European Commission

Social Situation Monitor

Inflationary shocks, economic aggregates and households' green transition: A causal machine-learning analysis using mixed-frequency data

Milos Ciganovic, Federico d'Amario & Massimiliano Tancioni

ICF

HIVA

EUROPEAN COMMISSION

Directorate-General for Employment, Social Affairs & Inclusion
Directorate F— Employment and Social Governance, Analysis
Unit F4 — Analysis and Statistics

Contact: Karolina Gralek
E-mail: Karolina.gralek@ec.europa.eu

European Commission
B-1049 Brussels

Inflationary shocks, economic aggregates and households' green transition: A causal machine-learning analysis using mixed-frequency data

Milos Ciganovic
Corresponding author
Sapienza University of Rome
email: milos.ciganovic@uniroma1.it

Federico D'Amario
Corresponding author
Sapienza University of Rome
email: federico.damario@uniroma1.it

Massimiliano Tancioni
Corresponding author
Sapienza University of Rome
email: massimiliano.tancioni@uniroma1.it

October 2023

Manuscript completed in October 2023

Luxembourg: Publications Office of the European Union, 2024

PDF ISBN 978-92-68-18106-5 ISSN 2811-6798 doi 10.2767/005050 KE-CF-24-010-EN-N

TABLE OF CONTENTS

INTRODUCTION

In times characterised by substantial increases in price levels (see Figure 1), the interrelations between inflationary shocks, economic activity, unemployment, and household green transition investments are key to understanding the extent to which a changed macroeconomic setting might impact the economy and society.

Figure 1 Evolution of overall Harmonised Index of Consumer Prices (HICP) and energy inflation, European Union (EU), Germany, France and Italy, December 1997 to August 2023

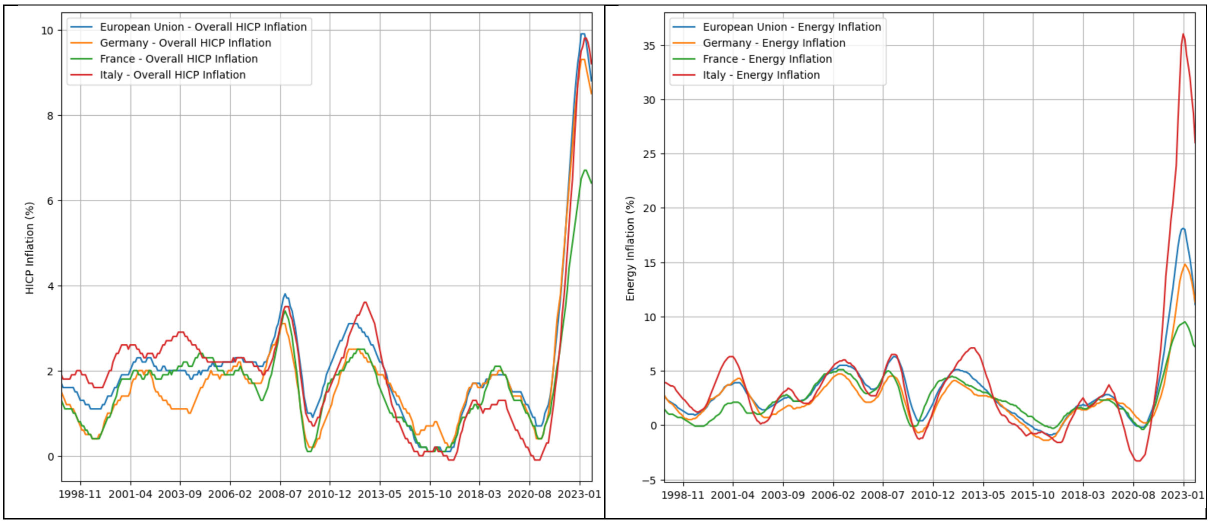

Notes: Overall HICP inflation (left subplot) and energy inflation (right subplot).
Source: Authors' elaboration, based on Eurostat data.

When cost of living increases and nominal wage dynamics remain relatively weak, households' real disposable income decreases, potentially resulting in social and economic instability and lower willingness to participate in energy-saving housing infrastructure policies. However, it can also lead to a potential increase in willingness to participate in such policies to reduce energy costs.

