

Price elasticities and demand-side real rigidities in micro data and in macro models*

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

Abstract

High degrees of demand-side real rigidities are able to generate the large monetary non-neutrality found in aggregate data. This paper provides micro-based evidence on the key parameters governing this rigidity using European homescan data. We find strong evidence for demand-side real rigidity, which is, however, significantly lower than that normally assumed in macro models. In a menu-cost model calibrated to our estimates, we show that our estimates are associated with reasonable values for production-side parameters, but they are not able to generate the degree of monetary non-neutrality observed in macro data.

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

Keywords: Demand curve, price elasticity, super-elasticity, price setting, real rigidities, 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, Maarten Dossche, 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. This paper previously circulated under the title “Microeconomic evidence on demand-side real rigidity and implications for monetary non-neutrality”.

[†]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, CEPR, and KOF ETH Zürich.

1 Introduction

The specification of consumer preferences and the resulting demand curves play a key role in the outcomes of economic models in many different fields such as industrial organization, trade, macroeconomics and monetary economics. The key parameters that describe the shape of these demand curves are the price elasticity, i.e., the percentage change in demand in response to a one percent change in a good's price, and the super-elasticity, which corresponds to the price elasticity of the price elasticity (Klenow and Willis, 2016).¹

While the most standard macroeconomic model assumes a constant elasticity of demand, following Dixit and Stiglitz (1977), demand curves with positive super-elasticities, following (Kimball, 1995), have been used in many recent models, since they introduce a demand-side real rigidity that can reconcile the simultaneous existence of a fairly high degree of price flexibility at the micro level and substantial monetary non-neutrality at the macro level. In addition, these models match the persistence that is often present in aggregate data.² If firms face demand curves with non-constant price elasticity, the response of demand to price changes is asymmetric: a reduction in the relative price increases relative demand for a good by less than a relative price increase of the same size reduces demand. Figure 1 plots the demand and profit functions associated with different parameters for price elasticity and super-elasticity: to illustrate the main mechanism, we refer to the demand specification employed by Gopinath and Itskhoki (2010) and the counterfactual with constant demand elasticity (CES) for this specification (red lines). When the super-elasticity is positive, demand and profits lie below the levels of the counterfactual when the relative price deviates from one, because firms have stronger incentives to keep prices closer to the aggregate

¹Demand curves firms are faced with shape firms' responses to cost shocks and other supply-side decisions and the elasticity of demand and the curvature of the demand curve (related to the super-elasticity) are important statistics that are used for many comparative static questions in the fields mentioned above (see Mrázová and Neary, 2017).

²See, e.g., the literature cited in the notes to Figure 1. Recently, Lindé and Trabandt (2019) show that also the missing deflation in the aftermath of the Great Recession can be replicated in a macroeconomic model with demand-side real rigidity.

price level if the super-elasticity is positive, compared to the counterfactual. Therefore, desired markups co-move negatively with relative prices, which makes firms’ profits more sensitive to the prices of competitors and thereby gives rise to strategic complementarities in price-setting. Overall, compared to CES demand, positive super-elasticities cause prices to have a more sluggish response to both idiosyncratic and aggregate shocks.

Macroeconomic models embracing Kimball-type preferences have used a broad range of values for the calibration of the super-elasticity parameter, which is illustrated in Figure 1.³ The demand curves exhibit fairly different behavior on the side of consumers in response to price changes, which causes firms to have different profit functions. These differences are not surprising given that the values of the individual models are generally derived to match macro moments against the backdrop of different theoretical frameworks and time periods.


We provide empirical estimates for the size of price elasticities and super-elasticities for a wide range of products. We employ a rich homescanner dataset that includes data on 76 goods categories in three European countries: Belgium, Germany and the Netherlands. One advantage of using homescan data is that we can directly track consumers’ shopping behavior over time with different retailers, and we can control for important consumer characteristics, such as income or age, which arguably influence the shape and location of demand curves. Applying a flexible discrete choice model, we use these data to provide micro-based evidence on price elasticities and the associated super-elasticities. Compared to previous micro-based contributions, our data coverage is considerably broader.⁴

Our results support previous micro-based empirical evidence and thus underpin the criticisms of Chari et al. (2000) and Klenow and Willis (2016) that the values for

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

⁴Goldberg and Hellerstein (2013), e.g., use Dominick’s *Finer Foods* data and estimate the super-elasticity of the goods category “beer” at 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) and Verhelst and Van Den Poel (2012) find super-elasticities in the range of 4 for goods with an elasticity of 3 or larger.

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



Notes: Figure 1 plots the demand functions and the resulting profit functions implied by the values of the elasticities and super-elasticities of demand resulting from our estimates and 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 (2016). These papers include Bergin and Feenstra (2000), Chari et al. (2000), Coenen et al. (2007), Eichenbaum and Fisher (2005), Gopinath and Itskhoki (2010), Kimball (1995), Klenow and Willis (2016) and Woodford (2003). 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 paranthesis after the author name(s) indicate the value of the elasticity and super-elasticity assumed. The shaded area covers the range of demand and profit functions that are compatible with our empirical estimates of the elasticity and super-elasticity. To determine this range, we used the elasticity values associated with the 10% and 90% of the empirical distribution of this parameter. The super-elasticity corresponding to these elasticity values were obtained using the empirical relationship documented in Appendix C.1.

super-elasticities employed in most of the macro literature imply highly implausible behavior. The shaded area of Figure 1 represents the range of demand and profit functions that are compatible with the empirical distribution of our estimates.⁵ Our estimates suggest that many of the specifications with large super-elasticities considered in the literature belong to rather extreme cases, that are close to or even clearly above the 90th percentile of the estimated super-elasticities. More moderate values such as those assumed in Bergin and Feenstra (2000) and Gopinath and Itskhoki (2010) are closer to

⁵The borderline cases considered are the ones corresponding to the 10th (1.41) and 90th percentiles (11.45) of the empirical distribution for the elasticity estimates.

the majority of our empirical estimates. Our results, moreover, suggest fairly moderate elasticity values well below 10. By simulating a standard menu cost model calibrated on the basis of the obtained micro evidence on elasticity and super-elasticity values, we show that the demand-side real rigidities implied by the estimates are rather small. Demand side real-rigidities and nominal rigidities alone are not large enough to match both micro and macro facts on prices. Moreover, we show that the high degree of monetary non-neutrality resulting from the assumptions of (unrealistically) large super-elasticity values would require highly implausible assumptions regarding production-side parameters.

The rest of this paper is structured as follows: in Section 2 we describe our dataset and provide some descriptive statistics. Section 3 outlines the approach which we use to estimate demand curves, whereas Section 4 presents the empirical results. Section 5 outlines 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 that has not been used in the macro literature before. The data have been made available by AiMark (Advanced International Marketing Knowledge).⁶ The database is comparable to the IRI household panel and the AC Nielsen homescan data for the US.⁷

⁶AiMark is a nonprofit 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.

⁷For example, Coibion et al. (2013) use IRI data to examine the cyclical properties of prices. Kaplan and Schulhofer-Wohl (2017) and Stroebel and Vavra (2018) use AC Nielsen homescan data to examine price dynamics at the household and regional level, respectively.

2.1 Description of the dataset

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 provided with a scanning technology which it uses to scan all the products purchased at retail outlets, including all major supermarket chains (such as Rewe and Aldi in Germany or Albert Heijn and C1000 in the Netherlands), drugstores, small corner shops and internet stores.

For each product bought, the household scans the bar-code, which uniquely identifies the product via the Global Trade Item Number (GTIN)⁸ and enters the number and associated price for this product into the scanning device.⁹ The dataset contains a description of each product and a classification system of the goods into more aggregate product categories.¹⁰ The household also provides the name of the retailer where it bought each product. The products belong to the categories of fast-moving consumer 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 on the data employed in the estimation.¹¹ As mentioned above, we used available scanner data for Belgium, Germany, and the Netherlands. The sample period is from 2005 to 2008. The number of households is 1,746 for Belgium, 11,631 for Germany, and 4,030 for the Netherlands. We observe 21.8

⁸The GTIN-13 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).

⁹In case the product does not have a bar-code, the household enters this information manually.

¹⁰We constructed a common classification scheme for the products in our dataset (for all countries) that is based on the scheme 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 of country representatives of GfK and extensive documentation on the different classification schemes, to which we had access at the data providers' offices.

¹¹See Appendix A for the sampling procedure.

million transactions in the dataset, ranging from 2.1 million for the Netherlands to 10.4 million for the Netherlands. These observations include the purchases of 191,334 unique products. The products come from 76 different categories (such as beer or butter: see Table E.1 in the Appendix for a list of all categories included in the estimation).

Table 1: *Database - Overview*

Country	Households	Unique products	Product Categories	Transactions
Belgium	1,746	49,808	31	2,057,254
Germany	11,631	72,617	20	9,376,238
Netherlands	4,030	68,909	25	10,373,400
Total	17,407	191,334	76	21,806,892

Notes: Entries in the column “Households” report the number of different households in the estimation sample. The “Products” column provides the number of unique GTINs that are included in our estimation sample. The column “Categories” contains the number of product categories included in the sample, whereas the column “Transactions” reports the number of purchases we observe.

For the econometric analysis, we first split the data into 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 that would be available in the same retailer at a given shopping trip of a household. We therefore construct alternative prices as described in Appendix 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, overall, comparable to those obtained for the US IRI homescan dataset reported in Coibion et al. (2013). We find an average frequency of price changes (fpc) of 17.1% on a monthly basis, which is a bit lower than the 23.8%

reported for the US data (Table 2).¹² These figures are also comparable to the frequencies of price changes in the CPI and PPI micro data for the euro area. Álvarez et al. (2006) report a frequency of 28% for unprocessed food and 14% for processed food in the CPI and 12% for nondurable, nonfood items in the PPI. Most of the products in our dataset fall into one of these nondurable goods categories, which suggests that the mean frequency of 17.1% and the median frequency of 15.6% seem reasonable. Excluding sales, we find a mean (median) frequency of 15% (12.6%), suggesting that the sales in our dataset do not change the frequency of price changes substantially.¹³ The share of price increases in all price changes ($fraction^{up}$) reported in the evidence for the euro area collected in Dhyne et al. (2005) is 54%, which is close to our mean and median.¹⁴

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

	Including sales		Excluding sales	
	Mean	Median	Mean	Median
fpc	17.11	15.65	14.97	12.62
$fraction^{up}$	52.62	51.49	52.99	51.54
$size$	-0.44	0.03	-0.42	0.04
$size^{abs}$	17.05	14.54	17.43	14.74
$size^{up}$	16.73	13.49	17.02	13.43
$size^{down}$	-17.25	-15.46	-17.74	-15.7

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

The average absolute size of price changes, $size^{abs}$, is 17.1% in our data and thus a bit smaller than the sizes of price changes for the US IRI data reported in Coibion et al. (2013). Their average price increase (decrease) is 20.6% (22.4%), whereas the sizes of the

¹²See Coibion et al. (2013), Table A1 in their Appendix.

