Price Setting on the Two Sides of the Atlantic - Evidence from Supermarket Scanner Data[☆]

Peter Karadi[a], Juergen Amann[b], Javier Sánchez Bachiller[c], Pascal Seiler[d], Jesse Wursten[e]

[a] *corresponding author, European Central Bank, DG Research, CEPR.*

[b] *OECD.*

[c] *University of Luxembourg.*

[d] *ETH Zurich, KOF Swiss Economic Institute.*

[e] *Faculty of Economics and Business, KU Leuven.*

Abstract

We compare supermarket price settings in the US and the euro area and assess their impact on food inflation. We introduce a novel scanner dataset of Germany, the Netherlands, France, and Italy (EA4) and contrast it with an equivalent dataset from the US. We argue, first, that product-level volatility is higher in the US, and it is responsible for the more flexible price setting there. Second, we provide direct evidence of state dependence in price setting. We find that it raises price flexibility by around 25% in both regions through the endogenous selection of large price changes. Both facts are well-represented by a mildly state-dependent price-setting model, and they jointly explain over a third of the difference in food inflation volatility between the US and the euro area, and also around a third of the difference between the inflation response to the Covid shock in Germany and Italy.

[☆]All opinions expressed are personal and do not necessarily represent the views of the European Central Bank or the Eurosystem. We would like to thank IRi for making the data available. All estimates and analyses in this paper, based on data provided by IRi, are by the authors and not by IRi. We thank the editor, Yuriy Gorodnichenko, an anonymous referee, Luca Dedola, Lukas Henkel, Ryan Kim, Oleksiy Kryvtsov, Raphael Schoenle, participants of the ECB-Cleveland Fed Inflation Drivers and Dynamics 2022 Conference and seminars at Brandeis University and the European Central Bank for their comments and Ivan Borna Brcina, Matthias Gnewuch, and Marta Oliva Riera for their excellent research assistance.

Email addresses: peter.karadi@ecb.europa.eu (Peter Karadi), juergen.amann@oecd.org (Juergen Amann), javier.bachiller@uni.lu (Javier Sánchez Bachiller), seiler@kof.ethz.ch (Pascal Seiler), jesse.wursten@kuleuven.be (Jesse Wursten)

1. Introduction

Food inflation is more volatile in the US than in the euro area and, correspondingly, responded more forcefully to the recent Covid-19 pandemic (see Figure 1). Price setting in the food-retail sector has macroeconomic relevance because food consumption accounts for around one-fifth of consumption in both regions, and the salience of grocery prices makes them influence households' aggregate inflation expectations (**?**). Previous research has established that price flexibility depends both on the frequency of repricing (*how many* prices change) and the extent of state dependence in price setting (*which* prices change) (**???**). In this paper, we use new store-level scanner data from the euro area and a corresponding dataset from the US to carefully measure these two features of supermarket price setting, which assess their impact on the differences in food-inflation volatility.

We find that both the higher frequency and the stronger state dependence of price changes contribute to the higher flexibility of supermarket inflation in the US relative to the euro area. We argue that the driving force behind both factors is higher cross-sectional volatility in the US. Larger product-level fluctuations both (i) force retailers to adjust prices more frequently and (ii) raise price misalignments, which increase the selection of large price changes. Our conclusions have implications for both model selection and policy.

The paper introduces a novel store-level scanner dataset acquired from the marketing company IRi by the European Central Bank in the context of the Price-setting Microdata Analysis Network (PRISMA). The dataset covers Germany, the Netherlands, France, and Italy (EA4), over five years between 2013-2017.[1] The dataset records weekly prices of over 2 million products in over 37 thousand stores in a spatially representative sample. We contrast it to evidence obtained from the US IRi Academic Dataset, an analogous weekly panel over the period 2001-2012 of over 200 thousand products in over 3000 stores covering the 50 most relevant US markets.

We use the datasets to characterize key features of price settings in the US and the euro area. First, we contrast standard moments about the repricing frequency and the size distribution of price changes. We filter out temporary sales (**??**), which account for the majority of price changes, but contribute only marginally to fluctuations in inflation at regular business cycle frequencies. In line with previous evidence, we find that sales-filtered reference prices change infrequently, and the average absolute size of price changes are large in both regions (**???**). Previous research has also concluded that one needs price-adjustment frictions and

[1] For the analysis of the Covid shock, we use an auxiliary dataset, which covers the period between mid-February to mid-May in 2019 and 2020 in Germany and Italy for a subset of the stores. For details, see Section 6.

Figure 1: Food and non-alcoholic beverage inflation in the US and euro area, COICOP 01, harmonized prices, year-on-year

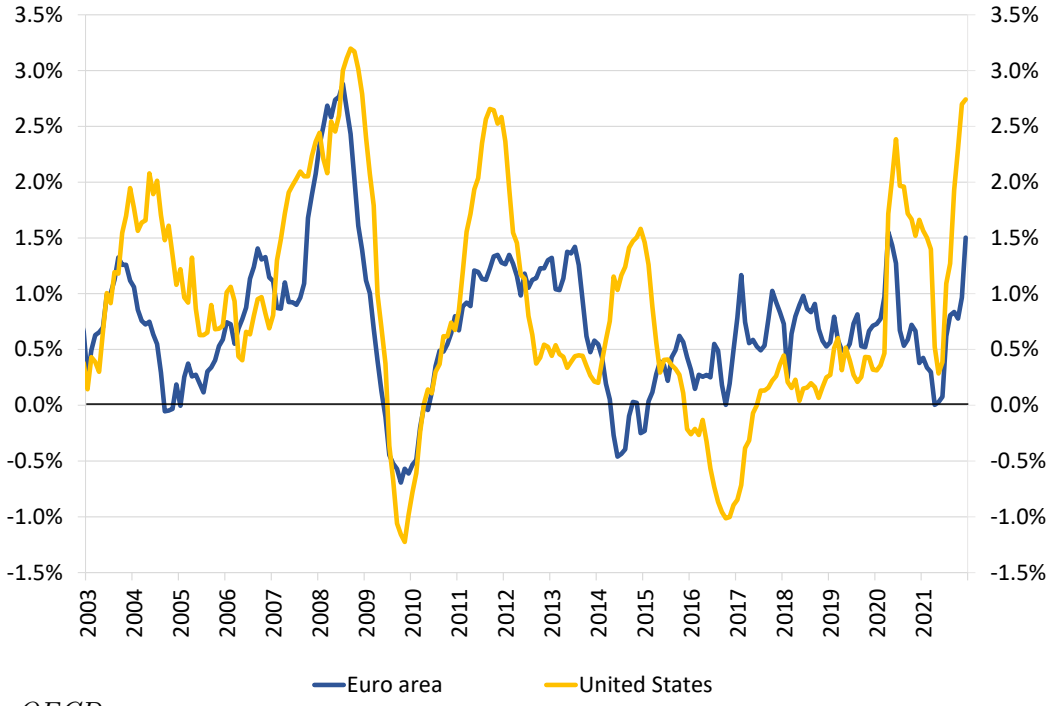

Source: OECD

Notes: The figure shows the evolution of year-on-year food and non-alcoholic beverage inflation in the US and euro area between 2003-2021. The series show clear co-movement over most of the period (correlation: 59%), and the US inflation shows higher volatility than euro area inflation (standard deviations: US: 0.95%, EA4: 0.64%).

large product-level shocks to account for the low frequency and the large magnitude of price changes (**?**). We show that the frequency of price changes (implying price changes once every 7 months in the US versus 10.5 months in the EA4) is substantially higher in the US. The higher frequency is accompanied by a higher average size of price changes (15% in the US versus 9% in EA4). This indicates that the product-level volatility is larger in the US relative to the euro area.

Second, we measure the extent of state dependence in price setting in the two regions. State dependence determines the endogenous selection of large price changes and can raise the volatility of inflation. We use the unparalleled cross-sectional granularity of the data to generate data moments, which are directly informative about state dependence. In particular,

we create a proxy for price-misalignments as the distance of a (log) price from the average price of the same product in those competitor stores that changed their prices in the same month. The average price of price-adjusting stores reveals the optimal reset price in a wide class of models (**????**). To assess the extent of state dependence, we measure both the probability of price adjustment as a function of the misalignment (adjustment hazard) and the density of misalignments following the framework of **?**. We find that state dependence is higher in the US than in the euro area. It raises aggregate price flexibility by around 25% in both regions leading to a larger absolute impact in the US, where price flexibility due to frequency is already higher. This implies that popular price-setting models that ignore state dependence (**?**) can underestimate price flexibility by roughly a third. Notably, the key difference between the extent of state dependence is driven by the more dispersed density of price misalignments, which is strongly influenced by the already established higher volatility of product-level shocks. Our conclusions about the state dependence of price changes are supported by additional data moments. Specifically, the kurtosis of standardized price changes, which decreases with higher state dependence in a wide class of models (**?**), is low in both regions and lower in the US than in the euro area. Furthermore, the duration-hazard of reference price changes is increasing in both regions in line with state dependence in price setting, after we suitably control for unobserved heterogeneity.

Next, we conduct a structural analysis of the price-setting moments, which confirms that higher product-level volatility is one of the key underlying causes of differences in price-setting and food-inflation volatility across the two regions. We use the state-of-the-art state-dependent price-setting model of **?** to estimate three underlying structural parameters affecting price setting: (i) the magnitude of price-adjustment (menu) costs, (ii) the standard deviation of idiosyncratic shocks, and (iii) the magnitude of information-acquisition costs, which determines the level of state dependence in the model between the time-dependent (**?**) and the fixed-menu-cost (**?**) extremes. The most notable difference between the US and euro area is the higher volatility of idiosyncratic shocks in the US, while both the price-adjustment and the information-acquisition costs are quite similar in the two regions. The model can account for over a third of the observed difference in food-inflation volatility between the US and euro area.

Finally, we provide evidence on the response to aggregate shocks by assessing supermarket prices in Germany and Italy during the first wave of the Covid-19 pandemic. The shock raised supermarket demand (by restricting access to food away from home) with limited impact on costs (supermarkets were essential sectors sheltered from the lockdowns). The comparison is interesting because the frequency of price change is substantially lower in Germany than in Italy. The difference is mainly driven by the lower idiosyncratic volatility in Germany. We document that the inflation response was lower in Germany, and the model calibrated to the

price-setting moments of the two countries can explain around a third of the difference.

Related literature: The paper is related to different strands of the literature. We contribute to the literature that compares price settings in the euro area and the US by introducing a new supermarket scanner dataset and contrasting key price-setting moments, like the frequency and the size of price changes. **?** compare price settings in the two regions using microdata underlying the Consumer Price Index. They confirm that the frequency and the size of (sales-filtered) price changes are larger in the US not only in the processed food sector, as in our sample, but also in the whole economy, albeit at a somewhat smaller degree.

We also contribute to estimating the extent of state dependence in price setting. We calculate moments that are directly informative about state dependence, like the generalized- and the duration hazard functions, utilizing the high granularity of the scanner data. We find that the generalized hazard function, which expresses the probability of price changes as a function of price misalignment, is upward-sloping both in the US and in the euro area in line with state dependence. To proxy price misalignments, we use distance from competitors' reset prices (**?**), which is a valid proxy in a wide range of price-setting models. Our results confirm previous results, which use distance from competitors' prices on more restrictive samples (**??**), and are consistent with complementary estimates using distance from an estimated cost measure (**??**). We show, furthermore, that the duration hazard, which measures the probability of a price change as a function of the age of the price, is upward-sloping in both regions when we use sales-filtered reference prices and control for unobserved heterogeneity. Upward sloping duration hazard is in line with state-dependent pricing models (see, for example, **??**). Our evidence is different from **????**, which find the hazard decreasing, but in line with **?**, which find it non-decreasing for most disaggregated product-groups.

We assess the implications of our evidence by estimating key structural parameters of a state-of-the-art price-setting model (**?**) in both regions. The model features fixed (menu) costs of price adjustment (**?**), product-level technology shocks (**?**), and information frictions, which allow it to capture the infrequent and large price adjustments and state-dependence in line with our evidence. As **?**, **?** and **?**, we find that state dependence raises the flexibility of the price level in both regions. We argue that higher volatility of product-level shocks in the US is the key reason behind cross-country differences in price setting and food-inflation volatility. This result is related to **?**, who argues that variation of idiosyncratic volatility over time (as opposed to across countries) implies time-varying price flexibility over the business cycle in a related framework. Higher cross-sectional volatility in the US versus the euro area has also been documented using various alternative measures. Using stock returns, **?** and **?** document higher idiosyncratic volatility in the US than in major euro area countries (see also **?**). Relatedly, **?** find higher firm-level employment-growth volatility in the US than in

selected euro area countries (see also **?**).

The paper is structured as follows. Section 2 describes the data. Section 3 describes conventional moments of price changes in the two regions, including frequency, size, and higher-order dispersion measures. Section 4 presents more complex moments, including the generalized (price gap) and the duration (price age) hazard functions, and quantifies the level of state dependence in the two regions. Section 5 conducts structural analysis, and Section 6 contrasts the price-setting response to the Covid shock in Germany and Italy. Section 7 concludes.

2. Data

This section introduces the novel euro area dataset and shows its key features together with its US counterpart. We also present the data-cleaning steps we implemented to improve the informativeness of the data for the analysis of price setting.

2.1. Data coverage

The data covers 4 euro area countries: Germany, the Netherlands, France, and Italy between 2013-2017 and the US between 2001-2012.[2] The datasets are weekly panels of total revenues (TR_{psw}) and units sold (Q_{psw}) for each product p in store s in week w. We refer to a product in a store as an item. Unit-value prices of each item are calculated as revenues over units sold ($P^{uv}_{psw} = TR_{psw}/Q_{psw}$).

2.1.1. Product coverage

The granularity of the datasets is unsurpassable: they include *all* products sold in each store in the sample.[3] The products are identified with their unique barcodes (EANs in the euro area and UPCs in the US).[4] The number of unique products ranges from around 390.000 to

[2]Even though the US and EA4 datasets do not overlap, this does not hinder the comparison of the key moments used in our analysis, as they are fairly stable over our sample period (see, for example, Figure E.12 in the appendix on the frequency of reference-price changes).

[3]The US sample only includes products within a selected 30 broad product categories: beer, blades, carbonated beverages, cigarettes, coffee, cereal, deodorant, diapers, facial tissue, frankfurters, frozen dinner, frozen pizza, household cleaner, laundry detergent, butter, mayonnaise, milk, mustard&ketchup, peanut butter, paper towels, photography supplies, razors, salty snacks, shampoo, spaghetti sauce, sugar substitutes, toilet tissue, toothbrush, toothpaste, yogurt.

[4]The EANs of *private-label* products are masked to protect the identity of the supermarket chain. We exclude private-label products from the analysis in France, where the revenues and quantities sold of all private-label EANs are aggregated by store, confounding the evolution of item-level prices.

Table 1: Data coverage

	US	DE	FR	IT	NL
Time series	2001-2012		2013-2017		
# products (td)	205	370	423	698	392
# stores (td)	3.3	10.3	5.9	14.3	6.6
# observations (bn)	2.7	13.79	10.02	10.96	7.67
# 2-digit ZIPs	51	97	93	93	94
# chains	147	17	43	435	29
% in HICP/CPI	19.6	18.5	23.3	23.4	20.7
annual exp. (bn €/$)	6.2	32.8	56.19	42.16	30.02

776.000 in the euro area and over 200.000 in the US (see Table 1).

Products sold in supermarkets include food, alcoholic- and non-alcoholic beverages, personal-care products, as well as goods for household maintenance. They cover around 20% of the consumer baskets (see Table 1). The expenditure distribution in the IRi samples approximates quite well the true consumption pattern of households across major product categories (see Appendix A).

We conduct the analysis below using a subsample for each country to ease the computational burden. Specifically, we select a 5% random sample of EANs in each EA4 country, and a 25% random sample of UPCs from the US.[5] The random choice of products ensures that the sample is representative. We include all the stores and time periods in the subsample, where and when the selected products were sold in positive quantities.

2.1.2. Store coverage

The datasets are representative of the brick-and-mortar sale of participating supermarket chains. The participating chains include regular and discounter supermarkets as well as drug stores.[6] The store IDs are masked to protect the identity of the supermarkets, but they are unique over time, which allows us to track the price spells of items over time.[7] In the euro

[5]The US sample includes fewer products and stores (see Table 1). Choosing a relatively larger subsample makes the number of items in the US sample the same order of magnitude as in the euro area countries.

[6]The datasets exclude 'hard' discounters like Lidl, Aldi, or Walmart.