In today's uncertain economic landscape, the impact of inflation is a critical area of investigation. Factors such as sharp increases in commodity prices, sectoral supply constraints, and expansionary policies have been identified as contributing to the magnitude and persistence of the recent surge in inflation (Blanchard and Bernanke, 2023; Gagliardone and Gertler, 2023). The sudden increase in raw materials and energy prices observed since the end of 2021 was initially attributed to the confluence of COVID-19-related bottlenecks (supply component), strong policy interventions (fiscal and monetary), and subsequent reopening of economic and social activities (demand component). Higher production costs did not reduce firms' margins, with some institutional actors citing increases in profit rates as additional components of the price surge. The Russian war of aggression against Ukraine disrupted existing natural gas supply arrangements, fuelling inflationary pressures and uncertainty, both for future price evolution and its macroeconomic impacts. Professional forecasters underestimated the inflationary peak (Aliaj et al., 2023), increasing uncertainty in inflation expectations. This has become a crucial amplifying factor for inflation itself and its transmission mechanisms to the real sector (Cascaldi-Garcia et al., 2023). Despite these macroeconomic challenges, wholesale electricity prices in the EU declined in the first half of 2023. This was aided by the addition of new renewable power capacity, which has helped to mitigate the scale of previous price increases, reducing pressure on inflation. The introduction of the EU natural gas price cap from 15 February 2023 contributed to reducing the volatility created by Russia in the gas market.

Irrespective of the multifaceted origins of the recent price surge, there are three contrasting effects of inflation on the real sector and the ongoing green transition process. Firstly, families across various European countries might be incentivised to take advantage of green policies for structural renovations to reduce gas consumption and transition toward alternative energy sources (Long, 1993; Opoku, 2020).

Secondly, the sharp rise in prices could increase savings by delaying investment and private consumption. This effect may be complicated by the labour and skills shortages in sectors crucial to the green transition, possibly delaying project implementation (Howard, 1978; Callen and Thimann, 1997) and increasing the cost of skilled labour.

Thirdly, households in the lower part of the income distribution may have reduced investments and savings (and/or increases in indebtedness) to sustain the necessary consumption expenditure, perhaps compounded by a worsening labour market. To date, the labour market has been robust (high employment and low unemployment levels), but this may change if prices continue to increase over a prolonged period (Hille and Möbius, 2019; Coibion et al., 2020).

This research note investigates the causal effects of inflationary shocks on household economic indicators, unemployment rates, and other critical labour market indicators for three countries: Germany, Italy and France. These are the three major EU economies and, together, represent nearly half of the EU population. The variables of interest are: unemployment rate, employment rate, part-time employment as a percentage of employment, temporary contracts as a percentage of employment, gross household investment rate, gross household saving rate, and households' consumption expenditure. The research note also analyses the causal effects of inflationary shocks on the take-up of energy-saving housing infrastructure policy (EIP) in Italy, based on the number of requests for the Italian Superbonus 110%.[1] It does not provide such detailed analysis for France and Germany, given the lack of publicly available official data on energy-saving policies implemented. All of the labour market-related target variables are broken down by the main sex and age cohorts.

[1] Tax benefit governed by Article 119 of Legislative Decree No 34/2020 (Relaunch Decree), which consists of a 110% deduction of expenses incurred starting from 1 July 2020 for the implementation of specific interventions aimed at energy efficiency and static consolidation or reducing the seismic risk of buildings. Subsidised interventions include the installation of photovoltaic systems and infrastructure for charging electric vehicles in buildings.

In the absence of a widely accepted framework for understanding how economic shocks spread, this study employs cutting-edge machine learning techniques to analyse the data – specifically, so-called Double Debiased Machine Learning (DML) (Chernozhukov et al., 2018).

This method is designed to account (and control) for the possibility that a multitude of factors could jointly affect economic outcomes following unexpected shocks of different natures. Compared to more standard alternatives, it ensures a more accurate assessment of the direct impact of inflation on key economic indicators, without relying on contentious assumptions about cause and effect. 'Causal' machine learning integrates the power of automated learning algorithms and the concept of causality, taking advantage of the ability to select the most informative control variables from high-dimensional datasets.

The knowledge garnered from this study could be useful to policymakers designing and implementing evidence-based policy interventions to overcome challenges encountered by households in transitioning to eco-friendly practices. These interventions can concurrently address the adverse impact of inflation on unemployment and on the economy. The findings also have the potential to inform the development of policies targeting the specific financial constraints faced by households and businesses. Ultimately, it contributes to the wider discourse on the interrelation between inflation, green transition and unemployment in the context of a sustainable and inclusive economic recovery.

The remainder of the paper is organised as follows: Section 2 contains the literature review, Section 3 presents the data definition and issues, variables considered, issue of mixed frequencies, and ragged edges. Section 4 presents the estimation methodology. Section 5 summarises some selected results, while Section 6 provides some suggestions for further research. Appendix A provides information on the Superbonus dataset, Appendix B presents further literature on the methodology, and Appendix C gives details on the methodology employed.