¹³This result might of course be due to the fact that we can identify sales only via a v-shaped sales filter, which classifies all price changes that are exactly reverted in the next month, as a sale. We do not have a sales flag in the dataset.

¹⁴The distributions of the sizes of price changes and the frequencies of price changes are shown in Figure F.1 in the Appendix.

increases (decreases) in our data are 16.7% (17.3%).¹⁵

The key relationship we analyse in the following sections is the relationship between changes in demand and changes in prices. In Figure F.2 in the Appendix, we show the correlation between changes in prices and changes in quantities purchased at the product (GTIN) level. The correlation is negative and significant at the one percent level, but the coefficient is very small, suggesting that a one percent increase in the price is associated with only a small decline in the quantities purchased. Even though this correlation is not yet an estimation of a demand curve, in particular, the causality and the relative prices are not taken into account, it is already indicative of the later findings that demand elasticities, even though they are much higher than what the correlation suggests, are nevertheless, not very large.

3 Empirical method

In this section, we briefly describe the econometric model used. We estimate demand elasticities using a nested multinomial logit with random coefficients. Employing a nested framework enables us to depict a consumer’s decision of buying in a given product category and, conditional on choosing to buy in a category, which product to buy within the category. The model thus combines the decision to substitute across product categories and within product categories. Using a multinomial logit setup allows us to model a consumer’s choice among J alternative product brands within a product category, and the random coefficients

¹⁵In general, the statistics on the size of price changes in homescan data are larger than the statistics for CPI microdata for the euro area, as reported in Dhyne et al. (2005). These authors report an average price change of 15 to 16% for unprocessed food and of 7 to 8% for processed food in the euro area. One explanation for this result might be that many small price changes in CPI data are erratic, as shown in Eichenbaum et al. (2014): these authors find that many small price changes in CPI micro data are due to measurement issues, for example because of product replacements and quality changes, 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.

make it possible to capture potential random taste variations across consumers, which make price elasticities vary over brands, product categories, and consumers. Applying a discrete choice model to estimate demand elasticities has the advantage that the demand system derived from such a model is consistent with a wide class of CES- and non-CES utility functions, including Kimball-type preference specifications.¹⁶ In addition, we use the control function approach to control for price endogeneity.

3.1 Discrete choice specification

The model consists of two nests. In the upper nest, a household chooses whether or not to buy in a given category c . Conditional on choosing that category, in the lower nest, the household decides on the brand within the category. For example, a household chooses to buy an item in the category “pasta”, and within this category, it chooses to buy a certain brand, for example “Barilla”. We first describe the multinomial choice model for the lower nest and then the discrete choice model for the upper nest.

The conditional utility of household i purchasing brand j , conditional on buying in category $c(j)$, at shopping occasion t in the *lower* nest is given by

$$U_{ijt|c(j)} = \beta_{ij} - \alpha_i p_{ijt} + \delta x_{ijt} + \varepsilon_{ijt}, \quad (1)$$

where β_j is a vector of brand-specific effects and p_{ijt} is the purchase price at which household i buys brand j . Additional controls in x_{ijt} are the household’s income and average age (interacted with brand-specific dummies), brand-region fixed effects, and a

¹⁶McFadden and Train (2000) show that choice probabilities from a random utility model can be approximated very closely by a multinomial logit model. See (Ch. 6 Train, 2009, for an overview). Anderson et al. (1987) and Anderson et al. (1992) document that the demand system derived from a nested logit model is also generated by a CES utility function. This result is generalized to non-CES utility functions, including Kimball-type preferences, in Thisse and Ushchev (2016).

proxy for loyalty.¹⁷ We allow α and β_j to vary across households and model heterogeneity by assuming normal mixing distributions for β_j and a lognormal distribution for α . This assumption 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} .^{18, 19} The choice in the lower nest is a multinomial model since it represents the choice among five alternatives: product brand j , one of the other three product brands in the same category, and the outside brand.

Utility from purchasing in category $c(j)$ in the upper nest is given by

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

where w_{ict} counts the number of weeks since the household last purchased in the category and $IncV_{ict}$ is the inclusive value from (1), $\ln \sum_j [\hat{\beta}_{ij} - \hat{\alpha}_i p_{jt} + \hat{\gamma}_i loyal_{ijt} + \hat{\delta} x_{ijt}]$, which is the utility derived from the choice options in the lower nest in category $c(j)$. The parameter ρ describes the need to buy some products regularly (for example, if a household last bought a product in the category “toilet paper” many weeks ago, the probability of buying a product from this category should increase), while Ψ captures the effect of the utility

¹⁷Region 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, there are the three regions: Brussels Capital Region, Flemish Region and Walloon Region. For Germany, there are 16 German Landers. For the Netherlands, there are four regions: north, east, west and south Netherlands. Loyalty is proxied by a count variable, which represents the number of purchases of the same brand in the same category in the past and is included to control for motives to buy a given product that are unrelated to a brand’s current price (for example, due to habits). See, for example, Gordon et al. (2012), who apply the nested multinomial logit model to estimate demand elasticities.

¹⁸We index the shopping occasion by t , which does not refer to a well-defined constant frequency, but to the date when household i shops and buys a product in category c .

¹⁹Actually, the assumption in the multinomial logit model is that the difference between the two error terms of choice j and another choice option m in the same product category are logistically distributed. For identification, the scale of utility is irrelevant, since only differences between the utility obtained from buying brand j and the utility from buying the other brands in the same category matter. Therefore, the error term ε_{ijt} is assumed to be iid extreme value because the difference between two extreme value variables is distributed logistically. The mean utility from the outside option is not separately identified; it is therefore set to zero. See Appendix B for more details.

household i obtains from the choices in category c . The error term ν_{ict} is assumed to be distributed iid logistic. The choice in the upper nest is therefore a logit model, with a zero/one decision to buy in a given category $c(j)$. We use a maximum likelihood estimator, which constrains Ψ to lie between zero and one.

To control for endogeneity in prices, we use the control function approach (Petrin and Train, 2010). The main idea is to obtain a measure of the variation in prices that is endogenous and to include it directly as a control variable in the main estimation, so that it captures the variation in the unobserved factor that is not independent of the endogenous variable. The first step in implementing this methodology is to regress the observed price on the mean price of the same brand in other regions within the same country.²⁰ The second step is to retain the residual, which reflects the component of prices that is correlated with demand shocks, from this equation and to include it as a variable in equation (1). This approach allows us to estimate price elasticities of demand and super-elasticities of demand in a very flexible manner while controlling for price endogeneity.²¹

To estimate the model, we apply a three-step approach. First, we estimate control functions and retain the residual $\hat{\mu}_{jt}$. Second, we estimate the mixed logit for the lower nest with $\hat{\mu}_{jt}$ and x_{ijt} in equation (1). Third, we estimate the logit for the *upper* nest using the inclusive value calculated from the results obtained in the second step. Below, we describe how we then calculate the elasticities and super-elasticities for each household within a category.

²⁰Underlying this ‘Hausman’-instrument is the assumption that prices are correlated across regions due to common marginal costs but demand shocks average out in the mean price of the same brand in other regions. See Hausman (1996) and Nevo (2001).

²¹An alternative way to estimate the model would be to use the BLP (Berry et al., 1995) approach as in Nakamura and Zerom (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, because they are needed to estimate product-specific constants, which should remove the endogeneity from the error term. If market shares are not observed precisely, the control function approach is more reliable. See Petrin and Train (2010) for a discussion.

3.2 Demand elasticities and super-elasticities

Demand elasticities and super-elasticities can be derived from the nested mixed logit model in a straightforward manner (see Appendix B.3 for derivations). The demand elasticity for the upper nest is given by $\theta_{ijc(j)}^u = \alpha_i P_{ij|c(j)}(1 - P_{ic(j)})p_{ijt}$ and the elasticity for the lower nest is $\theta_{ijc(j)}^l = \frac{\alpha_i}{\Psi} p_{ijt}(1 - P_{ij|c(j)})$. $P_{ij|c(j)}$ denotes the predicted probability that household i chooses brand j , conditional on buying a product in category $c(j)$. $P_{ic(j)}$ is the predicted probability from the upper nest of choosing category $c(j)$. Because the probability of choosing brand j in category $c(j)$ can be written as the product of the marginal probability $P_{ic(j)}$ and the conditional probability $P_{ij|c(j)}$, the total elasticity is just the sum of the upper and the lower elasticities $\theta_{ijc(j)} = \theta_{ijc(j)}^u + \theta_{ijc(j)}^l$.

One advantage of using 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.²² The super-elasticity of the upper nest is $\epsilon_{ijc(j)}^u = 1 + \theta_{ijc(j)}^l - \Psi \alpha_i P_{ij|c(j)} P_{ic(j)} p_{ijt}$ and for the lower nest, it is $\epsilon_{ijc(j)}^l = 1 - \frac{\alpha_i}{\Psi} p_{ijt} P_{ij|c(j)}$. The total super-elasticity is the sum of the two and thus is given by $\epsilon_{ijc(j)} = \epsilon_{ijc(j)}^u + \epsilon_{ijc(j)}^l$.