[7]To guard the identity of the stores, store information is only included in our sample if there is a sufficient number of stores (for example, at least three in France) by geographical area and store type. In most cases

area countries, our dataset includes 75 percent of the IRi stores. In two countries (Germany and Italy), some supermarket chains only share a representative sample of their stores with IRi. We 'upweight' sample stores using projection weights created using information about the population of stores by geographic unit and store type, which is also part of the dataset (see Appendix B for details).

The euro area datasets are spatially representative in each country. The datasets include the location of the stores up to the first two digits of their ZIP code. The 2-digit ZIP areas partition the countries into around 100 regions (see Table 1). The US dataset covers 50 urban markets across the US. These markets approximately correspond to 50 Metropolitan Statistical Areas (MSA) out of the 384 MSAs in the mainland US in 2010 and cover 73% of the US population.[8]

2.2. Data cleaning

The focus of our analysis is reference prices (**??**), and we conduct a series of filtering steps to obtain them.

First, we estimate posted prices from weekly unit-value prices. We conduct two filtering steps. First, we filter out same-direction consecutive changes. We do this to minimize the impact of mid-week price changes. Intuitively, a mid-week price increase raises the average price only partially in the initial week and passes through fully only during the second week. Second, we round prices upwards to the nearest cent. We do this to mitigate the impact of buyer-specific discounts (see Appendix C for details).

Next, we construct weekly reference prices (P^f_{psw}) as 13-week running modal prices (**?**). Reference prices capture persistent changes in prices and disregard temporary sales that are completely reversed within weeks. By focusing on reference prices, we capture an overwhelming share of fluctuations in supermarket inflation at business cycle frequencies, and we filter out a large share of high-frequency variation caused by temporary sales (see Appendix D.1 for details). Temporary sales also account for a sizable fraction of the frequency of posted-price changes, almost two-thirds in most countries (see Table D.7). Despite an ongoing debate whether sales are an active adjustment margin for retailers to aggregate fluctuations (**???**), there is a widespread consensus that most adjustment at business cycle frequencies is achieved through reference prices. Previous research has also documented that sales inflation does not

(in France and the US, for example), store information is withdrawn in these cases. In other cases (in Italy, for example), the geographical granularity becomes coarser (1-digit as opposed to 2-digit ZIP areas).

[8]Therefore, even though the US sample is not spatially representative, it covers the most populous areas providing a relevant sample of supermarkets across urban areas.

respond significantly or responds only marginally to aggregate shocks (**???**). This justifies our focus on the behavior of reference prices in the rest of the paper.

Last, we transform weekly data to a monthly frequency. This facilitates comparison with monthly micro price data underlying the official price indices, helps us concentrate on more persistent price adjustments that are more relevant at business cycle frequencies, and overcome some of the weaknesses of the data, including the sizable share of missing data due to zero sales. We define the monthly item price as the (highest) mode of the item price over the month.[9]

3. Key moments of price changes

In this section, we characterize key features of reference price changes in supermarkets across the 4 euro area countries and the US. We focus on conventional moments, including frequency, size, and kurtosis of price changes, which were found to be relevant by the theoretical literature to influence the flexibility of the aggregate price level.

Table 2: Key moments of reference-price changes, expenditure-weighted

Moments	**US**	**EA4**	**DE**	**FR**	**IT**	**NL**
Frequency	13.8%	9.5%	5.3%	15.6%	9.6%	10.3%
Size	15.2%	9.3%	10.5%	4.5%	11.4%	9.2%
Kurtosis	2.65	3.2	2.9	3.8	3.32	2.71

Note: The table presents the frequency and the average absolute size of reference-price changes and the kurtosis of the standardized reference-price changes, all weighted by monthly expenditure weights.

[9] Using the mode guarantees to choose one of the weekly reference prices, so the time aggregation does not introduce artificial prices. This would happen if one instead used the mean or calculated monthly unit prices. Picking the *highest* mode in case of multi-modality tilts the monthly prices towards the (more persistent) reference prices, which tend to be above the sales prices.

3.1. Frequency

The frequency of reference-price changes is a key indicator of price flexibility. As the first row of Table 2 shows, the average frequency[10] in EA4 supermarket prices are fairly low, only 9.5 percent monthly. This suggests that reference prices change infrequently, only once around every 12 months, on average. The low frequency indicates that supermarkets face price-adjustment frictions, which hinder them from adjusting prices flexibly to changes in costs. The price flexibility is higher in the US, where the frequency of reference price changes is 13.8 percent, implying an average duration of 7 months.[11]

There is a notable heterogeneity in frequency across euro area countries. In Italy and the Netherlands, the frequency is close to the EA4 average, but in Germany, it is particularly low at 5.3% (19 months average duration) and in France, it is particularly high at 15.6% (6.5 months duration), and thus even higher than in the US.

3.2. Size

The average absolute size of reference price changes is large: 9.3 percent in the EA4 countries, on average. Its magnitude way exceeds what could be explained by trend inflation or aggregate fluctuations, which are both small during our sample period. Instead, they indicate an important role for idiosyncratic, product-level shocks. The size of price changes is higher in the US, where it reaches 13.8 percent. The higher size accompanied by a higher frequency indicates a more volatile product-level environment in the US, a factor that we will analyze further in a structural framework in Section 5. The size of price changes also varies across euro area countries. The average size is particularly low in France, which, together with the high frequency, indicates lower-than-average price-setting frictions there.

The first panel of Figure 2 shows the histograms of the absolute price change distributions in both areas. The size of price changes in both regions is dispersed, with many small as well as large price changes. Smaller price changes are more frequent, and larger price changes are less

[10] All moments are weighed by annual expenditure weights. Formally, monthly frequency is calculated as $\xi_t = \frac{\sum_i \omega_{it} I_{it}}{\sum_i \omega_{it}}$ where I_{it} is an indicator that takes the value 1 if the reference price of item i (a product in a particular store) changed from month $t-1$ to month t and 0 otherwise. The weights ω_{it} are annual expenditure weights. Table 2 reports average monthly moments over the sample. Average absolute size and percentiles are calculated analogously. Kurtosis is calculated using the subsample of items with at least 5 reference price changes, standardized by the mean and standard deviations at the item level, and weighed by expenditure weights.

[11] The frequency is stable over time, so the issue of non-overlapping US-EA4 samples should not hinder the international comparison (see Figure E.12 in the appendix).

frequent in the EA4 relative to the US. Correspondingly, the dispersion of the price changes is smaller in the EA4 with an interquartile range of 9 percent, while it is 16 percent in the US.

Figure 2: Absolute and standardized reference-price-change distributions

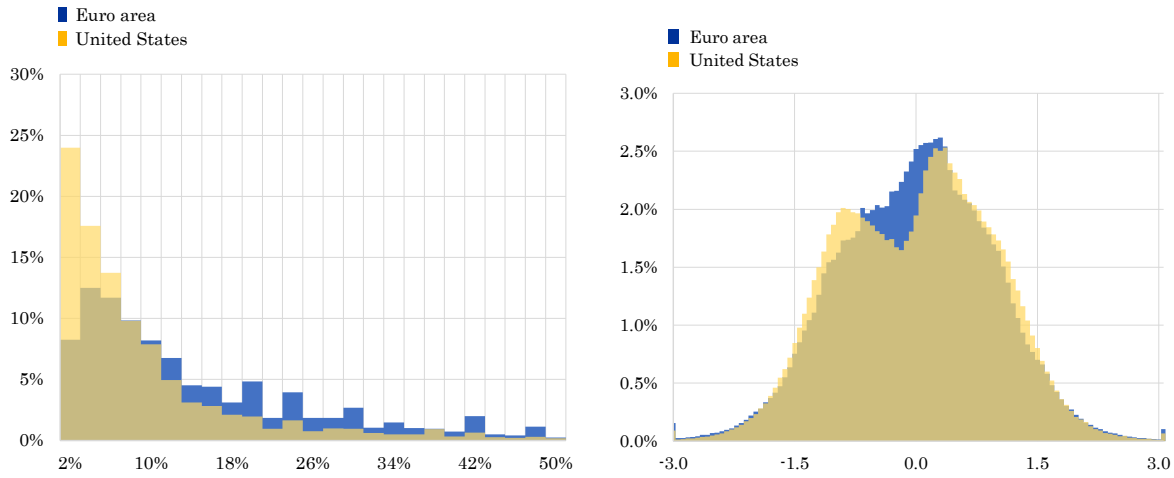

Absolute price changes | Standardized price changes

Note: The figure shows the absolute (left panel) and the standardized (right panel) reference-price-change distributions in both regions. The size is large and dispersed in both regions, and it is larger and more dispersed in the US than in the euro area. The shape of the standardized price change distribution is bimodal in the US, while it is unimodal in the euro area, with lower kurtosis in the US.

3.3. Higher-order moments

The shape of the price-change distribution can be informative about the extent of state dependence in price setting in a wide class of models (**?**). The second panel in Figure 2 shows the shape of reference-price change distribution in both regions. The reference-price changes are standardized at the product-store level to minimize the potential bias caused by cross-product heterogeneity in the mean or standard deviation of price changes. The figures indicate a kurtosis of 2.7 in the US and 3.2 in the euro area, around the kurtosis of 3 of the Gaussian distribution.

The distribution shows some pronounced bimodality in the US with some 'missing' mass close to zero, which is in line with the presence of fixed costs of price adjustment. At the same

time, the share of small reference-price changes stays high also in the US, much higher than models with strong state dependence would predict (**?**).

4. Evidence on state dependence - generalized and duration hazards

The conventional moments described in the previous section provide only indirect information about an important feature of price setting: the extent of its state dependence. Previous research has established that state dependence, which influences *which* prices adjust, can have as large an impact on the aggregate price flexibility as the frequency, which determines *how many* prices adjust. For example, in realistic models of price setting with strong state dependence (for example, **?**), the price level can respond almost completely flexibly to monetary policy shocks even though only a few prices adjust. The reason is that in these models, firms face a small fixed menu cost to change prices, so they find it optimal to adjust prices with large misalignments. When these prices change, they change by a lot, which can offset the impact of price rigidity and make the price level flexible. In this section, we present two sets of moments that are more directly informative about the extent of state dependence than conventional moments.

4.1. Generalized hazard

The first moment is the generalized (price gap) hazard function, which expresses the probability of price adjustment as a function of the price misalignments, or price gaps. The price gap is the distance between the posted price and the optimal 'reset' price the store would set if all price-adjustment frictions were temporarily absent. The gap, therefore, influences the strength of the product-level price-adjustment impetus: the larger it is, the further the price is from its optimal level, causing a potentially larger profit loss through either suboptimal demand (if the price is too high) or suboptimal markup (if the price is too low).

A key empirical challenge is that the optimal reset price is unobservable. As a proxy, we calculate the competitors' reset price (**?**). It is the average reference price[12] of the same products in those competing stores that also changed the price of the same product in the same month. The measure also controls for the permanent store- and category-level price differences caused by heterogeneity in amenities, geography, or market power. The proxy relies on the assumptions that (i) the price of the same good among price-changing competitors tracks well the evolution of the product's wholesale price and aggregate demand conditions, which are the primary drivers of the optimal reset price, (ii) differences in amenities and market

[12]By concentrating on reference prices, the measure controls for the impact of temporary sales.

power between stores cause permanent store- and category-level differences between prices, and (iii) chains follow national price-setting strategies (**?**), so local demand conditions have an insignificant impact on the optimal reset prices. We validate our proxy by showing that the size of the price change has a very tight, almost exactly one-to-one negative relationship with the price gap.

Formally, we formulate the competitor-reset-price gap x_{pst} for product p in store s in month t in three steps. First, we take the (logarithm of) the sales-filtered reference prices p^f_{pst}. Second, we calculate an unadjusted gap as $\tilde{x}_{pst} = p^f_{pst} - \bar{p}^f_{p(-s)t}$, where $\bar{p}^f_{p(-s)t}$ is the average reference-reset-price of the same product across those alternative stores that changed the price of the same product in month t. Third, we deal with the persistent heterogeneity across stores (i.e., chains, locations) by subtracting the average store- and category-level gap α_{cs} and reformulate the price gap as $x_{pst} = \tilde{x}_{pst} - \alpha_{cs}$, where product p belongs to category c.

Figure 3: Price-gap density and the size of non-zero price changes as a function of the price gap, and the generalized and duration hazards.

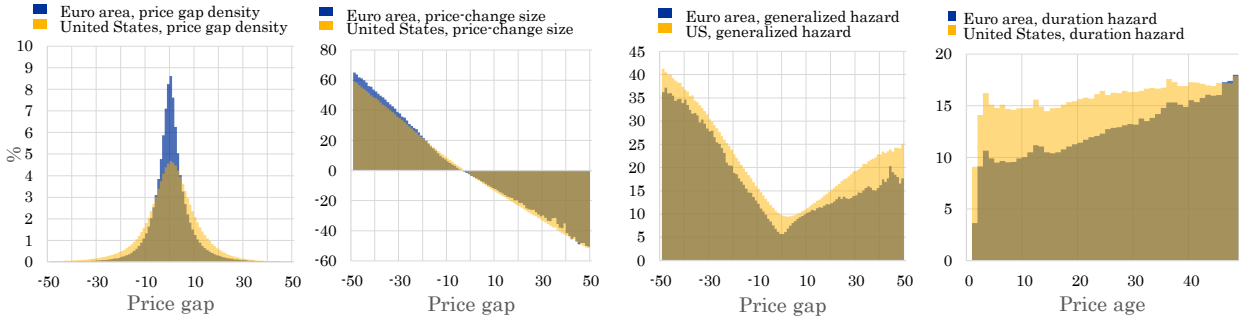

Note: The figures show the frequency of reference price changes (generalized hazard, left panel) and the average size of non-zero reference price changes (right panel) as a function of the price gap, and the density of the price gap (middle panel) in EA4 and the US. The V-shape hazard indicates the presence of state dependence in price setting, albeit at a moderate level in both regions. The density indicates wide dispersion of price gaps, higher on average in the US. The size figure validates the price gap measures showing a tight relationship between the gap and the eventual price-change size.

The first panel in Figure 3 shows the density of the price-gap distributions in the four euro area countries and the US. To arrive at the densities, we control for unobserved heterogeneity across items and the common impact of aggregate fluctuations by estimating item- and time-

fixed effects[13] in a panel regression of the form

$$x_{pst} = \alpha_{ps} + \alpha_t + \varepsilon_{pst}, \tag{1}$$

and calculating the share of normalized gaps ($x_{pst} - \hat{\alpha}_{ps} - \hat{\alpha}_t$) in the 101 unit-percentage-point ranges between -50.5 and 50.5 percent. We censor the normalized gaps at -50.5 and 50.5 percent.

The figure shows that the gaps are high, on average, and higher in the US (14%) than in the EA4 (10%). At the same time, the distribution is dispersed in both regions with a high mass of small gaps and a fat tail of large gaps. This is true even though we control for sales-related price changes as well as permanent differences between the store-specific prices.

We now assess the relationship between the price gap in period $t-1$ and the probability and average size of price adjustment in the following month t. We aim to estimate this relationship non-parametrically with a minimal set of structural assumptions. First, we allocate price gaps into 101 bins, each covering a unit percentage-point range between -50.5 and 50.5 percent. The indicator function $I_{pst-1}^{[x_{j-1},x_j)}$ for bin j takes the value 1 in case the gap $x_{pst} \in [x_{j-1}, x_j)$, and 0 otherwise. Second, we estimate a relationship coefficient (β_y^j) between the gap x and a variable of interest $y_{pst,t+1}$ (frequency or size) for each bin j using the following panel specification:

$$y_{pst,t+1} = \sum_{j=1}^{J} \beta_y^j I_{pst-1}^{[x_{j-1},x_j)} + \alpha_{ps} + \alpha_t + \varepsilon_{pst} \tag{2}$$

where α_{ps} are product-store-, and α_t are time-fixed effects. The fixed effects help us to control for unobserved heterogeneity across items and common co-movement caused by aggregate fluctuations. Third, we obtain the estimated relationship as a sum of two components. The first component is the β_y^j coefficients for $j = [1, 101]$. The second component is the average of the estimated fixed effects $\text{mean}_{ps}\hat{\alpha}_{ps} + \text{mean}_t\hat{\alpha}_t$ added to each bin j. Adding the second component makes sure that the weighted average across bins approximates the sample average of the variable of interest y.

The second panel in Figure 3 shows the average size of non-zero price changes as a function of the price gap in EA4 and the US. It is estimated following the above-described steps when

[13] An alternative specification would be to include item- and time-store fixed effects. The time-store fixed effects would control for store-wide synchronization of price changes (see, for example, **?**). The specification leads to virtually identical US hazard function estimates (not shown), suggesting that synchronization is present, but plays a quantitatively insignificant role.

the dependent variable is the non-zero reference-price changes $y_{pst,t+1} = \Delta p^f_{pst+1}|_{\Delta p^f \neq 0}$. The figures show a tight, negative, almost exactly one-to-one relationship between the gap and the average non-zero price changes in the subsequent month.[14] This validates our price-gap measure by showing that stores choose to close the gap, on average, when adjusting the price.

