DEUTSCHE BUNDESBANK
EUROSYSTEM

Discussion Paper

Deutsche Bundesbank
No 34/2023

Nowcasting consumer price inflation using high-frequency scanner data: Evidence from Germany

Günter W. Beck
(University of Siegen and Miggroprices)

Kai Carstensen
(Kiel University)

Jan-Oliver Menz
(Deutsche Bundesbank)

Richard Schnorrenberger
(Kiel University)

Elisabeth Wieland
(Deutsche Bundesbank)

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet http://www.bundesbank.de

ISBN 978–3–95729–969–7
ISSN 2941–7503

Non-technical summary

Research question

Forecasts of current-month inflation ("nowcasts") are of great importance for central banks and market participants, especially during turbulent times. Since official inflation figures are only available with a certain time lag and at a monthly frequency, we study how millions of highly granular and weekly household scanner data combined with novel machine learning (ML) techniques can help to improve the nowcasts of monthly German inflation in real time.

Contribution

We construct more than 180 weekly scanner-based price indices at the lowest aggregation level underlying official German inflation, such as those of butter and coffee beans. Then, we apply ML techniques to exploit this large set of scanner-based price indices to nowcast headline inflation and major subcomponents used in monetary policy analysis, such as processed and unprocessed food.

Results

Our weekly scanner-based price indices track their official counterparts extremely well. Within a mixed-frequency modeling framework, we demonstrate that these scanner-based price indices improve inflation nowcasts at this lowest aggregation level, notably as soon as after the first seven days of a month. Combining these data with ML techniques, we also obtain substantial predictive gains for the major policy-relevant subcomponents of inflation. Finally, by adding high-frequency information on energy costs and travel services, we construct nowcasts for headline inflation that are on par with, or that even outperform, survey-based inflation expectations, which are notoriously difficult to beat.

Nichttechnische Zusammenfassung

Fragestellung

Die Prognose der Inflation des laufenden Monats ("Nowcast") ist für Zentralbanken und Marktteilnehmer von großer Bedeutung, insbesondere in turbulenten Zeiten. Da amtliche Inflationszahlen nur mit einer gewissen Zeitverzögerung und auch nur monatlich verfügbar sind, untersuchen wir, wie Millionen wöchentlicher Haushalts-Scannerdaten in Kombination mit neuartigen Techniken des maschinellen Lernens (ML) dazu beitragen können, den Nowcast der monatlichen deutschen Inflationsrate in Echtzeit zu verbessern.

Beitrag

Zunächst erstellen wir mehr als 180 wöchentliche scannerbasierte Preisindizes auf der niedrigsten Aggregationsebene, die der offiziellen deutschen Inflation zugrunde liegt, beispielsweise für Butter und Kaffeebohnen. Dann wenden wir ML-Techniken an, um diesen großen Satz scannerbasierter Preisindizes für die Prognose der Gesamtinflationsrate und wichtiger Unterkomponenten wie verarbeitete und unverarbeitete Nahrungsmittel, die in der geldpolitischen Analyse verwendet werden, zu nutzen.

Ergebnisse

Unsere scannerbasierten Preisindizes bilden ihre amtlichen Gegenstücke sehr gut nach. Im Rahmen eines Mixed-Frequency-Modellierungsrahmens zeigen wir, dass unsere scannerbasierten Preisindizes die Inflations-Nowcasts auf dieser niedrigsten Aggregationsebene verbessern, insbesondere bereits nach den ersten sieben Tagen eines Monats. Durch die Kombination dieser Daten mit ML-Techniken erzielen wir auch erhebliche Vorhersagegewinne für die wichtigsten politikrelevanten Unterkomponenten der Inflationsrate. Unter der Hinzunahme hochfrequenter Informationen zu Energiekosten und Reisedienstleistungen erzeugen wir schließlich Nowcasts für die Gesamtinflationsrate, die den umfragebasierten Inflationserwartungen, die bekanntermaßen schwer zu schlagen sind, ebenbürtig sind oder diese sogar übertreffen.

Deutsche Bundesbank Discussion Paper No 34/2023

Nowcasting Consumer Price Inflation Using High-Frequency Scanner Data: Evidence from Germany*

Günter W. Beck
University of Siegen and Miggroprices

Kai Carstensen
Kiel University

Jan-Oliver Menz
Deutsche Bundesbank

Richard Schnorrenberger
Kiel University

Elisabeth Wieland
Deutsche Bundesbank

22nd November 2023

Abstract

We study how millions of highly granular and weekly household scanner data combined with novel machine learning techniques can help to improve the nowcast of monthly German inflation in real time. Our nowcasting exercise targets three hierarchy levels of the official consumer price index. First, we construct a large set of weekly scanner-based price indices at the lowest aggregation level underlying official German inflation, such as those of butter and coffee beans. We show that these indices track their official counterparts extremely well. Within a mixed-frequency modeling framework, we also demonstrate that these scanner-based price indices improve inflation nowcasts at this very narrow level, notably already after the first seven days of a month. Second, we apply shrinkage estimators to exploit the large set of scanner-based price indices in nowcasting product groups such as processed and unprocessed food. This yields substantial predictive gains compared to a time series benchmark model. Finally, we nowcast headline inflation. Adding high-frequency information on energy and travel services, we construct highly competitive nowcasting models that are on par with, or even outperform, survey-based inflation expectations that are notoriously difficult to beat.

Keywords: Inflation nowcasting, machine learning methods, scanner price data, mixed-frequency modeling.
JEL codes: E31; C55; E37; C53.

*Corresponding author: Günter Beck, Universität Siegen, Unteres Schloss 3, 57072 Siegen, Germany, tel.: +49 271 7403223. Email: guenter.beck@uni-siegen.de. We are very grateful to GfK for providing access to its household scanner data and to Alfred Dijs (Miggroprices), without whom the project would not have been possible. Special thanks also go to Muzammil Hussain, Liuxiu Liu and Thi Ngoc My Nguyen for their excellent data support. This paper was presented at the ECONDAT 2023 spring meeting, the 43rd International Symposium on Forecasting in Charlottesville (USA), the IAAE 2023 Annual Conference in Oslo and the Inflation: Drivers and Dynamics 2023 Conference hosted by the Federal Reserve Bank of Cleveland and the ECB in Frankfurt am Main. We thank all participants for providing useful comments and suggestions, notably Paweł Macias, Malte Knüppel, Livia Paranhos and Katja Heinisch for discussing this paper thoroughly, as well as Michele Lenza, Karol Szafranek and Damian Stelmasiak. The views expressed in this paper are those of the authors only and do not necessarily reflect those of the Deutsche Bundesbank or the Eurosystem.

1 Introduction

The economic shock induced by the COVID-19 pandemic posed unprecedented challenges to policymakers and triggered an enormous demand for reliable real-time information about the state of the economy, including inflation dynamics. A similarly strong need for a timely and continuous flow of information about ongoing price developments arose after Russia's invasion of Ukraine in February 2022, when an immense inflationary wave started to unfold. Whereas official macroeconomic statistics are only available with a certain time lag and at fixed intervals, non-traditional high-frequency data such as web scraping and scanner data in combination with machine learning (ML) techniques represent a promising toolkit for policymakers to monitor ongoing and potentially disrupting developments in real time and to make better-informed decisions in such situations (Tissot and de Beer, 2020, Doerr, Gambacorta, and Maria Serena, 2021).

The usefulness of real-time information is not constrained to extraordinary periods, however, as it allows more generally for a faster identification of the current state of the economy, thereby enabling a timelier and more targeted response by policymakers. It can also aid in quantifying the impact of policy measures more precisely (Buda, Carvalho, Corsetti, Duarte, Hansen, Moura, Ortiz, Rodrigo, Rodríguez Mora, and Alves da Silva, 2023) and possibly adjusting them more swiftly. Furthermore, in the absence of timely official data, there is a risk that economic agents, by basing their decisions on private data sources, may amplify idiosyncratic shocks. Having high-frequency data at hand, policymakers can dampen such effects via a regular data-driven communication strategy (see, e.g., Assenmacher, Glöckler, Holton, and Trautmann, 2021).

The aim of this paper is to analyze how the combination of non-standard high-frequency price data with state-of-the-art machine learning methods helps to nowcast inflation in real time. More specifically, we demonstrate that weekly household scanner data improve inflation nowcasts, both at a very high level of disaggregation and for major product groups and headline inflation. Our data stem from the household panel of the market research company GfK and contain daily purchases of fast-moving consumer goods (henceforth denoted as GFK:FMCG) at the barcode level, primarily covering food, beverages, personal and household care products. Using the comprehensive product descriptions available in the dataset, we transform the highly granular daily price information into weekly price indices that closely match the official price indices at the most disaggregate level used in the German consumption basket, the so-called COICOP-10 item level.[1] With the help of a recursive out-of-sample nowcast experiment, we then document that the application of mixed data sampling (MIDAS) and machine learning techniques to these data (combined with official monthly inflation series and some complementary high-frequency data) yields highly informative nowcasts as soon as after only seven days of a month. Not surprisingly, the nowcast accuracy increases with the number of days included. After 14 and particularly after 21 days, the headline inflation nowcasts even outperform standard surveys of market expectations that are notoriously difficult to beat.

[1] The COICOP-10 item level provides a considerably more detailed disaggregation than the COICOP-5 level used for the classification of euro area inflation.

Our nowcasting exercise proceeds in three steps that relate to the three hierarchy levels of the German consumption basket our project focuses on: highly disaggregate COICOP-10 items, product groups, and headline inflation. We start at the COICOP-10 level and construct weekly price indices from the granular GFK:FMCG data with the help of time-product dummy regressions used successfully both in the literature (de Haan, Hendriks, and Scholz, 2021) and at statistical offices (Eurostat, 2022). For each COICOP-10 item, we then specify a U-MIDAS model (Ghysels, Santa-Clara, and Valkanov, 2004) that uses the weekly index to predict, on days 7, 14, 21, and 28 of a month, its official counterpart, which is measured at a monthly frequency. We document that this approach reduces the nowcast error substantially relative to a univariate time series benchmark model. The advantage is particularly pronounced for COICOP-10 items classified as unprocessed fruit and vegetables and dairy products and fat, for which we achieve root mean squared error (RMSE) reductions in the range of 40%-60%. Importantly, large nowcasting gains accrue even if only the scanner data of the first seven days of a month are included.

In the second step, we focus on product groups such as unprocessed and processed food that are regularly monitored by policymakers and market participants and for which we have weekly scanner data information available. As these product groups consist of large numbers of COICOP-10 items, direct nowcasting models of group-specific inflation rates need to efficiently integrate a multitude of weekly GFK:FMCG price indices, which is why we resort to shrinkage estimators from the machine learning toolkit. We find that, compared to a standard time series benchmark model, the group-specific inflation nowcasts of both unprocessed food and processed food benefit considerably from adding the weekly information. Specifically, we document reductions in the relative RMSE of up to 25%. When considering more disaggregate product groups, the relative forecasting gains are particularly large for dairy products and fat (reduction in RMSE of roughly 45% to 55%), unprocessed fruit and vegetables (reduction of around 20% to almost 40%), processed meat and fish (reduction of more than 25% to almost 40%), and unprocessed meat, fish and eggs (reduction of nearly 20% to 25%). Again, the superiority of our ML-based approach becomes apparent as early as after day 7 of a month.

In the final step, we construct nowcasts of headline inflation. To this end, we split the German consumption basket into six components – unprocessed food, processed food, energy, package holidays, non-energy industrial goods (NEIG), and services – which we consider separately. Specifically, we fit a mixed-frequency machine learning model directly to the monthly inflation rate of each component, produce nowcasts, and aggregate the nowcasts to headline inflation by applying the official HICP weighting scheme. To the set of weekly predictors, we add price quotes for energy and package holidays, which are two of the most volatile and difficult-to-predict inflation components of the German HICP. The resulting headline inflation nowcasts consistently outperform not only a time series benchmark approach but also Bloomberg market expectations if at least information up to day 14 is included.

To study the merits of aggregate versus disaggregate inflation nowcasting, we supplement the direct machine learning models with a bottom-up nowcasting approach which works as follows. For COICOP-10 items matched by a weekly predictor, we apply the U-MIDAS model discussed in the first step. For the remaining COICOP-10 items, we fit a time series benchmark model. We then aggregate the item-level inflation nowcasts to headline inflation nowcasts by the official HICP weighting scheme. We document that this bottom-up approach is slightly outperformed by the direct machine learning approach in normal periods but dominates it in turbulent times. In addition, it has proved highly competitive compared to market expectations, even during the inflation hike of 2022. From this, we conclude that in terms of inflation nowcasting, direct machine learning models are difficult to beat in normal times but do not necessarily adapt quickly enough to large shocks. Overall, this suggests that there is not a single nowcasting method which uniformly outperforms all its competitors. Rather, it is the careful integration of informative high-frequency data into nowcasting models that makes a difference.

The outline of this paper is as follows. Section 2 discusses related research and emphasizes our contribution in some detail. Section 3 explains the high-frequency scanner data and how we derive price indices that mirror the official indices published by the statistical office. In Section 4, we describe our nowcasting strategy and in Section 5, we report the results. Robustness checks are presented in Section 6, while Section 7 concludes.

2 Literature review

Our paper relates to several strands of the literature. First, the use of scanner data for economic research can be dated back to as early as the 1970s. In an excellent survey article, Dubois, Griffith, and O'Connell (2022) report how both household and retail scanner data have been fruitfully exploited in research on firm and consumer behavior. Within the area of consumer prices, scanner data have been used to compute household-specific inflation (Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019), assesses inequality across countries (Beck and Jaravel, 2021), study price-setting strategies of firms (Butters, Sacks, and Seo, 2022; Karadi, Amann, Sánchez Bachiller, Seiler, and Wursten, 2023), estimate price elasticities of consumer demand (Beck and Lein, 2020), track the effects of the COVID-19-related lockdown on prices and product variety (Jaravel and O'Connell, 2020), and measure the cross-border effects on prices within the euro area (Messner, Rumler, and Strasser, 2023). In addition, scanner data supplemented with survey questionnaires (D'Acunto, Malmendier, Ospina, and Weber, 2021) and survey-based information treatments (Weber, Gorodnichenko, and Coibion, 2022) can help to reveal the manner in which the way inflation expectations are formed and affect spending plans. Our paper contributes to this literature by showing that household scanner data can also successfully be employed to nowcast headline inflation.

Second, our paper relates to a burgeoning body of literature that was spurred on the COVID-19 pandemic and seeks to construct high-frequency measures of existing low-frequency macroeconomic series. Examples include real-time indicators of house prices (Anenberg and Laufer, 2017), GDP (Eraslan and Götz, 2021) and – using individual bank account information – private consumption (Buda, Carvalho, Hansen, Ortiz, Rodrigo, and Rodríguez Mora, 2022). Regarding inflation, the billion prices project initiated by Cavallo and Rigobon (2016) has shown that online prices can successfully be used to build a high-frequency price index that closely mirrors the official headline inflation published by statistical offices. In addition, using household scanner data for the UK, Jaravel and O'Connell (2020) provide a measure of food price inflation that closely tracks the official number at least on an annual frequency during the COVID-19 period. Studies that draw on web-scraped data to examine price effects in this period include Watanabe (2020), Cavallo and Kryvtsov (2023) and Stelmasiak, Szafranek, Macias, and Błażejowska (2023). Non-standard, high-frequency data have also turned out to be useful in the study of price and consumption effects following natural disasters as, e.g., in Cavallo, Cavallo, and Rigobon (2014), Gagnon and López-Salido (2019) or Watanabe (2020). Furthermore, using household scanner data for the UK, Jaravel and O'Connell (2020) provide a measure of food price inflation that closely tracks the official number at least on an annual frequency, and Alvarez and Lein (2020) offer an online inflation measure for Switzerland by combining web-scraped prices with consumption weights derived from debit card transactions. We add to this literature by using scanner data to construct weekly inflation measures for unprocessed food, processed food, and non-durable goods, both on the aggregate level and for detailed subcomponents such as "butter", "coffee beans", and "sanitary cleaner". Since our data also contains information on quantities, we are able to use real-time weights in the construction of price indices. The availability of such detailed inflation series is particularly useful in times of crises, as we have illustrated with respect to the price development of selected food items following Russia's invasion of Ukraine (Beck, Carstensen, Menz, Schnorrenberger, and Wieland, 2022).

Third, we contribute to the literature on nowcasting key macroeconomic variables. Typically, research in this field has focused on providing monthly GDP estimates; however, fuelled by the COVID-19 pandemic and the war in Ukraine, there has been increased interest in producing high-frequency measures of monthly data, such as inflation. Some studies in this field use traditional data sources such as weekly gasoline or commodity prices (Modugno, 2013; Breitung and Roling, 2015; Knotek II and Zaman, 2017; Clark, Leonard, Marcellino, and Wegmüller, 2022; Aliaj, Ciganovic, and Tancioni, 2023) and report robust forecasting and nowcasting gains compared to econometric benchmark models or market expectations. Another branch of the literature uses web-scraped price data to predict aggregate and disaggregate food price inflation (Macias, Stelmasiak, and Szafranek, 2023; Powell, Nason, Elliott, Mayhew, Davies, and Winton, 2018) and headline inflation (Harchaoui and Janssen, 2018; Aparicio and Bertolotto, 2020), again documenting improved forecasting accuracy. We add to this literature by showing that scanner data is a very promising candidate for the nowcasting of inflation at both the aggregate and the disaggregate level.

Fourth, our paper relates to an earlier body of literature that addressed the question of whether it would pay off to forecast headline inflation by explicitly using subcomponents or even the full breakdown of the inflation rate as inputs (Hendry and Hubrich, 2011; Ibarra, 2012; Espasa and Mayo-Burgos, 2013; Bermingham and D'Agostino, 2013) Generally, the studies show that it does indeed help to take disaggregate information into account. Recently, this line of research has been taken up again by Joseph, Kalamara, Kapetanios, Potjagailo, and Chakraborty (2022) who use disaggregate inflation data combined with machine learning methods to forecast headline inflation in the UK. Related to this academic debate, central banks have always forecasted different components of the inflation rate, both for statistical reasons and for understanding the underlying price dynamics (Benalal, del Hoyo, Landau, Roma, and Skudelny, 2004; Capistrán, Constandse, and Ramos-Francia, 2010; Huwiler and Kaufmann, 2013; Giannone, Lenza, Momferatou, and Onorante, 2014; Ulgazi and Vertier, 2022). By using the German inflation rate which is the one with the most detailed breakdown worldwide, we show that combining disaggregate inflation nowcasts into an aggregate nowcast for headline inflation is a highly competitive approach.

From a methodological perspective, our paper relates to recent advances in machine learning that seek to improve inflation forecasts by exploiting large datasets. While this literature dates back at least to Stock and Watson (1999), simple univariate models have been found to be very difficult to beat (Atkeson and Ohanian, 2001; Stock and Watson, 2007). Recent results are more promising. In particular, Garcia, Medeiros, and Vasconcelos (2017), Medeiros, Vasconcelos, Veiga, and Zilberman (2021) and Babii, Ghysels, and Striaukas (2022) apply a multitude of machine learning techniques to large macroeconomic datasets and demonstrate that this can lead to notable forecasting gains (see also Paranhos, 2021; Li, Liao, and Quaedvlieg, 2022; Goulet Coulombe, Leroux, Stevanovic, and Surprenant, 2022; Hauzenberger, Huber, and Klieber, 2023). In a similar vein, Joseph et al. (2022), Botha, Burger, Kotzé, Rankin, and Steenkamp (2022), and Barkan, Benchimol, Caspi, Cohen, Hammer, and Koenigstein (2022) show that inflation forecasts can also benefit from applying machine learning to large sets of highly disaggregate price indices. Our findings add to this body of research by revealing that machine learning tools provide an effective solution for handling a large set of disaggregate price series in a mixed-frequency setting. In particular, we show that it pays off to combine shrinkage methods with the weekly GFK:FMCG dataset to produce higher quality nowcasts for major product groups, such as for unprocessed and processed food, and headline inflation.

Finally, our approach of nowcasting inflation at a very disaggregate level with the help of machine learning models also helps to produce robust results during severe crises. As discussed by Bańbura, Leiva-Leon, and Menz (2023), adjusting standard forecasting models to cope with the effects of large shocks is no mean feat. As we will show, modeling inflation at a disaggregate level helps to take into account very specific policy measures taken during the pandemic and the recent energy price hike, both of which only affect a limited number of disaggregate price series.

3 Data

We base our nowcasts on a weekly dataset of consumer prices from three different sources. Most importantly, we use household scanner data for daily purchases of fast-moving consumer goods, mainly consisting of food items and non-durable goods. In addition, we collect weekly energy prices from the European Commission and daily transaction data for travel services from a private travel booking system provider. We combine this high-frequency dataset with monthly inflation rates compiled by the German statistical office from 2003 until 2022. This section provides details about the data sources and the data transformations necessary to construct high-frequency price indices that match the official HICP series as closely as possible.[2]

3.1 Fast-moving consumer goods (GFK:FMCG)

The most comprehensive part of our dataset consists of home scan purchases of fast-moving consumer goods (FMCG) collected from private households by the market research company GfK. It records, at a daily frequency, the purchases of around 30,000 households which constitute a representative sample of the German household population, and mainly includes food, beverages, and personal care items. The dataset starts in January 2003 and covers around 200,000 different products and an average of around 30 million observations per year. Table (1) provides an illustration of the structure of the GFK:FMCG data using "butter" as an example. In addition to the price paid on a particular day and the barcode of a single product, the dataset contains detailed information about products and retailers.

Table 1: Illustrative example of household scanner data

GfK Household Panel							Mapping to COICOP Classification	
Household	Product Description	Barcode	Quantities	Sales	Retailer	Purchase Date	10-digit Code	Product-Category
1	Green Hill Butter 250g	400123123123	1	3.39 €	A	28.11.2022	0115100100	Butter
2	Lovely Butter 250g	400456456456	2	6.58 €	B	01.12.2022	0115100100	Butter
3	Lovely Butter 250g	400456456456	1	3.39 €	C	01.12.2022	0115100100	Butter
4	Green Hill Butter 250g	400123123123	1	3.29 €	B	02.12.2022	0115100100	Butter
5	Green Hill Butter 250g	400123123123	2	6.98 €	A	03.12.2022	0115100100	Butter
6	Sunny Sunflower oil 1l	100445566123	1	2.29 €	B	01.12.2022	0115400100	Sunflower oil
7	Blossom Sunflower oil 1l	100112233123	1	3.99 €	C	01.12.2022	0115400100	Sunflower oil
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

Note: Fictitious entries.