4 Empirical results

To summarize the estimates, we first calculate all elasticities for all categories, brands within each category, and households. Before we calculate the empirical statistics, we

²²This is because the model allows for heterogeneity in consumers' price sensitivity, which contributes to the curvature of demand. This model nests the CES with a super-elasticity of zero as a special case. See also (Hellerstein, 2008) and (Goldberg and Hellerstein, 2013). Appendix B.3 derives the super-elasticities.

trim the distribution of the estimates.²³ We then calculate an unweighted mean and two versions of a weighted mean: one in which we weight the elasticities by importance in total expenditures and one where we weight the elasticities by estimation precision. For the first measure of the weighted mean, labelled ‘Weighted Mean Expend.’ in Table 3, we use expenditure shares as weights. For the second measure of the weighted mean, labelled ‘Weighted Mean Variance’, we use the inverses of the estimated variances of the demand elasticity coefficients, α_i , as weights.²⁴

Figure 2: *Distribution of the elasticities and super-elasticities*



Notes: Panel a) shows the distribution of the estimated elasticities, θ , while panel c) shows the distribution of the estimated super-elasticities, ϵ . All estimates vary across households and brand-product category combinations. The distributions are shown for the trimmed sample.

The distributions of all estimates of the elasticity and super-elasticity for each category and household are shown in Panels A and B of Figure 2, respectively. The distribution of

²³We trim the distribution because we estimate a very large number of category- and household-specific coefficients, and there are a few outliers, which would otherwise contaminate the estimates of the mean. More specifically, for the super-elasticities and markup-elasticities, which are not constrained in the sign of the elasticity, we cut off the upper and lower 10% of all estimates. We trim demand elasticities by cutting off the upper 10% of the absolute value of the estimates and elasticities estimated below an absolute value of 1, because they are not in line with a model of monopolistic competition. The medians of the super-elasticities and markup-elasticities are unaffected by this symmetric trimming, while the median of the untrimmed distribution of demand elasticities would be even lower (at 2.07).

²⁴This weighing procedure is inspired by methods used in meta analyses that aim to provide an overall effect estimate (weighted average) across multiple studies analyzing the same relationship. The so-called fixed effect estimate is defined as the weighted average of an effect size (for example, a regression coefficient), which is weighted by the inverse of the estimated variance of the respective effect size. Rice et al. (2018) show that the fixed-effects estimate is a robust metric to quantify an overall average effect.

the 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 the price elasticities across categories and households. This heterogeneity 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 range between zero and five.

How large should the parameters for price elasticities and super-elasticities be in macroeconomic models? For the median, the price elasticity is rather small, at a value of 3.20 (Panel A of Table 3). The mean and weighted means are a bit higher, between 4.81 and 5.40. Many macroeconomic models assume a price elasticity of ten, which seems to be a relatively large value, given that more than 85% of our estimates are below ten. Not surprisingly, elasticities in the lower nest (the choice between brands within a product category) are higher than the elasticities in the upper nest (the choice across product categories): this is intuitive, since we would expect the elasticity of substitution to be higher for closer substitutes.

Even though 90% of the estimated values for the super-elasticities are positive, they are well below ten (Panel B of Table 3). The median estimate for the super-elasticities is 1.93, and the (weighted) means are slightly lower, between 1.44 and 1.59.²⁵ 95% of our estimates are below a value of five, suggesting that the super-elasticities are positive but small. Taking these pieces of evidence together, the data does not support assumptions of a super-elasticity parameter of ten or higher.

As shown in Gopinath and Itskhoki (2010), the elasticity of the markup μ is a function of the demand elasticity and the super-elasticity and can be expressed as $\frac{\partial \mu}{\partial \ln P}|_{P=1} = \frac{\epsilon}{\theta-1}$. This implies that the markup elasticity increases in the super-elasticity but decreases in

²⁵These results are close to the estimates for the beer market found in Goldberg and Hellerstein (2013), and somewhat lower than the estimates for the coffee market found in Nakamura and Zerom (2010) or for the European retailer found in Dossche et al. (2010).

Table 3: Descriptive Statistics Estimates

Panel A: Estimates of Demand Elasticities							
	Mean	Median	Weighted Mean		Percentile		Sign.
			Expend.	Variance	90 th	10 th	5%
Total elasticity	4.81	3.20	5.40	5.03	11.21	1.41	
Upper nest	0.42	0.29	0.40	0.46	1.02	0.07	100.00
Lower nest	4.27	2.78	4.89	4.41	10.19	1.16	96.4
Panel B: Estimates of Super-elasticities							
	Mean	Median	Weighted Mean		Percentile		
			Expend.	Variance	90 th	10 th	
Total super-elasticity	1.59	1.93	1.44	1.50	2.23	0.05	
Upper nest	-0.55	0.07	-0.69	-0.92	0.83	-3.58	
Lower nest	2.01	1.67	2.02	2.17	3.34	1.08	
Implied markup elasticity	0.99	0.68	0.87	0.51	2.61	-1.07	

Notes: Panel A shows the summary statistics for the estimated demand elasticities θ , and Panel B shows the estimates for the super-elasticities ϵ and (in the last row) the implied markup elasticities $\frac{\partial \mu}{\partial \ln P}|_{P=1}$. The first two columns show the mean and median values of the estimated elasticities. The third and fourth columns show the weighted means, which are weighted by expenditure share and by the inverse of the estimated variance of α_i . The fifth and sixth columns show the 90th and 10th percentile of the distribution of estimates, respectively. The last column shows the share of estimates that are significant (in the row ‘Upper nest’, we show the share of Ψ ’s that are significant at the 5% level and in the row ‘Lower nest’, we show the share of α_i ’s that are significant at the 5% level). Since the super-elasticities are a function of the same parameters as the demand elasticities, we do not repeat the significance levels in Panel B. The demand elasticities are multiplied by -1 , also before calculating the markup elasticities.

the demand elasticity. We therefore also calculate the implied markup elasticities for each estimated value of demand elasticity and super-elasticity.²⁶ Even though our estimated super-elasticities are small, the estimated demand elasticities are also smaller than what is often assumed in macroeconomic models. Therefore, the implied markup elasticities remain economically significant and range between 0.5 and 1 for the (weighted) mean and median, implying that firms adjust their markups by approximately one-half to one percent if a firm’s relative price changes by one percent (Panel B of Table 3).²⁷

²⁶Since there is a correlation between the estimates of the elasticity and the estimates of the super-elasticity, we calculate the markup elasticity for each estimate and provide descriptive statistics for the resulting distribution of markup elasticities.

²⁷These estimates are close to the markup elasticity estimates in Amiti et al. (2016), who report estimates in the range of 0.6 – 1.

5 A menu-cost model with Kimball-type preferences

In this section, we describe the theoretical model that we use for calibration to quantitatively assess the role of our estimated demand-side real rigidity for monetary non-neutrality. A detailed derivation is provided in Appendix D. The model we employ belongs to the class of menu-cost models in which price adjustments are costly, and it features a demand-side real rigidity.²⁸ For the purpose of our study, this model class is well suited because it allows us to quantitatively assess the size of the idiosyncratic shocks and/or menu costs that are needed to match major price facts when empirically plausible values for the demand elasticity and super-elasticity are employed.

5.1 Model setup

We assume that our economy is inhabited by a representative household, a continuum of firms and a monetary authority that controls the evolution of nominal GDP. The composite consumption good, C_t , that enters households' utility is created by the costless aggregation of a continuum of differentiated goods, $c_t(z)$, which are supplied by monopolistic firms. We implicitly define the composite consumption good, C_t , using an aggregator of the form $\int_0^1 \Upsilon\left(\frac{c(z)}{C}\right) dz = 1$, 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 (Kimball, 1995). Following Klenow and Willis (2016), we employ

$$\Upsilon(x) = 1 + (\bar{\theta} - 1) \exp\left(\frac{1}{\bar{\epsilon}}\right) \bar{\epsilon}^{(\frac{\bar{\theta}}{\bar{\epsilon}} - 1)} \left[\Gamma\left(\frac{\bar{\theta}}{\bar{\epsilon}}, \frac{1}{\bar{\epsilon}}\right) - \Gamma\left(\frac{\bar{\theta}}{\bar{\epsilon}}, \frac{x^{\frac{\bar{\epsilon}}{\bar{\theta}}}}{\bar{\epsilon}}\right) \right], \quad (3)$$

²⁸Other recent papers employing a menu-cost model framework include, Nakamura and Steinsson (2010), who we follow closely, Kryvtsov and Midrigan (2013), Nakamura et al. (2017), Klenow and Willis (2016), Klepacz (2017), Vavra (2013) and Gautier and Le Bihan (2018), among others.

with $x = \frac{c(z)}{C}$ and $\Gamma(u, z)$ denoting the incomplete gamma function. The resulting demand function is given by

$$c(z) = \left[1 - \bar{\epsilon} \ln \left(\frac{p(z)}{P} \right) \right]^{\frac{\bar{\theta}}{\bar{\epsilon}}} C, \quad (4)$$

where demand for good z depends positively upon overall consumption demand C and negatively upon the relative price of good z .²⁹ $\bar{\theta}$ and $\bar{\epsilon}$ determine the steady-state sizes of the elasticity and super-elasticity of demand, respectively.³⁰ 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 case where $\bar{\epsilon} > 0$, i.e., the NON-CES case, the price elasticity of demand, $\theta^{NON-CES}$, and the super-elasticity, $\epsilon^{NON-CES}$, are respectively given by

$$\theta^{NON-CES} = -\frac{\bar{\theta}}{1 - \bar{\epsilon} \ln \left(\frac{p(z)}{P} \right)} \quad \text{and} \quad \epsilon^{NON-CES} = \frac{\bar{\epsilon}}{1 - \bar{\epsilon} \ln \left(\frac{p(z)}{P} \right)}. \quad (5)$$

These expressions show that for $\bar{\epsilon} > 0$, the elasticity of demand increases in the relative price of good z and will be larger than that of 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.

Firms produce goods using labor, which is subject to idiosyncratic changes in labor productivity, and set prices to maximize profits. Changing prices is associated with a menu cost. The firms' objective function is given by $E_0 \sum_{t=0}^{\infty} D_{0,t} \Pi_t(z)$, where profits

²⁹In deriving equation (22), we made use of an approximation result obtained by Gopinath and Itskhoki (2010, Appendix). The CES case can be directly derived by employing the aggregator function $\Upsilon \left(\frac{c(z)}{C} \right) =$

$\left(\frac{c(z)}{C} \right)^{\frac{\bar{\theta}-1}{\bar{\theta}}}.$

³⁰See Equations (12) and (25) in the Appendix for the formal definitions of these two parameters.