We are now ready to turn to one of the key empirical moments we are interested in: the generalized hazard function, shown on the third panel of Figure 3. They are estimated for each region following the steps outlined above when the dependent variable is an indicator function that takes the value 1 in case the reference price of product p in store s changed in period t+1, and 0 otherwise $y_{pst,t+1} = I^f_{pst+1}$. The figures show clear evidence for state dependence in price setting in both regions: the probability of price adjustment clearly increases with the price gap as illustrated by the V-shape of the hazard functions. The slopes of the hazard functions, however, are moderate in both regions: the estimated probability of adjustment stays below 40% even for price gaps of 50 percent. The (density-weighted) average slope is 0.51 in EA4 and 0.38 in the US, the difference between the regions is caused by the larger slope at lower gaps, where the largest mass of price gaps are concentrated. Additionally, the probability of price change is strictly positive even at zero gaps and the hazard functions are asymmetric: the probability of adjustment is higher in case the item is below the competitors' reset price than when it is above. The height of the hazard function is larger in the US, in line with the higher frequency of price changes there, as we have already documented above.

4.2. Duration hazard

An alternative way of looking at state dependence is the duration hazard, which expresses the probability of price adjustment as a function of the months elapsed since the last price adjustment. In the presence of state dependence, the duration hazard is upward-sloping, because the probability of price change rises as the optimal price drifts further and further away from the posted price. The advantage of using granular scanner data to estimate the hazard function is that we can control for the cross-item heterogeneity, which can bias the slope estimate downward. We estimate the following panel regression

[14] The slope of the relationship is actually somewhat above 1. Such an outcome is consistent with the presence of concurrent item-level shocks, which hit after the gap was measured (remember, the gap is lagged by a month). If items with higher unobserved shocks are changed with a higher probability, which is in line with state dependence, the selection effect increases the average absolute size of price changes at each gap size. The fact that the relationship is only marginally steeper than 1 suggests that the measured gap plays a quantitatively much more important role than the unobserved concurrent shocks.

$$I_{pst,t+1} = \sum_{j=1}^{48} \beta^j I^j_{pst-1} + \alpha_{ps} + \alpha_t + \varepsilon_{pst}, \tag{3}$$

where the indicator function I^j_{pst-1} takes a value 1 if the reference price of product p in store s in month $t-1$ is j months old, and 0 otherwise. As with the generalized hazard, we add the average of the estimated item- and time-fixed effects to the β^j coefficients in order to make the weighted average of the coefficients approximate the frequency of reference-price changes.

The fourth panel of Figure 3 shows the results for EA4 and the US. It shows that the duration hazard is upward-sloping in both regions: the probability of adjustment increases with the age of the product. The slope of the adjustment hazard is higher in EA4 than in the US. In Appendix G, we show that controlling for both cross-item heterogeneity as well as sales-related price changes is essential to obtain upward-sloping hazards.

4.3. State dependence and price-level flexibility

In the previous section, we have argued that the V-shaped generalized hazard function and the upward-sloping duration hazard function are in line with state dependence in price setting. In this section, we quantify the extent of this state dependence. A natural measure of state dependence is how much it contributes to price flexibility, specifically to the price-level impact of a permanent money shock. To measure this, we follow the framework of **?**, who showed that under mild conditions, the generalized hazard function and the density provide sufficient information to quantify the contributions of the intensive and extensive margins of adjustment. We first describe the framework and explain how the relevant objects in the model relate to our empirical moments before turning to use it to decompose an aggregate money shock to adjustment margins.

In the price-setting framework of **?**, there is a continuum of firms each producing a single product i. Firms set the (log nominal) prices of their product (p_{it}) subject to price-adjustment friction. If these frictions were temporarily absent, the optimal price in period t would be p^*_{it}. The optimal price is driven by both aggregate and idiosyncratic factors $p^*_{it} = m_t + \nu_{it}$. For simplicity, we assume that shocks to both m_t and ν_{it} are permanent. The aggregate shock m_t shifts the optimal nominal price of all firms, whereas the idiosyncratic shock ν_{it} affects only firm i. The gap between the price and its optimal value $x_{it} = p_{it} - p^*_{it}$ is the relevant state variable and is sufficient to characterize each firm's price-setting choice. Assuming that the product i is sold in a continuum of stores, the average price set by price-changing stores reveals the optimal price p^*_{it}, in line with our empirical application.

The firms' price adjustment decision can be described by a generalized hazard function $\Lambda(x)$. The function takes values between 0 and 1, and its value expresses the probability of price adjustment for a firm with a price gap x. The hazard function is constant in the time-dependent **?** model: there, the probability of adjustment is independent of the price gap. At the other extreme, in the fixed menu cost model (**??**), the hazard function is a step function, which takes the value 0 when the gap is within the inaction band, and 1 otherwise. **?** shows that a continuum of intermediate hazard functions can arise when the menu cost is an i.i.d. random variable **?**, and when the firm is subject to rational inattention friction as in **?** (see also **?**).

In this economy, inflation can be expressed as

$$\pi = \int -x\Lambda(x)f(x)dx \tag{4}$$

where $f(x)$ is the density of price gaps across firms, and we suppressed subscripts for notational convenience. The expression is intuitive: the inverse price gap $(-x)$ is the size of the price adjustment, when it takes place, and the hazard is the probability of a price adjustment taking place. Their product summed across the gap distribution and weighted by the density of the gap is, therefore, equal to the inflation rate.

The question is how flexibly the inflation rate reacts to a small aggregate money increase m? **?** point out that the aggregate shock increases the optimal price of all firms, so it reduces the price gaps of each firm uniformly. The response to the aggregate shock can therefore be expressed as a derivative of the expression on the right-hand side of Equation (4) with respect to x, which implies

$$\frac{\partial \pi}{\partial m} = \underbrace{\int \Lambda(x)f(x)dx}_{\text{intensive}} + \underbrace{\int x\Lambda'(x)f(x)dx}_{\text{extensive}}, \tag{5}$$

where $\Lambda'(x)$ is the slope of the hazard function. The expression has two terms. The first term, which **?** dub the intensive margin results for each adjusting firm changing their prices by marginally more to incorporate the impact of the aggregate shock. Notably, it is exactly equal to the frequency of price adjustment, and this is the only margin that is active in the time-dependent **?** model, which has a constant hazard. The second term is the extensive-margin effect, which takes into account any shifts in the identity of price-adjusting firms. The slope of the hazard function appears in this expression because it measures the mass of new price adjusters as the aggregate shock shifts the price-gap density. The extensive margin is powerful if the new adjusters are primarily those with large price gaps. This tends to be the case with strongly state-dependent (S,s)-type menu cost models (**?**), where it is optimal to adjust prices with the largest gaps in the presence of fixed menu costs of price adjustment.

Our empirical estimates on the hazard function and the density of the price gap shown in Figure 3 allow us to conduct the **?** decomposition described by Equation (5). The intensive-margin effect is the average frequency, approximated here with the average of the hazard function weighted by the density at each bin. To obtain the extensive-margin effect, we first calculate the slope of the hazard function at each bin as the centered finite difference between subsequent bins. Second, we multiply the slope with the size of the misalignment and, third, we calculate a weighted average using the density weight of each bin.

The first row of Table 3 shows the overall impact effect of a permanent money shock in the euro area and in the US and in each four euro area countries. The second and the third rows decompose these into intensive-margin and extensive-margin effects (or state dependence). The fourth and fifth rows show the relative contributions of the two channels.

The table shows that the effect is higher in the US relative to EA4 because of both stronger intensive- and extensive-margin effects. The stronger intensive-margin effect is the consequence of a higher frequency of price changes. The stronger extensive-margin effect is the consequence of stronger state dependence, which is driven by two main factors: the slope of the hazard function and the average absolute size of gaps.[15] In the simple and realistic case of a symmetric and (piecewise) linear hazard function, the slope of the hazard is constant $|\Lambda'(x)| = \Lambda'$, where $\Lambda' > 0$ is a parameter. It is straightforward to see that in this case, the extensive-margin effect in Equation (5) is simply $\int x\Lambda'(x)f(x)dx = \Lambda' \int |x|f(x)dx$, which is the product of the slope of the hazard and the average absolute size of price gaps. The expression shows that the extensive-margin effect does not only depend on the slope of the hazard function but also the dispersion of the price gaps. Indeed, the extensive-margin effect is higher in the US even though the slope of the hazard function is somewhat higher in the EA4. The reason is that the higher dispersion of the price gaps more than compensates for the lower slope. Higher product-level volatility through its impact on higher price-gap dispersion, therefore, raises the state dependence of price setting. The relative contribution of the extensive-margin effect is around 25 percent in both the euro area and the US. This means that accounting for state dependence raises the price-level flexibility by around 33%=25%/(1-25%) relative to a time-dependent benchmark (**?**).

As Table 3 also shows, there is a sizable heterogeneity among euro area countries in the size of the overall impact effect. The heterogeneity is mainly driven by differences in the intensive-margin effect, determined by the frequency of price changes. The extent of state dependence among euro-area countries is similar, with the notable exception of France, where

[15] There is a potential third factor: the covariance of the slope and the gap, but this factor plays a marginal role in the realistic case of an approximately linear hazard function (see **?**).

it is substantially below average.

Table 3: Overall impact effect and absolute and relative contributions of adjustment margins

Margins	**US**	**EA4**	**DE**	**FR**	**IT**	**NL**
Overall impact effect	18.4%	11.6%	8.6%	15.3%	12.3%	14.0%
Intensive (absolute)	13.9%	8.8%	5.3%	13.1%	9.0%	10.7%
Extensive (absolute)	4.5%	2.7%	3.3%	2.2%	3.3%	3.3%
Intensive (relative)	75.4%	76.4%	61.8%	85.4%	73.1%	76.3%
Extensive (relative)	24.5%	23.6%	38.2%	14.6%	26.9%	23.7%

*Note: The table presents the overall impact effect of a marginal money shock and the absolute and relative contributions of the intensive- and extensive-margin effects (**?**). The table shows that stronger state dependence (extensive-margin effect) amplifies price flexibility in the US relative to EA4.*

5. Structural analysis

In this section, we interpret the evidence through the lens of a price-setting model (**?**). We ask which structural features drive the differences in price setting in the food retail sector between the US and the euro area and across euro area countries.

5.1. Structural model

The model we use is a quantitative price-setting model with price-adjustment costs and information frictions. It provides a micro foundation for the popular 'random menu cost' models (**??**) and includes the time-dependent **?** model and the fixed-menu-cost model of **?** as special cases.

We sketch the key features of the model here and direct the interested reader to the original paper for details and derivations. The paper generalizes the fixed menu cost model of **?**. There is a continuum of differentiated goods (i), which are sold in a market with monopolistic competition. This market structure gives the producer of each good market power to set prices at a markup above the marginal cost. The market power is determined by the elasticity of demand, which, in turn, is governed by the (constant) elasticity of substitution parameter ε.

The production requires labor, and the product-specific productivity is subject to idiosyncratic shocks. As argued by **?**, these shocks are necessary to explain the large absolute size of price changes. Specifically, productivity follows a random walk, with an idiosyncratic shock $z_t(i)$ with standard deviation σ_z ($A_t(i) = A_{t-1}(i) + z_t(i), z_t(i) \sim N(0, \sigma_z^2)$). All the relevant firm-level information is incorporated into the price gap, defined as the distance of its (log) price from its (log) optimal price $x_t(i) = p_t(i) - p_t^*(i)$.[16] In particular, its profit is a function of the price gap, and it is maximized when the price gap is zero. The price gap fluctuates as idiosyncratic shocks hit the optimal price, and the firm does not necessarily reset it to zero because the adjustment of the product price ($p_t(i)$) is costly.

The firms face two types of adjustment costs. First, as in **?**, the firm needs to pay a fixed (menu) cost κ in case it conducts a price review. After paying the cost, the firm obtains full information, thereby it learns its price gap and optimally closes it. Second, the firm needs to decide about the timing of its price review under imperfect information about the state of the economy, therefore about its price gap. The imperfect information is modeled as rational inattention, whereby the firm can obtain a costly signal $f(x)$ about the price gap, and the cost increases linearly with the informativeness (I) of the signal with a coefficient θ ($\theta I = -\theta E[\log f(x)]$)). **?** establishes two useful results. First, the optimal policy is described by a hazard function $\Lambda_t(x_t)$: a firm chooses to obtain a signal with probability $\Lambda_t(x_t)$ as a function of its price gap x_t and conducts a price review in case it receives a signal. Second, the functional form of the hazard function is well defined, it is (weakly) increasing with the (absolute value of the) price gap, and its slope depends on the information cost parameter θ. As the cost parameter $\theta = \infty$ increases without limit, the hazard function approaches a constant, which is the time-dependent **?** case, and the cost parameter is zero $\theta = 0$, the hazard function approaches a step function as in the fixed menu cost **?** case. In between these two extremes, the theoretical hazard function shares some key features of the empirical hazard functions shown in Section 4. In particular, it is increasing with higher absolute gaps, implies a positive hazard at zero gap, and is asymmetric with a higher probability of adjustment when prices are below the reset price.

We assess the dynamic impact of aggregate fluctuations in the model. We do this by approximating the aggregate equilibrium conditions up to a first order around the non-linear stationary equilibrium using the method proposed by **?**.[17] As in **?**, we assume that aggregate nominal expenditure equals the money supply $P_t Y_t = M_t$. Money growth follows an exogenous autoregressive process $g_{Mt} = \rho_M g_{Mt-1} + \varepsilon_{Mt}$, with $\varepsilon_{mt} \sim N(0, \sigma_m^2)$. Money shocks

[16] It is easy to see that the price gap can be equivalently expressed as the difference of the normalized price ($q_t(i)$) as defined in **?** from its optimum ($x_t(i) = p_t(i) - p_t^*(i) = q_t(i) - q_t^*$).

[17] We extend the code used in **?**. We thank Anton Nakov for posting his code.

can alternatively be interpreted as nominal expenditure shocks. We assume, furthermore, that the production of each product i is affected by an aggregate productivity factor A_t.[18] Aggregate productivity follows a first-order autoregressive process $A_t = \rho_A A_{t-1} + \varepsilon_{At}$, with $\varepsilon_{At} \sim N(0, \sigma_A^2)$.

5.2. Estimation

Our goal in this section is to identify the most relevant structural features, which account for the differences between the US and EA4 price setting and across euro area countries. We do this by matching the empirical moments obtained in previous sections to estimate key structural parameters in the model.

We calibrate some parameters to levels used in the literature following **?**, with one difference. We set the elasticity of the substitution parameter (ε) to 3. This is the parameter used by **?**, which implies markup levels relevant for supermarkets.[19] Furthermore, we calibrate the autoregressive coefficient of the money growth process (ρ_m) to 0.61 as **?**, and the autoregressive coefficient of the aggregate productivity ρ_A to 0.95, which is a standard value in the literature.

The three parameters we estimate are (i) the standard deviation of the idiosyncratic shocks (σ_z), which affects the volatility of the product-level environment, and the two parameters governing the price-adjustment costs: (ii) the review (menu) cost (κ) and (iii) the information cost (θ). We estimate these parameters by targeting three empirical moments with their simulated counterparts in the stationary equilibrium: the shape of the generalized hazard[20], and the frequency and size of the price changes.[21] We also check how the model matches some untargeted moments, like the duration hazard and the standardized price-change distribution.

[18] Formally, $y_t(i) = A_t A_t(i) h_t(i)^{1/\phi}$, where $y_t(i)$ is the (log) output of firm i, A_t is an aggregate, $A_t(i)$ is firm-specific productivity, $h_t(i)$ is firm-level labor, and ϕ is a parameter governing the extent of decreasing returns to scale of labor.

[19] The parameter is below that used by **?** (6). The lower parameter implies weaker competition and a flatter profit function, which helps us to match the consistently low slope of the empirical hazard function.

[20] The estimation algorithm minimizes the squared difference between the empirical and the theoretical hazard functions, weighted by the price-gap density. The empirical hazard, which is calculated as the probability of price change *next month* as a function of the current price gap, is matched with a simulated hazard, which similarly expresses the probability of price change in the following period as a function of the current gap.

