.
- (2011): *In Search of Real Rigidities*, Chicago: University of Chicago Press, vol. 25.
- GORDON, B. R., A. GOLDFARB, AND Y. LI (2012): “Does Price Elasticity Vary with Economic Growth? A Cross-Category Analysis,” *Journal of Marketing Research*, 50, 4–23.
- HAUSMAN, J. (1996): *Valuation of New Goods under Perfect and Imperfect Competition*, In: *The Economics of New Goods* ed. by T. Bresnahan and R. Gordon, University of Chicago Press, 307–342.
- HELLERSTEIN, R. (2008): ““Who Bears the Cost of a Change in the Exchange Rate? Pass-through Accounting for the Case of Beer,” *Journal of International Economics*, 76, 14–32.
- KIMBALL, M. (1995): “The Quantitative Analytics of the Basic Neomonetarist Model,” *Journal of Money, Credit and Banking*, 27, 1241–1277.
- KLENOW, P. J. AND J. L. WILLIS (2006): “Real Rigidities and Nominal Price Changes,” Federal Reserve Bank of Kansas City Working Paper No. RWP06-03.
- KRUSELL, P. AND A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867–896.
- LEVY, D., M. BERGEN, S. DUTTA, AND R. VENABLE (1997): “The Magnitude of Menu Costs: Direct Evidence From Large U. S. Supermarket Chains,” *Quarterly Journal of Economics*, 112, 791–825.

- MAĆKOWIAK, B. AND F. SMETS (2009): “On Implications of Microeconomic Price Data for Macroeconomic Models,” in *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*, ed. by J. Fuhrer, J. S. Little, Y. K. Kodrzycki, and G. P. Olivei, MIT Press, Cambridge, Massachusetts, 291–332.
- MIDRIGAN, V. (2011): “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, 79, 1139–1180.
- NAKAMURA, E. AND J. STEINSSON (2008): “Five Facts About Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 123.
- (2010): “Monetary Non-Neutrality in a Multisector Menu Cost Model,” *The Quarterly Journal of Economics*, 125, 961–1013.
- NAKAMURA, E. AND D. ZEROM (2010): “Accounting for Incomplete Pass-Through,” *The Review of Economic Studies*, 77, 1192–1230.
- NEVO, A. (2001): “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, 69, 307–342.
- PETRIN, A. AND K. TRAIN (2010): “A Control Function Approach to Endogeneity in Consumer Choice Models,” *Journal of Marketing Research*, 47, 3–13.
- TRAIN, K. (2009): *Discrete Choice Methods with Simulation*, Cambridge University Press, 2 ed.
- VERHELST, B. AND D. V. D. POEL (2012): “Implicit Contracts and Price Stickiness: Evidence from Customer-Level Scanner Data,” Working Papers of Faculty of Economics and Business Administration, Ghent University, Belgium 12/776, Ghent University, Faculty of Economics and Business Administration.
- WOODFORD, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press.

A Data preparation

This section describes how we get from the raw data to the final estimation data set.

In a first step, we group the data by categories and work with the data at the country level.

A list of categories by country is provided below

1. We drop households, where we did not observe any purchases at least one per year.
2. We replaced the household characteristics for income and age, where we are provided with ranges, with the median of the range. For Belgium, we are provided with 5 income categories: less than 496 EUR, 496–1239 EUR, 1240–1983 EUR, 1984–2726 EUR, and more than 2726 EUR. We do not have information on the consumers' age for Belgium. For Germany, we were provided with 16 income categories: less than 500 EUR, 500–749 EUR, 750–999 EUR, . . . , 3750–3999 EUR, and more than 4000 EUR. The 12 age categories for Germany are less than 19 years, 20–24, 52–29, . . . , 65–69, 70 and older. For the Netherlands, the data distinguishes 19 income categories: less than 7001 EUR, 700–900 EUR, 900–1100 EUR, . . . , 3900–4100 EUR, and more than 4100 EUR. The 11 age categories for the Netherlands are 12–19 years, 20–24, 25–29, . . . , 50–54, 55–64, 65–74, and more than 75 years.
3. We compute the price p used in the estimation as the price of a good per unit. For example, for a 500-grams pot of yogurt, which costs 99 cents, the price used in estimation is 0.198.
4. We removed outliers, defined as an observation, where the price p of a product is more than 200% larger than the average price of the identical product (GTIN).
5. Identify price changes by comparing prices of the same GTIN at different dates for the same retailer in two consecutive months. Calculate the size of price changes and the frequency of price changes per category and months. Calculate mean and median price adjustment size and frequency per category.
6. Nominal expenditure shares are calculated for each brand within a category within a country. We then rank the brands by expenditure share and use the top four brands plus a composite of all other brands (as the outside option) for estimation of the lower