LITERATURE REVIEW

A vast body of literature explores the economic impact of inflation. While this study does not aim to provide a comprehensive review of the literature, several important studies are described below, focusing on the specific economic variables in question.

Inflation's multifaceted influence on personal saving behaviour has been explored in various studies, beginning with Howard (1978) identifying the erosion of the real value of net liquid assets, thereby affecting individuals' willingness to save. The impact varies, depending on whether inflation is expected or unexpected, with evidence suggesting that inflation expectations may discourage personal saving behaviour in some contexts (e.g. Japan). This research note, by construction, addresses the impact of the unexpected inflation component, while also tackling expected inflation, by providing an analysis of the dynamic transmission mechanics.

Building on the theme of inflation's impact on personal finance, Lieb and Schuffels (2022) explored the reaction of consumer spending to anticipated inflation, analysing whether or not distinct aspects of households' financial positions determine how which consumption responds to expected inflation. The study highlighted two plausible channels. Firstly, a rise in inflation expectations can bolster consumption through direct increases in anticipated real wealth among households with nominal financial liabilities. Secondly, by impacting the real interest rate, expected inflation may interact with wealth only where households are not budget-constrained or have sufficient liquidity to adjust consumption in line with current real interest rates, facilitating a shift of funds between consumption and savings. This analysis complements the work by Lieb and Schuffels (2022) by providing an alternative approach to considering possible heterogeneous effects of inflation for different segments of the labour force.

In the context of Organisation for Economic Co-operation and Development (OECD) economies, the determinants of household saving have been a focal point of research. Opoku (2020) undertook a comprehensive examination, revealing that several factors – including inflation – have a significant negative long-term impact on household saving. Callen and Thimann (1997) analysed the influence of tax and social security systems, underscoring the importance of governmental structures in shaping household saving behaviour and enriching understandings of how inflationary factors may play a role. In Italy, Coibion et al. (2020) explored the causal relationship between inflation expectations and firms' economic decisions in the corporate sector, finding that increased inflation expectations lead to higher prices and credit utilisation, while simultaneously reducing employment levels, illustrating the far-reaching effects of inflation on the labour market.

Hille and Möbius (2019) conducted a nuanced analysis of the relationship between energy prices and employment, revealing that rising energy prices had no significant effect on net employment within manufacturing sectors, but significantly positive effects on the broader economy. A study analysing the dynamics of crude oil prices of Organization of Petroleum-Exporting Countries (OPEC) and non-OPEC countries gained insights into the relationship between energy prices and employment, highlighting competitive behaviours between OPEC and non-OPEC countries, and offering valuable perspectives on the broader economic implications of energy price fluctuations. In the context of energy conservation (i.e. decision and practice to use less energy), an econometric analysis by Long (1993) confirmed that rising energy prices create an incentive for households to invest in energy-saving measures. More recently, Gagliardone and Gertler (2023) showed that the surge in inflation is driven by a combination of oil price shocks and 'easy'' monetary policy, even after controlling for demand shocks and shocks to labour market tightness. The key mechanism explaining the strength of the transmission channel was the low elasticity of substitution between oil and labour, preventing adaptation in the productive sphere. This research note complements these studies, particularly Hille and Möbius (2019) and Long (1993), by addressing the specific roles of aggregate inflationary shocks and energy sector inflationary shocks.

References for the selected estimation procedure were drawn from very recent literature on causal machine learning. The field of causal machine learning has its roots in a seminal paper by Belloni et al. (2014), in which the authors discussed how advancements in data mining could be used to enhance the inference of structural model parameters. Their work inspired other innovations and enhanced the landscape of machine learning applied to causal inference.[2] The methodology used here draws significant inspiration from works on causal machine learning in general (Belloni et al., 2014; 2017) and Chernozhukov et al. (2018). The latter presented an innovative approach, DML, for estimating treatment and structural parameters using a variety (in principle, any) of machine learning algorithms. The method is particularly useful in complex environments with many variables to consider. The paper by Chernozhukov et al. (2018) is widely considered one of the most significant recent advancements in the field of causal inference.

[2] Farrell (2015), Belloni et al. (2017), Hartford et al. (2017).