[21] For internal consistency of our quantitative exercise, the frequency and the size measures we match here are derived from (unweighted, truncated at +-50%) generalized hazard and density estimates. In particular, frequency is measured as $\sum_j \Lambda_j f_j$, and size as $\sum_j |x_j| \Lambda_j f_j / \sum_j \Lambda_j f_j$, where Λ_j is the height of the generalized hazard, f_j is the relative share of products in the price-gap bin j, and x_j is its midpoint. These measures are not equal to the (weighted, untruncated) frequency (EA4: 5.5% vs 8.9%, US: 13.8% vs 14.1%) and size (EA4: 9.3% vs 8.0%, US: 15.2% vs 11.8%) measures reported in Sections 3.1 and 3.2, but they are close with comparable relative magnitudes.

5.3. Results

Figure 4 shows the match of the theoretical and empirical generalized hazards and densities for EA4 and the US. The fit is good for both the hazards and the densities, especially over the range where most of the mass concentrates, as indicated by the shaded areas. The distribution of the gaps in the euro area is more concentrated than in the US, which the theoretical distribution can only partially capture. Specifically, even though the theoretical distribution captures the lower standard deviation of the gap distribution (a moment, which is targeted through the size of price changes - the inverse price gap) in the EA4, it underestimates the kurtosis of the distribution. One reason for the high kurtosis can be the heteroskedasticity of price gaps across products. Indeed, the fit to the *standardized* price-change distribution (see the second row of Figure 4), which controls for this heterogeneity, is superior, and actually somewhat better for the euro area than for the US. The model is also reasonably good at matching the duration hazard, even though the moment was not directly targeted. In Appendix H, we show that a more realistic model, which assumes an asymmetric linear hazard function and leptokurtic idiosyncratic shock distribution (Laplace) in the euro area can match the hazard function, the shape of the gap distribution and the price-change distribution (including its unimodality in the euro area and bimodality in the US) better. However, it has essentially the same response to an aggregate money growth shock as the baseline model, suggesting that the fit of the baseline model is sufficiently good.

Table 4: Estimated parameters

Parameters	**US**	**EA4**	**DE**	**FR**	**IT**	**NL**
Review cost (κ)	21.6%	15.9%	22.0%	5.4%	20.1%	21.2%
Stdev. of idiosyncratic shocks (σ_z)	5.5%	3.0%	2.7%	2.1%	3.6%	3.9%
Information cost (θ)	0.65	0.65	0.4	0.54	0.58	1.07

Note: The table shows that state dependence is present, but mild in both regions (information frictions are high). Higher idiosyncratic-shock variation in the US plays a prominent role in explaining higher frequency and size of price changes.

Table 4 shows the estimated structural parameters for the euro area, the US, and the euro area countries. Several results are worth pointing out. First, the information cost parameters are finite, indicating the presence of state dependence in line with an increasing hazard function. Second, the information costs are sizable, indicating a mild state dependence, which is quantitatively closer to the time-dependent **?** model than the strongly state-dependent fixed-menu-cost **?** model. This is in line with flat hazard functions. Third, the information cost

Figure 4: Estimation, targeted and non-targeted moments

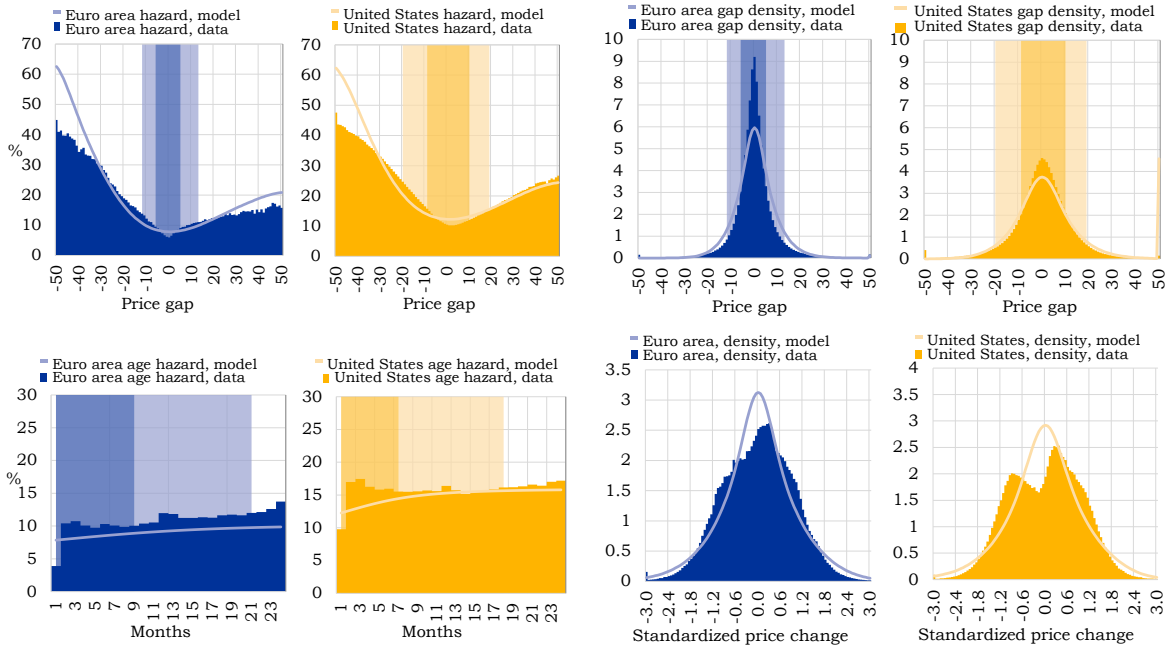

Note: The figures show the match of the simulated and the empirical generalized hazards and price-gap densities (matched moments, first row) and the duration hazard and price-change densities (unmatched moments, second row) in the euro area (EA) and the US. Shaded areas cover the 67% (darker) and 90% (lighter) mass of the corresponding densities.

parameters are actually the same across the regions, indicating similar levels of state dependence. Fourth, the estimated review (menu) cost parameters are somewhat higher in the US. Finally, fifth, the quantitatively most relevant structural reason for the differences across the regions is the distinct standard deviation of the idiosyncratic shocks. Figure 5 shows impulse responses to a money growth shock in models calibrated in the euro area and in the US and in a simulation of the euro area with US counterfactual idiosyncratic volatility. The figure shows that the inflation response is more flexible, and the output response is correspondingly larger and more persistent in the US, and the difference in idiosyncratic volatility is responsible for most of the difference. Therefore, the volatility of the product-level environment is higher in the US, which leads at the same time to (i) higher frequency of price changes, (ii) higher size of price changes, and (iii) more dispersed price gap distribution contributing to stronger aggregate responses both at the intensive and the extensive adjustment margins.

Figure 5: Simulated response to a money growth shock

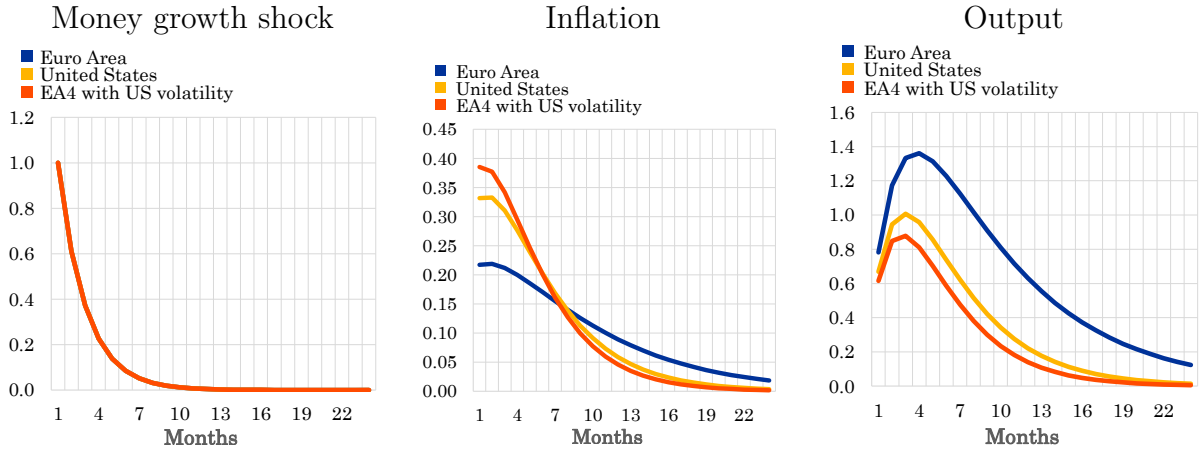

Note: The figures show the impulse response to a unit money growth shock in the model calibrated to the euro area, to the US, and the euro area with US idiosyncratic volatility. The figures show that the inflation response is weaker in the euro area, and the output response is correspondingly larger and more persistent. The difference is primarily driven by the higher idiosyncratic volatility in the US.

Among the euro area countries, the heterogeneity in the standard deviation of idiosyncratic shocks is also sizable and plays a key role in explaining the overall heterogeneity in price flexibility. The extent of price-adjustment frictions (κ) is similar across countries and comparable to the US, with a notable exception of France, where prices are estimated to be much more flexible. The information cost parameter varies somewhat, but within a limited range, implying mild state dependence in all euro area countries.

What share of the difference in the volatility in food inflation between the US and the EA4 is explained by the model? The question is important because even if the model can successfully capture differences in key features of price setting, different volatilities across countries can be the consequence of numerous unmodelled factors, like the heterogeneity in the nature and magnitude of aggregate or sectoral shocks. We assess the importance of differences in price setting through a simple exercise. First, we estimate the standard deviation of aggregate money shocks σ_m to match the standard deviation of year-on-year inflation in the US in the model and the data. Then we measure the predicted inflation volatility in the model if we feed the same money shocks to the model calibrated to the euro area. We repeat the same exercise with aggregate productivity shocks (σ_A). The first three columns of Table 5 show the results. As the second row shows, assuming money (or nominal expenditure) shocks, the model correctly predicts lower inflation volatility in the euro area, but it only

explains somewhat over a third of the difference observed in the data. As the third row shows, assuming productivity shocks, the model explains most of the difference in inflation volatility. It is outside of the scope of the paper to assess the importance of the two shocks in driving food inflation, but arguably both types of shocks are active. We conclude that differences in price-setting frictions, though not the only factor at play, can account for a relevant share of the observed difference in inflation volatility.

Table 5: Simulated inflation response

	Inflation volatility			Covid shock		
	US	**EA4**	**Difference**	**DE**	**IT**	**Difference**
Data	0.95%	0.64%	0.31%	0.54%	1.65%	1.11%
Model (money)	0.95%	0.83%	0.12%	0.92%	1.31%	0.39%
Model (tfp)	0.95%	0.66%	0.29%			

Note: The table shows the inflation volatility in the US and EA4 (columns 1-3) and the inflation response to the Covid shock in Germany and Italy (columns 4-6) in the data and in the model. Around one-third of the observed difference is accounted for by the model.

6. Price setting during Covid-19 in Germany and Italy

In this section, we analyze the price-setting response of German and Italian supermarkets to the first wave of the Covid-19 lockdowns. The shock had a large, persistent, and broadly similar effect on supermarket demand in both countries. Contrasting the response in the two countries is relevant because price setting is heterogeneous across the countries: the frequency of price changes is structurally higher in Italy than in Germany (see Table 2), and the extent of state dependence is similar across the countries (see Table 3). Price-setting models, therefore, predict higher flexibility of the Italian supermarket inflation, which we can test in the data.

6.1. Data

The analysis in this section uses an auxiliary dataset covering large German and Italian supermarkets over three months, encompassing the first wave of the Covid-19 pandemic from mid-February till mid-May in 2020. The dataset also covers the analogous period in 2019,

which we will use as the base period in our index calculations. The dataset covers 20 2-digit ZIP areas.[22]

Our analysis uses the 2013-2017 German and Italian pre-Covid sample as a benchmark to assess the significance of changes observed over the 2019-2020 Covid period. To minimize the impact of compositional shifts over time, we restrict our baseline sample to stores and products which appear with positive sales in both the first quarter in 2013 and the sample quarter in 2020. The majority of stores are such 'established' stores[23]. A sizable fraction of the products is such 'established' products[24].

6.2. Supermarkets and the first wave of the Covid-19 pandemic

The Covid-19 pandemic and the accompanying lockdown measures had a large and persistent impact on supermarket demand. During the lockdowns, access to food away from home was severely restricted, while supermarkets were deemed essential and sheltered from the impact of the lockdowns. The Italian government imposed a national lockdown on 9 March and gradually eased it only after mid-May. In Germany, a federal lockdown was introduced on 22 March and was gradually eased from early May. In both countries, supermarkets stayed open during the lockdowns, while alternative access to food and beverages was restricted, as restaurants, canteens, and bars were deemed unessential and closed.

Our data allow us to quantify the magnitude of the demand change because the scanner data records weekly expenditures at the store-produce level. We restrict attention to established products in established stores, which are the focus of our analysis. We measure year-on-year nominal expenditure growth as the 52-week change in overall expenditure on items sold in positive quantities both in the current and base weeks.

The first row of Figure 6 shows the evolution of nominal expenditure growth between mid-February to mid-May in German and Italian supermarkets. The figure shows that the expenditure growth significantly exceeded its long-term average. The increase was particularly pronounced during the weeks preceding the introduction of the lockdowns. The growth rate reached as much as 19-29 percent during this 'stock-up shock,' as households increased their home stock of non-perishable groceries for precautionary reasons. The expenditure growth during the lockdowns stayed persistently well above average. It stabilized by the end of

[22] The ZIP areas in the sample cover 16% and 40% of the population and a share of supermarket expenditures of 22% and 46% throughout 2013-2017 in Germany and Italy, respectively.

[23] 668 out of 815 stores in Germany and 1486 out of the 2387 stores in Italy.

[24] 57.000 out of 266.000 products in Germany and 83.800 out of 535.500 products in Italy with an expenditure share in Germany of 43.43% and in Italy of 42.43%.

Figure 6: Nominal expenditure growth and reference price inflation in supermarkets during the Covid-19 pandemic, year-on-year

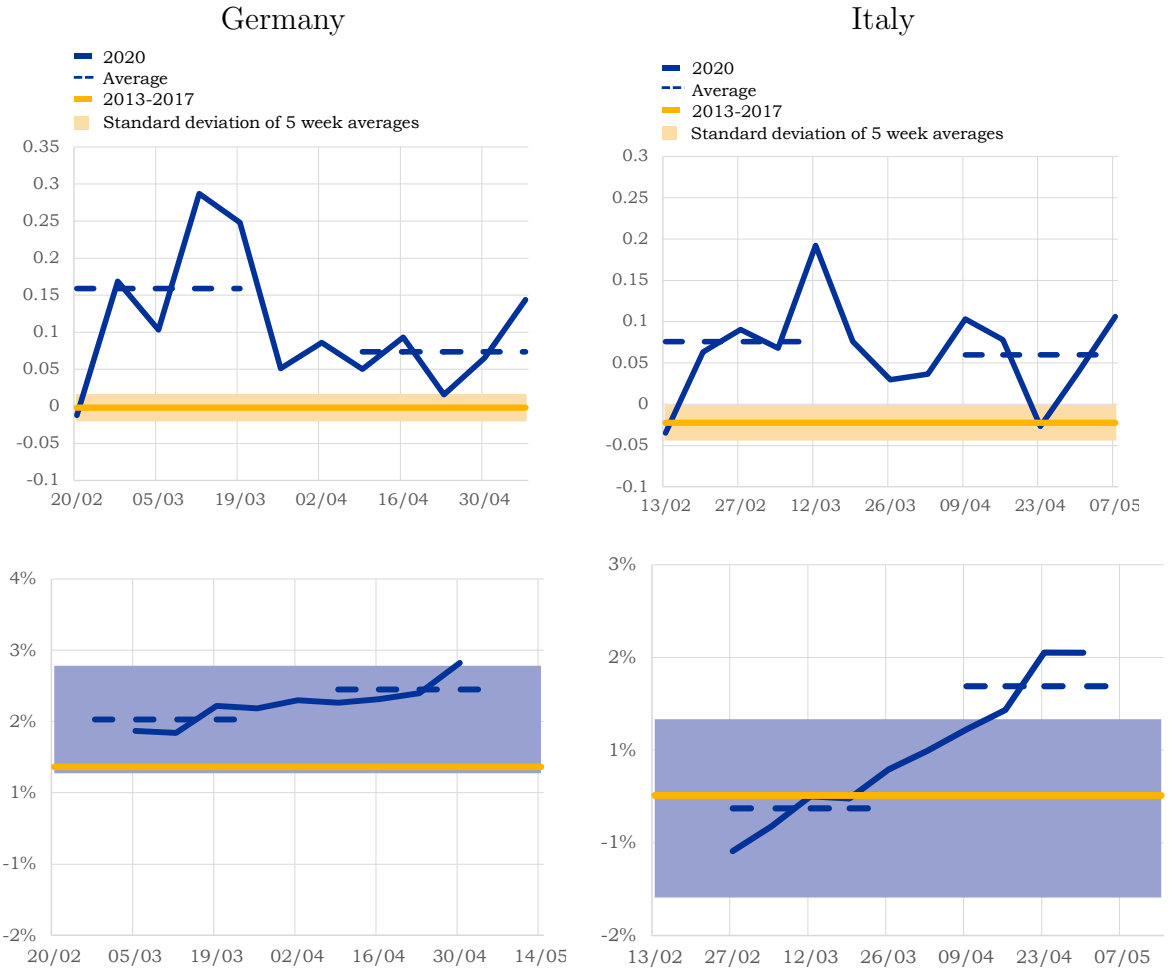

Note: The figure shows the weekly, year-on-year nominal expenditure growth (blue solid line) between mid-February and mid-May in 2020 in Germany and Italy (first row) and weekly reference price inflation (second row). It shows that the 5-week-average expenditure growth (blue-dashed line) exceeded the average long-term expenditure growth (yellow dashed lines) by more than a standard deviation in both Germany and Italy. It shows that the increase in the average 5-week-inflation over the quarter was smaller (0.54%) and within a +- one-standard-deviation band in Germany, while it was three times as large in Italy (1.65%) and clearly exceeded the band.

our sample at around 7.4 percent in Germany and at 6 percent in Italy, which significantly

exceeded the long-term nominal expenditure growth experienced over the 2013-2017 period.