$\Pi_t(z)$ are given by

$$\Pi_t(z) = p_t(z)y_t(z) - W_tL_t(z) - KW_tI_t(z). \quad (6)$$

$y_t(z)$ denotes firm's z sales, which are equal to $c_t(z)$. The firm's z revenue is represented by $p_t(z)y_t(z)$, labor costs by $W_tL_t(z)$, and $KW_tI_t(z)$ 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).³¹ We solve the firm's optimization problem using dynamic programming. The details of the numerical solution strategy are given in the appendix.

The monetary authority controls the path of nominal GDP according to the process

$$\ln Y_t^N = \mu + \ln Y_{t-1}^N + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2), \quad (7)$$

where aggregate nominal GDP, $Y_t^N = P_t C_t$, grows at a constant long-run rate μ and is subject to temporary shocks

5.2 Calibration

To examine the quantitative implications of our empirical findings for the plausibility of firms' price-setting parameters and 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. The parameters in this group comprise the subjective discount factor, β , the parameter of relative risk aversion, γ , the elasticity of labor supply, ψ , steady-state labor supply, L^* and the persistence parameter of idiosyncratic technology shocks, ρ . The values chosen for these coefficients are standard in the literature. We

³¹In the specification underlying our simulations, we additionally include a term capturing the fixed nominal costs of production. This step is motivated by the fact that our fairly low values for the elasticity would give rise to large markups and therefore firm profits, which we cannot observe in the national accounts.

employ a discount factor $\beta = 0.96^{1/12}$ implying an annual (steady-state) interest rate of approximately 4%. The parameter of relative risk aversion, γ , is set equal to 1.5, which corresponds to the median value estimated by Smets and Wouters (2003) (1.6) for the euro area. For the value of the elasticity of labor supply, ψ , we likewise choose a value conforming to the median value estimated by Smets and Wouters (2003). These two values influence the interaction between the real wage and spending and labor supply. We set $\bar{L} = 1/3$, i.e., the steady-state labor supply is assumed to correspond to one-third of the total time endowment. The persistence of the productivity shocks is assumed to be 0.7.³²

To calibrate the mean inflation rate and the standard deviation of money supply shocks, we use CPI data obtained from Eurostat over the period from 1999 to 2008 for Belgium, Germany and the Netherlands. The mean inflation rate, 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 the considered model specifications*

Parameter description	Value
Subjective discount factor, β	$0.96^{1/12}$
Relative risk aversion, γ	1.5
Elasticity of labor supply, ψ	1
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 across the specifications (Table 5, column 1). Reflecting the considerable heterogeneity in our estimates for the elasticity and, to

³²In our model, a higher shock persistence implies that firms will tend to adjust prices more frequently in response to shocks of similar sizes due to their higher associated longevity. As a result, smaller values for the menu costs and shock variances are needed to match the observed price characteristics. The main conclusions from our calibration exercise remains unchanged (see working paper version where the case of a shock persistence of 0.9 is considered).

a smaller degree, the super-elasticity, we consider a variety of scenarios for these two parameters. The cases examined represent permutations of the selected values (median, 10th percentile and 90th percentile) from the two distributions of the estimated coefficients (see Table 3).³³ For each elasticity value, we also consider the case of a super-elasticity of zero (CES case). As a comparison, we report preference specifications employed in the macro literature and illustrated in Figure 1 in the lower panel of Table 5.

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 Table 2.³⁴ The menu costs are reported as a fraction of steady-state revenue.

The coefficient values associated with our estimates reveal some clear patterns. In line with our intuition, the sizes of both parameters tend to increase with the size of the super-elasticity. Apart from the extreme low elasticity case, both menu costs and standard deviations of idiosyncratic productivity shocks are about 1.5 to 2 times higher for the larger super-elasticity values. Moreover, for small (large) steady-state elasticities the calibrated values for the menu costs tend to be small (large) whereas the standard deviation measures tend to be high (low).³⁵

The sizes of menu costs are largely plausible. For the CES cases, menu costs correspond to 0.40% (of total income) for the 10th percentile of the estimated elasticities, to 2.53% for the median elasticity and reach a maximum of 9.31% for the 90th percentile of the estimated elasticities. When the frequency of price changes is taken into account, this

³³Further results employing additional values of the estimated distributions including the weighted means are reported in Table D.1 of Appendix D.3.

³⁴We also conducted our simulation exercise using price statistics excluding sales. The results are qualitatively and quantitatively similar and are reported in the Appendix of the working paper version.

³⁵The intuition underlying these findings is as follows: The smaller the elasticity the flatter the demand and as a result the profit function becomes. This can plainly be seen in Figure 1 for these two curves associated with for our most extreme considered case (1.41). As a consequence, deviations from the optimal relative price associated with a moderate cost shocks do not lead to large reductions in profits and thus the firm will not adjust its price even for relatively small adjustment costs. To match the observed median price changes in the data, the model therefore requires large idiosyncratic shocks.

amounts to monthly adjustment costs of 0.06%, 0.40% and 1.46%, respectively.³⁶ Apart from the menu cost value obtained for the large-elasticity cases, the obtained figures are broadly in line with those provided by independent evidence on menu costs as reported, for example, by Levy et al. (1997). These authors find the costs of changing prices are approximately 0.7% of revenue. For positive values of the super-elasticity, the menu costs take values of approximately 1% for the majority of the considered cases and are thus likewise broadly consistent with the empirical evidence. Values of well above 1% are only obtained for combinations that include the largest elasticities.

Considering the cases from the macro literature, we can basically identify three groups in terms of the compatibility of obtained menu costs with empirical evidence. For (Bergin and Feenstra, 2000, BF), (Gopinath and Itskhoki, 2010, GI), and (Klenow and Willis, 2016, KW), we reasonable moderate menu costs (0.60%-1.47%, all taking into account the frequency of price changes), for (Woodford, 2003, W) and Eichenbaum and Fisher (2005) values are more than 3 times larger than the Levy et al. (1997) estimates, and for (Chari et al., 2000, CKM) and (Kimball, 1995, K) they are certainly unrealistic (with values of around 48%, they are more than 60 times larger than the Levy et al. (1997) estimates).

In regard to the standard deviations of idiosyncratic productivity shocks, we obtain numbers for the CES case of around 7.5%.³⁷ Interestingly, the figures do not generally increase dramatically for positive values of the super-elasticity. Apart from the lowest elasticity case and very few exceptions for all other cases, their sizes remain in a fairly narrow range around 10% and are thus well below the value of 28%, which is labeled as clearly unrealistic by Klenow and Willis (2016). The low-elasticity cases require large

³⁶The differences in obtained menu costs between the low- and high-elasticity cases reflect the fact that the profit function of a firm facing a very low elasticity of substitution is very flat implying that gains from price changes are small such that already very low menu costs prevent a firm from adjusting its price after a shock. An opposite argument applies to a firm facing a high elasticity of demand.

³⁷Gopinath et al. (2015) estimate a value of 13% for Southern Europe and Bachmann and Bayer (2014) of 9% for Germany, both in annual data. This result is likely due to the fairly high frequency and size of the price changes in our data.

idiosyncratic shocks to match the rather high frequency and size of price changes, since firms' incentive to adjust prices is otherwise very low.

Looking at the cases in the literature we find that the values are of similar plausible size for BF and to some extent for GI and W, are sizeably larger for KW and EF and they are clearly implausible for CKM and K.

5.3 Simulation results

Employing the calibrated parameter values, we next investigate the degree of monetary non-neutrality associated with the obtained elasticity and super-elasticity values. Following Nakamura and Steinsson (2010), we measure monetary non-neutrality using three 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 is to a nominal shock, the higher the value for CIR will be. Second, we compute the variance of aggregate consumption spending that results from model simulations where only aggregate nominal shocks hit the economy and compare it to the variance of detrended real consumption observed in our sample countries (*Var. of C explained*).³⁸ Third, we provide a measure of the persistence of the simulated real consumption series (*Pers. of C*). This number is obtained by fitting an AR(13) process to the artificially generated consumption series and summing its autoregressive coefficients.

Considering the simulation results for our measures of monetary non-neutrality we see that the obtained values are generally fairly moderate, independently of the preference specification. Certainly, one reason for this finding is that we match production-side

³⁸Unlike 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.

Table 5: Outcome comparisons for alternative elasticity-super-elasticity specifications

(El., super-el.)	Menu costs	Std.dev. id. shocks	CIR	Var. C expl.	Pers. of C
Panel A: Estimated values					
(3.20 0.00)	2.53	7.62	0.004	0.99	0.25
(3.20 1.93)	4.74	14.48	0.012	4.26	0.49
(3.20 2.23)	4.2	15.51	0.008	4.62	0.52
(3.20 0.05)	2.65	7.79	0.007	3.01	0.46
(1.41 0.00)	0.43	7.62	0.004	1.1	0.27
(1.41 1.93)	2.05	38.53	0.022	11.31	0.69
(1.41 2.23)*	2.32	41.45	0.053	14.09	0.72
(1.41 0.05)*	0.54	8.46	0.025	10.08	0.66
(11.21 0.00)*	9.31	7.52	0.001	0.13	0.09
(11.21 1.93)	15.14	8.97	0.004	3.55	0.47
(11.21 2.23)	15.3	9.36	0.008	3.91	0.5
(11.21 0.05)	14.45	7.68	0.008	3.94	0.47
Panel B: Specifications from the literature					
BF (3.00 1.33)	3.84	12.89	0.007	4.57	0.5
CMK (10.00 385)*	311.72	946.88	0.008	2.91	0.43
EF (11.00 10.00)	22.08	15.86	0.019	10.47	0.69
GI (5.00 4.00)	9.4	15.3	0.007	3.8	0.49
K (11.00 471)*	311.89	941.93	0.008	2.9	0.4
KW (5.00 10.00)	7.92	18.67	0.012	4.74	0.51
W (7.76 6.67)	15.5	15.65	0.012	5.4	0.55