nest.

7. To construct the Hausman instrument, we calculate the average price of the same brand in the same month for all NUTS regions excluding the region for each observation. We then regress the price p on monthly time dummies, region dummies (NUTS3 regions), brand-region fixed effects (NUTS1 regions), and the instrument. The residual of this regression is then used in the estimation.
8. To construct the loyalty variable *loyal*, we count how often a household has bought the same brand within the same category in the past.
9. The data contain all prices of all goods bought. We cannot directly observe prices of the alternatives. To construct the prices of the four alternatives, which are not chosen at a shopping occasion, we use the following procedure:
 - Search the data for the price of an alternative brand, bought within the same week, at the same retailer, in the same NUTS1 region.
 - If the search above was not successful, search the data for the price of an alternative brand, bought within the same month, at the same retailer, in the same NUTS1 region.
 - If the search above was not successful, search the data for the price of an alternative brand, bought within the same week, at the same retailer.
 - If the search above was not successful, search the data for the price of an alternative brand, bought within the same month, at the same retailer.
 - If the search above was not successful for any of the alternatives, drop the observation.
10. Estimate the lower nest for each category. Drop the category, if the price coefficient is not negative in an initial conditional logit specification, which we use as initial values for the Halton draws used for maximum simulated likelihood. We also drop the category, if the likelihood function did not converge after 25 iterations.
11. Construct observations for upper nest estimation, where a category was not chosen (outside option). We apply the following procedure to each category for each country:
 - Merge all observations (for all categories) for each household, which is included in

the category.

- Define an indicator variable, that is equal to one if the household chose the category at a shopping occasion and zero otherwise.
- Calculate the inclusive value for the observations with the indicator variable being unity.
- Construct the inclusive value, the control-function residual, and the variable *loyal* for the observations with the indicator variable being zero by following a similar procedure as above for the construction of prices of alternatives not chosen. In particular, we construct observations for price and the control-function residual by using the price and residual from the same brand for the same week, for the same retailer, the same region if observed. If that was not observed, we apply the same criteria, but use price and residual from the same month, instead of the same week. If both are not observed, we drop the observation in the estimation.

B Additional tables

Table 6: *Overview of goods categories included in estimation*

Belgium	Germany	Netherlands
Beer	Body powder	All-purpose cleaner
Bleach	Butter	Antrycide
Breakfast cereals	Canned instant meal	Bathroom polish
Buillons	Canned pickles	Cleansing tissue, dry
Butter	Ceramic glass cleaner	Dishwashing liquid
Canned instant meal	Dentifrice	Eye make-up
Canned Vegetables	Dried mushrooms	Foot care
Condensed milk /creamer	Eau de toilette, women's	Fresh bakery products
Cotton wool	Eiskonfekt	Hair Conditioning Products
Crispbread	Foils	Hairsprays
Curd cheese	Inceticide	Honey
Dessert sauce, wet	Liquors	Ketchup
Flavouring/Herbs	Metal cleaner	Stain remover
Frozen Dinners and Entrees	Pickled gherkin	Vinegar
Hairsprays	Poultry	
Inceticide	Shampoos	
Incontinence products	Sherbet powder	
Isotonic drinks	Sweet dishes	
Ketchup	Yogurt	
Laugen-/Cocktailgebaeck		
Mouthwash		
Packet Soup		
Seasoning and cocktail sauce		
Soap		
Soft spirit		
Sugar		
Sweet spreads		
Tea		
Tube cleaner		
Wine		

Notes: goods categories are sampled according to the description in the Appendix A. Some drop out of the estimation for one country and not for the other, because for some countries, there is enough data or the model converged after 25 iterations, whilst for others, there is not enough data or it did not converge.