6.3. Inflation response

How did prices respond to the persistent increase in demand? The second row of Figure 6 shows the evolution of reference prices in Germany and Italy (see Appendix I for the evolution of posted-price inflation). It shows that the increase between the first and the third month was substantially lower in Germany (0.54%) than in Italy (1.65%), albeit from a higher initial level. The increase in Germany stayed within a +- one standard deviation confidence band, while it clearly exceeded the band in Italy. This happened despite the similarity of the shock, the type of retailers, and the basket of products.

Can differences in price setting account for the differences in the inflation response? Previously we have shown that prices change less frequently in Germany (see Table 2), while the extent of state dependence is similar across the two countries (Table 3). The structural model calibrated to the two countries rationalized this outcome with comparable price-setting frictions, but substantially smaller product-level volatility and somewhat lower information frictions in Germany. Table 5 shows the simulated impact of a permanent 6% nominal expenditure shock on the 2-month cumulative inflation (equivalent to a change in the year-on-year inflation between the 3rd and the 1st month in the data) in the two models calibrated to Germany and Italy. The model reproduces the more sluggish inflation response in Germany and accounts for around one-third of the difference between the countries.

7. Conclusion

The paper contrasted price settings in the euro area and the US using a novel supermarket-scanner dataset in four euro area countries and in the US. It has found that the higher flexibility of food inflation in the US is driven both by the higher frequency of repricing and the stronger state dependence of price changes. It argues that the driving force behind both factors is a more volatile product-level environment in the US. The models were able to explain over a third of the differences in inflation volatility across the regions.

Our conclusions have implications for both model selection and policy. First, the evidence is in line with models with sizable nominal rigidities in both regions, which amplify the impact of monetary and fiscal policy on the real economy. The higher nominal rigidities in the euro area imply that, at least in the food sector, changes in nominal expenditure growth have a smaller impact on prices and a larger impact on quantities than in the US. Second, the evidence presented in the paper supports state dependence in price setting. Even though we find that the estimated magnitude of state dependence has a mild impact on price flexibility in response

to small aggregate shocks, state dependence necessarily implies that prices endogenously become more flexible after large aggregate shocks and higher trend inflation (**???**). Third, the sizable differences in the implied product-level volatility between the US and the euro area raise important questions for future research. Although in the simplest class of price-setting models, product-level volatility matters only insofar as it influences frequency and state dependence (**?**), in more complicated models, it can have an independent impact on price flexibility, as high product-level volatility can make retailers limit their attention to aggregate fluctuations (**?**), which could mitigate their responsiveness to aggregate shocks. Its key role in driving differences across regions also highlights the importance of further research to understand better the underlying sources of the product-level volatility, including whether they are the consequence of larger shocks or higher responsiveness to these shocks (**?**).

Online appendix

Appendix A. Official vs IRi expenditure distribution

Figure A.7 contrasts the expenditure distribution within the IRi sample across nine product categories in each country, and the corresponding expenditure share of the same category in the official HICP (EA4) and CPI (US) indexes.[25] Some categories, like bread, meat, fruits, and vegetables are somewhat underrepresented in the IRi supermarket samples, as these products are regularly purchased also from specialized stores. The match is less tight in the US sample, which only includes a selected set of product categories (for example, only processed sausages 'Frankfurters' as meat products).

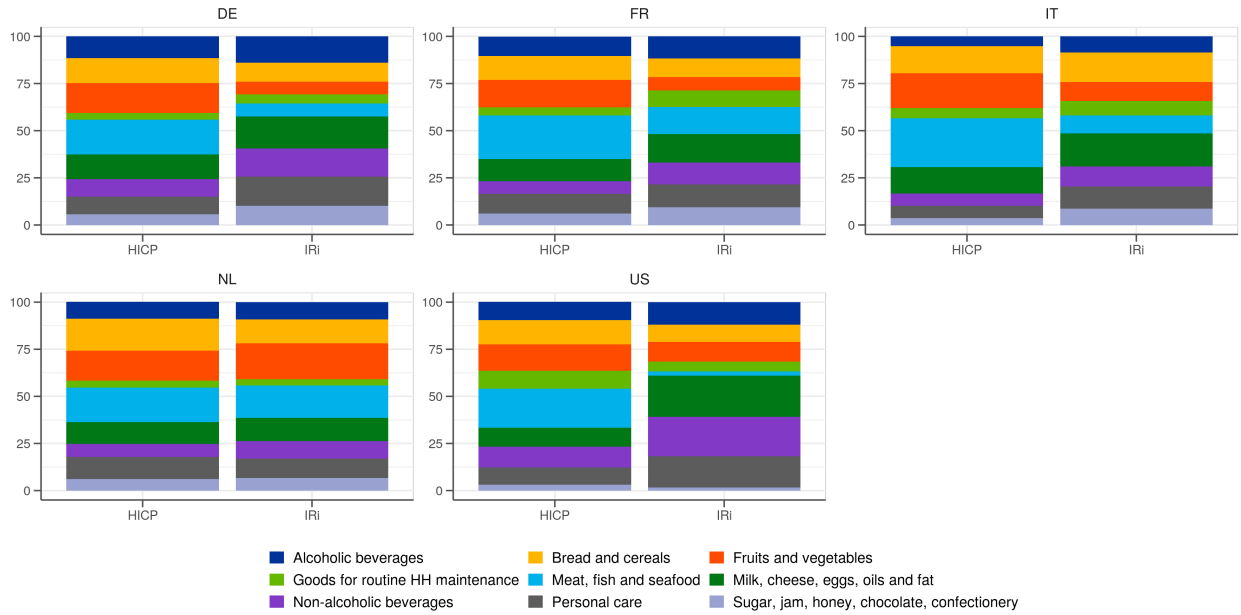

Figure A.7: Official vs IRi expenditure shares by category

[25] The nine categories are constructed to represent large, but still fairly homogeneous groups of products with a sizable share in our sample across all five countries. They are constructed as a suitable combination of 3-digit and 4-digit COICOP categories. We use categories (EA4) and subcategories (US) provided by IRi to allocate products into the nine categories.

Appendix B. Construction of store projection weights for Germany and Italy

In the German and Italian samples, some supermarket chains share a representative sample (as opposed to a full census as in the US, France, and the Netherlands) of their stores with IRi. To maintain the representativeness of our sample, we need to adjust the weight of the sample stores suitably.

In particular, we 'upweight' stores (s) by projection weight ν_s. The weights for the stores that appear as census are $\tilde{\nu}_{sw} = 4/3$, which offsets the impact of the dataset being a 75 percent share of the full sample.

To obtain projection weights for sample stores in Germany and Italy, we need to estimate the overall number of stores by store type. This data is part of the IRi dataset.[26] In Italy, the overall number of stores by store type is available annually at the end of each year: here, we use linear projections between end-of-year observations to obtain estimates of the weekly number of stores by store type $\tilde{N}_{Sw}$. In Germany, the number of stores by store type is only available at the end of 2017 (N_{ST}). Here we use the evolution of the stores in the census population in our sample to obtain estimates of the weekly number of stores by store type. First, we calculate average weekly entry (γ_S) and exit rates (δ_S) by census store type (S) over the 2013-2017 sample. Second, the estimated number of stores by store types are obtained by assuming constant entry and exit rates by store type $\tilde{N}_{Sw} = (1 + \gamma_S - \delta_S)\tilde{N}_{Sw-1}$, $\tilde{N}_{ST} = N_{ST}$.

The weights for sample stores are obtained as

$$\tilde{\nu}_{sw} = \frac{4}{3}\frac{\text{non-census population by type in week w}}{4/3 \cdot (\text{non-census sample by type in week w})}, \tag{B.1}$$

where (non-census population in type in week w) = (estimated aggregate population by type in week w ($\tilde{N}_{Sw}$))-4/3(the number of census stores by type in week w).

Finally, we normalize the projection weights to make sure they sum to unity each week:

$$\nu_{sw} = \frac{\tilde{\nu}_{sw}}{\sum_s \tilde{\nu}_{sw}}, \text{for each } w. \tag{B.2}$$

Appendix C. Posted-price filter

Unit-value prices do not necessarily reflect posted prices. There are two main reasons for this. First, mid-week price changes generate unit values that are in-between actual prices.

[26] The number of stores is available by store type and geographical area, but the latter we ignore in the current analysis.

Table C.6: Data-cleaning moments

	DE	FR	IT	NL	US
% same-direction changes	2.15	5.39	8.1	3.58	6.03
% also fractional	1.66	3.71	5.36	1.65	3.31
% fractional price	7.6	8.05	11.66	5.91	6.96
% below closest integer	68.93	53.83	59.48	62.33	58.95
% missing (obs)	43.91	42.09	46.58	42.97	38.49
% missing (exp)	55.37	46.88	59.12	38.9	55.5
% missing (exp >4w)	22.16	21.04	26.27	16.7	13.49

Note: 'Missing (exp)' refers to the expenditure share of products that record zero sales in a single (over four consecutive) week(s).

Second, coupons and other buyer-specific discounts can reduce the average revenue from a product below its posted price. We transform unit prices to estimated posted prices using the following filtering rules.[27]

First, to reduce the impact of mid-week price changes, we filter out same-direction consecutive price changes. A one-time mid-week permanent price change necessarily generates such same-direction consecutive unit-price changes. The mid-week price change increases the average weekly unit price only partially in the initial week, and pass-through fully only during the second week. Formally, if we observe two consecutive price increases ($I^{+}_{psw,w-1} > 0, I^{+}_{psw+1,w} > 0$) or decreases ($I^{-}_{psw,w-1} > 0, I^{-}_{psw+1,w} > 0$), we conclude that there was a mid-week price change during week w. We set the end-of-the-week posted price during this week as the unit-value price in the following week $P_{psw} = P^{uv}_{psw+1}$. As Table C.6 shows, 2-8 percent of the prices are affected by the same-direction filter. Out of these filtered prices, usually, over half are fractional (fractions of a cent). As fractional unit values cannot be posted prices, their presence strongly confirms mid-week price changes. Their high share suggests that the filter recovers the true posted prices in most cases. And even though some of the filtered same-direction price changes could have been true adjustments, filtering them out biases our results only marginally, especially at the monthly frequency, which is going to be our focus.

Second, to mitigate the impact of buyer-specific discounts, we round prices upwards to the nearest cent. Posted prices need to be integers in cent units. However, 6-12 percent of unit-value prices are fractional even after controlling for same-direction price changes (see Table

[27] There is a potential third reason: a within-week temporary price discount. These within-week price changes would be recorded as (smaller) changes in the weekly average price, potentially distorting the price-setting moments at the highest frequencies. As our focus is monthly frequency, we do not expect such changes to influence our conclusions.

C.6). As most of the deviations from the posted price result from discounts, we round the prices upwards. Indeed, the closest integer is higher than the price in over 60 percent of the fractional prices. A higher than 50 percent share is expected when the fractional prices are caused by discounts paid by a small fraction of the buyers. In cases when the discounts are paid by such a small fraction of the buyers that the unit-value price deviates from the posted price by at most a cent, our filter picks up the actual posted price. Even when the discounts reduce the average price by more than a single cent, upward rounding brings us closer to the posted price. However, there can still be many cases when the share of buyers paying a discount is large enough to reduce the unit-value price away from the posted price by more than a single cent. In these cases, the filter does not recover the actual posted price. Therefore, we show the robustness of our results below when we exclude fractional prices from the analysis.

The posted price also remains unobserved when there is no sale of the item in a particular week. Zero-sale weeks ($I^{\emptyset}_{psw}$) for existing items[28] are frequent in the data. In particular, 38-47 percent of the observations are missing (see Table C.6). Furthermore, the expenditure share of items with missing observations is also high, so the issue does not only affect rarely-sold unpopular products with a small expenditure share. In particular, the annual expenditure share of products with at least one missing observation over a year is 40-60 percent. The issue is somewhat less pressing if we realize that consecutive missing observations are usually short, much shorter than a month. In particular, the annual expenditure share of products with at least one case of 4 or more consecutive missing observations is between 15-25 percent. This is the relevant metric in our analysis, which focuses on monthly price developments: as monthly prices are missing only if weekly prices are missing for 4 weeks consecutively. The presence of a not insignificant fraction of missing prices is still can be considered a caveat of our dataset, and its potential impact needs to be carefully assessed in the analysis below.

The dataset requires careful treatment during the rare occasions when the product identifiers stop referring to the same product over time. This happens in the US sample in 2007:01, 2008:01, and 2012:01, when the identifiers of some private-label products get reassigned by IRi. We lack additional information about the rules followed during the reassignment, so, conservatively, we assume that new private-label products replaced old private-label products during these three months, and we do not link price spells of private-label products over these months. We treat similarly a subset of German beer- and beer-cocktail products in 2014:01, when their EAN got reassigned to refer to a crate instead of a bottle (which could

[28] We consider a product p in store s existing in week w if $M_{ps} \leq w \leq T_{ps}$, where M_{ps} is the date of entry (the first week when product p was sold in store s), and T_{ps} is the date of exit (the last week it was sold).

occasionally generate artificial 24-fold price increases): we treat them as separate products and do not link their prices over the 2014:01 period. Lastly, we drop from the Dutch dataset over the 2013-2014 period a subset of (overwhelmingly fresh) products, which had inconsistent unit treatment resulting in unreliable price development. In particular, we drop products with 'internal use' EANs and 'random weight' volume measurements over the 2013-2014 period. The internal use EANs are assigned by the stores to products packaged internally (e.g., fresh meat). The 'random weight' volume measurement implied a non-standardized unit treatment before 2014, which could have resulted in random unit variation over time if the store changed its reporting. To avoid artificial variation in our data, we drop these products from the analysis. The treatment impacts a small subset of the products (around 12 percent share of annual expenditures) over only two years and only in the Dutch data, so we expect it to have a marginal impact on our analysis.

Appendix D. Inflation

In this section, we construct an inflation index and compare its dynamics to the official food-at-home inflation subindices.

Our baseline inflation index is constructed as a geometric average of price changes weighted by their annual expenditures. Formally:

$$\Pi_t = \prod_{ps} \left(\frac{P_{pst}}{P_{pst-1}} \right)^{\omega_{pst-1,t}}, \tag{D.1}$$

where Π_t is the gross inflation rate in month t, and P_{pst} is the posted-price of product p in store s in month t. The weight is the annual expenditure on product p in store s as a share of the annual expenditures.

Formally, $\omega_{psy} = \sum_{t \in y} TR_{pst} / \sum_{ps} \sum_{t \in y} TR_{pst}$, where TR is the total revenue (nominal expenditure), and y is the year of month t.[29] The price index $P_t = \prod_{s=0}^{t} \Pi_s$ is simply a chained product of the inflation index.

[29] To arrive at monthly expenditures TR_{psm}, we transform weekly expenditures as

$$TR_{psm} = \frac{52}{12} \frac{\sum_{w \in m} \sum TR_{psw}}{\sum_{w \in m} 1}, \tag{D.2}$$

where the normalization controls for the number of weeks in the month (either 4 or 5). We first divide the sum of expenditures by the number of weeks in the particular month and then multiply it by the average number of weeks in the year. The index only considers items, which exist both in periods $t-1$ and t, therefore

Figure D.8: The year-on-year change of the IRi supermarket inflation index, and the official food and beverage subindexes.