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 in Panel A are taken from the estimation results as reported in Table 3. The values in Panel B are those used in (Bergin and Feenstra, 2000, BF), (Chari et al., 2000, CKM), (Eichenbaum and Fisher, 2005, EF), (Gopinath and Itskhoki, 2010, GI), (Kimball, 1995, K), (Klenow and Willis, 2016, KW) and (Woodford, 2003, W). 3) Menu costs 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. Generally, a difference of less than 0.00045 between model and empirical statistic is used as a matching criterion. Cases marked with an asterisk are characterized by very long-lasting convergence processes. In these cases, we therefore employed a convergence value of 0.0025 and a somewhat coarser grid. We are currently running programs applying finer convergence criteria and grids. 4) CIR (cumulative impulse response), Var. of C (variance of aggregate real consumption) and Pers. of C (persistence of consumption) represent the measures of monetary non-neutrality as described in the main text. 5) The variance of real consumption is obtained from the HP filtered quarterly real consumption series (summed over the three countries included in the data sample) from 1995 to 2008. Data source: Eurostat.

parameters such that the firms' frequency and size of price changes correspond to that observed in the data. Given that our firms adjust their prices fairly frequently and given that average price changes are fairly sizeable, we can expect firms to respond quite quickly and comprehensively to monetary shocks which will generally be dominated by idiosyncratic shocks which need to be large to match the observed price facts. The reported figures for the case of constant elasticities clearly illustrate that nominal rigidity alone can generate only very small degrees of monetary non-neutrality, confirming several results including those of Golosov and Lucas (2007) and Nakamura and Steinsson (2010), for example. The proportion of the variance in real consumption explained ranges from 0.13% to 1.1%. With higher super-elasticities, this share increases to 4.62% for the median (14.1%, 3.91% for the low and high elasticity cases, respectively). Our measures of monetary non-neutrality increase by a factor of 2 to 8. Persistence on the other hand tends to increase from values of around 0.25 in the CES case to numbers around 0.5 and 0.6 in the non-CES cases. As noted in the previous subsection, this increase in monetary non-neutrality is associated with higher menu costs and variances of the idiosyncratic productivity shocks, which in most cases, are largely compatible with the empirical evidence.

The responses for the model with demand-side real rigidity are larger on impact and more pronounced. This result is illustrated in Figure 3, which shows the cumulative impulse responses of real consumption to a nominal demand shock for the case of the median elasticity estimate (3.20%) without (in red) and with real real rigidity (in blue). The explicit consideration of the empirically documented real rigidity not only increases the effect of the 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 increases from only 2 months for the CES case to 6 months for the case with real rigidity, which is still much smaller than what one normally finds in the data.

Comparing the obtained values for the measures of monetary non-neutrality implied by

our estimates with those employed in the literature, we observe that - not surprisingly - those specifications with large and extreme preference parameters are associated with larger degrees of monetary non-neutrality. However, the differences are not as sizeable as one might have expected given the heterogeneity in preference parameter values. Given that these increases come at the cost of having to assume highly implausible values for the menu costs and the variance of idiosyncratic shocks, our findings do not support the use of most of these specifications in macro models.

Figure 3: *Response to a monetary policy shock*




Figure 3 plots the cumulative impulse response of real consumption expenditures to a (positive) shock to nominal aggregate demand both for the CES (left panel) and NON-CES (right panel) demand functions. The elasticity value corresponds to the estimated median value (3.20). The value for the super-elasticity is 1.93 (median estimate).

6 Conclusions

The findings of [Bils and Klenow \(2004\)](#) that micro prices are changed relatively often and that both large price increases and decreases frequently occur have challenged the previously dominating view among 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 included introducing real rigidities

resulting from non-constant 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 employ a new dataset on consumer retail transactions that 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 that values for the price elasticity parameters range between 3 and 5. 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 (see Nevo, 2000, for example). Similarly, we find that the super-elasticity parameters are much lower than the values used in macro models, with values in the range of 1 to 2. Together with the demand elasticity estimates, these imply a markup elasticity of approximately $0.6 - 1$.

To quantitatively assess the importance of demand-side real rigidity, we augment a model with only nominal rigidity and augment it with 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 approximately three to five because the model including demand-side real rigidity shows a degree of monetary non-neutrality that is three to five times larger than the pure nominal rigidity model. These calibrations imply plausible values for menu costs and idiosyncratic shock variances and can still match the observed frequency and size of price changes. This is particularly true when a high persistence in idiosyncratic shock processes is assumed. However, these demand-side real rigidities do not generate a lot of aggregate non-neutrality. Our results suggest that only approximately

5% 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 result 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

- 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,” *Journal of the European Economic Association*, 4, 575–584.
- AMITI, M., O. ITSKHOKI, AND J. KONINGS (2016): “International Shocks and Domestic Prices: How Large Are Strategic Complementarities?” Working Paper 22119, National Bureau of Economic Research.
- ANDERSON, S., A. DE PALMA, AND J. THISSE (1987): “The CES is a discrete choice model?” *Economics Letters*, 24, 139–140.
- ANDERSON, S. P., A. DE PALMA, AND J.-F. THISSE (1992): *Choice Theory of Product Differentiation*, MIT Press, Cambridge, MA.
- BACHMANN, R. AND C. BAYER (2014): “Investment Dispersion and the Business Cycle,” *American Economic Review*, 104, 1392–1416.
- 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.
- CHARI, V. V., P. J. KEHOE, AND E. R. MCGRATTAN (2000): “Sticky Price Models of the Business Cycle,” *Econometrica*, 68, 1151–1180.

- 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.
- DEATON, A. AND J. MUELLBAUER (1980): “An Almost Ideal Demand System,” *American Economic Review*, 70, 312–326.
- 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. VAN DEN POEL (2010): “The Kinked Demand Curve and Price Rigidity: Evidence from Scanner Data,” *Scandinavian Journal of Economics*, 112, 723–752.
- 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.
- GAUTIER, E. AND H. LE BIHAN (2018): “Shocks vs Menu Costs: Patterns of Price

Rigidity in an Estimated Multi-Sector Menu-Cost Model,” Working papers 682, Banque de France.

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.

GOPINATH, G., S. KALEMLI-OZCAN, L. KARABARBOUNIS, AND C. VILLEGAS-SANCHEZ (2015): “Capital Allocation and Productivity in South Europe,” Working Paper 21453, National Bureau of Economic Research.

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.

KAPLAN, G. AND S. SCHULHOFER-WOHL (2017): “Inflation at the household level,” *Journal of Monetary Economics*, 91, 19–38.

- 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 (2016): “Real rigidities and nominal price changes,” *Economica*, 83, 443–472.
- KLEPACZ, M. (2017): “Price setting and volatility: evidence from oil price volatility shocks,” Manuscript.
- KRUSELL, P. AND A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867–896.
- KRYVTSOV, O. AND V. MIDRIGAN (2013): “Inventories, markups, and real rigidities in menu cost models,” *The Review of Economic Studies*, 1, 249–276.
- 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.
- LINDÉ, J. AND M. TRABANDT (2019): “Resolving the Missing Deflation Puzzle,” CEPR Discussion Papers 13690, C.E.P.R. Discussion Papers.
- MCFADDEN, D. AND K. TRAIN (2000): “Mixed MNL Models for Discrete Response,” *Journal of Applied Econometrics*, 15, 447–470.
- MIDRIGAN, V. (2011): “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, 79, 1139–1180.
- MRÁZOVÁ, M. AND J. P. NEARY (2017): “Not So Demanding: Demand Structure and Firm Behavior,” *American Economic Review*, 107, 3835–74.

- NAKAMURA, E. AND J. STEINSSON (2010): “Monetary Non-Neutrality in a Multisector Menu Cost Model,” *The Quarterly Journal of Economics*, 125, 961–1013.
- NAKAMURA, E., J. STEINSSON, P. SUN, AND D. VILLAR (2017): “The elusive costs of inflation: Price dispersion during the U.S. Great Inflation,” Manuscript. Columbia University.
- NAKAMURA, E. AND D. ZEROM (2010): “Accounting for Incomplete Pass-Through,” *The Review of Economic Studies*, 77, 1192–1230.
- NEVO, A. (2000): “A Practitioner’s Guide to Estimation of Random-Coefficients Logit Models of Demand,” *Journal of Economics & Management Strategy*, 9, 513–548.
- (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.
- RICE, K., J. P. T. HIGGINS, AND T. LUMLEY (2018): “A re-evaluation of fixed effect(s) meta-analysis,” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 181, 205–227.
- SMETS, F. AND R. WOUTERS (2003): “An estimated stochastic dynamic general equilibrium model of the euro area,” *Journal of the European Economic Association*, 1, 1123–1175.
- STROEBEL, J. AND J. VAVRA (2018): “House Prices, Local Demand, and Retail Prices,” *Journal of Political Economy*, forthcoming.

- THISSE, J. AND P. USHCHEV (2016): “When Can A Demand System Be Described By A Multinomial Logit With Income Effect?” HSE Working papers WP BRP 139/EC/2016, National Research University Higher School of Economics.
- TRAIN, K. (2009): *Discrete Choice Methods with Simulation*, Cambridge University Press, 2 ed.
- VAVRA, J. (2013): “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation,” *The Quarterly Journal of Economics*, 129, 215–258.
- VERHELST, B. AND D. VAN DEN 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.

Supplementary Material to “Price elasticities and demand-side real rigidities in micro data and in macro models”

A	Data preparation	A2
B	Estimation methodology	A5
B.1	Multinomial logit	A5
B.2	Control function approach	A6
B.3	Derivation of elasticities and super-elasticities	A7
C	Empirical results: additional findings and robustness check	A9
C.1	Relationship between elasticities and super-elasticities	A9
C.2	Robustness of the empirical estimates	A9
D	A menu-cost model with Kimball-type preferences	A14
D.1	Model setup	A14
D.1.1	Households	A14
D.1.2	Firms	A18
D.1.3	Monetary policy	A19
D.2	Model solution	A20
D.3	Model outcomes - detailed results	A23
E	Additional tables	A24
F	Additional figures	A25

A Data preparation

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

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 for which we did not observe any purchases for 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-gram pot of yogurt, which costs 99 cents, the price used in the estimation is 0.198.
4. We removed outliers, defined as an observations where the price p of a product is more than 200% higher than the average price of the identical product (GTIN).
5. We identify price changes by comparing prices of the same GTIN at different dates for the same retailer for two consecutive months. We calculate the size of the price changes and the frequency of the price changes per category and months. We

- calculate the mean and median price adjustment size and frequency per category.
6. The 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 the 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. We estimate the lower nest for each category. We 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. We construct observations for the 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 that 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 for which the indicator variable is zero by following a similar procedure as above for the construction of the prices of alternatives not chosen. In particular, we construct observations for price and the control-function residual by using the price and residual for the same brand, for the same week, for the same retailer, and for the same region if observed. If that was not observed, we apply the same criteria but use the price and residual from the same month instead of the same week. If both are not observed, we drop the observation in the estimation.
 - Drop category if Ψ is estimated to be lower than 0.001.