Table 7: *Sample-specific preference and production-process parameters and simulation results (excluding sales)*

	Median elasticity		Weighted mean elasticity	
	Low persistence	High persistence	Low persistence	High persistence
CES				
Elasticity	3.58	3.58	2.08	2.08
Superelasticity	0	0	0	0
Menu costs	4.84	4.72	2.06	2.04
Std.dev. of id. shocks	8.83	6.9	8.86	6.82
<i>fpc</i>	13.08	13.12	13.1	13.08
<i>fraction^{up}</i>	55.98	54.98	54.52	54.13
<i>size^{abs}</i>	17.25	17.26	17.23	17.22
<i>size^{up}</i>	16.46	16.77	16.88	16.99
<i>size^{down}</i>	18.25	17.87	17.65	17.48
Model outcomes				
CIR	0.011	0.009	0.012	0.008
Var. of C	0.01	0.008	0.01	0.009
Var. of C: data	0.245	0.245	0.245	0.245
Var. of C explained	0.041	0.033	0.041	0.037
NON-CES				
Elasticity	3.58	3.58	2.08	2.08
Superelasticity	1.39	1.39	1.47	1.47
Menu costs	7.6	7.16	4.71	4.76
Std.dev. of id. shocks	13.76	10.56	20.82	16.25
<i>fpc</i>	13.14	13.08	13.09	13.1
<i>fraction^{up}</i>	54.22	54.17	52.54	52.85
<i>size^{abs}</i>	17.27	17.23	17.21	17.21
<i>size^{up}</i>	17	16.99	17.5	17.39
<i>size^{down}</i>	17.58	17.51	16.89	17
Model outcomes				
CIR	0.013	0.012	0.015	0.022
Var. of C	0.015	0.016	0.023	0.019
Var. of C: data	0.245	0.245	0.245	0.245
Var. of C explained	0.061	0.065	0.094	0.078


Notes: 1) Menu costs are given by the percentage share of steady-state revenue $\frac{\theta-1}{\theta} \frac{K}{Y^*}$. 2) The values for the elasticity and super-elasticity are taken from the estimation results as reported in Table 3. 3) Menu cost and standard deviation parameters are chosen so as to match the median frequency and size of price changes reported in rows *fpc* and *size^{abs}* of the upper panel (descriptive statistics including sales) of Table 2. 4) CIR (cumulative impulse response) and Var. of C (variance of aggregate real consumption) represent measures of monetary non-neutrality as described in the main text. The numbers for the variance of C are multiplied by 10^4 . (5) To compute the variance of real consumption we employ HP filtered quarterly real consumption (summed over the three countries included in the data sample) from 1995 to 2008. The source of the data is Eurostat.

C Additional figures


Figure 5: *Inflation in the sample and HICP inflation rates*



(a) Belgium




(b) Germany



(c) Netherlands


Notes: Monthly inflation rates in our data sample are calculated as the weighted (by expenditure share) average change in the price level of all products bought in the estimation sample. Monthly HICP inflation rates are taken from Eurostat and cover the same period as our sample (2005–2008).

Figure 6: *Sizes and frequencies of price changes in the data set*




Notes: the size of a price change is the percentage change from the previous month price of the identical product, given that a price for the previous month for the same household and retailer can be observed. The frequency is the monthly frequency of price changes at the product level, where a price spell is identified only if a price for the previous month for the same household and retailer can be observed. Analogously, only prices, where two prices are observed in two consecutive months and identical are not a price change. All prices, where we did not observe two consecutive months for the same household and retailer are dropped before calculating these figures.


Figure 7: Demand curves implied by elasticity estimates



(a) Median elasticity



(b) Weighted mean elasticity



(c) Mean elasticity

Notes: Demand curves implied by the various elasticity estimates reported in Table 3.