DATA ELABORATION

DATA COLLECTION

This study uses a comprehensive approach to estimate the causal impact of inflation on Italy, Germany and France. Highly detailed, country-specific datasets are created that include outcome (target) and treatment (shock) variables, as well as an extensive range of series used as controls for confounding factors (variables that may contemporaneously have a causal impact on both outcome and treatment variable, leading to endogeneity issues, i.e. biased estimates). Having data on confounding factors allows their use as 'controls' to overcome potential biases from their omission. Figure 2 presents a possible confounding factor which can be adopted in the investigation of the causal effect of education level on unemployment. Given that family wealth can affect both variables, it is considered a good control (confounding factor).

Figure 2 Example of confounding factor

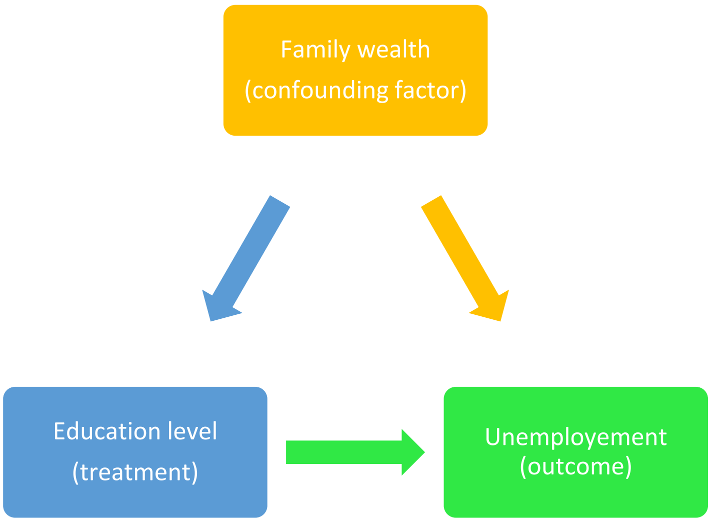

Source: Authors' elaboration.

The datasets are drawn from different sources, including national statistics institutes, national central banks, Eurostat and the European Central Bank (ECB). Alongside the outcome and treatment variables, a large number of possible confounding factors are selected. The datasets comprise 190 (Italy), 140 (Germany), and 217 (France) variables (outcome, treatment, controls).

The investigation is rooted in the study of a monthly dataset, stretching from 28 February 2011 to 31 December 2022. To further enrich the research and to shed light on the specific impact of inflation on the uptake of the Italian EIP Superbonus 110%, a specific weekly dataset[3] was also used for Italy, spanning 27 August 2021 to 31 December 2022. The use of high-frequency data is motivated by the need for sufficient sample information for the specific periods impacted by the policy targeted in the study. Exploring these different timeframes and datasets enables a more comprehensive picture of the intricate relationship between inflationary tendencies and the broader economic landscape. All variables, except those related to Superbonus 110%, are log-differentiated. Consequently, the estimated parameters can be interpreted as elasticities[4] between treatment and target variables, indicating the percentage increase in target variable following a 1% increase in the treatment variable.

Data limitations mean that this is inherently an aggregate, macro-level analysis. Reliance on very recent periods in which the inflationary shocks materialised implies an analysis deeply rooted in the high-frequency time dimension. With the lowest of the mixed-frequency data fixed at quarterly frequency, individual or distributional effects (whose data are available only annually and with a substantial time delay) are outside the scope of this analysis. Any future study of the distributional effects of inflationary shocks would be facilitated by the availability of more up-to-date data and new evidence on potential transmission channels.

A group of 15 target variables[5] is identified for each country for which the causal effect of shocks to HICP inflation and energy inflation is estimated. The variables chosen are:

- Unemployment rate;
- Employment rate;
- Part-time employment as percentage of employment;
- Temporary contracts as percentage of employment;
- Gross household investment rate;
- Gross household saving rate;
- Households' consumption expenditure.

In order to eliminate biases due to potential confounders, control variables were carefully selected, reflecting a multitude of aspects related to the overall economy of each country. More specifically, the variables chosen pertain to the labour market (including activity rate and labour cost index), as well as sentiment indexes encompassing construction, retail, household, and service confidence indicators. Financial variables such as debt securities and interest rates, as well as the prices of different energy items and industrial production indexes, are also included.

[3] See Table A1 in Appendix A for a detailed description of the dataset.

[4] The elasticity 'θ' between variables A and B is defined as the percentage change in A due to a percentage change in B; synthetically: $\theta = \frac{\%\Delta A}{\%\Delta B}$.

[5] For labour market-related target variables, the estimated causal effect also considers sex and age cohorts.