B Estimation methodology

This section gives a brief description of the discrete choice model underlying the elasticity estimates and the estimation using the control function approach.

B.1 Multinomial logit

We first show that the scale of utility is irrelevant for discrete choice estimation and that the assumption of an iid extreme value distributed error term results in a multinomial model with probabilities that take the form of logits. We then describe the choice probabilities in more detail.

To show how the assumption on an iid extreme value distributed error term is consistent with a logit probability, consider in a generic model the probability of choosing brand j over brand m in the same product category with J choice alternatives. This is modelled as the probability that the utility derived from product j is higher than the utility derived from product $m \neq j$, $Prob(U_j > U_m)$. Denote the observed part of utility by V and the unobserved error term by ε . Then, $Prob(V_j + \varepsilon_j > U_m + \varepsilon_m) = Prob(\varepsilon_m - \varepsilon_j < V_j - V_m)$. Let the joint density of the random vector $\varepsilon = (\varepsilon_1, \dots, \varepsilon_J)$, including ε_j and ε_m , be denoted by $f(\varepsilon)$: then, the probability of choosing brand j can be written as $P_j = \int_{\varepsilon} I(\varepsilon_m - \varepsilon_j < V_j - V_m) f(\varepsilon) d\varepsilon$, where I is an indicator function that is one if the expression in parentheses is true and zero otherwise. Since only differences in utility matter, the dimension of this integral can be reduced to $J - 1$ dimensions by writing the integral over error differences $\varepsilon_m - \varepsilon_j \equiv \tilde{\varepsilon}_{jm}$ as $P_j = \int_{\tilde{\varepsilon}_j} I(\tilde{\varepsilon}_{jm} < V_j - V_m) g(\tilde{\varepsilon}_j) d\tilde{\varepsilon}_j$, where $\tilde{\varepsilon}_j$ is an $J - 1$ dimensional vector over the error differences between products j and all other products $\neq j$. Furthermore, since the differences of extreme-value distributed variables have the form of logits, $g(\cdot)$ is a logistic distribution.³⁹

³⁹See Train (2009) and references therein.

Next, we show the derivation of the conditional and marginal probabilities from the upper and lower nest. The probability of household i choosing brand j within category $c(j)$ is the product of the conditional probability of choosing brand j within the category $c(j)$, $P_{ij|c(j)}$, conditional on an item in category $c(j)$ being chosen and the marginal probability of choosing an item in category $c(j)$, $P_{ic(j)}$. It is convenient to decompose the probabilities because the marginal and conditional probabilities take the form of logits,

$$\begin{aligned} P_{ic(j)} &= \frac{e^{W_{ic(j)} + \lambda_{c(j)} IncV_{ic(j)}}}{\sum_{\ell=1}^C e^{W_{i\ell} + \lambda_{\ell} IncV_{i\ell}}} \\ P_{ij|c(j)} &= \frac{e^{Y_{ij}/\lambda_{c(j)}}}{\sum_{k=1}^J e^{Y_{ik}/\lambda_{c(j)}}}, \end{aligned}$$

where $IncV_{ic} = \ln \sum_J e^{Y_{ij}/\lambda_{c(j)}}$ is the inclusive value, which is the expected utility household i receives from the J choice alternatives (brands) within category $c(j) \in C$. $\lambda_{c(j)}$ is the log-sum coefficient, and Y_{ij} represents the explanatory variables, which vary across alternatives within category $c(j)$ and include household characteristics. $W_{ic(j)}$ is a vector of variables that describe category $c(j)$ and does not vary across j .

B.2 Control function approach

This section of the Appendix describes the control function approach in more detail. Consider that regression $p_{jt} = \theta'_j Z_{jt} + \mu_{jt}$, where Z_{jt} is the mean price of the same brand in other NUTS1 regions within the same country (Hausman, 1996; Nevo, 2001). The control function is a variable that captures the conditional mean of the correlation between the observed variables and the error term. Suppose that 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

$$U_{ijt|y_{it}=1} = \beta_{ij} - \alpha_{ij}p_{jt} + \gamma_i w_{ijt} + \delta x_{ijt} + \lambda \mu_{jt} + \tilde{\varepsilon}_{ijt}^1 + \varepsilon_{ijt}^2. \quad (8)$$

Thus, we include, in addition to the variables discussed in the main text, the residual from the control function as an additional explanatory variable.

B.3 Derivation of elasticities and super-elasticities

The elasticity for the upper nest is derived from the probability of choosing category c in the upper nest,

$$\begin{aligned} P_{ic} &= \frac{e^{W_{ic} + \Psi IncV_{ic}}}{\sum_C e^{W_{ic} + \Psi IncV_{ic}}} \\ \frac{\partial IncV_{ic}}{\partial p_j} &= \frac{\partial (\ln \sum_J e^{Y_{ij}/\Psi})}{\partial p_j} \\ &= \alpha_i P_{ij|c(j)} / \Psi \end{aligned}$$

which we then use to calculate the upper nest elasticity $\theta_{ijc(j)}^u$

$$\begin{aligned} \frac{\partial P_{ic}}{\partial p_{ijt}} \frac{p_{ijt}}{P_{ic}} &= \frac{e^{\rho w_{ict} + \Psi IncV_{ict}}}{\sum_C e^{\rho w_{ict} + \Psi IncV_{ict}}} \frac{p_{ijt}}{P_{ic}} \\ &= \frac{e^{W_{ic} + \Psi IncV_{ic}} \alpha_i P_{ij|c(j)} \sum_C e^{\rho w_{ict} + \Psi IncV_{ict}} - e^{W_{ic} + \Psi IncV_{ic}} e^{W_{ic} + \Psi IncV_{ic}} \alpha_i P_{ij|c(j)}}{(\sum_C e^{\rho w_{ict} + \Psi IncV_{ict}})^2} \frac{p_{ijt}}{P_{ic}} \\ &= (\alpha_i P_{ij|c(j)} P_{ic} - P_{ic} P_{ic} \alpha_i P_{ij|c(j)}) \frac{p_{ijt}}{P_{ic}} \\ &= \alpha_i P_{ij|c(j)} (1 - P_{ic}) p_{ijt} \end{aligned}$$

Similarly, the elasticity for the lower nest is derived as

$$\begin{aligned}
P_{ij|c(j)} &= \frac{e^{(\beta_{ij} + \alpha_i p_{ijt} + \gamma_i w_{ijt} + \delta x_{it})/\Psi}}{\sum_J e^{(\beta_{ij} + \alpha_i p_{jt} + \gamma_i w_{ijt} + \delta x_{it})/\Psi}} \\
\frac{\partial P_{ij|c(j)}}{\partial p_{ijt}} \frac{p_{ijt}}{P_{ij|c(j)}} &= \frac{e^{(\beta_{ij} - \alpha_i p_{jt} + \gamma_i w_{ijt} + \delta x_{it})/\Psi} * (\alpha_i/\Psi) * (\sum_J e^{(\beta_{ij} + \alpha_i p_{jt} + \gamma_i w_{ijt} + \delta x_{it})/\Psi})}{\sum_J (e^{(\beta_{ij} + \alpha_i p_{jt} + \gamma_i w_{ijt} + \delta x_{it})/\Psi})^2} \frac{p_{ijt}}{P_{ij|c(j)}} \\
&\quad - \frac{e^{(\beta_{ij} + \alpha_i p_{jt} + \gamma_i w_{ijt} + \delta x_{it})/\Psi} * e^{(\beta_{ij} + \alpha_i p_{jt} + \gamma_i w_{ijt} + \delta x_{it})/\Psi} \alpha_i/\Psi}{\sum_J (e^{(\beta_{ij} + \alpha_i p_{jt} + \gamma_i w_{ijt} + \delta x_{it})/\Psi})^2} \frac{p_{ijt}}{P_{ij|c(j)}} \\
&= P_{ij|c(j)} (\alpha_i/\Psi) (1 - P_{ij|c(j)}) \frac{p_{ijt}}{P_{ij|c(j)}} \\
&= \frac{\alpha_i}{\Psi} p_{ijt} (1 - P_{ij|c(j)}).
\end{aligned}$$

The super elasticity for the upper nest is given by

$$\begin{aligned}
\frac{\partial \theta_{ijc(j)}^u}{\partial p_{ijt}} \frac{p_{ijt}}{\theta_{ijc(j)}^u} &= \alpha_i [P_{ij|c(j)} (1 - P_{ic}) + p_{ijt} \frac{\partial P_{ij|c(j)}}{\partial p_{ijt}} (1 - P_{ic}) - p_{ijt} P_{ij|c(j)} \frac{\partial P_{ic}}{\partial p_{ijt}}] \frac{p_{ijt}}{\theta_{ijc(j)}^u} \\
&= 1 + \Psi \alpha_i P_{ij|c(j)} (1 - P_{ic}) p_{ijt} \frac{\theta_{ijc(j)}^l}{\theta_{ijc(j)}^u} - \Psi \alpha_i P_{ij|c(j)} P_{ic} p_{ijt} \\
&= 1 + \theta_{ijc(j)}^l - \Psi \alpha_i P_{ij|c(j)} P_{ic} p_{ijt},
\end{aligned}$$

and that for the lower nest is given by

$$\begin{aligned}
\frac{\partial \theta_{ijc(j)}^l}{\partial p_{ijt}} \frac{p_{ijt}}{\theta_{ijc(j)}^l} &= \frac{\alpha_i}{\Psi} [(1 - P_{ij|c(j)}) - p_{ijt} \frac{\partial P_{ij|c(j)}}{\partial p_{ijt}}] \frac{p_{ijt}}{\theta_{ijc(j)}^l} \\
&= \frac{\alpha_i}{\Psi} [(1 - P_{ij|c(j)}) \frac{p_{ijt}}{\theta_{ijc(j)}^u} - p_{ijt} P_{ij|c(j)}] \\
&= 1 - \frac{\alpha_i}{\Psi} p_{ijt} P_{ij|c(j)=1}.
\end{aligned}$$

C Empirical results: additional findings and robustness check

C.1 Relationship between elasticities and super-elasticities

In section 5 below, we show that demand elasticity and super-elasticity have a negative correlation, conditional on their steady state values (see equations 25 and 26). We show that this also holds in our estimates. To illustrate the correlation, we regress the estimated super-elasticity, $\epsilon_{ijc(j)}$, on the estimated demand elasticity $\theta_{ijc(j)}$. The relation is negative, but not very large. An elasticity of 3.20 (5.40) is, according to the correlation, associated with a super-elasticity of 0.95 (0.65).⁴⁰

Table C.1: *Relation between elasticities and super-elasticities*

	(1) OLS	(2) MLS	(3) WLS Expend.	(4) WLS Variance
Elasticity	-0.103*** [0.001]	-0.124*** [0.001]	-0.093*** [0.001]	-0.087*** [0.001]
Constant	1.322*** [0.002]	1.346*** [0.003]	1.150*** [0.004]	1.339*** [0.002]
R^2	0.07	0.04	0.07	0.05

Notes: this table reports the results of regression $\epsilon_{ijc(j)} = a + b\theta_{ijc(j)} + e_{ijc(j)}$, which illustrates the correlation between demand elasticities and super-elasticities. The first column shows OLS estimates, the second Median Least Squares (MLS), the third weighted least squares (WLS), where the weights are expenditure shares and the fourth column reports WLS estimates, where the weights are based on the inverse estimated variance of α_i as a weight. Robust standard errors are in brackets, *** p<0.01, ** p<0.05, * p<0.1.