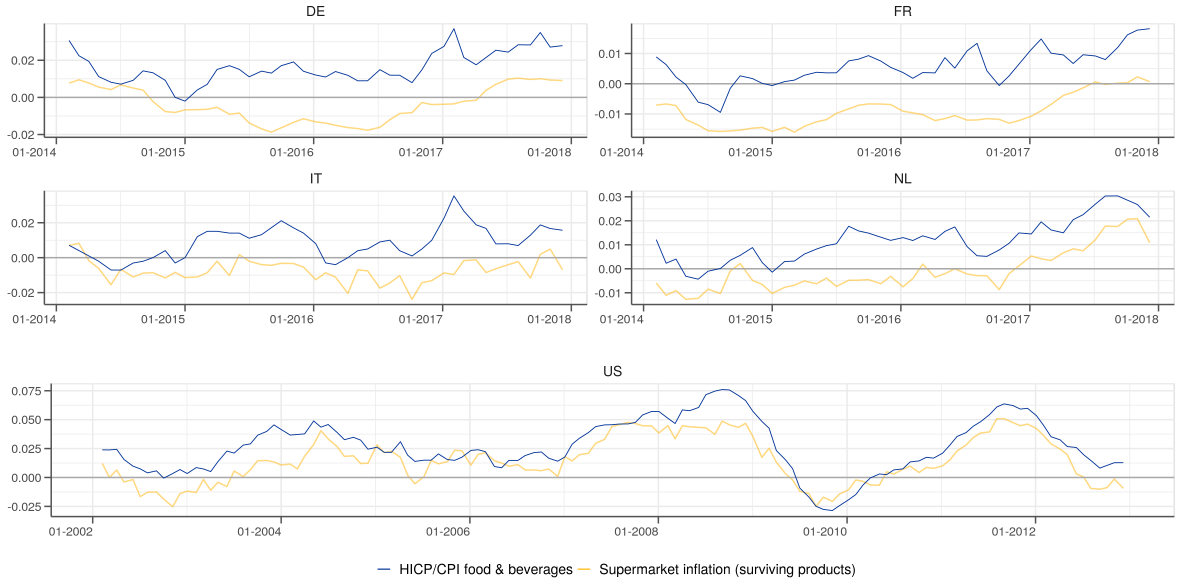

Note: The figure shows the evolution of the year-on-year IRi supermarket inflation and the official food and beverage subindexes. The comovement between the series is apparent, especially at low frequencies.

Our baseline inflation index captures the business cycle fluctuations in official food and beverage inflation reasonably well. Figure D.8 shows the evolution of the year-on-year inflation rates of the two series in each country. The comovement is apparent, especially at low frequencies.[30]

the actual expenditure weights used are

$$\omega_{pst-1,t} = \frac{I_{pst-1,t}\omega_{psy}}{\sum_{ps} I_{pst-1,t}\omega_{psy}},$$

where $I_{pst-1,t}$ is an indicator function that takes the value 1 if product p in store s exists in both months $t-1$ and t.

[30] At the same time, our inflation index underestimates the level of official inflation. As we detail in the appendix, the primary reason for this is that our index excludes the impact of new product introductions. These tend to have a small impact on inflation variability at business cycle frequencies (see also **?**), but can substantially raise the level of inflation.

In the upcoming analysis, we decompose the baseline index into key components to establish relevant stylized price-setting facts.

Appendix D.1. Temporary sales

A salient feature of price spells is their high-frequency variation. Prices regularly get reduced (or increased) temporarily for a couple of weeks, after which they tend to return to exactly the initial price. As we show momentarily, most price changes in our sample are due to such temporary sales. Previous research has established that the nature of high-frequency price changes is distinct from those of more persistent reference price changes (**???**). While reference prices are driven primarily by costs, sales are used as a marketing tool to trigger households to try out new products and stores, to gain the trade of bargain-hunter households, as well as a tool to fine-tune inventory. The high-frequency variation influences inflation dynamics differently than the evolution of reference prices, therefore, it is instructive to analyze them separately.

We employ the reference price filter of **?** to filter out temporary sales. As in **?**, we iteratively update the reference prices to align the reference-price change with the posted price change. We do this because a change to a new reference price is sometimes picked up with a delay by the algorithm (it takes a while till the new price becomes a mode within the rolling window). The algorithm corrects this by aligning the change in the reference price with the change in the posted price. As an additional step, we control for clearance and introductory sales in the first and last 5 weeks of the spell (**?**). We do this by carrying forward (backward) the reference price in the 6th week before the last (after the first) price of the spell. A key advantage of the reference-price filter over a more conventional regular-price filter that controls for V-shaped temporary price cuts (**?**) is that it also controls for temporary increases (spikes) in price spells. Such increases can be rationalized, for example, by inventory management: higher prices temporarily reduce demand and make sure the store does not run out of the product until a new delivery arrives. Spikes can account for as high as one-third of high-frequency price changes (**??**).

Figure D.9 decomposes our baseline inflation series into a reference price inflation series and a sales inflation series. The reference price inflation series is constructed analogously to our baseline series (Equation D.1) with reference prices replacing posted prices. Sales inflation is defined as the difference between posted-price inflation and reference-price inflation. The figure shows that sales-price inflation is subject to large high-frequency noise and it explains a small share of inflation variability at business cycle frequencies. This is one of the key reasons why we concentrate on reference-price changes in our analysis.

Even though sales inflation accounts for a small share of inflation variability at business cycle

Figure D.9: The year-on-year reference-price and sales inflation

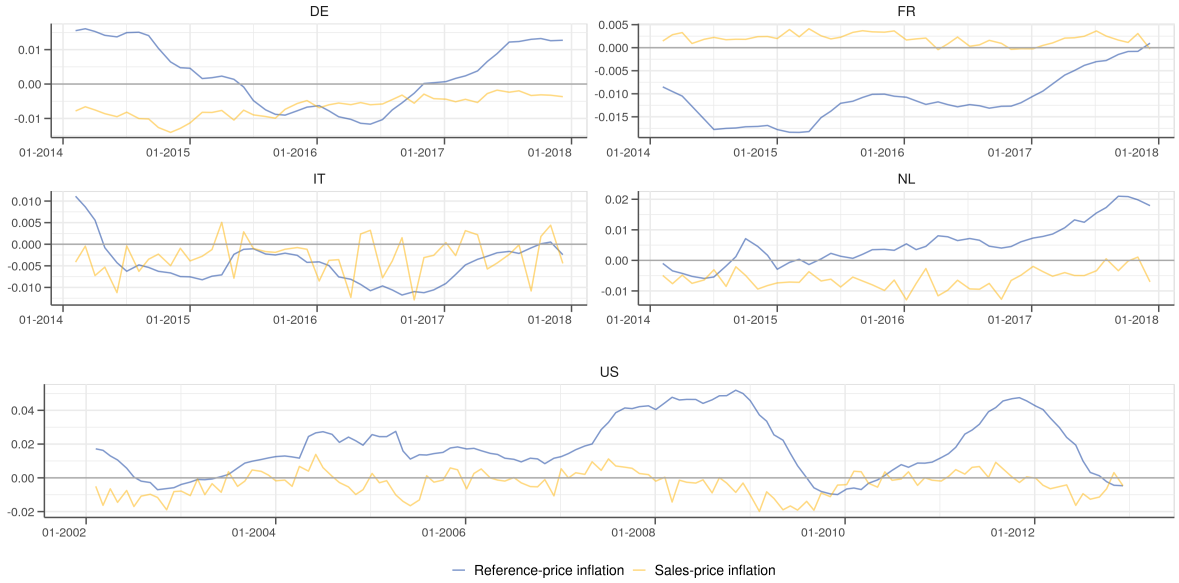

Note: The figure decomposes inflation into reference-price inflation and residual sales inflation. The figure shows that the reference-price inflation evolves smoothly and it accounts for most inflation variation at business cycle frequencies. Sales inflation, in contrast, varies little at business cycle frequencies, while it is subject to a sizable high-frequency noise.

frequencies, a large fraction of price changes are due to sales. To show this, we calculate the average monthly frequency of price changes for both posted- and reference prices. We weigh the item-level frequency with annual expenditure weights analogously to our baseline inflation index and take a simple average over time. Formally, the frequency (ξ_t) of monthly price changes is

$$\xi_t = \sum_s \sum_p \omega_{pst-1,t} I_{pst-1,t}, \tag{D.3}$$

where $I_{pst-1,t}$ is an indicator that takes the value 1 if the posted price of product p in store s in month t changed from the previous month, and 0 otherwise. The frequency of reference-price changes is calculated analogously with an indicator function that takes the value 1 in case of a reference-price change. Table D.7 shows the frequency of posted- and reference-price changes in the 4 euro area countries and the US in rows 1 and 2 and shows their ratio in row 3. The table shows that around 2/3 of price changes are due to sales and this share is fairly stable across countries.

Table D.7: Frequency of posted- and reference-price changes

Frequency (monthly, mean)	DE	FR	IT	NL	EA4	US
Posted	12.41%	25.47%	27.56%	24.77%	21.42%	39.35%
Reference	4.53%	15.27%	9.04%	10.06%	8.97%	13.34%
Ratio	2.74	1.67	3.05	2.46	2.39	2.95

Note: The table presents the frequency of posted- and reference price changes and their ratio. It shows that almost 2/3 of price changes are due to sales.

Appendix D.2. Features of the baseline inflation index

The construction of our baseline index resembles that of the HICP and CPI, which can also be characterized as chained, annual-expenditure-weighted price indices. One difference is that the weights are contemporaneous in our index, while HICP and CPI rely on lagged expenditures. The advantage of setting contemporaneous weights is that we do not need to restrict our analysis to products that exist also in the preceding year. This is a relevant advantage in the fast-moving supermarket-product category, where there is a sizable turnover between products. (Linking closely related products with distinct barcodes over time is out of the scope of this paper, while statistical offices put substantial effort into regularly replacing exiting with entering products after suitable quality adjustment.)

Annual weighting has multiple advantages over schemes that use more frequent weighing. First, because our focus is on price setting (as opposed to the measurement of the welfare-relevant inflation), annual weights minimize the impact of high-frequency quantity changes on the measurement of price dynamics, while still correctly tracking trend changes in the relative importance of different products. For example, it mitigates the seasonality of the index that more frequent weighting schemes would cause. This is particularly relevant among supermarket goods, where seasonal sales generate large seasonal variations in expenditures. Second, annual weighting mitigates a bias called chain drift, which impacts chained indices with high-frequency weighting schemes. Chain drift is present if the index does not return to 1 when the price returns to its initial level. One cause of the chain drift is the inventory behavior of households, who stock up during temporary promotions. As a result, the quantity of purchases declines below its initial level after it increased during the promotions. In the

presence of such dynamic behavior, even superlative chained indices (e.g., Tornqvist[31]) with high-frequency weights measure deflation even though the price returned to its initial level (**?**). Lower-frequency weights mitigate the impact of the chain drift by taking into account the longer planning horizon of the households (**?**) and bringing the index closer to fixed-base indexes, which are free of chain drift. To assess the remaining impact of chain drift on our index, we compare our index with the unchained year-on-year price index of existing products.

The results indicate that even annual weighting is not sufficient to completely eliminate chain drift from our baseline inflation measure. At the same time, the close to one correlation between the unchained and the chained series shows that the chain drift has an insignificant impact on the variation of measured inflation at business cycle frequencies, which is the main focus of our analysis.

Figure D.10: Chain drift: the year-on-year change of the baseline chained inflation index and the unchained 12-month inflation index.

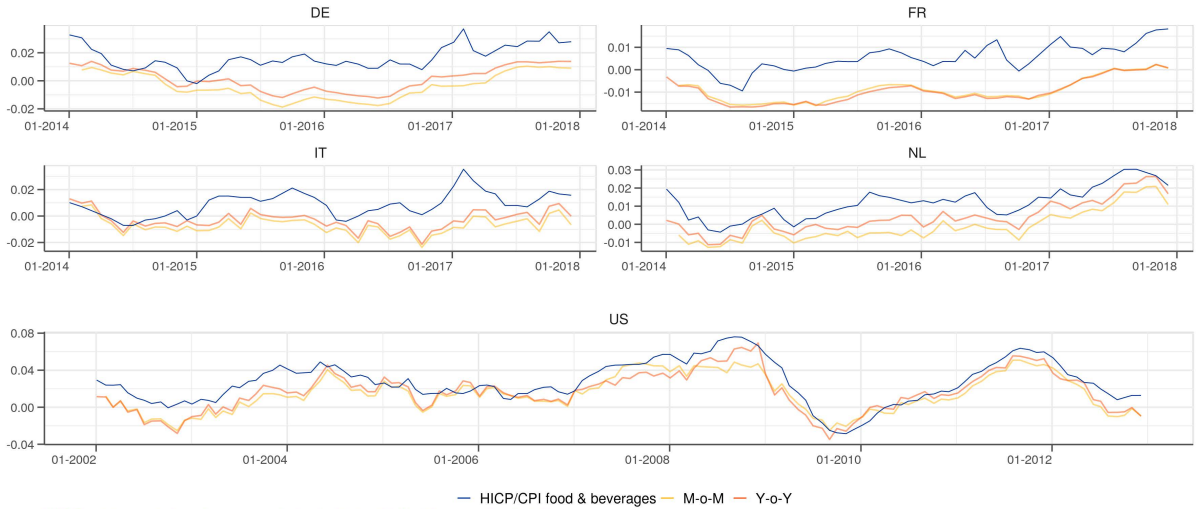

Note: The figure compares the year-on-year change of the baseline chained inflation index and the unchained 12-month inflation index. The level difference between the series indicates that our baseline inflation index suffers from some remaining chain drift, however, the close comovement of the series shows that the chain drift has an insignificant impact on the fluctuations of the series.

[31]The Tornqvist index between month t and $t-1$ equals to $\Pi_t = \prod_{ps} \left(\frac{P_{pst}}{P_{pst-1}} \right)^{(\omega_{pst}+\omega_{pst-1})/2}$, where ω_{pst} is the expenditure share of product p in store s in month t.

Table D.8: Chain drift: Average inflation of and correlation between the year-on-year change of the baseline chained inflation index and the unchained 12-month inflation index.

Series	moment	**DE**	**FR**	**IT**	**NL**	**US**
Inflation - chained 1m	average	-0.40	-0.92	-0.71	-0.10	1.47
Inflation - unchained 12m	average	0.13	-0.97	-0.31	0.39	1.67
1m-12m inflation	correlation	0.99	0.99	0.99	0.98	0.96

Appendix D.3. New-product introductions

A feature of the inflation index described by Equation (D.1) is that it includes the price change of existing products, but excludes the impact of product entry and exit on the price level. We justify our focus on the price setting of existing products by showing below that it is this component of the price level that is mostly responsible for its responsiveness to business-cycle fluctuations. The component of inflation caused by product entry and exit, which we call new-product inflation, in contrast, is broadly stable and insensitive to business-cycle fluctuations. Notably, this letter component, through the sizable difference between the price of the entering and exiting products, is what is responsible for the level of inflation.

The role of new-product inflation is not the focus of this paper. However, we find it useful to show some indicative evidence of its behavior by comparing the evolution of an 'all-product' inflation index to our baseline existing-product index. We measure the former as the monthly change in the price levels defined as the geometric average of prices weighted by the annual expenditure weights. Formally,

$$\Pi_t^{ap} = \frac{\prod_{ps} P_{pst}^{\omega_{pst}}}{\prod_{ps} P_{pst-1}^{\omega_{pst-1}}}. \tag{D.4}$$

where the weights ω_{pst} are annual expenditure-share of items that are present in month t. [32] The key difference between the existing-product and the all-product inflations is in the set of products they consider. The existing-product inflation only includes products that exist in both periods (the weights $\omega_{pst-1,t}$ are positive only for products that exist both in months $t-1$ and t), and the all-product price levels include all existing prices (the weights ω_{pst} are positive for all products existing in period t), therefore, the all-product inflation is taking account the price differences between exiting and entering products.

[32] Formally, the weights are given by $\omega_{pst} = I_{pst}\omega_{psy} / \sum_{ps} I_{pst}\omega_{psy}$, where I_{pst} is an indicator function that takes a value 1 if product p in store s is sold in positive quantities.