For the purpose of our study, this dataset offers several advantages. First, scanner prices stem from actual transactions and should thus closely co-move with official prices sampled from a representative product bundle by the statistical office. Second, information on prices and purchases are recorded by shopping day, i.e., it can be used to construct high-frequency series. The data is also well-suited for nowcasting purposes as weekly updates are already made available by GfK on the following Monday. Third, the dataset includes quantities, i.e., the number of units bought of a particular product. This information is important as we can use it to construct consumption weights that increase the correlation with the official price indices. Fourth, information is provided at the barcode level, which allows us to take into account

[2] A list of all data sources used is provided in Table A1 of Appendix A.1.

composition effects.[3] For example, if a particular brand of butter is purchased more frequently, it will receive a higher weight in the construction of the scanner-based price index of butter. While this effect might temporarily lead to deviations from the official price index, which is based on a consumption basket that is fixed in the short run, official weights are adjusted regularly and thus should converge with those derived from actual shopping behavior.

In our empirical exercise, we construct nowcasts both for headline inflation as well as for some product groups such as "unprocessed food", "processed food", or "non-durable goods', and, if possible, for the most detailed inflation series available. According to the COICOP system,[4] the German inflation rate can be broken down into different categories, ranging from goods, services and energy to more detailed components such as "vegetables" or "mineral water". In the euro area, the most detailed harmonized breakdown of the inflation rate is the so-called COICOP-5 level, whereas in Germany, inflation can be disaggregated further into the COICOP-10 level,[5] which, however, poses some challenges. First, the COICOP-10 series are compiled within the system of the national CPI and not the HICP. Hence, methodological differences between both concepts have to be taken into account when aggregating inflation nowcasts from the bottom up.[6] Second, the CPI is typically revised every five years with the introduction of a new base year by including new product groups and removing outdated ones. Therefore, some COICOP-10 series are only available on a shorter time period, which we try to overcome by extending them backwards using corresponding price indices from the previous base years.

Keeping this in mind, we match the products of the scanner data to the corresponding COICOP-10 items, making use both of the product categories and the detailed product descriptions included in the GfK household panel as well as the item descriptions contained in the HICP manual.

Based on the mapped scanner data, we proceed as follows. First, in each COICOP-10 component, the raw price $p_{i,d}$ of a given product i bought on a particular day d is defined as:

$$p_{i,d}^{raw} = \frac{sales_{i,d}}{unit_{i,d}}, \tag{1}$$

where $sales_{i,d}$ are the total expenses in euro for a given item and $unit_{i,d}$ denotes the number of items bought. By this, we obtain for each item and time period a sample

[3] Mostly, the barcode is given by the "Global Trade Identification Number" (GTIN), whereas GfK assigns a unique ID to products such as fresh food or private labels for which no GTIN is available.

[4] COICOP stands for "Classification of Individual Consumption by Purpose"; see Eurostat (2018) for details.

[5] For example, the COICOP-5 component "cheese and curd" is decomposed into the five COICOP-10 groups "hard cheese", "sliced cheese", "soft cheese", "curd" and "cream cheese". Indices computed at this lowest-available level are generally denoted as "elementary indices"; see IMF, ILO, OECD, Eurostat, UNECE, and The World Bank (2020), Chapter 1.

[6] The HICP is a chain-linked price index where weights are updated each year whereas the CPI is a fixed-base index where weights are updated only every five years. Moreover, the CPI includes prices for gambling and owner-occupied housing. Despite their methodological differences, headline rates of the HICP and CPI behave rather similarly over time.

of unit-value observations. Moreover, we omit outliers which are below and above the 1st and 99th percentiles of the price distribution at the COICOP-10 level. Due to the large number of households in the dataset, a specific item is often bought several times per day, in which case we compute the average price per item and day. Finally, we transform the data from daily into weekly frequency by defining four weeks per month such that the first week consists of day 1 until day 7, the second week from day 8 to 14, the third week from day 15 to 21 and the fourth week from day 22 to day 28. By using only 28 days of a given month, we tackle the problem of a shorter February and leap years.

We compile scanner-based price indices at the COICOP-10 level by running weighted time-product dummy (TPD) regressions. This method, proposed by Diewert (2005), is widely used in official price statistics to construct price indices from scanner or web-scraped data (de Haan et al., 2021; Eurostat, 2022). Specifically, for each week $t = 0, \ldots, T$ and product $i = 1, \ldots, N$, we fit the equation

$$\ln p_{i,t} = \beta^0 + \sum_{\tau=1}^{T} \delta^\tau d^\tau_{i,t} + \sum_{j=1}^{N-1} \gamma^j D^j_i + \varepsilon_{i,t}, \tag{2}$$

where $D^j_{i,t}$ represents a product dummy which takes the value 1 if $i = j$ (as identified by its barcode) and 0 otherwise, and $d^\tau_{i,t}$ denotes a time dummy which takes the value 1 if $t = \tau$ and 0 otherwise. Weights are given by the total expenses, $sales_{i,t}$, for a given product. This increases the price effect of popular products compared to those that are bought less frequently. Note that missing prices for a given item are treated in a similar way as in official price statistics: if an item is only temporarily missing, its last price is carried forward up to eight weeks, before it is replaced with another product.

For each week $t = 0, ..., T$, we estimate a price index from the exponential of the coefficient on the respective time dummy, such that:

$$I^{0,t}_{TPD} = 100 \times \exp\left(\hat{\delta}^t\right). \tag{3}$$

To mimic a real-time compilation of scanner-based price indices, we estimate equation (2) on a rolling window of 49 weeks, which covers at least one full year of scanner data (i.e., one month consisting of four weeks only). For example, the first estimation window covers period 1 to period 49 (providing a price index of the same length), the second estimation window period 2 to period 50, and so on. The linking of this sequence of 49-period price indices is performed in the spirit of a mean splice. By linking subsequent index values to the existing one, this yields a non-revisable real-time price index such that:

$$I^{0,t}_{TPD} = \prod_{k=t-\lambda}^{t-1} \left(I^{0,k}_{TPD} \times I^{k,t}_{[t-w+1,t]}\right)^{\frac{1}{\lambda}}, \tag{4}$$

where w denotes the window size (49 weeks) and λ is an overlapping linking period, which we set to eight weeks.[7] Hence, in our example above, the index in period 50 is obtained as a geometric average of the pairwise changes of the period-50 index value to the index values in periods 42 to 49 obtained from the second estimation window, each multiplied by the index value of the corresponding overlapping period from the first estimation window.

3.2 Energy and travel services

Unexpected price changes for energy and package holidays regularly contribute strongly to forecast errors of German headline inflation. Therefore, we also match these components with high-frequency information available in almost real time.

The energy component of the HICP consists of 14 price indices at the COICOP-10 level[8] that we try to match with the high-frequency price indicators discussed in the following. Most importantly, we use weekly price data from the Weekly Oil Bulletin (WOB) by the European Commission that has proven invaluable in previous work (Modugno, 2013; Aliaj et al., 2023). The WOB contains weekly information starting in 2005 about average fuel prices at the pump (Diesel, Supergrade petrol) and household-size deliveries of heating oil. The prices include duties and taxes and are thus much more closely related to consumer price indices than futures prices traded on financial markets, as used by Breitung and Roling (2015) and Knotek II and Zaman (2017). Nevertheless, we include the European Gas Spot Index (EGSI) as we have no other high-frequency information for the household gas supply.

We also use a few more daily energy price series from various sources as described in Table A1 in the appendix. All German gasoline stations have to report their intraday price changes of super and diesel fuels to the Market Transparency Unit for Fuels. We access this database via the data provider "Tankerkoenig" and take unweighted averages of all prices reported within a day. We also include daily measures of heating oil and wood pellet prices. We turn all daily series into weekly by taking an unweighted average of days 1-7 (week 1), 8-14 (week 2), 15-21 (week 3), and 22-28 (week 4).

Regarding travel services, we use high-frequency data from the travel booking system provider AMADEUS that include transaction prices for package holidays, which is defined as a combination of flight and accommodation services. Prices for inter-

[7] In the context of a rolling time window, a mean splice includes all overlapping periods as a linking period (Eurostat, 2022). However, in our weekly application, this yields 48 overlapping periods. We therefore opted to link over a shorter time period to ease the computational burden. We also computed alternative price indicators based on different index concepts and splicing methods. Overall, the TPD model with a splicing over eight weeks gave the best results in terms of in-sample correlations.

[8] These items are "Electricity" (COICOP no. 0451010000), "Natural gas, excl. share in the costs" (0452103000), "Share in the costs for gas central heating" (0452105100), "Liquefied gas, charging of a tank container" (0452200200), "Heating oil" (0453001100), "Share in the costs for oil central heating" (0453005100), "Coal briquettes" (0454100200), "Firewood, wood pellets or the like" (0454900100), "District heating" (0455002200), "Diesel fuel, cetane number below 60" (0722100100), "Diesel fuel, cetane number 60 and more" (0722100300), "Supergrade petrol, 95 octane" (0722201100), "Supergrade petrol, 98 octane and more" (0722204300), and "Liquefied petroleum gas" (0722301100).

national package holidays represent an important component of the German HICP because of their relatively large weight in the consumer basket and their relatively high volatility.[9] The AMADEUS dataset starts in 2012 and is available at a daily frequency.[10] We provide more details about the construction of the high-frequency index for package holidays in Section A.3 of the appendix. Basically, the index is compiled by constructing a weighted average of the prices for the most relevant travel destinations and that focuses on last-minute bookings up to 14 days before the travel date.

3.3 Descriptive statistics

Table 2 provides an overview of the HICP components for which we are able to construct a corresponding high-frequency price index using household scanner data for the full sample 2003 to 2022. In total, we cover about 12% of the basket underlying the German headline inflation rate, including the vast majority of unprocessed and processed food items.[11] With regard to non-energy industrial goods (NEIG), we match a significant proportion of non-durables but are unable to do so for most of the semi-durable and durable goods. Except for package holidays, our dataset does not comprise high-frequency prices for services which make up about 50% of the German consumer basket. However, service prices are either determined by administrative measures, such as insurances or tuition fees, or are rather sticky (see Gautier, Conflitti, Faber, Fabo, Fadejeva, Jouvanceau, Menz, Messner, Petroulas, Roldan-Blanco, Rumler, Santoro, Wieland, and Zimmer, 2023). Therefore, we do not expect the lack of high-frequency prices for services to impair our nowcasting results too strongly.

In Figure 2, we compare the year-over-year inflation rates of the HICP with its scanner-based counterparts by considering the product groups "food" and "non-durable goods".[12] Note that the HICP aggregates are derived from the underlying COICOP-10 series using only those series for which we also have high-frequency data available. While this distinction does not matter for food, energy and package holidays, it is relevant for non-durable goods given that we only cover about 40% of this aggregate.

Overall, the comovement between the weekly high-frequency and the monthly official inflation rates is very high. For food, non-durable consumer goods and energy, the correlation coefficient is about 0.9. Even if we use the HICP component for non-durable goods based on the full set of the underlying COICOP-10 series, the

[9] In 2023, package holidays make up 3.5% of the German inflation rate, compared to 0.2% in France and 0.5% in Italy.

[10] The dataset has been used by Henn, Islam, Schwind, and Wieland (2019) for price measurement and by Nagengast, Bursian, and Menz (2021) to estimate the exchange rate pass-through.

[11] Tobacco products are not available in the scanner data set since households are typically reluctant to reveal reliable information about their actual purchases.

[12] The graphs for "energy" and "package holidays" are shown in Figure A1 in the appendix. An online appendix to this paper (available upon request) plots the high-frequency inflation rates together with the official rates for all of the about 180 product groups which could be mapped to scanner data.

Table 2: Mapping between high-frequency price data and the German HICP

	HICP		Scanner data	
Component	COICOPs	Weight	COICOPs	Weight
Unprocessed food	38	2.4	30	2.0
Fruit	8	0.7	6	0.5
Vegetables	11	0.7	9	0.6
Meat & eggs	15	0.9	15	0.9
Fish	4	0.1	0	.
Processed food	142	11.1	116	8.1
Fruit	7	0.2	5	0.1
Vegetables	12	0.4	11	0.4
Meat	13	1.1	11	0.9
Fish	7	0.2	4	0.1
Bread & cereals	25	1.5	23	1.4
Dairy products & fat	18	1.5	14	1.4
Beverages	29	2.9	23	2.7
Other food products	28	1.2	25	1.0
Tobacco	3	2.1	0	.
NEIG	302	23.0	39	1.8
Non-durables	75	5.9	36	1.8
Semi-durables	139	8.7	3	0.1
Durables	88	8.4	0	.
Total HICP	482	36.5	185	11.9

Note: The table reports the number of COICOP-10 price indices and the associated weights of different HICP components in addition to the number of series and the share for which we have household scanner data available over the full sample 2003 to 2022. Weights refer to CPI expenditure shares of the base year 2020.

correlation still exceeds 0.5.[13]

[13] Correlations are computed by carrying forward the monthly inflation rates to each of the four weeks per month. This approach mirrors our forecasting setup more closely than aggregating the weekly series into monthly frequency in a first step.

Figure 1: HICP inflation and high-frequency counterparts

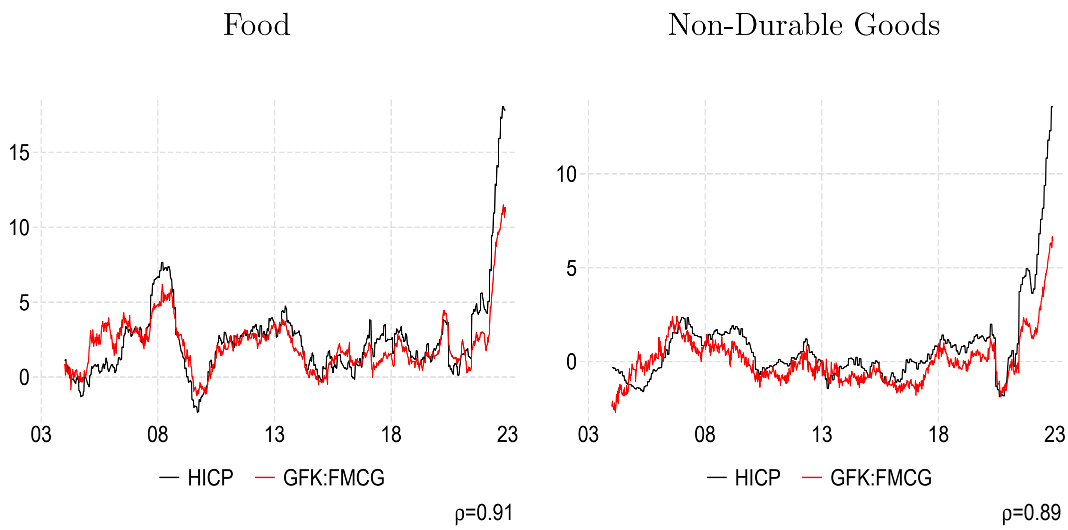

Note: The figure shows year-over-year inflation rates (% change) for HICP subcomponents aggregated using all of the corresponding COICOP 10-digit level series for which we have high-frequency data available. "HICP" refers to the aggregates using official COICOP 10-digit series, and "GFK:FMCG" refers to the scanner data for fast-moving consumer goods. ρ reports the correlation coefficient between both series.

It is worth noting that this high correlation of the aggregate inflation series masks some considerable heterogeneity at the COICOP-10 level. This is illustrated for a selected number of product groups in Figure A2 that highlight some general patterns.[14] First, as the indices for "cucumbers", "cherries" and "whole milk" illustrate, official inflation rates tend to exhibit larger fluctuations compared to their high-frequency counterpart. This could be due to the fact that the statistical office samples only a subset of products, whereas scanner-based indices use all products being bought by households, which might smooth out extreme price changes. On the other hand, we also observe that for other COICOP-10 components indices of high-frequency scanner data are more volatile than the official series. This is, e.g., the case for "multivitamin juice" or "salt". This could stem from sales or special offers that are more likely to be included in the high-frequency scanner dataset rather than the monthly HICP data. Finally, towards the end of the sample, official inflation series have increased much more than their high-frequency counterpart. This points to substitution effects that arise if households switch from expensive products to cheaper ones.

[14]The graphs for all COICOP-10 components are provided in an online appendix to this paper (available upon request).

4 Nowcasting strategy

Our nowcasting exercise proceeds in three steps. In the first step, we evaluate at the highly disaggregate COICOP-10 item level how well the weekly GFK:FMCG price indices predict their current month's official HICP counterparts. For that purpose, we employ the well-established MIDAS framework advocated by Ghysels et al. (2004) that takes into account the different frequencies of the data.[15] We compare this approach to an autoregressive benchmark with seasonal dummy variables. We are thereby able to quantify the informational advantage of adding weekly scanner data to nowcast monthly inflation at the most granular level.

In the second step, we nowcast the inflation rates of product groups of the HICP consumption basket that are closely monitored by central banks, market observers, and business cycle experts and for which we have disaggregate weekly GFK:FMCG price information available. The most important of these groups are unprocessed food, processed food, and non-energy industrial goods. As they consist of large numbers of disaggregates – for example, unprocessed food is an aggregate of 142 COICOP-10 items – we implement machine learning techniques like the elastic net and the sparse group LASSO to achieve dimensionality reduction via shrinkage and produce direct inflation nowcasts at the product group level in a high-dimensional mixed-frequency environment. Again, our benchmark is a standard autoregressive model augmented with seasonal dummies. The setup allows us to study the extent to which the combination of high-frequency external information and up-to-date machine learning techniques can improve inflation nowcasts.

In the third step, we construct nowcasts of headline inflation. Here, we face the obstacle that the GFK:FMCG data cover only a part of the HICP consumption basket. Therefore, we add high-frequency data on price developments of energy and package holidays, which for Germany are the two most volatile and difficult-to-predict inflation components omitted so far. Rather than attempting to directly relate headline inflation to the multitude of dissaggregate information compiled, we use the product-group approach also applied in the previous step. Specifically, we target the following six components that make up to headline inflation: unprocessed food, processed food, energy, package holidays, non-energy industrial goods (NEIG), and services. For each of these components, we fit a machine learning model that includes all relevant disaggregate weekly information. We then compute headline nowcasts from component nowcasts by applying the official HICP weighting scheme. We complement the machine learning strategy with a bottom-up approach that generates inflation nowcasts exclusively at the COICOP-10 level as described in the first step and subsequently aggregates them, again using the official HICP weights. In this setup, we predict COICOP-10 items not matched by high-frequency covariates with the help of the seasonal-dummy autoregressive benchmark model. We compare these two approaches to market expectations, which have been shown by Bańbura et al. (2023) to be a very challenging benchmark.

[15] Alternatively, we could have used a dynamic factor model or mixed-frequency VAR approach estimated by means of a Kalman filter or Bayesian methods (see, e.g., Modugno, 2013; Cimadomo, Giannone, Lenza, Monti, and Sokol, 2022); however, in our setting, the much simpler MIDAS framework is sufficient.

We conduct a recursive out-of-sample nowcast experiment that covers the period of January 2016 to December 2022. This evaluation sample is dictated by data availability as quite a few official COICOP-10 price indices start only in January 2015. Nevertheless, it includes both the COVID-19 crisis and the recent surge in inflation, arguably not only two of the most challenging periods for inflation forecasting in recent history, but also especially hard for statistical models that are fitted to the data in normal times and do not incorporate all the structural information professional forecasters took into account. The following paragraphs describe our nowcasting strategy in more detail.

4.1 Benchmark nowcasts

We define the month-over-month inflation rate as $\pi^M_{c,t} = 100 \times (P_{c,t}/P_{c,t-1} - 1)$, and the year-over-year inflation rate as $\pi^A_{c,t} = 100 \times (P_{c,t}/P_{c,t-12} - 1)$, where $P_{c,t}$ denotes the HICP index of the COICOP-10 item $c = 1, \ldots, 644$ in month t. Then, we compute model-based benchmark nowcasts from a seasonal dummy autoregressive (SD-AR) model of the following form:

$$\pi^M_{c,t} = \alpha_0 + \sum_{j=1}^{p} \rho_j \, \pi^M_{c,t-j} + \sum_{s=1}^{13} \gamma_{c,s} \, d_{c,s,t} + \varepsilon_{c,t}, \tag{5}$$

Since we work with data that are not adjusted for seasonal effects, we augment the AR model with 11 monthly dummies $d_{c,1,t}, \ldots, d_{c,11,t}$, defining December as the reference case. We also add an Easter dummy, $d_{c,12,t}$, and a Pentecost dummy, $d_{c,13,t}$, to capture the specifics of the German public holiday season, which in turn affect the consumption patterns of German households.[16] Throughout the paper, we use the year-over-year inflation rate as our target variable; hence, we transform the resulting nowcasts for the month-over-month rates accordingly. The model specification is, in each estimation step, guided by the Bayesian Information Criterion (BIC). This includes the autoregressive lag order $p \in \{1, \ldots, 12\}$ and the decision of whether or not to include the dummies.[17]

[16] The Easter dummy measures how many days of the two Easter weeks are in March and April, while the Pentecost dummy measures how many of the three Pentecost days (Saturday to Monday) are in May and June.

[17] We also use a local mean model that consists only of an intercept estimated from the last twelve observations to accommodate the specifics of some administered prices. In addition, for seven administered price indices of various medical and veterinary services, driving license fees and motor vehicle registration fees we do not fit a model to the month-over-month inflation rate as these indices are known to change only rarely and then stay constant for years. Instead, we use a random walk forecast of the level. In fact, we have evaluated more generally the forecast accuracy of AR models estimated in year-over-year rates as well as random walk models specified both in price levels and inflation rates, but have found that – with the exception of the seven cases just listed – the SD-AR model yields the best overall results for the COICOP-10 items, the subcomponents, and headline inflation.