C.2 Robustness of the empirical estimates

To provide a robustness analysis of our estimates, we use an alternative methodology that does not rely on distributional assumptions that are similar to the econometric model

⁴⁰We use the MLS for calculating the predicted super-elasticity for the median value and accordingly the WLS for predicting the super-elasticity for the weighted mean.

applied in the main part of this paper. To do so, we use an extension of the almost ideal demand system (Deaton and Muellbauer, 1980), derived in Dossche et al. (2010), to estimate parameters describing the elasticity and super-elasticity of demand. Using this model, we show that the main conclusion from this paper, that demand elasticities and super-elasticities are relatively low, is confirmed.

The extended almost ideal demand system model includes an additional term, the squared relative price. This additional term allows for greater flexibility for the elasticity and super-elasticity estimates, as shown by Dossche et al. (2010). The basic empirical model we employ is given by

$$s_{m,t,i} = \alpha_i + \sum_{h=1}^5 \gamma_{m,i,h} \ln p_{m,t,h} + \beta_i \ln \left(\frac{Exp_{m,t}}{P_{m,t}} \right) + \sum_{h=1}^5 \delta_{i,h} \left(\ln \left(\frac{p_{m,t,h}}{P_{m,t}} \right) \right)^2 \quad (9)$$

where the subindex m denotes goods categories, t indicates the time period (month) and i denotes the individual brands within the category. $P_{m,t}$ denotes the price index of goods category m , and $p_{m,t,i}$ represents the price of one unit (e.g., gram, ml, and pieces) of brand i within category m . $Exp_{m,t}$ represents the total expenditures in category m at date t , defined as $Exp_{m,t} \equiv \sum_i p_{m,t,i} q_{m,t,i}$ with $p_{m,t,i} q_{m,t,i}$ being the total nominal expenditures on brand i . The left-hand side variable $s_{m,t,i} = (p_{m,t,i} q_{m,t,i}) / X_{m,t}$ represents the share of expenditures on the single item i . When $\delta_{j,h} = 0$, equation (9) nests the standard AIDS model proposed by Deaton and Muellbauer (1980). The price index for category m , $P_{m,t}$ is defined as the Stone Price Index

$$\ln P_{m,t} = \sum_i s_{m,t,i} \ln p_{m,t,i}, \quad (10)$$

i.e., the weighted average price of goods belonging to category m .

Once one has obtained the estimates of the parameters in (9), one can obtain the elasticities and super-elasticities of demand. The steady-state elasticity, where relative

prices equal unity, is given by

$$\bar{\theta}_j = 1 - \frac{\gamma_{jj}}{s_j} + \beta_j \quad (11)$$

and the steady-state super-elasticity is given by

$$\bar{\epsilon}_j = \frac{1}{\bar{\theta}_j} \left((\bar{\theta}_j - 1)(\bar{\theta}_j - 1 - \beta_j) - \frac{2\delta_{jj}(1 - s_j)}{s_j} + 2(\delta_{jj} - s_j \sum_{i=1}^N \delta_{ji}) \right). \quad (12)$$

As for the discrete choice model in the body of the paper, we concentrate on the top four brands within each category. The first four of these correspond to the brands with the largest total expenditure shares within the given category m over all time periods t , whereas the fifth brand represents all other goods and is denoted by *other* (the outside good).

Following Dossche et al. (2010), we then use a Seemingly Unrelated Regression (SUR) to estimate the following regression for each product category m :

$$\begin{aligned} s_{m,t,i} = & \alpha_i + \sum_{h=1}^5 \gamma_{i,h} \ln p_{m,t,h} + \beta_i \ln \left(\frac{X_{m,t}}{P_{m,t}} \right) + \sum_{h=1}^5 \delta_{i,h} \left(\ln \left(\frac{p_{m,t,h}}{P_{m,t}} \right) \right)^2 + \\ & \sum_{t=1}^T \eta_{i,t} D_t + \sum_{m=1}^M \kappa_{i,m} D_m + \theta_i Z_{m,t,i} + \varepsilon_{i,m,t}, \end{aligned} \quad (13)$$

where D_t and D_m are time and market fixed effects, respectively, and Z represents the average household income and age to control for potential effects of demographic variables on the elasticity estimates. We impose the same homogeneity and symmetry conditions as in Dossche et al. (2010); that is $\sum_{h=1}^5 \gamma_{j,h} = 0$, $\gamma_{j,h} = \gamma_{h,j}$. Because of these conditions and because the expenditure shares within a category sum up to one by construction, we can drop one equation from the system for estimation. We drop the equation for the outside good *other*.


The estimates in this robustness exercise largely confirm the conclusions drawn from

Table C.2: *Descriptive statistics estimates robustness*

Panel A: Estimates of Demand Elasticities							
	Mean	Median	Weighted Mean		Percentile		Sign.
			Expend.	Variance	90 th	10 th	5%
Total elasticity	1.37	1.20	1.44	1.10	4.46	1.01	60.07
Panel B: Estimates of Super-elasticities							
	Mean	Median	Weighted Mean		Percentile		Sign.
			Expend.	Variance	90 th	10 th	
Total super-elasticity	0.81	0.31	1.36	0.30	2.38	-0.61	56.08
Panel C: Estimates of Markup Elasticities							
Implied markup elasticity	12.38	1.04	5.20	3.56	4.46	-10.17	

Notes: Panel A shows the summary statistics of the estimated demand elasticities θ , and Panel B shows those for the estimates of the super-elasticities ϵ . Panel C shows the implied markup elasticity $\frac{\partial \mu}{\partial \ln P}|_{P=1} = \frac{\epsilon}{\theta-1}$. The first two columns show the mean and median values of the estimated elasticities. The third and fourth columns show the weighted means, weighted by expenditure share and by the inverse of the estimated variance of α_i . The fifth and sixth columns show the 90th and 10th percentiles of the distribution of estimates, respectively. The last column shows the share of estimates that are significant (in Panel A, the share of γ 's that are significant at the 10% level and in Panel B, the share of δ 's that are significant at the 10% level).

Figure C.1: *Distribution of elasticities and super-elasticities in the demand system estimates*



Notes: Panel a) shows the distribution of the estimated elasticities using the almost ideal demand system-type model described above, θ , while panel c) shows the estimates' super-elasticities, ϵ . All estimates vary over brand-product category combinations. The distributions are shown for the trimmed sample.

the baseline results in the body of the paper. The results for this robustness exercise are reported in Table C.2.⁴¹ The median (weighted mean) of demand elasticity estimated at 1.2 (1.44), and the 90th (10th) quantile of the distribution is 4.46 (1.01). The estimates are thus somewhat smaller than the ones reported in the discrete choice model above. They nevertheless confirm our main conclusion that the estimated elasticities are rather low compared to what is often assumed in the literature. This effect is even more severe in the estimates reported here. The estimated super-elasticities are 0.31 for the median and 1.36 for the weighted mean. The super-elasticity estimates are similar to those in the discrete choice model: they are low in the median and on average and are distributed more tightly, where 90% of the estimates are below 2.5.

⁴¹Because the demand system is defined across each product category and over time, we cannot report the upper and lower nests for the nested logit model in the body of the paper.

D A menu-cost model with Kimball-type preferences

In this section, we describe in more detail the theoretical model that we use for calibration 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).

D.1 Model setup

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 as a special case) rather than assuming a roundabout production setup.

D.1.1 Households

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

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{1}{1-\gamma} C_t^{1-\gamma} - \frac{\omega}{1+\psi} L_t^{1+\psi} \right], \quad (14)$$

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 (18) below.

Households choose composite consumption and labor to maximize (14) subject to the following budget constraint

$$P_t C_t + D_{t,t+1} B_{t+1} \leq B_t + W_t L_t + \int_0^1 \Pi_t(z) dz \quad t = 0, 1, \dots \quad (15)$$

This equation requires that aggregate consumption expenditures C_t and investment in financial assets $D_{t,t+1} B_{t+1}$ cannot be higher than the 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 asset that pays off one unit in period $t + 1$.