The all-product inflation is the relevant inflation index, if the exiting products are all replaced by similar quality entering products. This is admittedly a strong assumption. There are both 'true' product exits as well as 'true' product introductions without a matching entry or exit respectively in our dataset. But the replacements constitute arguably the majority of exits and entries.[33] These could include pure barcode changes (when the exact same product is reintroduced with a different barcode), changes in packaging, volume, or minor flavor/color/form upgrades. The true entries and the replacements require different treatment by the price index. With some reasonable assumptions about the utility function (including how quality affects demand), **?** and **?** show how the quality of true introductions can be assessed by relying on their market share relative to existing products. Intuitively, if their quality is high relative to their price (which is observable) their relative market share (which is also observable) is also going to be high. The same techniques, however, are not applicable in case a producer, which influences the supply of both the old and the new products, replaces an old product with a new one. In this case, the market share is influenced by the producer's choices and is not informative about the quality of the new product. There is arguably a lot of product replacement among supermarket goods, where what is changing is the packaging and the price, but not the quality of the good. In these cases, the all-product inflation is the valid index.

Table D.9: Level and comovement of key inflation indexes

Series	moment	**DE**	**FR**	**IT**	**NL**	**US**
All-product inflation	average	1.51	1.39	0.37	2.78	2.25
Surviving-product inflation	average	-0.40	-0.92	-0.71	-0.10	1.47
Official food and beverage inflation	average	1.63	0.56	0.91	1.21	2.72
All-Surviving inflation	correlation	0.93	0.74	0.71	0.52	0.96
Official-Surviving inflation	correlation	0.49	0.71	0.32	0.87	0.89

Figure D.11 presents the all-product inflation rates for all countries together with the relevant official inflation subindex. Table D.9 shows the average year-on-year inflation rates. The figures/tables show that the all-product inflation rates gets closer to the level of the HICP inflation than our baseline index and it captures the variation of the official indices well across different countries. This suggests that the all-product inflation, despite its simplicity,

[33] **?** using the US IRi Academic Dataset find, consistently with our statement, that 'product line extensions, such as flavor or form upgrades or novelty and seasonal items, are much more prevalent than the introduction of new brands.'

Figure D.11: New product introductions: The year-on-year change of the baseline inflation index, the all-product index, and the HICP/CPI subindexes.

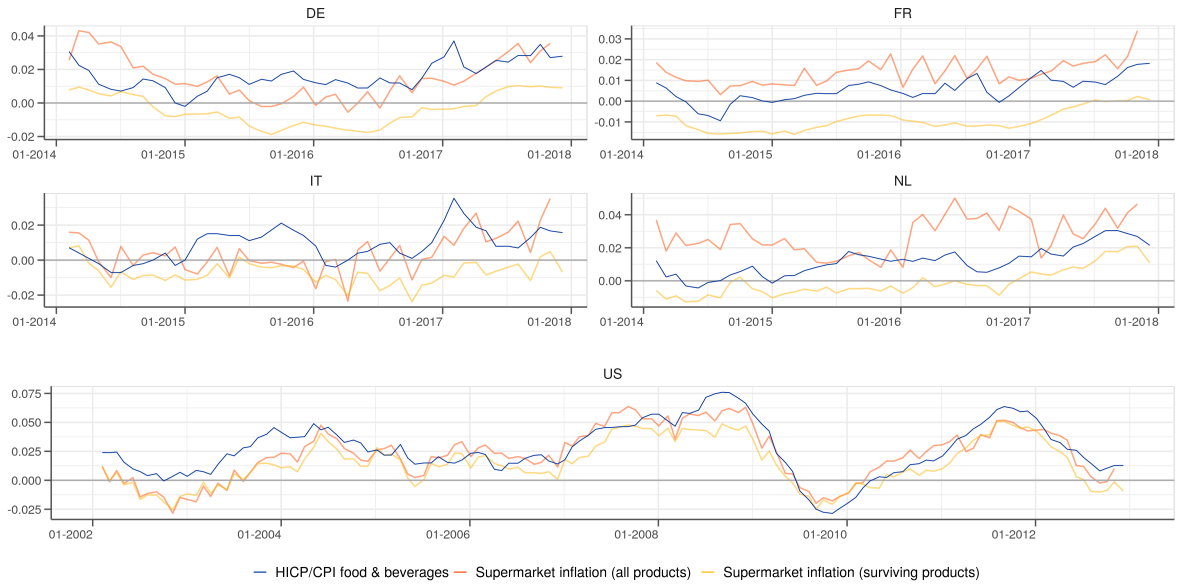

Note:

captures most of the information inherent in official price indices, which are based on much more careful judgment of product replacement and quality adjustment.

Appendix E. Time variation of EA4 and US frequency

Figure E.12: Frequency of reference-price changes, EA4 vs US

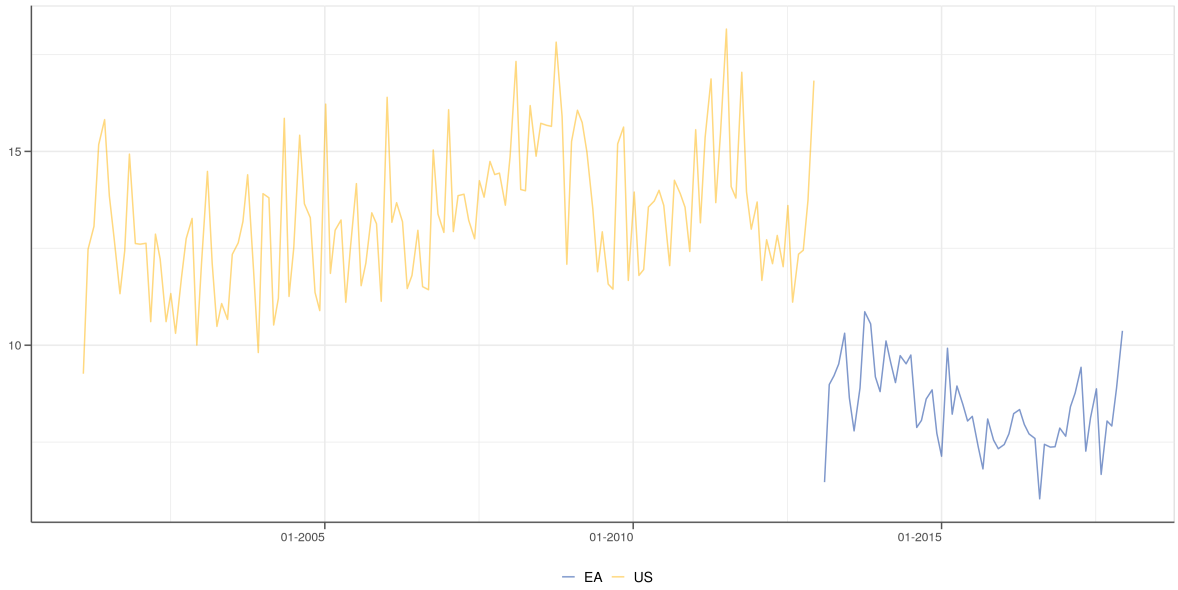

Note: The shows the evolution of the frequency of reference-price changes in the US (2001-2012) and in EA4 (2013-2017). The figure shows that the frequency is robustly lower, implying higher price rigidity in EA4 relative to the US.

Appendix F. Heterogeneity across euro area countries

Figure F.13: Generalized hazard functions and price-gap densities and duration hazard functions across euro area countries and the US

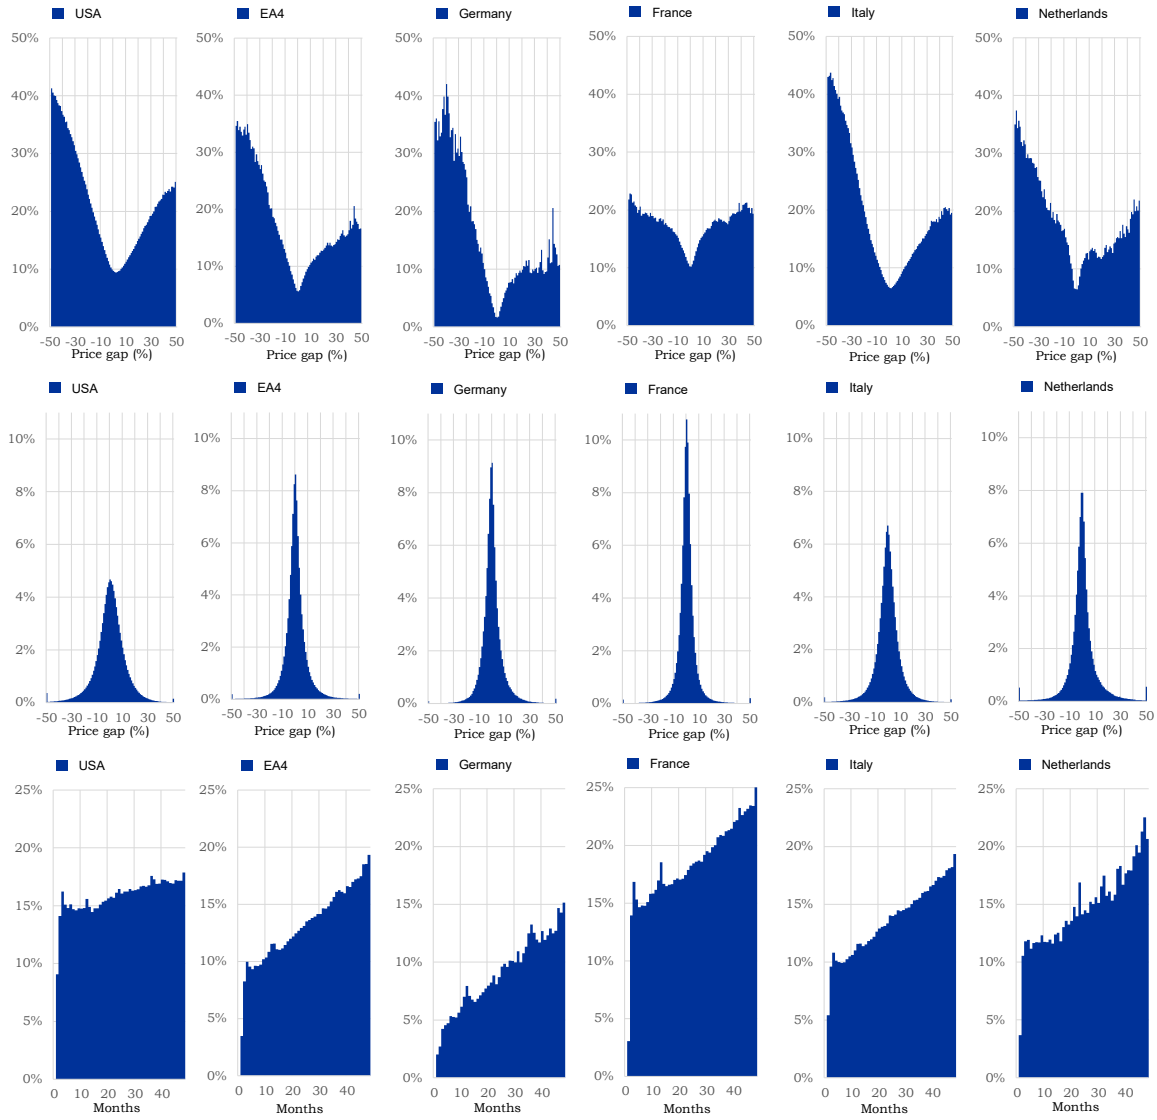

Note: The figure shows the generalized hazard functions (first row), the price-gap densities (second row), and the duration hazard functions (third row) across four euro area countries and the US. The figures show evidence for moderate state dependence in all countries, with notable heterogeneity across countries.

Appendix G. Duration hazard, sales and cross-sectional heterogeneity

Controlling for both cross-item heterogeneity as well as sales-related price changes is essential to obtain upward-sloping duration hazards. The left and right panels of Figure G.14 show, respectively, our estimates without controlling for fixed effects in Equation (3) and using posted-price changes, as opposed to reference price changes. The figures show that both factors bias the estimated slope downward, so much so that in both cases we would erroneously conclude that the hazard function is downward sloping.

Figure G.14: Duration hazard functions without controlling for cross-item heterogeneity and sales

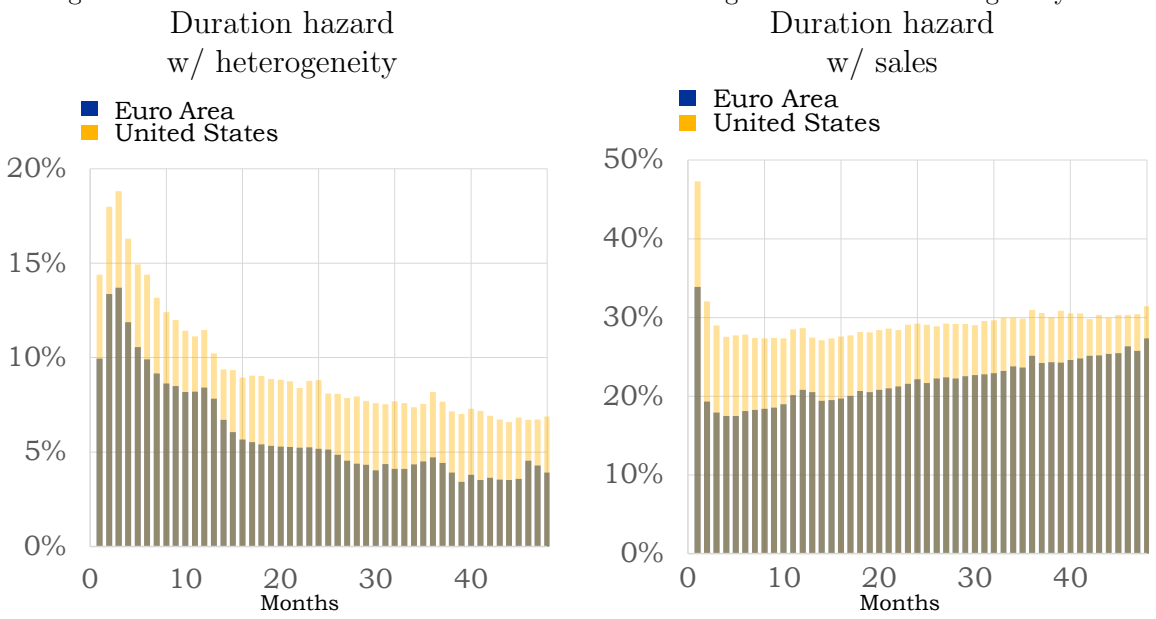

Note: The figures show the duration hazards without controlling for cross-item heterogeneity (left panel) and including sales-related price changes (right panel). The figures show that controlling for both factors is important to conclude that the hazard function is upward-sloping.

Appendix H. Robustness of structural model estimates

Appendix H.1. Asymmetric linear hazard

In this section, we present a price-setting model, which captures the shape of the generalized hazard function and the price-gap densities, and, consequently, also the price-change distributions better than the baseline **?** model. We ask the impulse responses of our baseline model are robust to these modifications.

We assume that the firm adjusts its price subject to a following hazard function

$$\Lambda(x) = \begin{cases} \omega + \alpha^{+} x & \text{if } x \geq 0 \\ \omega - \alpha^{-} x & \text{if } x < 0 \end{cases} \tag{H.1}$$

where ω is the height of the hazard function at zero gap, α^{-} is the (inverse) slope of the hazard for negative gaps and α^{+} is the slope for positive gaps.

We assume that the idiosyncratic shocks follow a Laplace distribution in the euro area $\varepsilon_t \sim L(0, \sigma_\varepsilon)$, where L is the Laplace distribution. We keep all the other features of the baseline model unchanged.

Figure H.15 shows the fit of the generalized hazard and the price gap densities of the model. The fit improves substantially relative to the baseline **?** model. The framework is even able to capture the bimodality of the US price-change distribution, which our baseline model was unable to achieve. However, as Figure H.16 shows, the improvement is quantitatively inconsequential, as it implies an insignificant change relative to the baseline as a response to a money growth shock.