4.2 Step 1: nowcasting item-level inflation rates

For each COICOP-10 item c for which we have weekly GFK:FMCG data available, we estimate a MIDAS model of the monthly HICP inflation rate, $\pi^M_{c,t}$, to account for the mixed-frequency environment. We use the baseline specification

$$\phi(L)\,\pi^M_{c,t+h} = \alpha_{0,h} + \beta_{c,h}\,B(L^{1/m};\theta)\,x^{(m)}_{c,t} + \sum_{i=s}^{13}\gamma_{c,s}\,d_{c,s,t+h} + \varepsilon_{c,t+h}, \tag{6}$$

where the subscript $t = 1,\ldots,T$ denotes the monthly time index and the superscript m denotes the high-frequency ratio within a month. The predictors include the weekly GFK:FMCG inflation rate, $x^{(m)}_{c,t}$, sampled four times more frequently than the target variable, and the set of monthly dummy variables, $d_{c,1,t},\ldots,d_{c,13,t}$, defined above. To account for temporal dependence, we specify an autoregressive polynomial, $\phi(L)$, with lag order $\{1, 12\}$ and a distributed lag polynomial, $B(L^{1/m};\theta)$, that aggregates the m high-frequency lags to the common frequency t.

For each month, we use high-frequency information only from its first four weeks, so that we have a fixed week-to-month ratio of $m = 4$ as required by the MIDAS model. The first week spans the seven initial days of a given month t, the second week includes days 8 to 14, and so on up to day 28. This strategy controls for the problem of overlapping calendar weeks across consecutive months and the heterogeneous number of days in different months. It is also in line with the typical convention in official price statistics according to which prices are primarily collected during the first three to at most four weeks of a month, implying that prices observed after the 28th day of a month hardly enter the HICP.

To integrate weekly observations into a monthly model, we use the following notation. We denote the four weekly observations by $x_t, x_{t-\frac{1}{4}}, x_{t-\frac{2}{4}}, x_{t-\frac{3}{4}}$, where x_t denotes the high-frequency inflation rate of the 4th week in month t over the same week of month $t-1$, $x_{t-\frac{1}{4}}$ refers to the inflation rate measured in the 3rd week and so on up to the 1st week of t.[18] Since the choice of $m = 4$ with a single predictor does not lead to a proliferation of parameters in equation (6), we implement the unrestricted MIDAS approach (hereafter *U-MIDAS*) with OLS estimation (see Foroni, Marcellino, and Schumacher, 2015; Ghysels and Marcellino, 2018, for more details).[19]

How do we deal with the ragged-edge feature of our mixed-frequency dataset? In each month t, we use the four information sets available on days 7, 14, 21 and 28 to nowcast the inflation rate $\pi^M_{c,t}$ of some item c which is matched by GFK:FMCG data. It is important to keep in mind that the official COICOP-10 inflation rates are only released with a two-week delay following the end of the reference month. Hence, nowcasts made on day 7 of month t use official inflation rates up to month $t-2$ and an estimate of month $t-1$ derived from the weekly data which we denote by $\hat{\pi}^M_{c,t-1}$. They also include the first weekly observation $x_{c,t-\frac{3}{4}}$ and fill in the remaining

[18] See Section B.1 of Appendix B for a more formal description of the baseline model (6) and its matrix notation.

[19] In our setting, the U-MIDAS approach delivers results that are roughly equivalent to nonlinear choices of the high-frequency aggregation scheme via $B(L^{1/m};\theta)$, such as the exponential Almon lag specification.

three weeks with this latest observation available. Nowcasts made on day 14 use the official inflation rate of the previous month $\pi^M_{c,t-1}$, which at that day has just been published, and add the first two weekly observations $x_{c,t-\frac{3}{4}}$ and $x_{c,t-\frac{2}{4}}$. Again, they use the latter to fill in the two missing weeks. This updating scheme is repeated on days 21 and 28. We list the information sets in the following:

Day 7: $\left(\hat{\pi}^M_{c,t-1}, x_{c,t-\frac{3}{4}}, x_{c,t-\frac{3}{4}}, x_{c,t-\frac{3}{4}}, x_{c,t-\frac{3}{4}}\right)$

Day 14: $\left(\pi^M_{c,t-1}, x_{c,t-\frac{3}{4}}, x_{c,t-\frac{2}{4}}, x_{c,t-\frac{2}{4}}, x_{c,t-\frac{2}{4}}\right)$

Day 21: $\left(\pi^M_{c,t-1}, x_{c,t-\frac{3}{4}}, x_{c,t-\frac{2}{4}}, x_{c,t-\frac{1}{4}}, x_{c,t-\frac{1}{4}}\right)$

Day 28: $\left(\pi^M_{c,t-1}, \underbrace{x_{c,t-\frac{3}{4}}}_{\text{day 7}}, \underbrace{x_{c,t-\frac{2}{4}}}_{\text{day 14}}, \underbrace{x_{c,t-\frac{1}{4}}}_{\text{day 21}}, \underbrace{x_{c,t}}_{\text{day 28}}\right)$

Note that this within-month random-walk update does not require estimation of any parameter and works well in our sample.

4.3 Step 2: nowcasting product group-specific inflation

Policymakers, professional forecasters, and market participants regularly monitor the inflation rates of important product groups. Given the restricted coverage of the GFK:FMCG data, we focus on the high-level product groups unprocessed food, processed food and non-energy industrial goods, which receive considerable attention in the Eurosystem, and the more disaggregated low-level product groups unprocessed fruit and vegetables; unprocessed meat, fish and eggs; processed fruit and vegetables; processed meat, fish and eggs; bread and cereals; dairy products and fat; beverages and other food products; and non-durables.

A natural starting point for nowcasting at the product-group level is to treat all COICOP-10 series belonging to that group as relevant predictors. As the U-MIDAS setting is not suited to handling such a large set of predictors, we follow the related literature which successfully applied shrinkage estimators in such data-rich environments (see, for instance, Garcia et al., 2017; Medeiros et al., 2021; Joseph et al., 2022).

We try two modeling approaches. Our first approach avoids mixed frequencies by aggregating the weekly GFK:FMCG indicators, $x^{(m)}_{c,t}$, to the monthly frequency, which yields a single $x^{(M)}_{c,t}$. Then, we apply standard shrinkage methods to estimate nowcasting models of the group-specific target inflation rates.[20] Specifically, we use the least absolute shrinkage and selection operator (LASSO), the ridge and the elastic net estimator.[21] We proceed as follows. Suppose we want to nowcast, on day 14 of

[20] Applying penalized U-MIDAS regressions to the large set of predictors defined at the weekly frequency (four weekly series for each predictor) is also feasible; however, this approach does not recognize serial dependence across high-frequency lags and thereby may be subject to random selection. Zhao and Yu (2006) show that LASSO selects the true model consistently if and (almost) only if the irrelevant covariates are not highly correlated with the predictors in the true model ("irrepresentable condition").

[21] We use the elastic net without tuning the relative weights of the L1 and L2 norms. Instead, we

month T, the official inflation rate, $\pi^M_{g,T}$, of product group g. We first regress $\pi^M_{g,t}$ on the full set of contemporaneous predictors $x^{(M)}_{c,t}$ belonging to group g using the sample $t = 1, \ldots, T-1$ of all monthly data available on that day. This yields a vector of estimated parameters $\hat{b}_g$. We then construct the day 14 estimates of $x^{(M)}_{c,T}$, $\forall c \in g$, as described in the previous step, substitute them on the right-hand side of the regression equation and compute a nowcast based on the estimated parameters $\hat{b}_g$.

The second approach applies the sparse-group LASSO (sg-LASSO) estimator proposed by Babii et al. (2022) and regresses the group-specific target inflation rate, $\pi^M_{g,t}$, directly on the large set of weekly GFK:FMCG inflation rates using orthogonal Legendre polynomials as the aggregation scheme. This approach has the advantage in that it performs shrinkage in a mixed-frequency rather than a low-frequency setting by recognizing serial dependence across different high-frequency lags, also taking into account the time series nature of the data.[22]

The tuning parameters of the aforementioned machine learning tools are determined in a data-driven manner using cross-validation to obtain optimal prediction performance.[23] Finally, to evaluate the nowcast precision of these machine learning approaches based on weekly GFK:FMCG information, we fit SD-AR benchmark models directly to the group-specific target inflation rates, $\pi^M_{g,t}$, from which we construct time series forecasts.

4.4 Step 3: nowcasting headline inflation

To nowcast headline inflation, we split it into the following six components: unprocessed food, processed food, non-energy industrial goods, energy, package holidays, and services. We nowcast each part separately and construct a headline nowcast by using the official HICP weighting scheme.

Bottom-up U-MIDAS approach. Our first nowcasting approach follows a bottom-up strategy. We estimate one nowcasting model for each COICOP-10 item and aggregate the item nowcasts to the six components of headline inflation using the official HICP weights. For those items that are matched by weekly GFK:FMCG data, we use the U-MIDAS model presented in Section 4.2. For all other items, except for energy and package holidays, we use the SD-AR model described in Section 4.1. This

impose equal weights.

[22]We use a Legendre polynomial of degree $L = 0$ which attributes equal weights to all high-frequency lags and delivered similar results compared to other choices of L but at a lower computational cost (see Section 6). By contrast, the $L = 1$ polynomial leads to an increasing linear function and thereby favors more distant lags, $L = 2$ features higher weights to very recent and more distant lags, and so on. See Section B.2 of Appendix B for a more formal description of the machine learning methods.

[23]We tune the hyperparameters of sg-LASSO via expanding cross-validation splitting the in-sample data into $k = 5$ folds and tests on the $k + 1^{\text{th}}$ fold so that it accounts for the time series nature of the data, although it only uses the end of the sample as the test set. Cross-validation of the standard shrinkage methods (LASSO, ridge and elastic net) also uses a training split of $k = 5$ folds but hereby assumes independent and identically distributed samples, which is also valid in a time series context provided the models yield uncorrelated errors (Bergmeir, Hyndman, and Koo, 2018). For a review of these cross-validation methods, see Goulet Coulombe et al. (2022).

implies that the nowcasts for almost all service items are based on the SD-AR model as we do not have any high-frequency indicator for services.

The energy component consists of 14 COICOP-10 items relating to household energy and fuels. Since the weekly energy prices discussed in Section 3.2 are not always good counterparts to these items, we proceed as follows. We first select, at each recursion of the nowcasting experiment, the weekly series most strongly correlated with the COICOP-10 item at hand from a relevant subset. We then run a U-MIDAS model including the selected series as predictor.[24]

Package holidays typically receive a weight of around 3%-4% in the overall HICP consumption basket (except in the years 2021 and 2022, when the COVID-19 pandemic reduced it to roughly 1%). Nevertheless, this component accounts for relevant fluctuations in headline inflation due to its high volatility, which is, unsurprisingly, dominated by a strong seasonal pattern. It consists of two COICOP-10 items, domestic and international package holidays, whereas international travels correspond to more than 95% of the total index. We model both series separately and account for the methodological change from January 2019[25], which revised them backward until 2015. Hence, we keep the real-time perspective of the exercise by producing nowcasts for the non-revised series until December 2018 – using June 2012 as the starting point of the sample due to a structural break in the seasonal pattern – while the revised series becomes the target as of January 2019 with a break-free sample starting in January 2015. Package domestic holidays are modeled with an AR structure with lags $\{1, 2, 6\}$ on the month-over-month growth rates augmented with seasonal dummies. For package international holidays, we fit the log level on the set of seasonal dummies and weekly AMADEUS indices aggregated to the monthly frequency – more precisely, "last-minute" bookings (see Section 3.2).[26] Moreover, we correct the nowcasts of this model because time-invariant dummies cannot account for changes in the price level over time. Specifically, we adjust each nowcast by the nowcast error observed on average in the previous two months. Even though this model is specified solely in monthly data, we summarize it under the heading "U-MIDAS approach" to keep the labeling simple.

Direct machine learning approach. Our second nowcasting approach uses machine learning and directly targets the six components of headline inflation. Given the fairly good performance of the LASSO estimator within the horse race presented

[24]Specifically, we use the following subsets to select from for each of the 14 COICOP-10 items: For "Electricity", "Natural gas, excl. share in the costs", "Share in the costs for gas central heating", and "District heating", we do not have well-matching weekly information available, which is why we select from the series provided by the WOB, the European Gas Spot Index (EGSI), and, for "District heating", a lagged moving average of heating oil prices. For "Liquefied gas, charging of a tank container" we select from LPG and heating oil prices. For "Heating oil", "Share in the costs for oil central heating", and "Coal briquettes", we select from heating oil prices provided by the WOB and Heizoel24. For "Firewood, wood pellets or the like", we use the series on pellet prices. For "Diesel fuel, cetane number below 60", "Diesel fuel, cetane number 60 and more", "Supergrade petrol, 95 octane", "Supergrade petrol, 98 octane and more" we select from the direct counterparts provided by the WOB and Tankerkoenig. For "Liquefied petroleum gas", we use the LPG prices provided by the WOB.

[25]For more details, see Box 5 in ECB Economic Bulletin Issue 2, 2019.

[26]The results are very similar when we use the first principal component of the AMADEUS indices instead of their average.

in Section 4.3, we replicate this ML strategy for unprocessed food, processed food, non-energy industrial goods and energy. Thereby, we proceed by first aggregating the weekly indicators of GFK:FMCG and ENERGY to the monthly frequency and specifying a LASSO regression of each target month-over-month inflation rate on the corresponding predictors. For package holidays, we directly model the COICOP-3 component that combines domestic and international packages, but using different models for both the non-revised series up to 2018 and the revised series as of 2019. In the former case, we implement an AR with lags $\{1, 4\}$ on the month-over-month growth rates augmented with seasonal dummies and monthly rates of the AMADEUS series. For the revised series, we replicate the log level regression, introduced above. As for service inflation, we cannot fit a direct forecasting model as we do not have any relevant high-frequency indicator available. Therefore, we resort to the same bottom-up nowcasts derived from the SD-AR model used in the bottom-up nowcasting approach above.

Benchmark approach. For the six components of headline inflation, we use aggregated nowcasts of the SD-AR model applied to each item at the COICOP-10 level as the benchmark. For headline inflation, we also use market expectations provided by Bloomberg and Consensus Economics as the benchmark. Each month, these data providers ask economists mainly working in private banks to report their best estimate for headline inflation in Germany, along with many other macroeconomic variables. The Bloomberg survey generally takes place within the third week of a given month and is available since 2015. Consensus Economics has only recently started to ask survey participants to also provide an estimate for inflation in the current month. The survey is conducted within the second week of each month and data is available since March 2021.

In Figure 2, we summarize the timeline of the different data releases of market expectations and HICP releases between September 2022 and November 2022. At the end of each of the four weeks in each month, we produce a new nowcast based on incoming high-frequency data. Note that in the first week of September, the nowcasts need to use our previous estimate for inflation, since the official numbers for the COICOP-10 series have not yet been published. The statistical office publishes a first estimate for headline inflation at the end of the corresponding month, for example, on September 29th, followed by the final numbers in the subsequent month, in this case, on October 13th. Consensus Economics has surveyed its participants on September 12th, and Bloomberg between September 22th to September 26th.

4.5 Real-time information set on non-standard policy measures, annual updates of HICP weights and imputed prices in times of crisis

Our evaluation period has been marked by extraordinary events as the outbreak of the COVID-19 pandemic in 2020 and Russia's invasion of Ukraine in 2022. These events brought with them a set of non-standard policy measures which also affected consumer prices, such as a temporary VAT cut in 2020 and one-time emergency aid measures for gas, heating and public transport. Notably, all these policy measures were communicated by the government well before their introduction and therefore

Figure 2: Timeline of market expectations and data releases

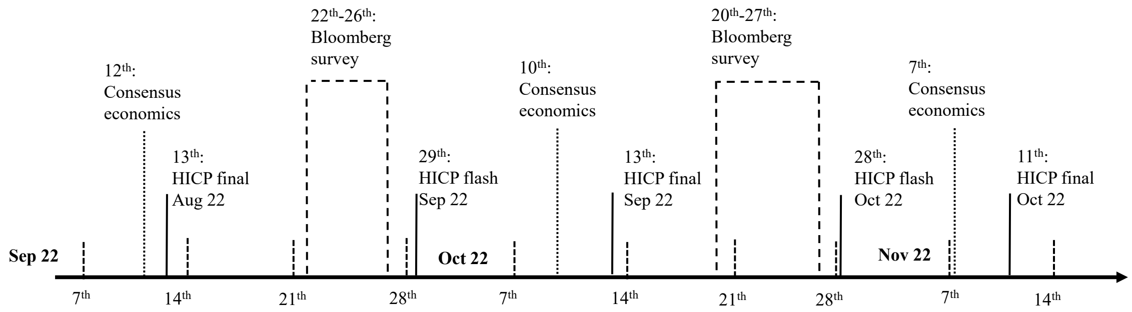

Note: The figure shows, for the months September 2022 until October 2022, the four weeks ending at days 7, 4, 21, and 28, at the end of which we produce a nowcast based on new high-frequency data. In addition, it shows the data releases for inflation and survey periods of market expectations.

included in professional forecasters' information set at that time. Hence, for a proper comparison with market-based inflation expectations, we include *a priori* information on the introduction and reversal of a given policy measure in our nowcasting models. More precisely, we apply ex-ante assumptions when nowcasting the months 2020M7 and 2021M1 (introduction and reversal of temporary VAT cut), 2022M7 and 2022M9 (introduction and ending of heavily reduced public transport tickets) as well as 2022M12 (one-time emergency aid for gas and heating).[27] If we exclude these *a priori* information on policy measures, our results for headline inflation would be obscured by larger nowcasting errors in those components where we do not have high-frequency price information (e.g. services and durable goods).

Finally, the COVID-19 pandemic also hampered price collection by statistical offices considerably. This is notably true for the prices of travel-related services which had to be estimated during lockdown periods, typically by using the month-over-month inflation rate of the previous (non-pandemic) year to reflect the seasonal pattern of these prices. Since this imputation procedure for package holidays was publicly known at that time, we take it into account when nowcasting the HICP subcomponent of package holidays.[28] Moreover, disaggregate HICP weights are updated at the beginning of a year, but typically only published with the final January HICP figures in February. Hence, in our nowcasting models, we use the previous year's weight information when nowcasting the headline rate for January.

[27] Appendix A.4 provides more details on the specific policy measures and how they are accounted for within our econometric framework.

[28] In Germany, the prices for package holidays were imputed during the period 2020M4-2020M6 and 2020M9-2021M5. See Eurostat's "Information on imputations made related to COVID-19".

5 Results

We present the results of our recursive out-of-sample nowcast experiment in three steps. First, we demonstrate the usefulness of including weekly GFK:FMCG data in nowcasting models of HICP inflation at the COICOP-10 level, i.e., the most disaggregate available COICOP level. Then, we show that combining these data with machine learning methods yields superior forecasts at more aggregate product-group levels. Finally, we document that even headline inflation nowcasts benefit considerably when this approach is used, outperforming market expectations in most periods.

5.1 Results of the item-level inflation nowcasts

Table 3 reports the root mean squared error (RMSE) of the monthly inflation nowcasts by the U-MIDAS model relative to the benchmark SD-AR model for selected COICOP-10 items for which weekly GFK:FMCG data are available, while absolute RMSE values for all items are presented in Table C4 of Appendix C. We group the results into panels of product groups such as unprocessed fruit and vegetables, processed meat and fish, and dairy products and fat. Within each panel, the columns refer to the within-month information sets of day 7, 14, 21, and 28. For readability, each cell is colored in a heatmap style where darker colors indicate a lower relative RMSE and hence a better nowcasting performance of the U-MIDAS model compared to the SD-AR benchmark.

For many products and nowcasting days, feeding the weekly GFK:FMCG information in a U-MIDAS model reduces the nowcast error substantially relative to the benchmark. Consider, for example, the RMSE of nowcasting the official year-over-year inflation rate for sweet peppers, which is the first item in the upper left panel. Day 7 nowcasts with the U-MIDAS model yield a relative RMSE of 0.51, cutting the nowcast error almost by half. To understand this result, recall that on day 7 of a month t, the latest official information available pertains to month $t-2$. Hence, the SD-AR benchmark model effectively needs to generate a two-step ahead forecast. By contrast, the U-MIDAS model uses the GFK:FMCG data, which are complete for month $t-1$ and even include the first week of month t. This information advantage can be expected to produce a more accurate nowcast as long as the GFK:FMCG index is reasonably correlated with its official counterpart.

On day 14, the relative RMSE of the U-MIDAS model for sweet peppers increases slightly to approximately 0.56. It remains around this level on days 21 and 28. Why is the superiority of incorporating GFK:FMCG data mitigated compared to day 7? On days 14, 21, and 28 of a month t, the official HICP of sweet peppers for month $t-1$ has already been published. The benchmark model thus only needs to generate a one-step ahead forecast so that the informational advantage of using GFK:FMCG data declines in this respect. At the same time, the U-MIDAS model can exploit additional weekly GFK:FMCG information of month t which may result in a more stable signal concerning this month. However, for sweet peppers – and for many other items – the effect of expanding the information set to days 22 and 28 does not lead to (notably) better nowcasts, which is probably related to the typical practice in official price statistics of recording many prices around the middle of a month.

Across product groups, GFK:FMCG information improves the nowcasts at the item level particularly strongly for unprocessed fruit and vegetables as well as dairy products and fat, for which RMSE reductions in the range of 40%-60% can be achieved, also indicated by statistically significant results of the Diebold and Mariano (1995) test of equal predictive accuracy in many of these cases. The U-MIDAS model also exhibits very good nowcasting properties for many items of the other product groups.

What if we directly match the monthly GFK:FMCG inflation series to their official HICP counterparts using OLS? This basic forecasting method (hereafter *OLS match*) avoids using the MIDAS model and constructs the nowcasts by aggregating the available weekly GFK:FMCG information to the monthly frequency. Table C3 in Appendix C replicates the relative RMSE figures and reveals that nowcasting accuracy can be slightly improved across items that already perform exceptionally well using U-MIDAS. This applies to most items within the category unprocessed fruit and vegetables and dairy products. By contrast, U-MIDAS still yields smaller RMSE values across product groups that cannot surpass nowcasting gains of 50%. Therefore, improvements in relationship to the SD-AR benchmark are mostly stemming from precise weekly GFK:FMCG information.

So what drives nowcasting success at the item level? The most important factor is a close match between the GFK:FMCG and official price indices. Figure 3 displays the predictive gain of using the GFK:FMCG price indicators as a function of the in-sample fit with their official counterparts. Evidently, a higher correlation typically goes hand in hand with a smaller RMSE compared to the benchmark model. This negative relation also holds across different nowcasting days. On average, including the GFK:FMCG information starts to reduce the nowcast RMSE once the correlation between the month-over-month rates exceeds 0.4.