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

$$D_{t,T} = \beta^{T-t} E_t \left[\left(\frac{C_T}{C_t} \right)^{-\gamma} \frac{P_t}{P_T} \right] \quad (16)$$

$$\frac{W_t}{P_t} = \omega L_t^\psi C_t^\gamma. \quad (17)$$

Equation (16) represents the standard intertemporal Euler equation linking consumption growth to the real interest rate while equation (17) 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 (2016), we implicitly define the composite

consumption good, C_t , using an aggregator of the form

$$\int_0^1 \Upsilon \left(\frac{c(z)}{C} \right) dz = 1, \quad (18)$$

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

$$\Upsilon \left(\frac{c(z)}{C} \right) = \left(\frac{c(z)}{C} \right)^{\frac{\bar{\theta}-1}{\bar{\theta}}} \quad (19)$$

which amounts to assuming 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 (2016) 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

$$\Upsilon(x) = 1 + (\bar{\theta} - 1) \exp \left(\frac{1}{\bar{\epsilon}} \right)^{\bar{\epsilon}(\frac{\bar{\theta}}{\bar{\epsilon}} - 1)} \left[\Gamma \left(\frac{\bar{\theta}}{\bar{\epsilon}}, \frac{1}{\bar{\epsilon}} \right) - \Gamma \left(\frac{\bar{\theta}}{\bar{\epsilon}}, \frac{x^{\frac{\bar{\epsilon}}{\bar{\theta}}}}{\bar{\epsilon}} \right) \right], \quad (20)$$

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

Given an optimal decision about overall consumption expenditure, C , households 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

$$c(z) = \left(\frac{p(z)}{P} \right)^{-\bar{\theta}} C, \quad (21)$$

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

$$c(z) = \left[1 - \bar{\epsilon} \ln \left(\frac{p(z)}{P} \right) \right]^{\frac{\bar{\theta}}{\bar{\epsilon}}} C, \quad (22)$$

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, respectively. When $\bar{\epsilon} = 0$, this demand function reduces to the CES case.⁴²

The major difference between the setups of Dixit and Stiglitz (1977) and Klenow and Willis (2016) 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

$$\theta(p) = \frac{\partial \ln c(z)}{\partial \ln p} \quad (23)$$

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

$$\epsilon(p) = \frac{\partial \ln \theta(p)}{\partial \ln p}, \quad (24)$$

where we allow both the elasticity and the super-elasticity of demand to be non-constant and depend on the relative price of a good.⁴³

⁴²In deriving equation (22), we made use of an approximation result obtained by Gopinath and Itskhoki (2010, Appendix).

⁴³Dossche 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).

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

$$\theta^{NON-CES} = -\frac{\bar{\theta}}{1 - \bar{\epsilon} \ln \left(\frac{p(z)}{P} \right)}, \quad (25)$$

whereas the super-elasticity of demand is given by

$$\epsilon^{NON-CES} = \frac{\bar{\epsilon}}{1 - \bar{\epsilon} \ln \left(\frac{p(z)}{P} \right)}. \quad (26)$$

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

$$(\theta^*)^{NON-CES} = \bar{\theta} \text{ and } (\epsilon^*)^{NON-CES} = \bar{\epsilon}. \quad (27)$$

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 higher than it is 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.

D.1.2 Firms

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

$$y_t(z) = A_t(z)L_t(z), \quad (28)$$

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

$$\ln A_t(z) = \rho \ln A_{t-1}(z) + \epsilon_t(z), \quad \epsilon_t(z) \sim N(0, \sigma_{\epsilon,t}^2). \quad (29)$$

Firms aim to maximize the discounted value of expected profits,

$$E_0 \sum_{t=0}^{\infty} D_{0,t} \Pi_t(z), \quad (30)$$

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

$$\Pi_t(z) = p_t(z)y_t(z) - W_t L_t(z) - K W_t I_t(z) - P_t U. \quad (31)$$

$y_t(z)$ 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, $K W_t I_t(z)$ 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.

D.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

$$\ln Y_t^N = \mu + \ln Y_{t-1}^N + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2). \quad (32)$$

D.2 Model solution

An equilibrium in this model is a set of policy rules for the endogenous variables that is consistent with the households' 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 with the corresponding demand function in real terms as follows:

$$\begin{aligned} \Pi_t^R(z) = & \left(\frac{p_t(z)}{P_t} \right) \mathcal{F} \left(\frac{p_t(z)}{P_t} \right) \frac{Y_t^N}{P_t} - \omega L_t^\psi C_t^\gamma \left(\frac{1}{A_t(z)} \right) \mathcal{F} \left(\frac{p_t(z)}{P_t} \right) \frac{Y_t^N}{P_t} - \\ & - K \omega L_t^\psi \left(\frac{Y_t^N}{P_t} \right)^\gamma I_t(z) - U, \end{aligned} \quad (33)$$

where $\mathcal{F}(\cdot)$ corresponds to equation (21) in the CES-Case and to equation (22) in the NON-CES-Case.

We solve the firm's optimization problem using dynamic programming.⁴⁴ 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 $\frac{p_{t-1}(z)}{P_t}$.

⁴⁴To 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

$$V\left(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{Y_t^N}{P_t}\right) = \max \left\{ V^{NC}\left(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{Y_t^N}{P_t}\right), V^C\left(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{Y_t^N}{P_t}\right) \right\}, \quad (34)$$

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:

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

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

$$\frac{P_t}{P_{t-1}} = \Gamma\left(\frac{Y_t^N}{P_{t-1}}\right) \quad (35)$$

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 and then 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).

D.3 Model outcomes - detailed results

Table D.1: *Detailed outcomes for alternative elasticity-super-elasticity specifications*

(El., super-el.)	Menu costs	Std.dev. id. shocks	CIR	Var. C expl.	Pers. of C
(4.81 0.00)	4.28	7.61	0.004	0.84	0.25
(4.81 1.59)	6.29	10.84	0.011	4.41	0.51
(4.81 1.93)	6.66	11.53	0.009	4.47	0.49
(4.81 1.44)	6.12	10.53	0.011	4.34	0.5
(4.81 2.23)	6.94	12.13	0.008	4.58	0.52
(4.81 0.05)	4.61	7.7	0.007	3.02	0.42
(3.20 0.00)	2.53	7.62	0.004	0.99	0.25
(3.20 1.59)	4.4	13.22	0.009	4.44	0.52
(3.20 1.93)	4.74	14.48	0.012	4.26	0.49
(3.20 1.44)	4.2	12.56	0.008	4.41	0.51
(3.20 2.23)	4.2	15.51	0.008	4.62	0.52
(3.20 0.05)	2.65	7.79	0.007	3.01	0.46
(5.40 0.00)	4.9	7.53	0.003	0.78	0.23
(5.40 1.59)	6.96	10.4	0.011	4.36	0.48
(5.40 1.93)	7.33	11.01	0.01	4.4	0.5
(5.40 1.44)	6.79	10.12	0.01	4.2	0.51
(5.40 2.23)	7.66	11.63	0.009	4.44	0.5
(5.40 0.05)	5.41	7.71	0.007	3.03	0.45
(1.41 0.00)	0.43	7.62	0.004	1.1	0.27
(1.41 1.59)	1.89	34.25	0.023	8.66	0.63
(1.41 1.93)	2.05	38.53	0.022	11.31	0.69
(1.41 1.44)	1.82	32.26	0.022	7.52	0.62
(1.41 2.23)*	2.32	41.45	0.053	14.09	0.72
(1.41 0.05)*	0.54	8.46	0.025	10.08	0.66
(11.21 0.00)*	9.31	7.52	0.001	0.13	0.09
(11.21 1.59)	14.93	8.71	0.008	3.69	0.49
(11.21 1.93)	15.14	8.97	0.004	3.55	0.47
(11.21 1.44)	14.91	8.81	0.005	3.56	0.48
(11.21 2.23)	15.3	9.36	0.008	3.91	0.5
(11.21 0.05)	14.45	7.68	0.008	3.94	0.47

Notes: 1) Menu costs are given by the percentage share of steady-state revenue $\frac{\theta-1}{\theta} K$. 2) The values for the elasticity and super-elasticity are taken from the estimation results as reported in Table 3. 3) Menu costs 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. Generally, a difference of less than 0.00045 between model and empirical statistic is used as a matching criterion. Cases marked with an asterisk are characterized by very long-lasting convergence processes. In the cases marked with an asterisk we therefore employed a convergence value of 0.0025 and a somewhat coarser grid. We are currently running programs applying finer convergence criteria and grids. 4) CIR (cumulative impulse response), Var. of C (variance of aggregate real consumption) and Pers. of C (persistence of consumption) represent the measures of monetary non-neutrality as described in the main text. 5) The variance of real consumption is obtained from the HP filtered quarterly real consumption series (summed over the three countries included in the data sample) from 1995 to 2008. Data source: Eurostat.

E Additional tables


Table E.1: *Overview of goods categories included in the estimation*

Belgium	Germany	Netherlands
Beer	Body powder	All-purpose cleaner
Bleach	Butter	Antrycide
Breakfast cereals	Canned instant meal	Bathroom polish
Buillons	Canned pickles	Bleach blocks
Butter	Ceramic glass cleaner	Chocolate bars
Canned instant meal	Dentifrice	Chocolate marshmallows
Canned vegetables	Dried mushrooms	Christmas pastries
Condensed milk /creamers	Eau de toilette, women's	Cleansing tissue
Cotton pads	Eiskonfekt	Dishwashing liquid
Crispy bread	Foils	Dry instant meals
Curd cheese	Inceticide	Eye makeup
Dessert sauce, liquid	Liquors	Foot care
Flavouring/herbs	Metal cleaner	Fresh bakery products
Frozen dinners and Entrees	Pickled gherkin	Hair conditioning products
Hairsprays	Poultry	Hairsprays
Insecticide	Shampoo	Hand dishwashing liquid
Incontinence products	Sherbet powder	Honey
Isotonic drinks	Sweet dishes	Ketchup
Ketchup	Wine, fruit	Mayonnaise
Lye bisquits	Whisky	Metal cleaner
Mouthwash		Mustard
Packet soup		Stain remover
Seasoning and cocktail sauce		Toilet paper, wet
Shampoo		Vegetable oils
Soap		Vinegar
Soft spirit		
Sugar		
Sweet spreads		
Tea		
Tube cleaner		
Wine		

Notes: The goods categories are sampled according to the description provided in 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, while for others, there is not enough data or it did not converge.


F Additional figures

Figure F.1: *Size and frequency of price changes in the dataset*



Notes: The size of a price change is the percentage change from the previous month's price of the identical product, given that the 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. Similarly, a price change is noted when two prices are observed in two consecutive months and they are identical. All prices for which we did not observe two consecutive months for the same household and retailer are dropped before calculating these figures.

Figure F.2: *Correlation between changes in prices and changes in quantity purchased*



Notes: This figure reports the log-change in prices of a unique product from one year to the next on the horizontal axis and the log-change in purchased quantity of the same product from one year to the next on the vertical axis. To construct annual data, we average all prices for a GTIN over a year and we summed all quantities over a year.