Figure H.15: Estimation, targeted and non-targeted moments, asymmetric linear hazard

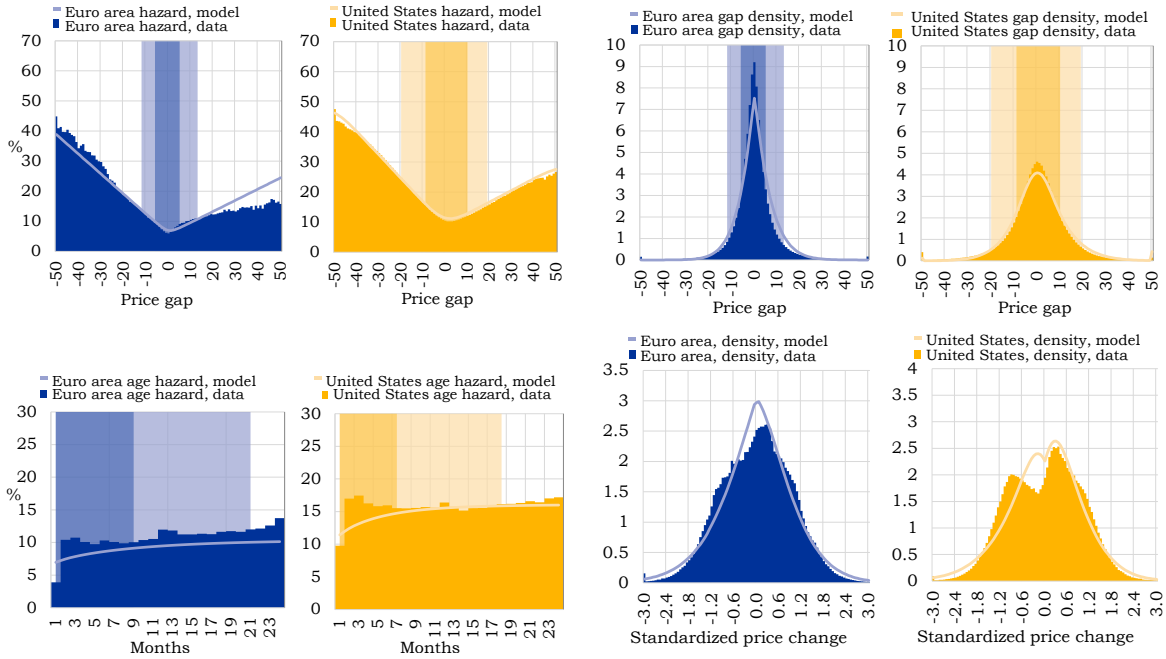

Note: The figures show the match of the simulated and the empirical generalized hazards and price-gap densities (matched moments, first row) and the duration hazard and price-change densities (unmatched moments, second row) in the euro area (EA4) and the US in the model with asymmetric linear hazard. The model matches the generalized hazard and the price-gap densities better than the baseline model. The model is even able to capture the bimodality of the price-change distribution in the US. Shaded areas cover the 67% (darker) and 90% (lighter) mass of the corresponding densities.

Table H.10: Estimated parameters, asymmetric linear hazard

Parameters	**US**	**EA4**
Stdev. of idiosyncratic shocks (σ_z)	5.3%	4.1%
Hazard at 0 gap (ω)	8.9%	5.9%
Slope ($x < 0, \alpha^-$)	0.78	0.66
Slope ($x \geq 0, \alpha^+$)	0.39	0.37

Note: The table shows the parameters of the model with asymmetric linear hazard. The EA4 calibration assumes idiosyncratic shocks with a Laplace distribution (US calibration keeps assuming a Gaussian distribution).

Figure H.16: Simulated response to a money growth shock in the model with asymmetric linear hazard

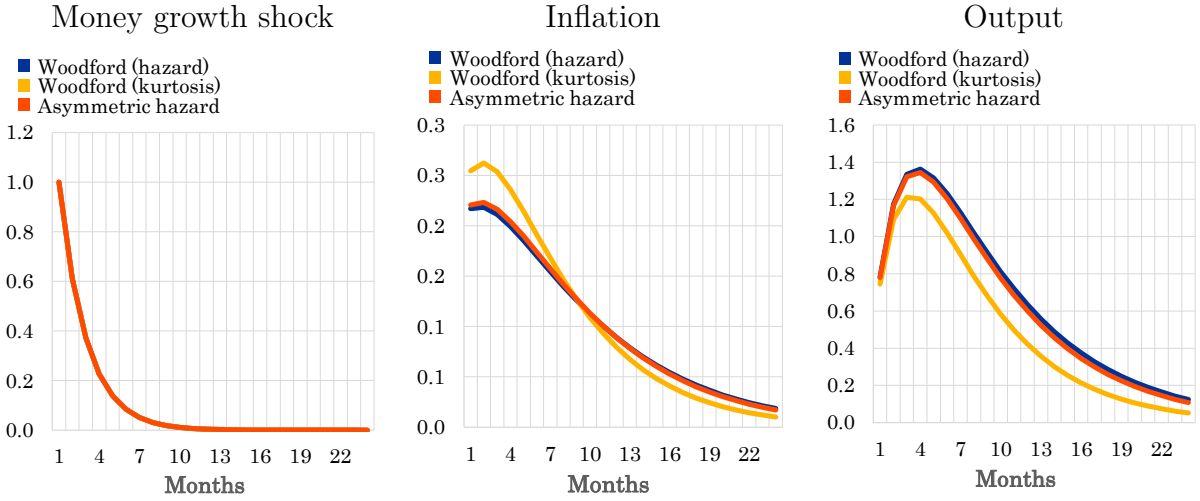

Note: The figures show that even though the model with asymmetric linear hazard and a Laplace distribution (in EA4) captures the shape of the hazard, the gap density, and the price-change distribution somewhat better, the improved fit has only a marginal impact on the impulse responses.

Appendix H.2. Targeting the kurtosis

This section presents an exercise, where we target the shape of the standardized price-change distribution along with the frequency and the size of the price changes, as in **?**. Not surprisingly, the model's fit with the price change distribution improves, at the cost of a worse fit with the hazard function. The calibration strategy implies a somewhat higher state dependence (lower information cost - α - parameters) and more flexible inflation response (see Figure H.16) than the baseline model. At the same time, the idiosyncratic shock volatility stays the key underlying factor driving the predicted differences in inflation volatility across regions, as in our baseline calibration.

Figure H.17: Estimation, kurtosis targeted

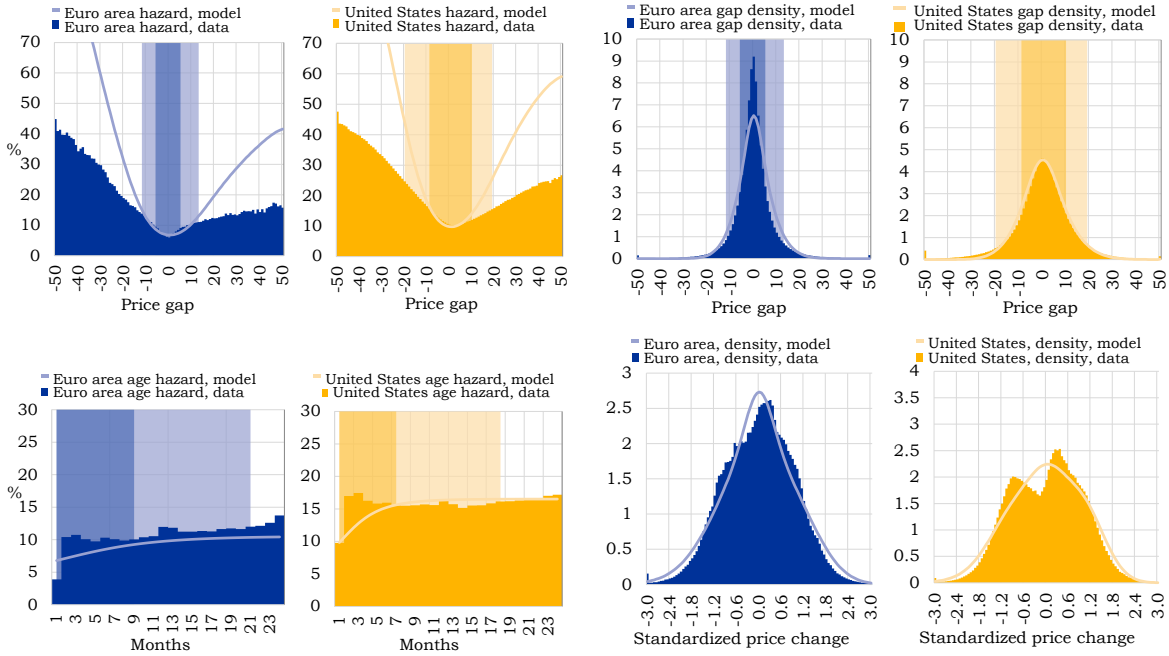

Note: The figures show the fit of the simulated and the empirical generalized hazards and price-gap densities (first row) and the duration hazard and price-change densities (second row) in the euro area (EA4) and the US in the model where the kurtosis of the price-change distribution is targeted (instead of the hazard function). The model matches the price-change distribution better at the cost of a worse fit of the hazard function. Shaded areas cover the 67% (darker) and 90% (lighter) mass of the corresponding densities.

Table H.11: Estimated parameters, kurtosis targeted

Parameters	**US**	**EA4**
Menu cost (κ)	11.3%	10.2%
Stdev. of idiosyncratic shocks (σ_z)	5.1%	2.9%
Information cost (α)	0.10	0.18

Note: The table shows the parameters of the model, where the kurtosis is targeted.

Appendix I. Supermarket inflation during Covid

We measure inflation using the year-on-year change in the Tornqvist price index with quarterly expenditure weights. The Tornqvist index is a superlative price index with desirable welfare-theoretical properties[34]. Quarterly expenditure weights reduce the impact of high-frequency variation in the composition of products both due to seasonal factors and temporary sales. Additionally, concentrating on year-on-year indexes minimizes the impact of seasonal variation as well as the potential impact of the 'chain drift,' which can be present with higher-frequency indexes relying on scanner data (**?**).

Formally, we calculate headline and reference-price inflation as

$$\pi_w^i = \sum_{ps} \gamma_{psw} (\log P_{psw}^i - \log P_{psw-52}^i), \tag{I.1}$$

where P_{psw}^i is the posted- $(i = h)$ and the reference price $i = f$ of product p in store s in week w and the weights are

$$\gamma_{psw} = \frac{I_{psw,w-52}(\omega_{psq-4} + \omega_{psq})/2}{\sum_{ps} I_{psw,w-52}(\omega_{psq-4} + \omega_{psq})/2}, \tag{I.2}$$

where $I_{psw,w-52}$ is an indicator function that takes the value 1 if product p in store s is sold in strictly positive quantities in both w and $w - 52$ and 0 otherwise[35], and ω_{psq} is the quarterly expenditure share of product p in store s in quarter q.[36]

Figure I.18 shows the weekly, year-on-year supermarket inflation in Germany and Italy between mid-February and mid-May. We concentrate on the 5-week-average inflation, which smooths out some high-frequency variability of the weekly series. The 5-week-average inflation started at around its long-term average in 2020 in both Germany and Italy and increased throughout the quarter in both countries. The increase was higher and clearly exceeded a one-standard-deviation band[37] in Italy (1.89 percentage points), while it was smaller and stayed

[34] It is the second-order approximation of the welfare-relevant price index under an arbitrary homothetic utility function.

[35] We match weeks with previous-year weeks based on their distance from the Easter week, the strongest seasonal factor over the mid-February-mid-May period we have data for in 2019 and 2020.

[36] The ensuing supermarket inflation rates in both countries co-move with the respective HICP food and beverages subindices. The correlation coefficients of the monthly inflations are 43% in Germany and 54% in Italy. The level of supermarket inflation is below the HICP subindices. The main reason is that we concentrate on surviving products and ignore the impact of new product introductions, which generate a major share of trend inflation.

[37] The band shows the standard deviation of 5-week-inflation rates over the first two quarters of the years between 2013-2017.

Figure I.18: Supermarket inflation during the first wave of the Covid-19 pandemic, year-on-year

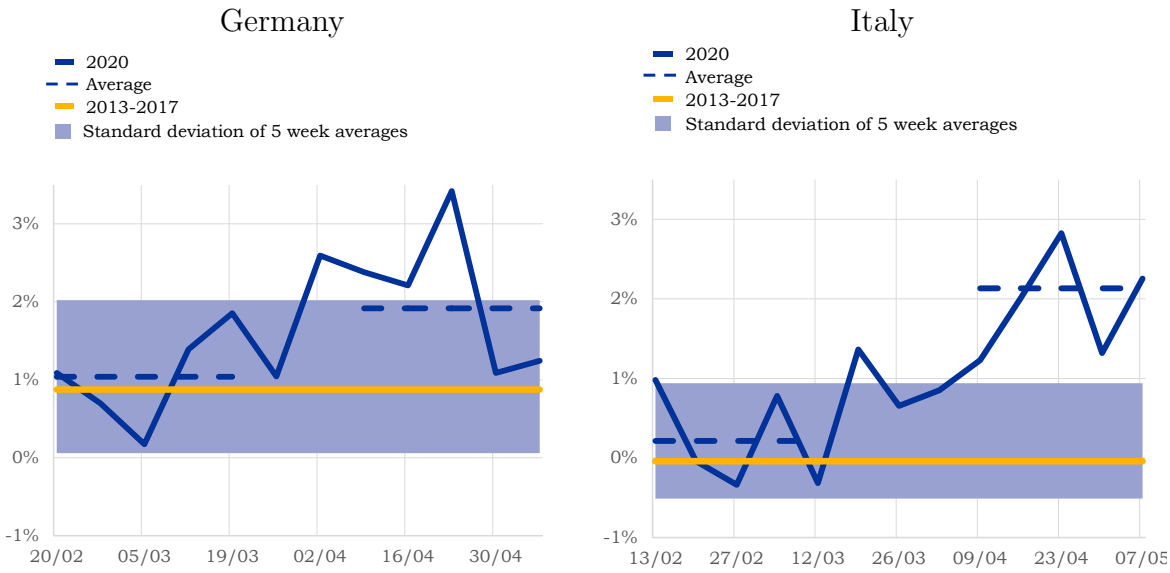

Note: The figure shows the weekly, year-on-year supermarket inflation (blue solid line) between mid-February and mid-May in 2020 in Germany and Italy. It shows that the average inflation in the first 5 weeks (blue-dashed line) stayed close to the average inflation rate during the first two quarters of 2013-2017 (yellow-dashed lines). Over the quarter, the 5-week-average inflation (blue dashed line) increased sizeably in both Germany (0.95%) and Italy (1.89%). The change stayed within a +- one-standard-deviation band in Germany but exceeded it in Italy.

within a one-standard-deviation band in Germany (0.95 percentage points). The increases are comparable to the HICP food-and-beverage subindexes between February and May in Italy (1.96 percentage points) and Germany (0.79 percentage points).

Figure I.19 shows the real expenditure growth in supermarkets in Germany and Italy. It is defined as the difference between nominal expenditure growth and the inflation rate. The figure shows that most of the nominal expenditure growth translated to real expenditure growth. This is consistent with the observation that prices responded sluggishly to the Covid shock.

Figure I.19: Real expenditure growth in supermarkets during the Covid-19 pandemic, year-on-year

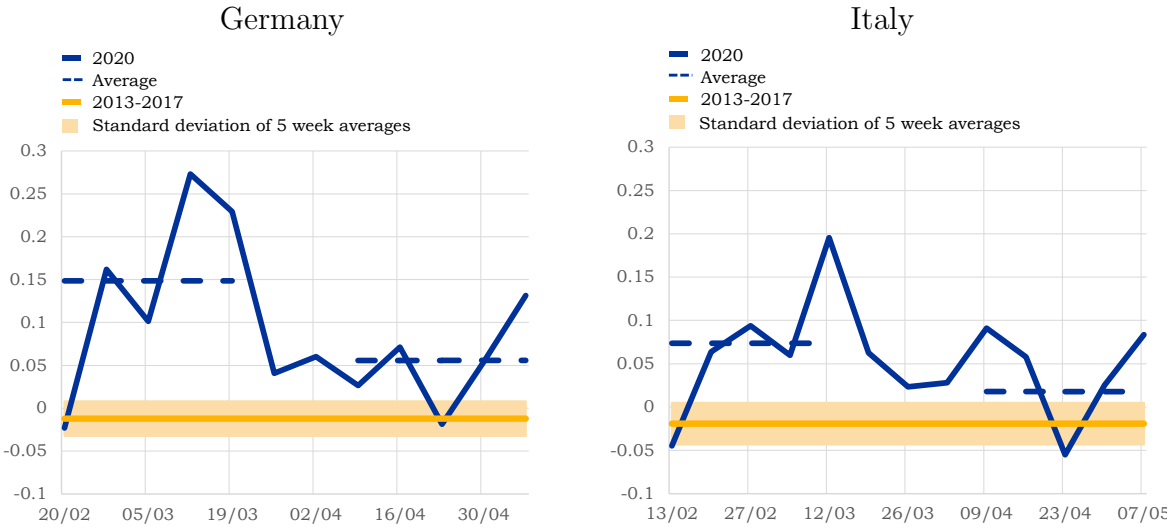

Note: The figure shows the weekly, year-on-year real expenditure growth (blue solid line) between mid-February and mid-May in 2020 in Germany and Italy. It shows that the 5-week-average expenditure growth (blue-dashed line) exceeded the average long-term expenditure growth (yellow dashed lines) by more than a standard deviation in both Germany and Italy. The expenditure growth was particularly high in the weeks preceding the lockdowns ('stock-up' shock), but stayed persistently high also during the lockdowns.