Figure 3 also shows some differences between high-level product groups. Processed food items (denoted by red dots) exhibit correlations between the GFK:FMCG and official price indices on the full range between zero and almost 1. They appear to be most representative of the relationship between correlation and nowcast gain just discussed. Most items of unprocessed food (denoted by green dots) show a strong or very strong correlation, but the nowcast benefit is more heterogeneous. By contrast, non-energy industrial goods (blue dots) mostly exhibit moderate to low correlations between the GFK:FMCG and official price indices; thus, including the scanner data in many cases does not pay off. These products are typically characterized by a greater product variety than food items, which makes it more difficult to fit the official price indices.

Table 3: RMSE for FMCG product-level inflation: U-MIDAS relative to the SD-AR benchmark

Unprocessed fruit and vegetables

	day 7	day 14	day 21	day 28
Sweet peppers	0.515**	0.562**	0.555**	0.567**
Butterhead Lettuce	0.555**	0.628**	0.582**	0.63**
Onion or garlic	0.578**	0.714**	0.678**	0.662**
Carrots	0.547**	0.718**	0.727**	0.727**
Grapes	0.75**	0.72**	0.693**	0.716**
Lambs lettuce	0.755*	0.714*	0.714*	0.715*
Cauliflower or cabbage	0.783**	0.758**	0.754**	0.741**

Processed fruit and vegetables

	day 7	day 14	day 21	day 28
Frozen chips or the like	0.522	0.673	0.652	0.64
Apple sauce	0.802	0.806	0.785	0.784
Tinned mushrooms	0.803	0.834	0.852	0.849
Potatoes	0.813	0.974	0.929	0.916
Potato crisps	0.924	0.907	0.918	0.913
Asparagus	0.94	0.913	0.932	0.926
Tinned pineapple	0.923	0.952	0.931	0.944

Unprocessed meat, fish and eggs

	day 7	day 14	day 21	day 28
Eggs	0.666	0.792	0.662	0.659
Fresh poultry	0.791	0.808	0.803	0.808
Minced pork	0.853	0.855	0.849	0.85
Minced beef	0.781	0.885	0.886	0.867
Beef for boiling	0.83	0.871	0.876	0.876
Pork chop or cutlet	0.881	0.902	0.902	0.88
Smoked pork chop	0.881	0.92	0.894	0.905

Processed meat and fish

	day 7	day 14	day 21	day 28
Ham or bacon	0.622	0.673	0.675	0.675
Salami or sausage	0.676	0.727*	0.72*	0.72*
Liver sausage	0.705	0.771	0.763*	0.769*
Fish fingers	0.751	0.777	0.755*	0.785
Lyoner pork sausage	0.733	0.791	0.771	0.781
Fried sausage	0.777	0.837	0.816	0.81
Prepared minced meat	0.807	0.863	0.848	0.837

Beverages and others

	day 7	day 14	day 21	day 28
Baking powder	0.544	0.675	0.685	0.679
Orange juice	0.713	0.76	0.763*	0.766*
Sugar	0.699	0.892	0.848	0.85
Apple juice	0.746	0.918	0.867	0.848
Blancmange powder	0.846	0.871*	0.858*	0.869
Food for infants	0.833**	0.868**	0.884*	0.889*
Ketchup	0.931	0.863	0.853	0.859

Bread and cereals

	day 7	day 14	day 21	day 28
Flour	0.582	0.557	0.565	0.557
Pasta	0.73	0.769*	0.774*	0.796
Rice	0.758	0.792*	0.806	0.809*
Ready to bake rolls	0.746	0.817*	0.823*	0.826*
Crisp bread	0.779	0.846*	0.838*	0.869*
Cornflakes and muesli	0.832	0.847	0.862	0.853
Frozen cake, tart or pie	0.795	0.872	0.867	0.875

Non-durables

	day 7	day 14	day 21	day 28
Aluminium foil	0.72	0.768*	0.752*	0.749*
Toilet tissue	0.8	0.8	0.796	0.803
Hair shampoo	0.725**	0.827**	0.857**	0.83**
Deo spray or deo roll on	0.768*	0.873	0.859*	0.851*
Hair spray or gel	0.839	0.85*	0.846*	0.858
Paper handkerchiefs	0.924	0.832	0.83	0.834
Toothpaste	0.866	0.887*	0.889*	0.885**

Dairy products and fat

	day 7	day 14	day 21	day 28
Whole milk	0.39	0.437	0.416	0.42
Condensed milk	0.359	0.442	0.437	0.436
Low fat milk	0.412	0.427	0.425	0.428
Cream	0.424*	0.461	0.455	0.457
Butter	0.433**	0.498**	0.457**	0.458**
Curd	0.539	0.638	0.622	0.612
Sliced cheese	0.593	0.68	0.678	0.68

Colour scale: 1, 0.9, 0.8, 0.7, 0.6, 0.5, 0.4, 0.3, 0.2

Sources: GfK household panel; own calculations.
Notes: The figure shows heatmaps of RMSE values for the U-MIDAS model relative to the SD-AR benchmark at nowcasting days 7, 14, 21 and 28 for the best-performing COICOP-10 items within selected FMCG product groups. Results for the Diebold and Mariano (1995) test in the event of outperformance relative to the SD-AR model are indicated by the symbols * (5% level) and ** (1% level).

Figure 3: Predictive gain of the FMCG data as a function of the in-sample fit with official counterparts

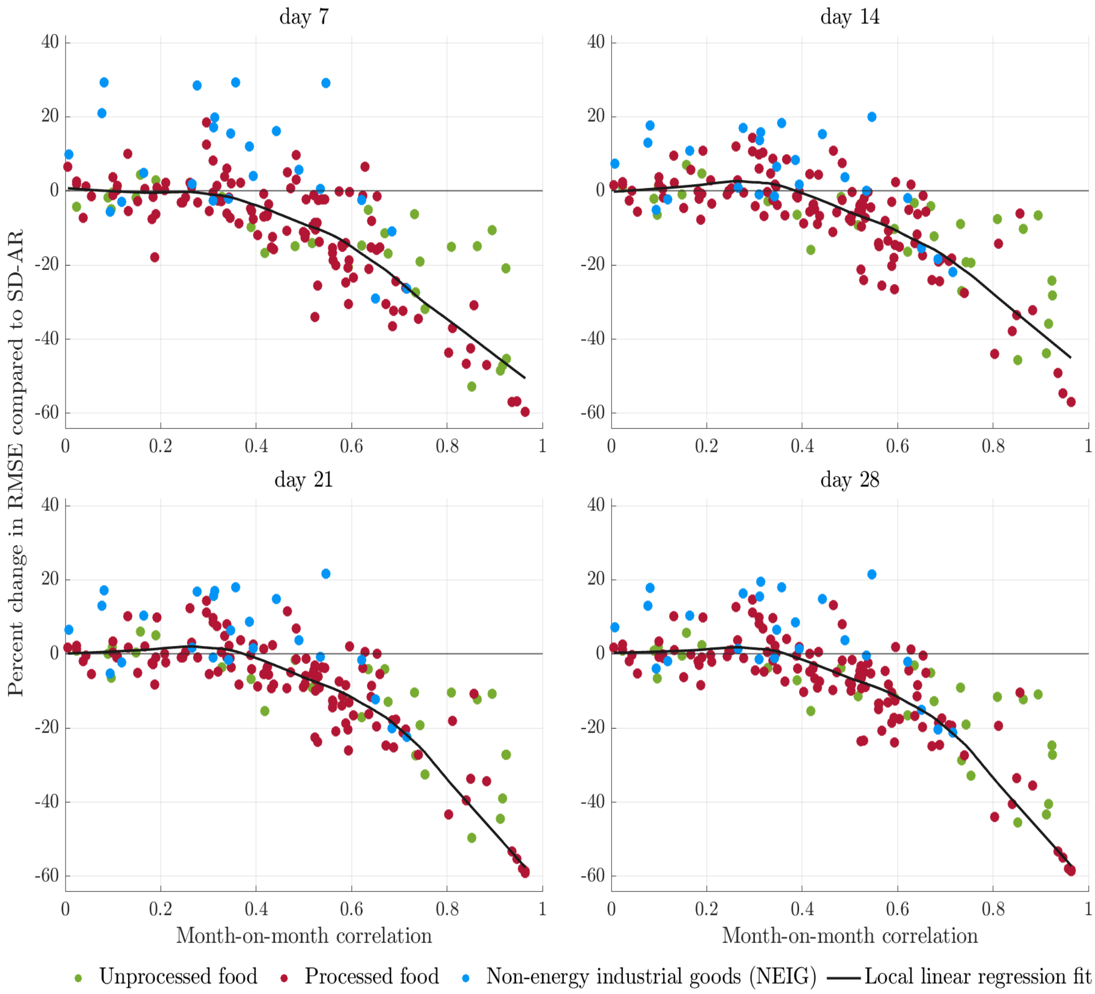

Sources: GfK household panel; own calculations.
Note: For each FMCG item at the COICOP-10 level, the figure shows the percentage change in RMSE of the U-MIDAS nowcasts (on days 7, 14, 21 and 28) compared to SD-AR as a function of the fit between GFK:FMCG indicators and their official counterpart based on correlations using month-over-month rates. Outliers at the 1st and 99th percentiles of the RMSE changes are removed.

5.2 Results of the product group-specific inflation nowcasts

The inflation rates of the three high-level product groups unprocessed food, processed food and non-energy industrial goods (NEIG) receive considerable attention in the Eurosystem. As described in Section 4.3, we use machine learning shrinkage methods to estimate direct nowcasting models for these product groups (and some subgroups) that successfully include the large set of regressors we have available, namely the underlying weekly GFK:FMCG price indicators at the COICOP-10 level. We compare these models to SD-AR benchmark models fitted to the group-specific inflation rates.

Table 4 displays the RMSEs of the group-specific shrinkage models relative to their benchmark. The top panel refers to the three high-level product groups. With regard to unprocessed and processed food, the shrinkage models clearly outperform the benchmark with reductions in the RMSE between 15% and 25% on all nowcasting days. By contrast, for non-energy industrial goods (NEIG) the benchmark dominates. This outcome likely reflects the different coverage rates of the GFK:FMCG data across product groups. As shown in Table 2, 30 out of 38 COICOP-10 items of unprocessed food and 116 out of 142 COICOP-10 items of processed food are matched, but only 39 out of the 302 NEIG with semi-durables and durables almost lacking completely. In addition, even the relatively few matched NEIG items, mostly non-durables, do not correlate very strongly with their HICP counterparts, as reported in Figure 3.

The bottom panel in Table 4 shows the results for the eight more disaggregated low-level product groups. In the large majority of cases, it pays off considerably to use shrinkage models that include weekly GFK:FMCG data. The advantage is particularly large for dairy products and fat (reduction in RMSE of roughly 45% to 55%), unprocessed fruit and vegetables (reduction of around 20% to almost 40%), processed meat and fish (reduction of more than 25% to almost 40%), and unprocessed meat, fish and eggs (reduction of nearly 20% to 25%). For processed fruit and vegetables, bread and cereals, and beverages and others, the nowcasting gains are more muted, but it is still generally beneficial to use shrinkage models. Only in the case of non-durables is there no clear difference to the benchmark, which likely again reflects the low correlation of the GFK:FMCG items at the COICOP-10 level with their HICP counterparts.

The weekly flow of information affects the nowcasting performance in a way very similar to the underlying COICOP-10 items discussed in the previous section. Most importantly, the information available at day 7 of a given month already turns out to be highly valuable. This likely reflects the fact that at day 7 of a month t, the benchmark model includes only official inflation rates of month $t-2$, while the shrinkage approaches use the full GFK:FMCG data of month $t-1$ and the first week of month t. The additional information exploited at day 14 typically further improves the nowcasts in absolute terms (see Figure C3 of Appendix C for the absolute forecast errors over time by nowcasting day), whereas this is not always the case relative to the benchmark, which on that day includes the official inflation rates of month $t-1$. Finally, the additional information gained in weeks 3 and 4 of a month is of minor quantitative importance.

Concerning the different shrinkage approaches, the nowcasting results do not favor a single method. The general conclusion is that it is important to include the weekly GFK:FMCG dataset and make it usable in an appropriate way. To this end, standard shrinkage methods (LASSO, ridge and elastic net) work generally as well as the sg-LASSO approach, which performs variable selection in a mixed-frequency setting and fully accounts for the time series nature of the dataset. Nevertheless, for processed food, which is the product group with by far the largest number of underlying COICOP-10 items that we match with GFK:FMCG data, the sg-LASSO is clearly superior. This may indicate that this approach is especially promising when it comes to very high-dimensional estimation and nowcasting settings.

Table 4: RMSE for FMCG product-group inflation: Various shrinkage methods relative to the SD-AR benchmark

Unprocessed food

	day 7	day 14	day 21	day 28
LASSO	0.788*	0.788**	0.755**	0.741**
Ridge	0.838*	0.832**	0.814**	0.805**
Elastic Net	0.796*	0.784**	0.754**	0.755**
sg-LASSO	0.788*	0.775**	0.778**	0.794**

Processed food

	day 7	day 14	day 21	day 28
LASSO	0.846	0.861	0.805*	0.81*
Ridge	0.865	0.875	0.824*	0.811*
Elastic Net	0.855	0.854	0.811*	0.82*
sg-LASSO	0.829	0.779	0.774*	0.781*

Non-energy industrial goods

	day 7	day 14	day 21	day 28
LASSO	1.093	1.075	1.081	1.083
Ridge	1.134	1.081	1.076	1.069
Elastic Net	1.088	1.084	1.073	1.095
sg-LASSO	1.069	1.073	1.08	1.081

(a) High-level product groups

Unproc. fruit and vegetables

	day 7	day 14	day 21	day 28
LASSO	0.658*	0.651**	0.623**	0.62**
Ridge	0.788*	0.81**	0.794**	0.79**
Elastic Net	0.658*	0.656**	0.626**	0.63**
sg-LASSO	0.722*	0.651**	0.643**	0.668**

Unproc. meat, fish and eggs

	day 7	day 14	day 21	day 28
LASSO	0.751	0.812	0.801	0.796
Ridge	0.763	0.812	0.808	0.79
Elastic Net	0.761	0.817	0.818	0.794
sg-LASSO	0.743	0.812	0.802	0.784

Processed fruit and vegetables

	day 7	day 14	day 21	day 28
LASSO	0.785	0.942	0.916	0.899
Ridge	0.83	0.939	0.911	0.88
Elastic Net	0.784	0.928	0.902	0.881
sg-LASSO	0.779	0.94	0.902	0.878

Processed meat and fish

	day 7	day 14	day 21	day 28
LASSO	0.675	0.723	0.675	0.638
Ridge	0.673	0.73	0.678	0.651
Elastic Net	0.674	0.728	0.665	0.631
sg-LASSO	0.622	0.711	0.68	0.668

Bread and cereals

	day 7	day 14	day 21	day 28
LASSO	1.004	0.912	0.856	0.827
Ridge	1.061	0.965	0.885	0.856
Elastic Net	0.999	0.914	0.864	0.828
sg-LASSO	0.967	0.82	0.829	0.841

Dairy products and fat

	day 7	day 14	day 21	day 28
LASSO	0.463*	0.521*	0.455*	0.442*
Ridge	0.563*	0.595*	0.534*	0.503*
Elastic Net	0.484*	0.52*	0.45*	0.446*
sg-LASSO	0.47*	0.531*	0.514*	0.507*

Beverages and others

	day 7	day 14	day 21	day 28
LASSO	0.963	0.953	0.93	0.905
Ridge	0.949	0.957	0.926	0.919
Elastic Net	0.953	0.952	0.929	0.917
sg-LASSO	0.933	0.908	0.914	0.917

Non-durables

	day 7	day 14	day 21	day 28
LASSO	1.076	1.026	0.995	0.977
Ridge	1.064	1.041	1.001	0.981
Elastic Net	1.079	1.028	0.994	0.983
sg-LASSO	1.048	1.013	1.004	1.004

(b) Product groups

Sources: GfK household panel; own calculations.
Notes: The figure shows heatmaps of RMSEs for nowcasts based on (i) shrinkage methods (LASSO, ridge and elastic net) and (ii) sg-LASSO relative to the SD-AR benchmark for FMCG higher-level components and subcomponents. Results for the Diebold and Mariano (1995) test in the event of outperformance relative to the benchmark are indicated by the symbols * (5% level) and ** (1% level).

5.3 Results of the headline inflation nowcasts

Table 5 shows the RMSE of the bottom-up U-MIDAS approach and the direct machine learning approach relative to the benchmark for headline inflation and its six components. Let us first focus on the bottom-up U-MIDAS approach. Recall that it uses U-MIDAS models to nowcast each COICOP-10 item that is matched by weekly GFK:FMCG, energy price, or international package holiday data, and the SD-AR model for the remaining items before it aggregates all these nowcasts with the help of the official HICP weights. By contrast, the benchmark model fits the SD-AR model to all COICOP-10 items, while it uses the same aggregation strategy. Hence, the relative RMSE tells us in a clear way how beneficial the inclusion of weekly external data for aggregate inflation forecasting is.

From the left panel of Table 5 we infer that the weekly external data make a considerable difference. On day 7, the informational advantage relates to unprocessed food, processed food, and energy, reducing the RMSE by roughly 30%, 25% and 47% respectively, whereas the SD-AR model is sufficient to nowcast package holidays, non-energy industrial goods, and services (for which we do not have any weekly indicators). This translates to a reduction in the RMSE of headline inflation by 18% compared to the benchmark. The informational advantage increases to levels close to 32% for headline inflation as more information accumulates over days 14, 21, and 28, mainly because nowcasts of the energy component and package holidays become more accurate. In fact, using the OLS match approach with monthly aggregated estimates of price changes – see Table C5 in Appendix C – leads to slightly improved performance for most components and headline inflation.

The right panel of Table 5 shows how the machine learning approach performs. Recall that it directly models the inflation rates of the six components as a function of all underlying weekly data and then aggregates these six nowcasts to headline inflation. Except for unprocessed food, the nowcasts generally deteriorate when compared to the bottom-up U-MIDAS both for the components (especially processed food) and to a smaller extent, for headline inflation. This result suggests that – in a setting where official aggregation schemes are known and easy to implement – there is no systematic advantage of employing machine learning models to estimate data-dependent aggregation weights. While the latter may, in theory, better reflect the dependence structure of the target series and its underlying predictors, we conjecture that in our case, the time series used to estimate them are neither long nor stable enough to outweigh the associated increase in estimation and thus nowcasting variance.

Is this disadvantage of the direct machine learning approach related to specific periods? To shed light on this issue, Figure 4 displays the squared forecast errors of the six components of headline inflation over time, where we focus on day 28. Note that we multiply these squared forecast errors with their COICOP weights, which we use to aggregate the component-wise forecasts to the headline forecast, in order to obtain an impression of their overall relevance – keeping in mind, of course, that the weighted sum of the six squared forecast errors (which we show in the figure) is not equal to the square of the weighted sum of the six forecast errors (which amounts to the headline forecast error). Nevertheless, comparing the top and bottom panels shows that the energy shocks of the year 2022 are the main source of forecast er-

Table 5: RMSE of headline inflation and its components: bottom-up U-MIDAS and direct machine learning approaches relative to the benchmark approach

	Bottom-up U-MIDAS				Direct ML			
	day 7	day 14	day 21	day 28	day 7	day 14	day 21	day 28
Unprocessed food	0.697^{*}	0.736^{*}	0.725^{*}	0.722^{*}	0.803	0.738^{**}	0.717^{**}	0.707^{**}
Processed food	0.743	0.767^{*}	0.763^{*}	0.761^{*}	0.985	0.918	0.886	0.858
Energy	0.532^{**}	0.36^{*}	0.323^{*}	0.324^{*}	0.637^{*}	0.423^{*}	0.463	0.355^{*}
Package holidays	1.102	1.064	1.009	0.941^{*}	1.056	1.058	1.035	0.998
NEIG	1.005	0.997	0.997	0.997	1.126	1.118	1.124	1.119
Services	1	1	1	1	1	1	1	1
Headline	0.829	0.721	0.701	0.681	0.821	0.744	0.746	0.7

Sources: GfK household panel; European Commission's Weekly Oil Bulletin; AMADEUS; own calculations.
Note: The figure shows heatmaps of RMSEs for nowcasts based on (i) the bottom-up U-MIDAS approach with aggregation via HICP weights and (ii) the direct machine learning relative to the benchmark approach, which is a bottom-up nowcast based on SD-AR models fitted at the COICOP-10 level. Results for the Diebold and Mariano (1995) test in the event of outperformance relative to the benchmark are indicated by the symbols * (5% level) and ** (1% level).

rors for the U-MIDAS approach, but even more so for the direct machine learning approach. On a smaller scale but still relevant are the forecast errors relating to non-energy industrial goods – where the direct machine learning in particular exhibits weaknesses during the pandemic and thereafter – while package holidays account for the majority of the errors throughout normal times like those prior to the pandemic, especially around 2018. [29] These findings indicate that estimating data-dependent aggregation weights is particularly detrimental in turbulent times.

Finally, we benchmark our model-based nowcasting results with expectations surveyed by Bloomberg and Consensus Economics among market participants. To that end, we report the cumulative sum of the loss differential of our model-based nowcasts versus the median survey expectations. Using the squared forecast error as our loss measure, we calculate the differential as

$$D_{t,ij} = -\sum_{\tau=1}^{t}(e^2_{\tau,M_i} - e^2_{\tau,S_j}), \qquad t = 1, \dots, T, \tag{7}$$

where e_{t,M_i} denotes the nowcast error of model M_i (either U-MIDAS or direct machine learning) and e_{t,S_j} denotes the nowcast error of market survey S_j (either Bloomberg or Consensus Economics). A positive value of $D_{t,ij}$ indicates that model i outperforms survey j while negative values imply the opposite.

Figure 5 shows the evolution of the loss differentials over time for nowcast days 14, 21, and 28 which are competitive to the Bloomberg survey.[30] The bottom-up

[29] See Figure C3 in Appendix C for a period-wise illustration of these forecast errors for GFK:FCMG groups.

[30] We leave out loss differentials for day 7 because the Bloomberg survey outperforms any model based on this very limited information set by far during the post-pandemic sample, compressing the scale of the axis measuring the loss differential so that the differences between days 14, 21, and 28 become difficult to distinguish. Figures that include day 7 are available upon request.

Figure 4: Contribution of HICP components to squared headline forecast errors

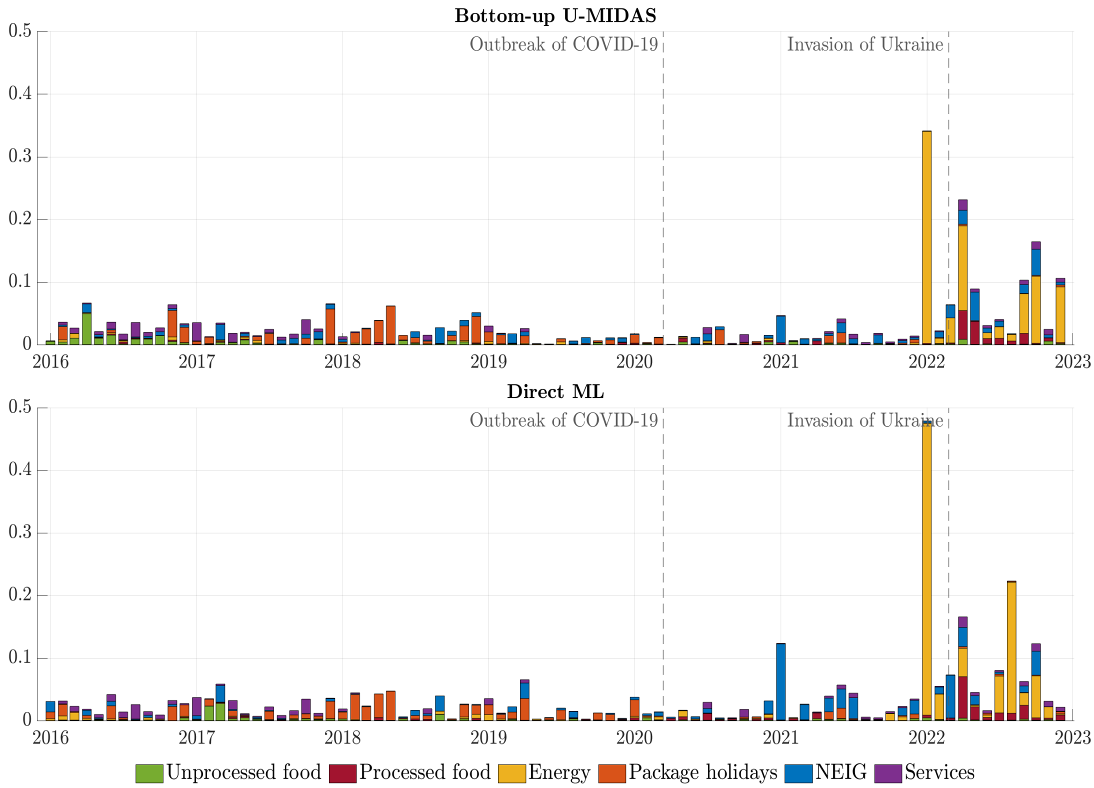

Sources: GfK household panel; European Commission's Weekly Oil Bulletin; AMADEUS; own calculations.
Notes: The figure shows the weighted squared forecast errors of the six components of headline inflation on day 28 of a month for the bottom-up U-MIDAS approach (top panel) and the direct machine learning approach (bottom panel).

U-MIDAS is roughly on par with Bloomberg expectations until the pandemic hits, while the direct machine learning approach even outperforms consistently during this period if at least high-frequency information up to day 21 is used. In fact, the modest underperformance of the bottom-up U-MIDAS compared to direct ML until 2020 stems primarily from larger nowcast errors for unprocessed food (see Figure 4). Following the pandemic period, forecasting gains in relation to the Bloomberg benchmark accumulate gradually and consistently over time so that our modeling strategies clearly outperform if at least information up to day 14 is included for the bottom-up U-MIDAS and day 28 for the direct ML.

Surprisingly, in January 2021, our simple ex-ante assumption for the VAT impact (see Section 4.5) led to a considerable jump in the performance of the bottom-up U-MIDAS compared to Bloomberg expectations. Moreover, the rising inflation scenario following the pandemic in 2021 reveals itself to be a challenging period for the direct machine learning approach in terms of keeping its cumulative advantage, especially using only the information set of days 14 and 21. Similarly, Russia's invasion of Ukraine favors the Bloomberg survey on impact irrespective of the information set; nevertheless, the bottom-up U-MIDAS and the direct ML on day 28 quickly catch up with previous advantage levels. This means that efficiently exploiting daily and weekly external information from energy markets allows us to maintain the edge in relation to Bloomberg expectations throughout 2022. Across information sets, marginal gains in loss differential mostly persist when information up to days 21 and 28 are included. Overall, these findings suggest that estimating data-dependent aggregation weights based on ML methods improves the nowcasting precision of headline inflation in normal times but can be detrimental in unstable environments.

The survey results published by Consensus Economics have only been reported since March 2021, which precludes a comparison over the whole nowcasting sample. Nevertheless, in the 20 months available, the bottom-up U-MIDAS and the direct machine learning approaches strongly outperform the survey expectations, as shown in Figure C4 in Appendix C.

Figure 5: Cumulative sum of the squared forecast error differentials: models versus Bloomberg survey-based expectations

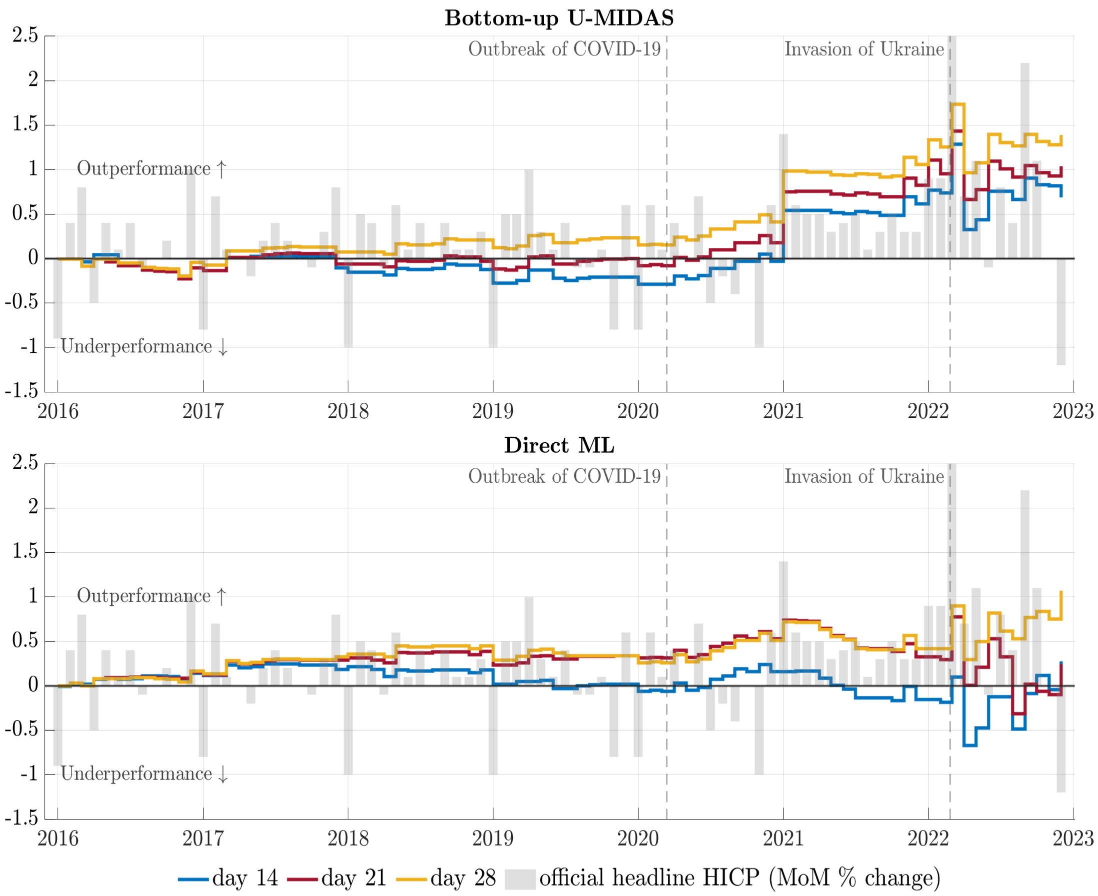

Sources: GfK household panel; European Commission's Weekly Oil Bulletin; AMADEUS; Bloomberg survey; own calculations.
Notes: The figure shows, on the left axis, the cumulative sum of the squared forecast error differential of the bottom-up U-MIDAS approach (top panel) and the direct machine learning in relation (bottom panel), respectively, in comparison to Bloomberg survey-based expectations on days 14, 21 and 28. The gray bars represent official month-over-month percentage changes in headline inflation.

6 Robustness analysis

We evaluate the robustness of the results reported in the previous section with respect to three aspects. First, we focus on the data input and investigate whether alternative methods of compiling GFK:FMCG price indices at the COICOP-10 level lead to different outcomes when compared to the (weighted) TPD regression described in Section 3.1. Second, given the high weekly volatility of some COICOP-10 item indices, we assess whether their predictive content improves by smoothing out volatile time series. Finally, we focus on the modeling strategy: (i) the robustness of the baseline U-MIDAS setting to alternative model specifications, extending the information set by non-contemporaneous weekly inflation rates and quantity indices; and (ii) the stability of product group-specific results to different folds of the cross-validation procedure in machine learning tools and different degrees of the sg-LASSO Legendre polynomial.

6.1 Alternative methods of compiling COICOP-10 price indices

In our baseline setting, we compute weekly price indices from the granular GFK:FMCG data employing a (weighted) TPD regression. In doing so, we follow the practice of statistical agencies worldwide that rely on this method due to its good results with respect to in-sample fit and nowcasting. To examine the extent to which this also holds for our GFK:FMCG data, we have implemented various alternative methods of computing weekly price indices.

We start by transforming daily prices into weekly aggregates using the arithmetic mean, the geometric mean and the median price. We do this both for the raw price data and the data excluding outliers. We consider several choices for the product sample underlying the computation of each price index constructed following the above-mentioned approaches. Due to the constant replacement and entry of new products in the GFK:FMCG dataset, price series can be highly volatile, with sudden spikes and jumps. Official price statistics address this issue by regularly selecting a basket of goods and services that is kept fixed for a specific period of time. Hence, we follow a similar approach by considering only prices of products that are available in each month of the entire sample since 2003, or within each month of the preceding two years. Next, we follow a standard time-dependent rule in official price collection and include only prices of products that have been bought between the 12th and the 18th of each month. Additionally, we focus on goods sold in discount shops, as these shops hold significant pricing power in Germany, causing other shops to adjust their prices accordingly. Finally, we employ a top-seller approach by selecting only those products with the highest market share that cover 50% of the market in the previous two years prior to the current year.

In total, combining the alternative aggregation methods and selection strategies described above, we end up with more than 50 different price indices for each COICOP-10 item. Using only "pre-sample" data dating back to before January 2016, the start of our nowcasting experiment, we then compare the month-over-month inflation correlations of the various methods with their official counterparts. This allows us to rank them and determine the best method based on the highest in-sample correlation for each COICOP-10 series.

Next, we use the respectively best method for each COICOP-10 item to replicate the analysis in Section 5.1. The results in Figure C5 of Appendix C show that the predictive properties found using the TPD-based GFK:FMCG indicators are, across all nowcast horizons, generally not affected when implementing the best pre-sample method by COICOP-10 instead. In fact, the comparison reveals a slight decay in the predictive ability of the best method compared to the TPD approach, especially in cases where the baseline U-MIDAS combined with TPD indicators substantially outperforms the SD-AR benchmark. Overall, the findings support the use of the TPD method as the baseline method since the in-sample fit can only be marginally improved by alternative methods while delivering the best solution from an out-of-sample perspective.

6.2 Smoothing out volatile inflation series

As described in Section 3.3, some COICOP-10 price indices can be highly volatile, which is a common property exhibited by high-frequency data. To filter out short-term noise of volatile GFK:FMCG price series and then reassess their predictive properties, we consider moving average filters using one to four weeks of past data. Figure C6 in Appendix C plots the changes in terms of nowcast RMSE when applying the four-week moving average smoother as a function of the weekly volatility level of each COICOP-10 item. It turns out that the predictive ability of smoothed inflation series worsens on average for the most volatile group, namely unprocessed food items, on days 7 and 14. By contrast, as we approach end-of-month nowcast horizons, changes in RMSEs become less pronounced. These findings indicate that it is difficult to reduce the noise in GFK:FMCG prices without affecting the signal that is related to official inflation dynamics.

6.3 Alternative U-MIDAS specifications and ML hyperparameter choices

In the first robustness check, we enrich the U-MIDAS information set in (6) by including non-contemporaneous GFK:FMCG weekly inflation rates from periods $t-1$ and $t-2$. The results are summarized in Figure C7 in Appendix C. They show that the RMSE is generally stable across COICOP-10 items and nowcasting days after accounting for the weekly inflation rates of $t-1$. Modest RMSE improvements in the range of 10% to 15% can be mostly attributed to items in which U-MIDAS already outperforms, to some extent, the SD-AR benchmark in the baseline scenario, especially when nowcasting on day 7. Unsurprisingly, the same pattern can be identified when additionally accounting for the high-frequency inflation lags of period $t-2$. Hence, incorporating past high-frequency information only leads to minor predictive gains, and mostly does so for processed food and NEIG items. This suggests that the autoregressive component in (6) sufficiently accounts for recent price dynamics while contemporaneous GFK:FMCG information constitutes the key signal for well-performing nowcasts. Moreover, a higher number of distributed lags in the U-MIDAS setting goes hand in hand with an increased nowcasting variance.

The second robustness check investigates whether GFK:FMCG quantity indices improve the quality of disaggregate inflation nowcasts beyond contemporaneous information of price indices. To this end, we add the month-over-month quantity indices at the weekly frequency, and their lags, as regressors in model (6). The RMSE practically remains unchanged across all COICOP-10 series, indicating that the inclusion of quantity information does not enhance the predictability of disaggregate inflation rates beyond what is already conveyed by GFK:FMCG price indices. As an exception, results for the COICOP-10 series "flour" display a noteworthy improvement when nowcasting days 22 and 28. This most likely reflects that the start of the war in Ukraine had strong effects on the average quantity-price relationship of "flour".

With the third robustness check, we examine whether the product group-specific nowcasting results discussed in Section 5.2 hold irrespective of our hyperparameter choices for the estimation of shrinkage methods. We start by increasing the degree of the sg-LASSO Legendre polynomial from $L = 0$ to $L = 1, 2$. The findings suggest a slight improvement in the precision of the nowcast in only a small number of cases where statistically significant outperformance is already achieved by sg-LASSO with $L = 0$ compared to SD-AR. Hence, it represents the optimal choice given that it promotes a higher dimensionality reduction and carries a smaller computational burden when estimating sg-LASSO coefficients. Similarly, we test for different folds of the cross-validation considering the grid set $k \in \{5, 10, 15, 20, 25\}$. Overall, these choices favor similar tuning parameters and model architectures, thus not altering the results for group-specific targets.

7 Conclusion

The recent decade has witnessed a burst in granular high-frequency information stemming from all parts of the economy. To be useful for policymakers and society at large, this vast amount of data needs to be processed with care and fed into appropriate models.

This paper demonstrates how pairing millions of household scanner data with state-of-the-art machine learning techniques yields highly competitive real-time inflation nowcasts for Germany, both at a very disaggregate level as well as for major product groups and headline inflation. The guiding principle of our approach is to use the economic structure inherent in the construction of official price indices to organize and condense the information carried by granular purchase decisions at the household level before we open up the machine learning toolkit. This strategy is reflected in the three steps of our analysis: we start at the most disaggregate level possible, proceed to an intermediate level of product groups, and finally turn to headline inflation.

In the first step, we construct a set of more than 180 weekly price indices at the COICOP-10 level from the granular scanner data. The virtue of this approach is that the mapping from the detailed product descriptions available in the scanner data to the official consumption basket underlying German inflation is straightforward as the COICOP-10 level includes items such as “butter”, “coffee beans”, and “sanitary cleaner”. We show that the scanner-based indices obtained in this way track their official counterparts well, especially for food items (with an average correlation of 0.9 for year-over-year inflation rates). When fed into a MIDAS model, they also improve disaggregate monthly inflation nowcasts, notably as soon as after the first seven days of a month.

In the second step, we turn to nowcasting product groups like unprocessed and processed food, taking all available weekly scanner-based indices into account. To this end, we apply shrinkage estimators to cope with the high-dimensional predictor set. Relative to a time series benchmark model, we obtain substantial predictive gains of up to 25% in terms of relative RMSE.

In the final step, we nowcast headline inflation. Once again, we use the COICOP architecture to do so. Specifically, we split headline inflation into six components which exhibit very heterogeneous time series properties and for which we can come up with different types of weekly predictors. We demonstrate that this approach yields highly competitive nowcasting models that are on par with, or even outperform, survey-based expectations, which are notoriously difficult to beat.

In summary, our strategy to combine fixed economic structures like the COICOP classification system with flexible machine learning tools turned out to provide us with accurate inflation nowcasts at different levels of aggregation. Our approach thereby exploits the virtue of granular data to provide an understanding of the disaggregate dynamics underlying overall inflation, whilst at the same time yielding valuable high-frequency real-time information about price developments of aggregates closely monitored by policymakers and market participants.

Looking ahead, given the considerable value added of high-frequency scanner data for inflation nowcasting documented in this paper, there is an urgent need to identify and exploit high-frequency information concerning those parts of the consumption basket underlying German inflation that are not covered by scanner data on fast-moving consumer goods. These include services, clothing and footwear, and the large range of items typically referred to as slow-moving consumer goods such as furniture, household appliances and the like.

References

Aliaj, T., M. Ciganovic, and M. Tancioni (2023). Nowcasting Inflation with Lasso-Regularized Vector Autoregressions and Mixed Frequency Data. *Journal of Forecasting 42*(3), 464–480.

Alvarez, S. E. and S. M. Lein (2020). Tracking Inflation on a Daily Basis. *Swiss Journal of Economics and Statistics 156*(18).

Anenberg, E. and S. Laufer (2017). A More Timely House Price Index. *The Review of Economics and Statistics 99*(4), 722–734.

Aparicio, D. and M. I. Bertolotto (2020). Forecasting Inflation with Online Prices. *International Journal of Forecasting 36*(2), 232–247.

Assenmacher, K., G. Glöckler, S. Holton, and P. Trautmann (2021). Clear, Consistent and Engaging: ECB Monetary Policy Communication in a Changing World. *ECB Occasional Paper 274*, 1–91.

Atkeson, A. and L. E. Ohanian (2001). Are Phillips Curves Useful for Forecasting Inflation? *Federal Reserve Bank of Minneapolis Quarterly Review 25*(1), 2–11.

Babii, A., E. Ghysels, and J. Striaukas (2022). Machine Learning Time Series Regressions with an Application to Nowcasting. *Journal of Business & Economic Statistics 40*(3), 1094–1106.

Bańbura, M., D. Leiva-Leon, and J.-O. Menz (2023). Do Inflation Expectations Improve Model-Based Inflation Forecasts? *ECB Discussion Paper 2604*, 1–47.

Barkan, O., J. Benchimol, I. Caspi, E. Cohen, A. Hammer, and N. Koenigstein (2022). Forecasting CPI Inflation Components with Hierarchical Recurrent Neural Networks. *International Journal of Forecasting 39*, 1145–1162.

Beck, G., K. Carstensen, J.-O. Menz, R. Schnorrenberger, and E. Wieland (2022). Real-Time Food Price Inflation in Germany in Light of the Russian Invasion of Ukraine. *VOXEU Column, link*.

Beck, G. and X. Jaravel (2021). Prices and Global Inequality: New Evidence from Worldwide Scanner Data. *Discussion Paper*, 1–50.

Beck, G. W. and S. M. Lein (2020). Price Elasticities and Demand-Side Real Rigidities in Micro Data and in Macro Models. *Journal of Monetary Economics 115*, 200–212.

Benalal, N., J. L. D. del Hoyo, B. Landau, M. Roma, and F. Skudelny (2004). To Aggregate or not to Aggregate? Euro Area Inflation Forecasting. *ECB Discussion Paper 374*, 1–67.

Bergmeir, C., R. J. Hyndman, and B. Koo (2018). A note on the validity of cross-validation for evaluating autoregressive time series prediction. *Computational Statistics & Data Analysis 120*, 70–83.

Bermingham, C. and A. D'Agostino (2013). Understanding and Forecasting Aggregate and Disaggregate Price Dynamics. *Empirical Economics 46*, 765–799.

Botha, B., R. Burger, K. Kotzé, N. Rankin, and D. Steenkamp (2022). Big Data Forecasting of South African Inflation. *Empirical Economics*.

Breitung, J. and C. Roling (2015). Forecasting Inflation Rates Using Daily Data: A Nonparametric MIDAS Approach. *Journal of Forecasting 34*(7), 588–603.

Buda, G., V. Carvalho, G. Corsetti, J. Duarte, S. Hansen, A. Moura, A. Ortiz, T. Rodrigo, J. Rodríguez Mora, and G. Alves da Silva (2023). Short and Variable Lags. *Cambridge Working Paper in Economics 2321*, 1–61.

Buda, G., V. Carvalho, S. Hansen, A. Ortiz, T. Rodrigo, and J. Rodríguez Mora (2022). National Accounts in a World of Naturally Occurring Data: A Proof of Concept for Consumption. *Cambridge Working Paper in Economics 2244*, 1–72.

Butters, A., D. Sacks, and B. Seo (2022). How Do National Firms Respond to Local Cost Shocks? *American Economic Review 112*(5), 1737–1772.

Capistrán, C., C. Constandse, and M. Ramos-Francia (2010). Multi-Horizon Inflation Forecasts using Disaggregated Data. *Economic Modelling 27*, 666–677.

Cavallo, A., E. Cavallo, and R. Rigobon (2014). Prices and Supply Disruptions during Natural Disasters. *Review of Income and Wealth 60*, 449–471.

Cavallo, A. and O. Kryvtsov (2023). What can Stockouts Tell us About Inflation? Evidence from Online Micro Data. *Journal of International Economics*, 103769.

Cavallo, A. and R. Rigobon (2016). The Billion Prices Project: Using Online Prices for Measurement and Research. *Journal of Economic Perspectives 30*(2), 151–178.

Cimadomo, J., D. Giannone, M. Lenza, F. Monti, and A. Sokol (2022). Nowcasting with Large Bayesian Vector Autoregressions. *Journal of Econometrics 231*, 500–519.

Clark, T. E., S. Leonard, M. Marcellino, and P. Wegmüller (2022). Weekly Nowcasting US Inflation with Enhanced Random Forests. *Mimeo*.

de Haan, J., R. Hendriks, and M. Scholz (2021). Price Measurement Using Scanner Data: Time-Product Dummy Versus Time Dummy Hedonic Indexes. *Review of Income and Wealth 67*(2), 394–417.

Deutsche Bundesbank (2020). The Fiscal Stimulus Package Announced by the Coalition Parties. *Bundesbank Monthly Report Juni*, 28–29.

Diebold, F. X. and R. S. Mariano (1995). Comparing Predictive Accuracy. *Journal of Business & Economic Statistics 13*(3), 253–263.

Dietrich, A., M. Eiglsperger, J. Mehrhoff, and E. Wieland (2021). Chain linking over December and methodological changes in the HICP: view from a central bank perspective. *ECB Statistics Paper Series 40*.

Diewert, E. (2005). Weighted Country Product Dummy Variable Regressions and Index Number Formulae. *Review of Income and Wealth 51*(4), 561–570.

Doerr, S., L. Gambacorta, and J. Maria Serena (2021). Big Data and Machine Learning in Central Banking. *BSI Working Paper 930*, 1–26.

Dubois, P., R. Griffith, and M. O'Connell (2022). The Use of Scanner Data for Economics Research. *Annual Review of Economics 14*, 723–745.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber (2021). Exposure to Grocery Prices and Inflation Expectations. *Journal of Political Economy 129*(5), 1615–1639.

Eraslan, S. and T. Götz (2021). An Unconventional Weekly Economic Activity Index for Germany. *Economics Letters 204*, 1–4.

Espasa, A. and I. Mayo-Burgos (2013). Forecasting Aggregates and Disaggregates with Common Features. *International Journal of Forecasting 29*, 718–732.

Eurostat (2018). Harmonised index of Consumer Prices (HICP). Methodological Manual. *Publication Office of the European Union, Luxembourg*.

Eurostat (2022). Guide on Multilateral Methods in the Harmonised Index on Consumer Prices (HICP) - 2022 edition.

Foroni, C., M. Marcellino, and C. Schumacher (2015). Unrestricted Mixed Data Sampling (MIDAS): MIDAS Regressions with Unrestricted Lag Polynomials. *Journal of the Royal Statistical Society 178*(1), 57–82.

Gagnon, E. and D. López-Salido (2019). Small Price Responses to Large Demand Shocks. *Journal of the European Economic Association 18*(2), 792–828.

Garcia, M. G., M. C. Medeiros, and G. F. Vasconcelos (2017). Real-Time Inflation Forecasting with High-Dimensional Models: The Case of Brazil. *International Journal of Forecasting 33*(3), 679–693.

Gautier, E., C. Conflitti, R. P. Faber, B. Fabo, L. Fadejeva, V. Jouvanceau, J.-O. Menz, T. Messner, P. Petroulas, P. Roldan-Blanco, F. Rumler, S. Santoro, E. Wieland, and H. Zimmer (2023). New Facts on Consumer Price Rigidity in the Euro Area. *American Economic Journal Macroeconomics*.

Ghysels, E. and M. Marcellino (2018). *Applied Economic Forecasting Using Time Series Methods*. Oxford University Press.

Ghysels, E., P. Santa-Clara, and R. Valkanov (2004). The MIDAS Touch: Mixed Data Sampling Regression Models.

Giannone, D., M. Lenza, D. Momferatou, and L. Onorante (2014). Short-Term Inflation Projections: A Bayesian Vector Autoregressive Approach. *International Journal of Forecasting 30*, 635–644.

Goulet Coulombe, P., M. Leroux, D. Stevanovic, and S. Surprenant (2022). How is Machine Learning Useful for Macroeconomic Forecasting? *Journal of Applied Econometrics 37*(5), 920–964.

Harchaoui, T. and R. Janssen (2018). How Can Big Data Enhance the Timeliness of Official Statistics?: The Case of the U.S. Consumer Price Index. *International Journal of Forecasting 34*(2), 225–234.

Hauzenberger, N., F. Huber, and K. Klieber (2023). Real-Time Inflation Forecasting Using Non-Linear Dimension Reduction Techniques. *International Journal of Forecasting 39*(2), 901–921.

Hendry, D. and K. Hubrich (2011). Combining Disaggregate Forecasts or Combining Disaggregate Information to Forecast an Aggregate. *Journal of Business & Economic Statistics 29*(2), 216–227.

Henn, K., C.-G. Islam, P. Schwind, and E. Wieland (2019). Measuring Price Dynamics of Package Holidays with Transaction Data. *EURONA 2/2019*, 95–132.

Huwiler, M. and D. Kaufmann (2013). Combining Disaggregate Forecasts for Inflation: The SNB's ARIMA Model. *Swiss National Bank Economic Studies 7*, 1–32.

Ibarra, R. (2012). Do Disaggregated CPI Data Improve the Accuracy of Inflation Forecasts? *Economic Modelling 29*, 1305–1313.

IMF, ILO, OECD, Eurostat, UNECE, and The World Bank (2020). Consumer Price Index Manual: Concepts and Methods. *IMF*.

Jaravel, X. (2019). The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector. *Quarterly Journal of Economics 134*(2), 715–783.

Jaravel, X. and M. O'Connell (2020). Real-Time Price Indices: Inflation Spike and Falling Product Variety During the Great Lockdown. *Journal of Public Economics 191*, 104–270.

Joseph, A., E. Kalamara, G. Kapetanios, G. Potjagailo, and C. Chakraborty (2022). Forecasting UK Inflation Bottom Up. *Bank of England Staff Working Paper 915*, 1–38.

Kaplan, G. and S. Schulhofer-Wohl (2017). Inflation at the Household Level. *Journal of Monetary Economics 91*, 19–38.

Karadi, P., J. Amann, J. Sánchez Bachiller, P. Seiler, and J. Wursten (2023). Price Setting on the two Sides of the Atlantic - Evidence from Supermarket Scanner Data. *Journal of Monetary Economics*.

Knotek II, E. S. and S. Zaman (2017). Nowcasting U.S. Headline and Core Inflation. *Journal of Money, Credit and Banking 49*(5), 931–968.

Li, J., Z. Liao, and R. Quaedvlieg (2022). Conditional Superior Predictive Ability. *The Review of Economic Studies 89*(2), 843–875.

Macias, P., D. Stelmasiak, and K. Szafranek (2023). Nowcasting Food Inflation with a Massive Amount of Online Prices. *International Journal of Forecasting 39*(2), 809–826.

Medeiros, M. C., G. F. Vasconcelos, Á. Veiga, and E. Zilberman (2021). Forecasting Inflation in a Data-Rich Environment: The Benefits of Machine Learning Methods. *Journal of Business & Economic Statistics 39*(1), 98–119.

Messner, T., F. Rumler, and G. Strasser (2023). Cross-Country Price and Inflation Dispersion: Retail Network or National Border? *ECB Discussion Paper 2276*, 1–56.

Modugno, M. (2013). Now-casting Inflation Using High Frequency Data. *International Journal of Forecasting 29*(4), 664–675.

Nagengast, A., D. Bursian, and J.-O. Menz (2021). Dynamic Pricing and Exchange Rate Pass-Through: Evidence from Transaction-Level Data. *European Economic Review 133*, 1–29.

Paranhos, L. (2021). Predicting Inflation with Neural Networks. *Warwick Economics Research Paper 1344*.

Powell, B., G. Nason, D. Elliott, M. Mayhew, J. Davies, and J. Winton (2018). Tracking and Modelling Prices Using Web-Scraped Price Microdata: Towards Automated Daily Consumer Price Index Forecasting. *Journal of the Royal Statistical Society 181*(3), 737–756.

Stelmasiak, D., K. Szafranek, P. Macias, and A. Błażejowska (2023). Online Food Prices and Shocks to Product Availability since Covid-19. *Applied Economics Letters*, 1–7.

Stock, J. H. and M. W. Watson (1999). Forecasting Inflation. *Journal of Monetary Economics 44*(2), 293–335.

Stock, J. H. and M. W. Watson (2007). Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit and Banking 39*(1), 3–33.

Tissot, B. and B. de Beer (2020). Implications of COVID-19 for Official Statistics: A Central Banking Perspective. *BIS IFC Working Paper 20*, 1–23.

Ulgazi, Y. and P. Vertier (2022). Forecasting Inflation in France: An Update of MAPI. *Banque de France Discussion Paper 869*, 1–33.

Watanabe, T. (2020). The Responses of Consumption and Prices in Japan to the COVID-19 Crisis and the Tohoku Earthquake. *CJEB Working Papers - Columbia Business School 373*, 1–16.

Weber, M., Y. Gorodnichenko, and O. Coibion (2022). The Expected, Perceived, and Realized Inflation of U.S. Households Before and During the COVID19 Pandemic. *NBER Working Paper 29640*.

Zhao, P. and B. Yu (2006). On Model Selection Consistency of Lasso. *The Journal of Machine Learning Research 7*, 2541–2563.

A Supplementary information on data, descriptive statistics and data construction

A.1 Overview of data sources

Table A1: Data sources

Variable	Source	Description
	Consumer Prices	
Headline inflation	DESTATIS	Harmonised index of consumer prices (HICP), unadjusted data.
10-digit series	DESTATIS	German national consumer price index (CPI), unadjusted data.
	Fast-moving consumer goods (GFK:FMCG)	
Prices	GFK	Scanner data recorded by private households on daily purchases.
	Travel services	
Prices	AMADEUS	Transaction data on daily bookings of package holidays by German travelers.
	Energy prices	
Euro super	Weekly Oil Bulletin	Average Monday pump price including duties and taxes
Diesel	Weekly Oil Bulletin	Average Monday pump price including duties and taxes
LPG motor fuel	Weekly Oil Bulletin	Average Monday pump price including duties and taxes
Heating oil	Weekly Oil Bulletin	Average price for deliveries of 2,000 to 5,000 liters incl. duties/taxes
Euro super	Tankerkoenig	Real-time prices of all gasoline stations in Germany
Diesel	Tankerkoenig	Real-time prices of all gasoline stations in Germany
Heating oil	Heizoel 24	Average daily heating oil price of participating delivery firms
Wood pellets	Holzpellets.net	Average daily wood pellets price of participating delivery firms
European Gas Spot Index (EGSI)	European Energy Exchange AG	Volume-weighted average price of all spot transactions concluded on the trading day. Values before October 2021 are mean values of the gas prices from the two former market areas GASPOOL and NetConnect Germany.
	Market expectations	
HICP nowcast	BLOOMBERG	Survey among market participants and professional forecasters.
HICP nowcast	CONSENSUS	Survey among market participants and professional forecasters.

Note: *DESTATIS*: Federal Statistical Office of Germany, *GFK*: Growth from Knowledge, *AMADEUS*: AMADEUS IT group, *HAVER*: Haver Analytics

A.2 Additional descriptive statistics of high-frequency data

Table A2 summarizes some basic characteristics of the GFK:FMCG data. Overall, we observe an average price of 1.7 euro, with the highest prices being paid for durables, and the lowest for dairy products and fat. Over the full sample, we observe about 100,000 distinct products by month distributed fairly equally across HICP subcomponents. The only exceptions are unprocessed fish and durables, where product coverage is relatively low; however, in more recent years, product variety has increased for these products, too. On average, a product stays in the sample for about 190 days, whereas dairy products have the longest lifetime with about 407 days compared to 47 days for durable goods. About 43% of all purchases are made in discount shops, with the largest discounter share attributable to unprocessed fish, followed by unprocessed meat and fruit. The availability of consumption data in discount shops is very important for nowcasting official inflation in Germany, since large discounters typically act as price leaders in the market. Finally, about 74.4% of household purchases are spend on processed food, in particular beverages and dairy products and fats, followed by non-energy industrial goods (15.5%) and unprocessed food (9.6%).

Table A2: Summary statistics of GFK:FMCG

	Prices							
Component	mean	median	10^{th}	90^{th}	#ID	#T	#discount	#expenses
Unprocessed food	2.2	1.8	0.9	3.8	5,853	150	0.53	9.6
Fruit	1.7	1.7	0.9	2.8	1,124	110	0.56	1.8
Vegetables	1.4	1.3	0.8	2.1	1,594	210	0.46	1.8
Meat & eggs	2.9	2.3	1.1	5.0	3,024	146	0.56	5.9
Fish	3.1	2.7	1.8	5.0	111	166	0.66	0.2
Processed food	1.5	1.1	0.5	2.7	66,459	271	0.46	74.4
Fruit	1.7	1.5	0.7	3.0	1,941	256	0.44	1.8
Vegetables	1.4	1.1	0.6	2.2	4,102	245	0.43	3.8
Meat	1.7	1.5	0.9	2.8	8,436	300	0.53	9.1
Fish	2.3	1.8	0.8	4.0	2,119	218	0.48	2.3
Bread & cereals	1.4	1.2	0.6	2.5	10,399	271	0.48	9.5
Dairy products & fat	1.1	0.9	0.4	2.0	10,167	407	0.46	16.4
Beverages	1.8	1.0	0.3	4.2	12,303	236	0.42	18.0
Other food products	1.4	1.1	0.5	2.6	16,991	226	0.46	13.5
Tobacco								
Non-energy industrial goods	2.3	1.5	0.5	4.5	24,509	87	0.27	15.5
Non-durables	2.3	1.5	0.5	4.5	23,752	87	0.27	14.9
Semi-durables	3.5	1.9	0.7	8.7	674	105	0.22	0.4
Durables	13.9	10.0	5.0	23.3	83	47	0.10	0.2
Total HICP	1.7	1.3	0.5	3.0	96,551	190	0.43	100

Note: This table reports the mean, the median, and the 10th and 90th percentile of the daily scanner prices. Column #ID shows the average number of unique products per month, #T the average lifetime of products in days, #discount the share of products bought in discount shops and #expenses the share of consumption purchases for each product category.

Finally, Figure A1 displays our supplementary high-frequency indicators on energy and package holidays, together with their official counterpart. For energy, comovement at the monthly frequency is nearly perfect, with a correlation coefficient of 0.93. The correlation for package holidays is also fairly high, at 0.45 if one adjusts the official inflation rate for a large jump in 2016 that reflects a statistical break due to the current HICP chain-linking practice (Dietrich, Eiglsperger, Mehrhoff, and Wieland, 2021).

Figure A1: HICP inflation and high-frequency counterparts

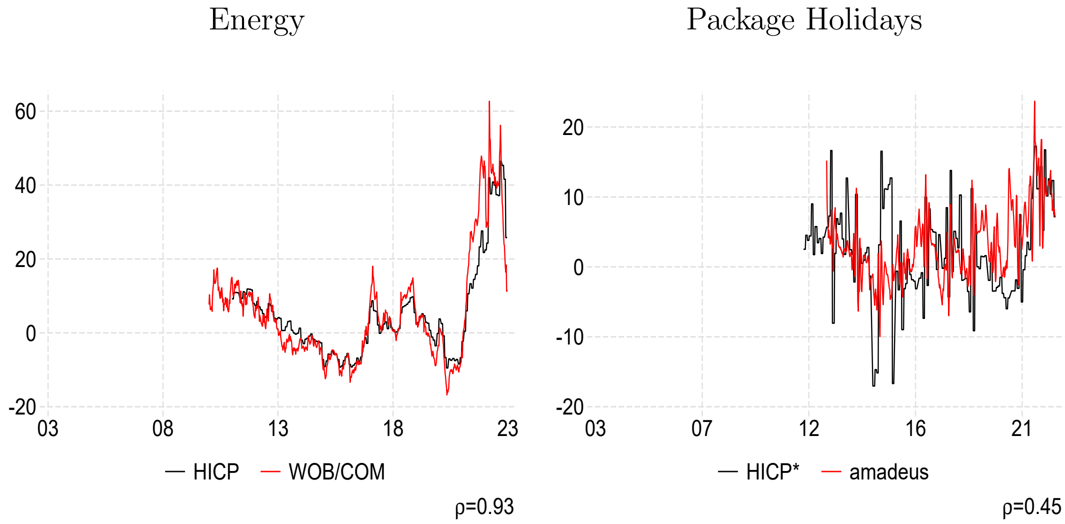

Note: The figure shows year-over-year inflation rates (% change) for HICP subcomponents aggregated using all of the corresponding COICOP 10-digit level series for which we have high-frequency data available. "HICP" refers to the aggregates using official COICOP 10-digit series, "WOB" refers to the data from the weekly oil bulletin and from commodity prices, and "amadeus" refers to the transaction data for package holidays. Note that the official inflation rate for package holidays has been adjusted for a large jump in 2015 caused by a statistical break due to chain-linking. ρ reports the correlation coefficient between both series.

Figure A2: Selected inflation rates of COICOP-10 items and weekly scanner-based counterparts

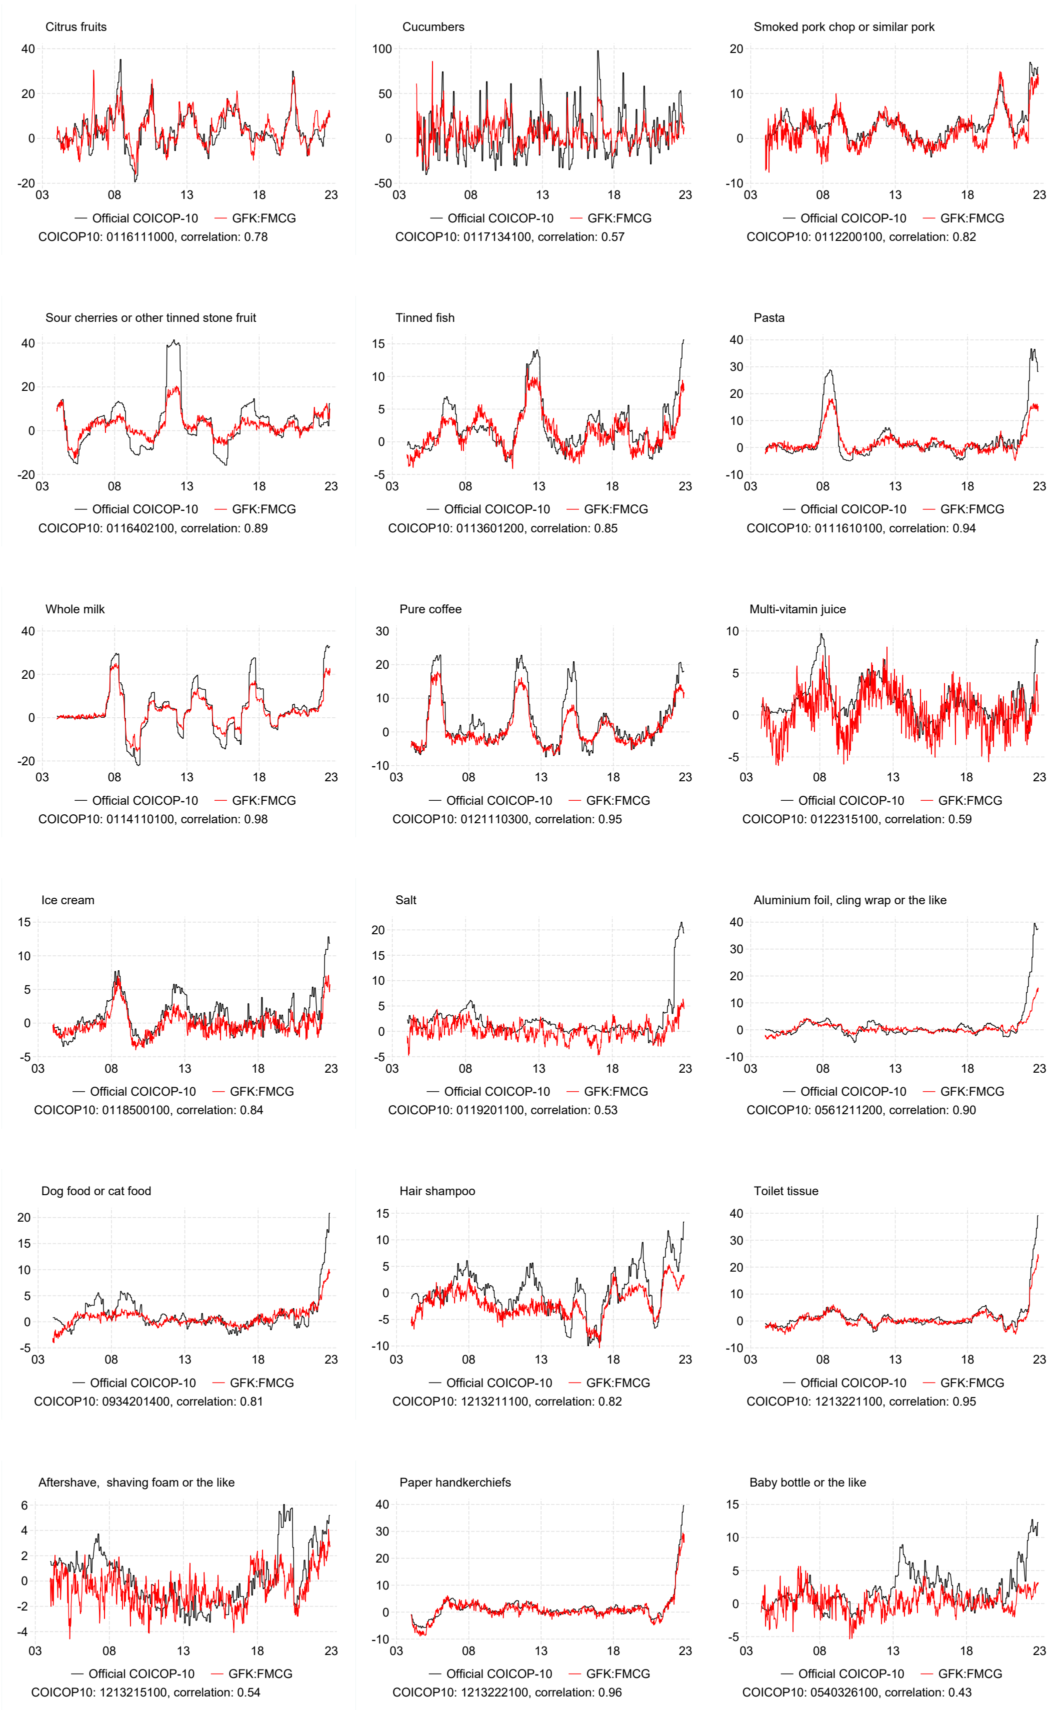

A.3 Construction of transaction-based price indicators for package holidays

The German HICP subindex on package holidays mainly consists of international flight holidays. Importantly, prices for package holidays enter the HICP on the day of travel rather than the day of booking, contrary to other products such as food, where prices are recorded on the day the price is recorded in-store. Therefore, in the case of package holidays, we can exploit advance booking for forecasting purposes, since we are already able to observe the bulk of bookings for a future package holiday in a given month well in advance.

We aggregate the AMADEUS micro dataset on package holidays in various stages. First, we omit all bookings in the dataset that have been canceled by the traveler or the travel agencies and select only those holiday regions which have been entered into the official price index up to 2022 (Turkey, Spain including Balearic Islands and Canary Islands, Greece, Dominican Republic and Egypt).[31] Second, since micro prices refer to the overall price of a booking, we perform a quantity adjustment and define the price by traveler and travel days:

$$p_{i,t}^{raw,amadeus} = sales\_exp_{i,t}^{amadeus} \frac{1}{N_{i,t}} \frac{1}{D_{i,t}} \tag{A1}$$

In addition, we compute the advance booking in days as the difference between the day of booking and the day of travel. To correct for outliers, we omit prices which are smaller than 20 euro and larger than 600 euro per traveler and day and with advance booking of more than one year. Third, similar to official price statistics, we categorize the bookings by advance booking to differentiate between different pricing schemes of early (last-minute) bookings, regular and early bookings. We opt for five different booking categories (bookings up to 14 days before departure, bookings between 15-30, 31-90, 91-180 and 181-270 days, as well as more than 270 days before departure), whereas last-minute bookings convey the most helpful signal to nowcast our target series. Finally, to meet the structure of the weekly nowcasts, we aggregate daily prices into weeks 1 to 4, but in addition, each week now consists of price indices categorized by five different categories of booking advances.

A.4 Modeling the impact of policy measures during our evaluation period 2016-2022

To keep our nowcasting models on a competitive level with expert opinions, we extend their information set by *a priori* knowledge on three non-standard policy measures implemented by the German government which affected consumer prices. All these measures were announced well before taking effect and were hence known to professional forecasters when they were nowcasting for the corresponding month.

First, to mitigate the negative economic effects of the COVID-19 pandemic, the

[31] The German price index for package holidays includes additional price representatives such as domestic package holidays, cruises and city trips. However, the bulk of the index represents those five regions above. See Henn et al. (2019) for a description of the underlying HICP methodology.

German government announced on 4 June 2020 that the VAT rate would temporarily be cut in the period July to December 2020. Specifically, it decreased the regular VAT rate (which applies to about 65% of prices collected in the HICP) from 19% to 16%, and the reduced rate that mainly applies to food (excluding beverages), newspapers and books from 7% to 5%. We implement this *a priori* information in our nowcasting models as follows. First, all COICOP 10-digit items are classified according to their VAT category (regular, reduced, or VAT-free). Next, while the dates and the extent of the VAT changes were announced beforehand and were thus known to forecasters, the pass-through to actual prices was not (see, for example, Deutsche Bundesbank, 2020). We thus include in our forecasts an ex-ante pass-through of 50% applied uniformly to all products, which mimics a Bayesian forecaster with a flat prior and symmetric loss function. This ex-ante VAT impact is subtracted from the target variable prior to model estimation and then added to the model's forecast. To account for the VAT changes ex-post, we fit a dummy variable that takes the value -1 in July 2020 and +1 in January 2021.

Second, we consider a reduction in the price of public transport tickets, which was passed by the German government on 21 May 2022 as a response to soaring energy prices after Russia's invasion of Ukraine.[32] Known as the 9 euro ticket, it allowed consumers to travel by local and regional public transport for 9 euros per month from June to August 2022. It reduced public transport prices considerably and in a foreseeable manner. We thus assume that forecasters were able to predict the corresponding changes in the affected price indices in June and September 2022.[33]

Third, the German government announced on 19 November 2022 a one-off emergency aid package for natural gas and heating taking effect in December 2022 ("December aid").[34] The government assumed the December installment of the households' contracts with their natural gas and district heating suppliers. Due to the complex and heterogeneous contract designs in Germany, it was unclear by how much this measure would affect the relevant HICP price indices. Similarly to our strategy for the temporary VAT cut, we apply a simple forecasting rule that a professional forecaster may have followed in real time. Specifically, we distinguish between homeowners and tenants. According to the latest EU-SILC survey, 46.7% of German households own their house or apartment. These households typically have a contract with a supplier of gas or district heating and thus directly benefited from the "December aid". The remaining households are tenants who typically pay their household energy via the landlord, from whom they receive an annual energy bill. Hence, they typically did not benefit (in December 2022) from the "December aid". We thus assume that in December 2022, the price for natural gas supply (COICOP 0452103000) and district heating (COICOP 0455002200) was zero for homeowners and unreduced for tenants. This implies an overall reduction by 46.7% for the December 2022 prices of natural gas supply and district heating.

[32]See note by Deutscher Bundestag.

[33]The affected COICOP-10 categories are "Train journey, short-distance" (0731111100), "Transport association, single or day ticket for adults" (0735011000), "Transport association, season ticket for apprentices" (0735013100), and "Transport association, monthly ticket for adults" (0735015000).

[34]See note by Deutsche Bundesregierung.

B Supplementary information on the econometric methodology

B.1 Nowcasting item-level inflation

To set the scene, let us assume that the latest official data on inflation of a COICOP-10 item c has been released for a given month t. For convenience of notation, however, we stop referring to the item c. Next, conditional on high-frequency data available up to period t and a pre-sample observation π_0^M, our baseline U-MIDAS (6), neglecting seasonal dummies and autoregressive lags greater than one, can be estimated via OLS following the matrix representation

$$\begin{bmatrix} \pi_1^M \\ \pi_2^M \\ \vdots \\ \pi_t^M \end{bmatrix} = \begin{bmatrix} 1 & \pi_0^M & x_1^{(m)} & x_{1-1/m}^{(m)} & x_{1-2/m}^{(m)} & x_{1-3/m}^{(m)} \\ 1 & \pi_1^M & x_2^{(m)} & x_{2-1/m}^{(m)} & x_{2-2/m}^{(m)} & x_{2-3/m}^{(m)} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \\ 1 & \pi_{t-1}^M & \underbrace{x_t^{(m)}}_{\text{4th week}} & \underbrace{x_{t-1/m}^{(m)}}_{\text{3rd week}} & \underbrace{x_{t-2/m}^{(m)}}_{\text{2nd week}} & \underbrace{x_{t-3/m}^{(m)}}_{\text{1st week}} \end{bmatrix} \begin{bmatrix} \alpha_0 \\ \rho_1 \\ b_1 \\ b_2 \\ b_3 \\ b_4 \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_t \end{bmatrix} \tag{B2}$$

Note that the matrix representation (B2) makes explicit the transformation of the high-frequency predictor $x_{t-k/m}^{(m)}$ into m low-frequency vectors $(x_{1-k/m}^{(m)}, \ldots, x_{t-k/m}^{(m)})'$, for $k = 0, \ldots, m-1$. Hence, model (6) is estimated in the low-frequency dimension, but our nowcasts can be updated each time high-frequency increments become available after t, whereas random walk forecasts of the most recent high-frequency information are used to deal with the "ragged-edge problem".

Finally, note that $B(L^{1/m}; \theta) = \sum_{k=0}^{K} B(k; \theta) L^{k/m}$ is a polynomial of length $(K+1)$ in the $L^{1/m}$ operator with $L^{k/m} x_t^{(m)} = x_{t-k/m}^{(m)}$. Despite we assume $K = m-1$ for simplicity, leading to a static model in the high-frequency component for $h = 0$, one might include high-frequency distributed lags $K > m-1$ that span over past low-frequency periods (see robustness section 6).

B.2 Nowcasting product group-specific inflation

To nowcast product groups covered by the GFK:FMCG dataset such as unprocessed food, processed food and non-energy industrial goods, let us assume a set of monthly aggregated inflation indicators $\mathbf{x}_t = (x_{1t}, \ldots, x_{qt})'$ such that $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_t)'$, where q denotes an abundant number of COICOP-10 series belonging to a given product group. Hence, conditional on official inflation data available at t, we model our higher-level product group $\pi^{\mathbf{M}} = (\pi_{g,1}^M, \ldots, \pi_{g,t}^M)'$ as a function of $\mathbf{X} = (\iota, \mathbf{x})$ using standard shrinkage methods such as LASSO, ridge and elastic net regression, where ι is a t-dimensional vector of ones. The hybrid elastic net estimator solves the following penalized least squares problem:

$$\hat{\beta} = \min_{\hat{\beta}} ||\pi^{\mathbf{M}} - \mathbf{X}\beta||^2 + \lambda \left(\alpha \, |\beta|_1 + \frac{(1-\alpha)}{2} \, ||\beta||^2 \right), \tag{B3}$$

where $\alpha \in (0,1]$ is a weight parameter that interpolates between LASSO ($\alpha = 1$) and ridge regression (as $\alpha \to 0$) while the regularization parameter λ controls the amount of shrinkage in β.[35] Hence, the idea is to shrink coefficients $b_{g,c}$ to or towards zero if the c-th COICOP-10 series is not relevant. Finally, we construct the monthly aggregated estimate of $\mathbf{x}$ using the available contemporaneous weekly information at the time of the nowcast and compute it based on the estimated parameters $\hat{b}_g$.

As a second class of models, we construct a nowcast for $\pi^M_{g,t+h}$ using the sg-LASSO-MIDAS framework (Babii et al., 2022, see) that handles high-dimensional mixed-frequency prediction problems. Let the matrix of covariates now be defined as:

$$\mathbf{X} = (\iota, \mathbf{X_1^{(m)}}W, \ldots, \mathbf{X_q^{(m)}}W), \tag{B4}$$

where $\mathbf{X_j^{(m)}} = (X^{(m)}_{c1}, \ldots, X^{(m)}_{ct})'$ is a $t \times m$ matrix of the c-th high-frequency covariate and W denotes a predetermined $m \times L$ matrix of weights based on Legendre polynomials of degree L that aggregate over the high-frequency lags. Then, the sg-LASSO estimator solves the penalized least squares problem:

$$\hat{\beta} = \min_{\hat{\beta}} ||\pi^{\mathbf{M}} - \mathbf{X}\beta||^2 + 2\lambda \left(\alpha \, |\beta|_1 + (1-\alpha) \, ||\beta||_{2,1}\right), \tag{B5}$$

where $||\beta||_{2,1} = \sum_{G \in \mathcal{G}} |\beta_G|_2$ is the group LASSO norm for a group structure $\mathcal{G}$ that hereby constitutes all high-frequency lags of a single covariate. Thus, in this case, $\alpha \in [0,1]$ determines the relative importance of LASSO sparsity and the group structure.[36] This implies that sg-LASSO promotes sparsity between and within COICOP-10 items, allowing us not only to select the relevant COICOP-10 series but also the appropriate lag structure of each item.

[35] Both tuning parameters λ and α are determined via expanding cross-validation.

[36] Note that $\alpha = 0$ leads to the group LASSO estimator, which is reminiscent of the elastic net regressor.

C Supplementary results

C.1 Nowcasting of product-level, group-specific and headline inflation

Table C3: RMSE for FMCG product-level inflation: OLS-match relative to the SD-AR benchmark

Unprocessed fruit and vegetables

	day 7	day 14	day 21	day 28
Sweet peppers	0.484**	0.496**	0.414**	0.431**
Butterhead Lettuce	0.532**	0.5**	0.452**	0.475**
Carrots	0.602**	0.722**	0.647**	0.602**
Grapes	0.672**	0.625**	0.627**	0.654**
Onion or garlic	0.587**	0.747*	0.682**	0.644**
Lambs lettuce	0.746*	0.733*	0.728*	0.736*
Cauliflower or cabbage	0.765**	0.722**	0.729**	0.755**

Processed fruit and vegetables

	day 7	day 14	day 21	day 28
Frozen chips or the like	0.543	0.601	0.51	0.454
Apple sauce	0.781	0.806	0.77	0.762
Tinned mushrooms	0.827	0.866	0.859	0.847
Potatoes	0.774**	0.938	0.944	0.911
Asparagus	0.914	0.898	0.901	0.894
Potato crisps	0.923	0.887*	0.901*	0.909
Frozen spinach	0.94	0.921	0.9	0.895

Unprocessed meat, fish and eggs

	day 7	day 14	day 21	day 28
Eggs	0.728*	0.787	0.738	0.709
Fresh poultry	0.842	0.842	0.825	0.808
Minced pork	0.84	0.834	0.827	0.819
Minced beef	0.778	0.872	0.863	0.858
Beef for boiling	0.828	0.876	0.877	0.877
Smoked pork chop	0.869	0.889	0.878	0.869
Pork chop or cutlet	0.913	0.921	0.918	0.902

Processed meat and fish

	day 7	day 14	day 21	day 28
Ham or bacon	0.65	0.691	0.669	0.655
Salami or sausage	0.711	0.759	0.716	0.695
Liver sausage	0.693	0.758	0.729	0.706
Lyoner pork sausage	0.779	0.833	0.788	0.767
Fish fingers	0.826	0.851	0.804	0.789
Prepared minced meat	0.813	0.87	0.834	0.804
Cold meat	0.857	0.868	0.855	0.848

Beverages and others

	day 7	day 14	day 21	day 28
Baking powder	0.55	0.686	0.682	0.679
Sugar	0.674	0.804	0.702	0.67
Apple juice	0.693	0.807	0.798	0.784
Orange juice	0.744	0.811*	0.776*	0.778
Ketchup	0.908	0.836	0.825	0.814
Pure coffee	0.844*	0.891	0.86*	0.857*
Food for infants	0.828**	0.886**	0.879**	0.866**

Bread and cereals

	day 7	day 14	day 21	day 28
Flour	0.594	0.561	0.55	0.554
Pasta	0.721	0.774*	0.761*	0.767*
Ready to bake rolls	0.735	0.829*	0.8**	0.778**
Rice	0.808	0.819*	0.817	0.81*
Cornflakes and muesli	0.852	0.856	0.831	0.821
Crisp bread	0.823	0.878	0.858	0.861*
Pizza or quiches	0.831	0.897	0.878	0.87

Non-durables

	day 7	day 14	day 21	day 28
Aluminium foil	0.776	0.817	0.797*	0.786*
Toilet tissue	0.82	0.844	0.813	0.792
Hair shampoo	0.744**	0.867**	0.848**	0.832**
Deo spray or deo roll on	0.812*	0.87*	0.85**	0.837**
Hair spray or gel	0.866	0.884*	0.858**	0.854**
Paper handkerchiefs	0.963	0.865	0.835	0.812
Toothpaste	0.888	0.909*	0.899**	0.888**

Dairy products and fat

	day 7	day 14	day 21	day 28
Whole milk	0.354*	0.369	0.316	0.3
Low fat milk	0.363	0.36	0.325	0.302
Condensed milk	0.332	0.376	0.356	0.361
Cream	0.355*	0.367*	0.354*	0.358*
Butter	0.353**	0.394**	0.377**	0.401**
Curd	0.356	0.518	0.468	0.499
Sliced cheese	0.582	0.648	0.606	0.588

Sources: GfK household panel; own calculations.
Notes: The figure shows heatmaps of RMSE values for the OLS match model relative to the SD-AR benchmark at nowcasting days 7, 14, 21 and 28 for the best-performing COICOP-10 items within selected FMCG product groups. Results for the Diebold and Mariano (1995) test in the event of outperformance relative to the SD-AR model are indicated by the symbols * (5% level) and ** (1% level).

Table C4: Absolute RMSE of the U-MIDAS model for FMCG product-level inflation

Unprocessed food

Product	COICOP ID	day 7	day 14	day 21	day 28	Product	COICOP ID	day 7	day 14	day 21	day 28
Beef for boiling	0112101100	2.61	1.57	1.58	1.58	Citrus fruits	0116111000	3.92	2.63	2.57	2.58
Roulade or loin of beef	0112107100	2.33	1.80	1.80	1.79	Bananas	0116120100	2.54	2.03	2.02	2.00
Beef	0112107200	2.04	1.32	1.32	1.32	Apples	0116130200	3.54	2.19	2.12	2.09
Minced beef	0112107300	3.32	2.20	2.20	2.15	Pears	0116140100	4.16	3.02	3.02	2.99
Veal	0112109100	1.14	1.02	1.01	1.01	Grapes	0116165100	7.98	5.12	4.93	5.09
Smoked pork chop	0112200100	1.91	1.29	1.25	1.27	Kiwis. melons or the like	0116170000	3.53	2.57	2.56	2.55
Minced pork	0112200200	3.48	2.22	2.21	2.21	Butterhead Lettuce	0117111100	12.49	10.27	9.52	10.30
Roast pork	0112200300	2.30	1.58	1.57	1.57	Lambs lettuce	0117119000	7.09	4.87	4.87	4.88
Pork chop or cutlet	0112200500	2.42	1.75	1.75	1.71	Cauliflower or cabbage	0117121000	7.79	6.04	6.01	5.91
Lamb	0112300100	1.61	0.95	0.97	0.98	Tomatoes	0117131100	10.71	8.52	8.14	8.12
Fresh poultry	0112410100	2.36	1.41	1.40	1.41	Sweet peppers	0117133100	6.73	5.69	5.62	5.74
Frozen poultry	0112410200	2.66	1.66	1.65	1.66	Onion or garlic	0117141100	5.50	3.73	3.55	3.46
Rabbit or game meat	0112500100	1.25	0.91	0.92	0.89	Mushrooms	0117142100	1.76	1.45	1.47	1.46
Liver or other edible offal	0112600100	1.16	0.93	0.93	0.93	Carrots	0117145100	3.93	3.02	3.06	3.06
Eggs	0114701100	1.74	1.20	1.00	0.99	Aspargus or the like	0117149100	7.26	5.53	5.44	5.52

Processed food

Product	COICOP ID	day 7	day 14	day 21	day 28	Product	COICOP ID	day 7	day 14	day 21	day 28
Rice	0111101100	1.42	0.98	1.00	1.00	Frozen vegetables	0117209100	1.36	0.87	0.87	0.88
Flour	0111201100	3.72	2.31	2.34	2.31	Dried vegetables	0117310200	1.17	0.72	0.73	0.72
Semolina	0111203100	2.79	1.60	1.58	1.60	Tinned gherkins	0117321100	1.69	1.04	1.03	1.03
White bread	0111311100	1.06	0.68	0.68	0.68	Tinned sauerkraut	0117323100	1.87	1.43	1.39	1.41
Rye bread or brown bread	0111312100	0.92	0.57	0.57	0.57	Tinned mushrooms	0117324100	1.25	0.77	0.79	0.79
Granary bread	0111313200	1.24	0.77	0.78	0.78	Tinned peas	0117325100	1.63	0.94	0.92	0.92
Ready to bake rolls	0111320200	1.31	0.77	0.78	0.78	Asparagus	0117328400	3.21	2.20	2.25	2.23
Fresh bread rolls	0111320300	1.58	1.02	1.01	0.99	Potatoes	0117401300	5.24	3.91	3.73	3.67
Sponge flan case	0111421100	1.25	0.74	0.74	0.73	Frozen chips or the like	0117402100	1.88	1.55	1.51	1.48
Frozen cake. tart or pie	0111423100	1.37	0.93	0.93	0.94	Potato crisps	0117500200	2.11	1.47	1.48	1.48
Fresh cake. tart or pie	0111424300	1.16	0.69	0.69	0.69	Sugar	0118100100	4.91	4.38	4.16	4.17
Biscuits	0111431200	1.88	1.51	1.51	1.52	Marmalade. jam or jelly	0118201100	1.91	1.60	1.59	1.60
Muffins or waffles	0111433100	1.51	0.86	0.85	0.85	Honey	0118203100	1.44	0.93	0.90	0.91
Crisp bread	0111442100	2.13	1.34	1.33	1.38	Cocoa based spread	0118205100	1.30	1.29	1.27	1.27
Toasted bread	0111444200	1.55	1.17	1.17	1.17	Slab of chocolate	0118301100	2.31	2.11	2.09	2.09
Rusk	0111446100	2.38	1.34	1.36	1.35	Chocolate	0118309100	1.48	1.17	1.17	1.17
Savoury biscuits	0111450100	2.05	1.70	1.70	1.67	Filled chocolates	0118401100	0.78	0.66	0.66	0.66
Pizza or quiches	0111500100	2.13	1.49	1.49	1.49	Boiled sweets	0118405100	0.83	0.66	0.67	0.67
Pasta	0111610100	2.18	1.34	1.34	1.38	Ice cream	0118500100	1.47	1.23	1.21	1.22
Pasta preparations	0111621200	2.51	1.94	1.93	1.93	Sweetener	0118601100	3.66	1.82	1.74	1.76
Oatflakes	0111701100	2.16	1.42	1.41	1.42	Vinegar	0119101100	1.43	0.89	0.89	0.89
Cornflakes and muesli	0111703100	0.98	0.81	0.82	0.81	Mustard	0119102100	2.05	1.19	1.18	1.18
Cake mix	0111801100	1.73	1.11	1.11	1.11	Ketchup	0119103200	2.95	1.70	1.68	1.69
Salami or sausage	0112710200	1.22	0.81	0.80	0.80	Sauce mix	0119103300	1.84	1.56	1.52	1.51
Ham or bacon	0112710300	1.09	0.74	0.75	0.75	Mayonnaise	0119104100	1.96	1.11	1.12	1.11
Lyoner pork sausage	0112721100	2.09	1.48	1.45	1.46	Salt	0119201100	2.56	1.57	1.57	1.57
Fried sausage	0112721200	1.67	1.17	1.14	1.13	Spices	0119203100	0.79	0.50	0.51	0.51
Cold meat	0112721300	2.02	1.46	1.46	1.47	Powdered infant milk	0119302100	0.82	0.52	0.53	0.53
Liver sausage	0112722100	1.42	1.01	1.00	1.01	Food for infants	0119303100	1.18	0.70	0.72	0.72
Tinned sausage	0112723100	1.57	0.91	0.91	0.91	Meat ready meal	0119406100	0.93	0.78	0.78	0.78
Meat based speciality salad	0112801100	1.61	0.92	0.92	0.92	Instant soup	0119911100	1.94	1.50	1.49	1.48
Frozen meat	0112805100	1.69	1.08	1.09	1.09	Tinned soup	0119913100	2.04	1.55	1.55	1.54
Meat-based ready meal	0112807200	1.14	0.73	0.73	0.72	Baking powder	0119930100	3.22	1.71	1.73	1.72
Prepared minced meat	0112808200	2.30	1.54	1.51	1.50	Blancmange powder	0119940100	1.89	1.14	1.12	1.14
Smoked fish	0113500100	1.82	1.24	1.21	1.21	Vitamin tablets or the like	0119990200	0.88	0.62	0.62	0.62
Tinned fish	0113601200	1.25	0.93	0.93	0.91	Pure coffee	0121110300	1.95	1.60	1.56	1.56
Fish marinade	0113602100	1.35	0.82	0.84	0.84	Instant coffee	0121121100	1.40	1.16	1.11	1.11
Fish fingers	0113603000	2.14	1.33	1.29	1.34	Black tea or green tea	0121201100	0.52	0.41	0.41	0.41
Whole milk	0114110100	1.75	1.20	1.15	1.16	Fruit tea or herbal tea	0121203100	0.94	0.79	0.79	0.79
Low fat milk	0114210100	2.02	1.31	1.30	1.31	Cocoa powder	0121300100	1.18	0.87	0.87	0.87
Condensed milk	0114300100	1.45	1.04	1.03	1.03	Sparkling mineral water	0122100100	1.22	0.89	0.88	0.88
Yoghurt	0114400200	1.91	1.45	1.54	1.54	Still mineral water	0122100200	1.18	0.93	0.92	0.92
Hard cheese	0114501100	2.22	1.38	1.39	1.36	Cola drink	0122211100	1.58	1.32	1.34	1.34
Sliced cheese	0114502100	2.35	1.73	1.72	1.73	Soft drink	0122219100	1.76	1.26	1.27	1.27
Soft cheese	0114503100	1.14	0.90	0.89	0.88	Apple juice	0122311100	1.73	1.35	1.28	1.25
Curd	0114507100	4.09	2.96	2.89	2.84	Orange juice	0122312200	1.46	1.03	1.03	1.04
Cream	0114601100	2.79	1.91	1.89	1.90	Multi vitamin juice	0122315100	1.25	0.97	0.96	0.97
Milk based dessert	0114604100	2.20	1.67	1.65	1.66	Vegetable juice	0122320300	1.40	0.79	0.80	0.82
Butter	0115100100	3.68	2.58	2.36	2.37	Liqueur	0211110100	0.59	0.45	0.44	0.45
Margarine	0115201100	2.65	2.14	2.07	2.05	Whisky	0211120100	0.75	0.60	0.60	0.59
Vegetable fat	0115209100	2.54	1.67	1.63	1.63	Brandy or cognac	0211130100	0.64	0.54	0.52	0.51
Sunflower oil	0115400100	6.72	3.26	3.18	3.25	Other spirits	0211140100	0.53	0.42	0.41	0.40
Dried fruit	0116301100	1.68	1.03	1.03	1.03	Red wine or rose wine	0212110200	0.75	0.52	0.52	0.52
Peanuts or trail mix	0116303100	0.95	0.82	0.81	0.81	White wine	0212120100	0.64	0.41	0.41	0.41
Apple sauce	0116401100	2.16	1.35	1.31	1.31	Sparkling wine	0212140100	1.29	1.02	1.03	1.03
Sour cherries	0116402100	2.15	1.81	1.78	1.75	Pils. dark or lager beer	0213100100	1.45	1.14	1.14	1.14
Tinned pineapple	0116403100	2.11	1.45	1.42	1.44	Wheat beer or Altbier	0213200100	1.03	0.87	0.87	0.87
Frozen spinach	0117201100	1.70	1.22	1.22	1.21	Non-alcoholic beer	0213300100	1.23	0.95	0.95	0.95

NEIG

Product	COICOP ID	day 7	day 14	day 21	day 28	Product	COICOP ID	day 7	day 14	day 21	day 28
Baby bottle or the like	0540326100	1.39	0.98	0.98	0.99	Hair spray or gel	1213212100	2.65	1.67	1.66	1.68
Heavy duty detergent	0561101100	1.37	0.91	0.91	0.91	Toothpaste	1213213100	1.36	0.85	0.85	0.85
Mild detergent	0561101200	1.84	1.17	1.17	1.17	Mouthwash or dental floss	1213214100	1.26	0.79	0.79	0.80
Fabric softener or starch	0561101300	1.91	1.17	1.17	1.17	Shaving foam	1213215100	1.07	0.75	0.75	0.75
Dishwashing detergent	0561103100	1.63	1.01	1.00	1.01	Toilet soap	1213216100	2.29	1.44	1.44	1.45
Sanitary cleaner	0561105100	1.39	0.84	0.85	0.85	Shower gel or foam	1213217100	1.50	1.06	1.07	1.07
Glass or furniture cleaner	0561105200	0.84	0.52	0.52	0.52	Toilet tissue	1213221100	1.72	0.98	0.98	0.98
All purpose cleaners	0561105300	1.56	0.97	0.97	0.97	Paper handkerchiefs	1213222100	1.75	0.95	0.94	0.95
Shoe polish	0561107100	1.74	1.10	1.09	1.10	Nappies for babies	1213223200	1.15	0.73	0.73	0.73
Filter paper	0561211100	1.04	0.64	0.63	0.64	Tampons or facial tissues	1213229100	1.67	1.03	1.03	1.03
Aluminium foil	0561211200	2.05	1.06	1.04	1.04	Perfume	1213231100	1.22	1.06	1.05	1.06
Candles	0561241100	1.73	1.16	1.16	1.15	Lipstick or lip care	1213232100	1.21	1.14	1.15	1.14
Scrubbing brushes or brooms	0561291000	0.71	0.47	0.47	0.47	Nail varnish	1213232200	1.14	1.00	0.99	0.99
Melissengeist tonic	0611032100	0.56	0.41	0.42	0.42	Make up	1213232300	1.31	1.18	1.18	1.19
Bird food	0934201200	1.24	0.81	0.81	0.81	Kajal pencil or mascara	1213232400	1.33	1.27	1.31	1.31
Dog food or cat food	0934201400	1.32	0.84	0.85	0.86	Hand cream	1213233100	1.16	0.93	0.93	0.93
Cat litter or bird sand	0934209100	1.10	0.67	0.68	0.68	Day cream or night cream	1213233200	0.95	0.82	0.82	0.82
Non electric toothbrush	1213105200	1.11	0.76	0.76	0.76	Baby cream	1213233300	1.23	0.77	0.77	0.78
Razor blades	1213105300	1.12	0.81	0.81	0.81	Deo spray or deo roll on	1213240100	1.57	1.20	1.18	1.17
Hair shampoo	1213211100	1.56	1.20	1.25	1.21						

Sources: GfK household panel; own calculations.

Figure C3: Absolute forecast errors over time: tracking the nowcasting performance of GFK:FMCG product groups

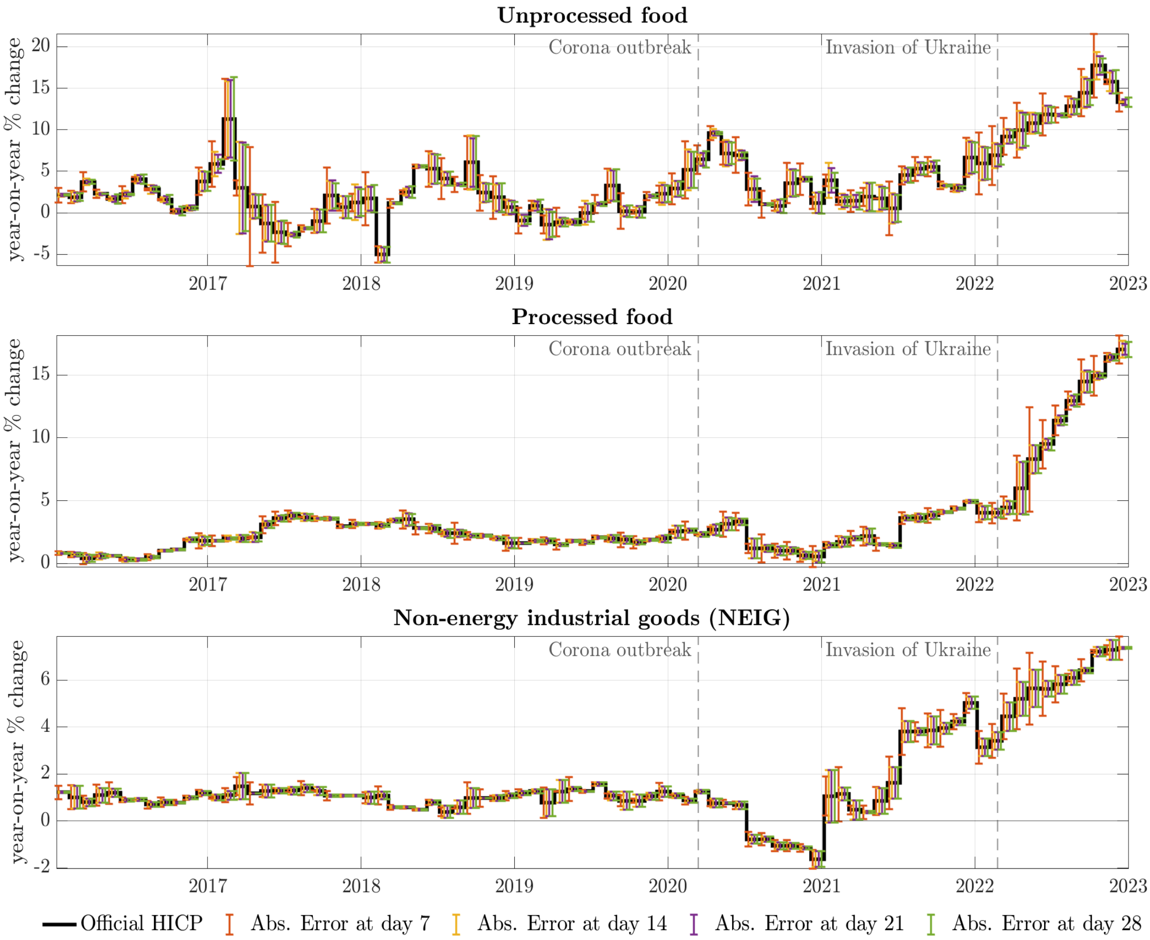

Sources: GfK household panel; own calculations.
Note: The figure shows the evolution over time of the official inflation rates and the absolute forecast errors of the best-performing models on days 7, 14, 21 and 28.

Table C5: RMSE of headline inflation and its components: bottom-up OLS match approach relative to the benchmark approach

Bottom-up OLS-match

	day 7	day 14	day 21	day 28
Unprocessed food	0.677*	0.695**	0.671**	0.667**
Processed food	0.844	0.844	0.799*	0.768*
Energy	0.432**	0.347*	0.323*	0.342*
Package holidays	1.102	1.064	1.009	0.941*
NEIG	1.014	1.003	1.001	0.998
Services	1	1	1	1
Headline	0.797	0.727	0.683	0.674

Sources: GfK household panel; European Commission's Weekly Oil Bulletin; AMADEUS; own calculations.

Note: The figure shows heatmaps of RMSEs for nowcasts based on the bottom-up OLS match approach with aggregation via HICP weights relative to the benchmark approach, which is a bottom-up nowcast based on SD-AR models fitted at the COICOP-10 level. Results for the Diebold and Mariano (1995) test in the event of outperformance relative to the benchmark are indicated by the symbols * (5% level) and ** (1% level).

Figure C4: Cumulative sum of the squared forecast error differentials: models versus Consensus market expectations

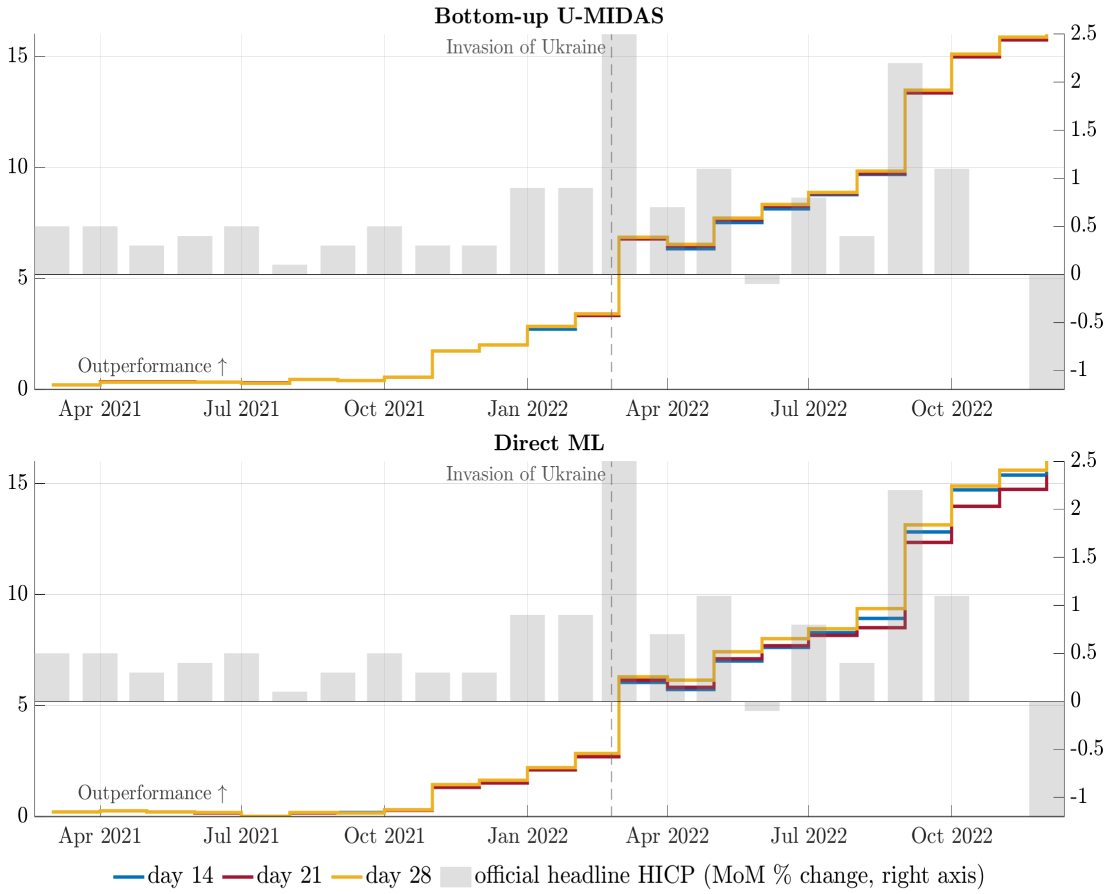

Sources: GfK household panel; European Commission's Weekly Oil Bulletin ; AMADEUS; Consensus survey; own calculations.
Notes: The figure shows, on the left axis, the cumulative sum of the squared forecast error differential of the bottom-up U-MIDAS approach (top panel) and the direct machine learning in relation (bottom panel), respectively, in comparison to Consensus market expectations on days 14, 21 and 28. The gray bars (right axis) represent official month-over-month percentage changes in headline inflation.

C.2 Robustness checks

Figure C5: RMSE change of implementing the best compiled COICOP-10 price indices

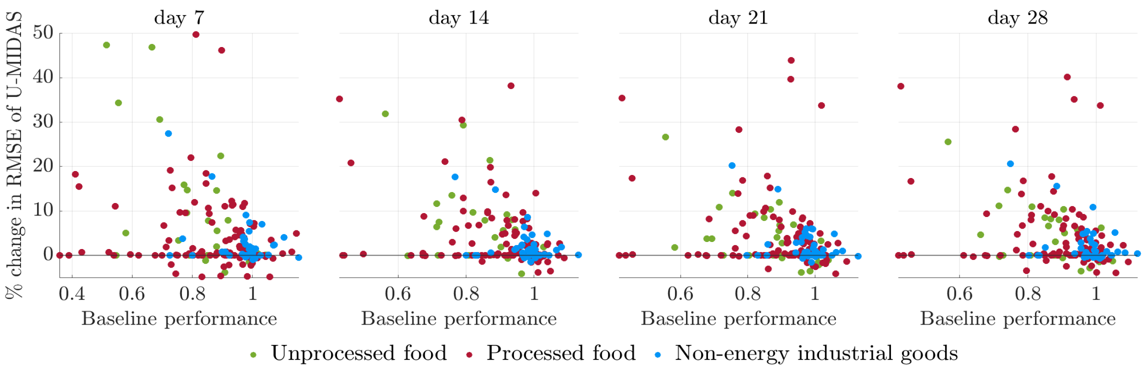

Sources: GfK household panel; own calculations.
Note: For each GFK:FMCG COICOP-10 item, the figure shows the percentage change in RMSE of implementing the best-compiled price index (based on month-over-month inflation correlations), compared to the baseline U-MIDAS (6) as a function of its relative performance to SD-AR.

Figure C6: RMSE change of a moving average smoother as a function of the volatility level

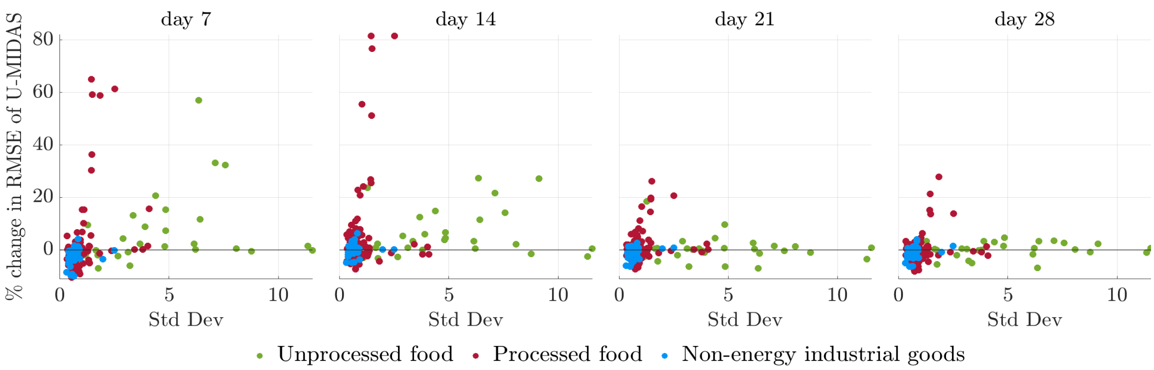

Sources: GfK household panel; own calculations.
Note: For each GFK:FMCG COICOP-10 item, the figure shows the percentage change in RMSE of applying a four-week moving average smoother, compared to the baseline U-MIDAS (6) as a function of the weekly volatility levels (standard deviation of GFK:FMCG month-over-month inflation rates).

Figure C7: RMSE change of additional high-frequency distributed lags

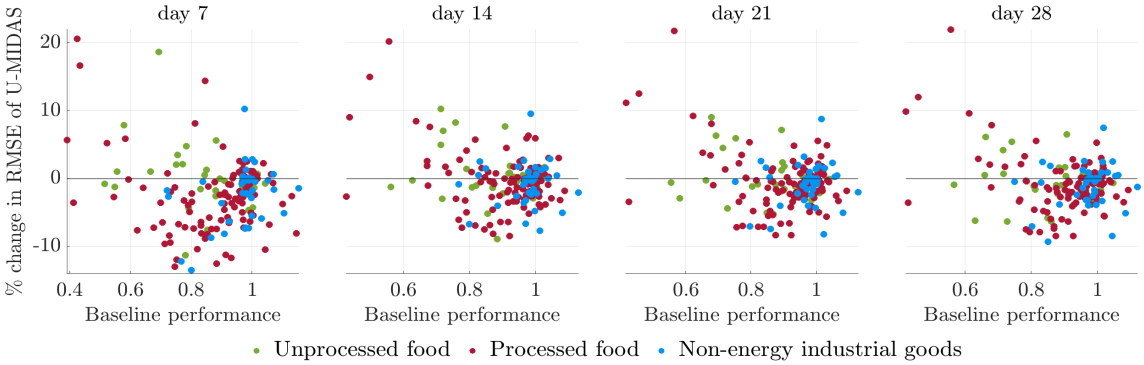

(a) Contemporaneous + high-frequency inflation rates at $t-1$

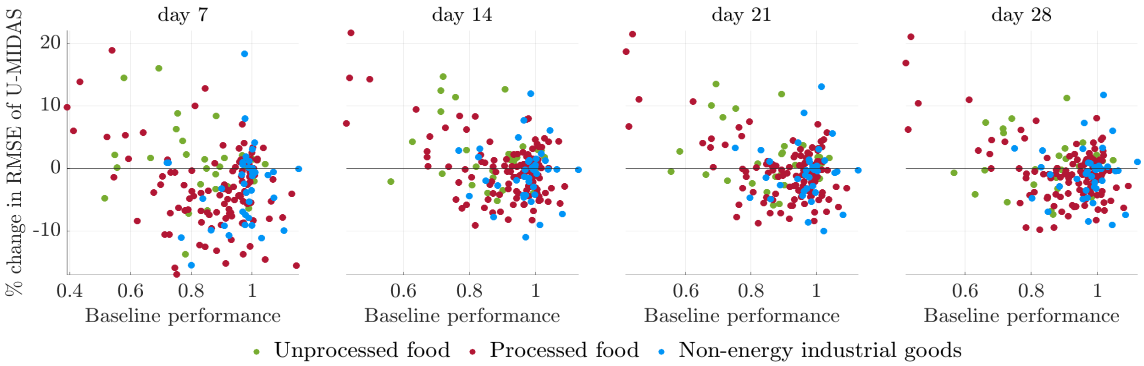

(b) Contemporaneous + high-frequency inflation rates at $t-1$ and $t-2$

Sources: GfK household panel; own calculations.
Note: For each GFK:FMCG COICOP-10 item, the figure shows the percentage change in RMSE of adding past high-frequency lags of GFK:FMCG indicators compared to the baseline static U-MIDAS (6) as a function of its relative performance to SD-AR.