A specialised weekly dataset for Italy determines the causal effect of inflation on EIP. This step is necessitated by the scarcity of official time series (either limited or not publicly available) on the progress of these green policies. A four-step approach is used to overcome this challenge:

1. Official statistics on national energy-saving housing infrastructure policies are collected, where available;[6]
2. Confounding factors and treatment variables are selected from the general dataset;
3. Weekly proxies (Google Trends data and weakly official data) are gathered to update the official time series from the previous steps;[7]
4. A weekly dataset is created by imputing missing values using the Kalman Filter (KF) and expectation maximisation (EM) algorithms and time interpolation methods, using proxies collected in step 3 as high-frequency drivers of unobservables (state variables).[8]

Table A1 in Appendix A provides an overview of the dataset composition for the Superbonus policy in Italy, alongside corresponding weekly state variables used for the purposes of imputation.

Although useful for handling missing data, imputation comes with certain caveats. For example, it might not be able to consider extreme values as missing observations when the algorithm is biased by overfitting (i.e. giving accurate predictions for a subset of data used in the estimation but not for other (new) data) to true low-frequency series. It can also affect data variability. Missing values that are replaced may increase or reduce the variability of the series and have an impact on statistical analysis and model performance evaluation. Addressing these issues is crucial for accurate results.

GOOGLE TRENDS

Google Trends provides insights into the popularity of search terms by displaying the frequency with which they are entered into Google's search engine, relative to the total search volume on the site during a specific period. These frequency data are presented at different temporal resolutions, depending on the timeframe in question. For example, daily data are available for time periods of fewer than nine months, while weekly data are presented for periods ranging between nine months and five years. Monthly data are provided for timeframes exceeding five years.

[6] EIP data for Italian Superbonus 110% were collected from the Italian National Agency for New Technologies, Energy and Sustainable Economic Development (ENEA), a public research institution operating in the fields of energy, environment, and new technologies to support competitiveness and sustainable development policies, https://www.enea.it/it

[7] Google Trends used here can be found in the weekly proxy column in Table A1 of Appendix A.

[8] The Kalman EM filter is particularly suitable for providing information-based estimates of (high frequency) missing data within (low frequency) observables. It is based on a state-space relation, where states are defined by high frequency information (such as Google Trends) used to fill the gaps of the low frequency information.

MISSING VALUES IMPUTATION

Generating a vast dataset with a multitude of control variables is a crucial element in the application of DML to estimate a causal relationship. This entails merging several time series, obtained at various frequencies, into a monthly dataset for general target variables and a weekly frequency dataset for EIPs. However, the use of mixed-frequency datasets can lead to missing values within observations detected at different frequencies, and missing values at the end of the samples ('ragged edges'). These irregularities must be addressed before inputting the data into the model to obtain an estimate of causal relations. The most commonly applied method for handling these issues is the KF, often combined with the EM algorithm.[9]

Ideally, an optimal imputation method should be capable of recovering true observations of a dataset after their removal during evaluation. However, this is not feasible in practice, as true observations in weekly frequencies are unavailable for comparison. Instead, the second-best solution must be employed, comparing actual and imputed observations, which are one time-step before them. The theory is that if the imputation mechanism works correctly, the transition from one observation to the next should be smooth (minimising internal variability) and imputed observations surrounding an actual value should be close to it.

In practice, the unobserved EIP weekly measures[10] are estimated using data from lower frequencies – monthly and quarterly data. The latent observation Vector Autoregression method (L-VAR) is employed (as introduced in Cimadomo et al. (2022) and recently applied in Aliaj et al. (2023)). This methodology treats the missing weekly observations as latent processes (i.e. unobserved or hidden variables inferred based on available data), which can be estimated using the KF. This constructs a systematic weekly dataset that can account for an irregular flow of information in the time dimension, such as mixed frequencies with ragged edges. In order to verify the empirical reliability of the high-frequency missing data imputation process, Pearson correlation coefficients are calculated between the lower frequency outcome variable and its weekly proxy (see Table A1 in Appendix A), used for imputation (see Table 1). The results indicate that the cumulative sum of Google Trend searches for 'Superbonus' shows a strong positive correlation with all official time series pertaining to the policy.

[9] Bańbura et al., (2010); Alternative imputation methods were examined: the performance of various techniques was evaluated, including mean imputation, decision trees, random forests, forest-like mechanisms, Bayesian ridge, k-nearest neighbours, nuclear norm minimisation, and the KF combined with the EM algorithm. Given space constraints, this research note outlines only the general approach taken to determine the optimal technique for the analysis.

[10] High frequency data are necessary for sufficient density of information around the periods impacted by the policy in question.

Table 1 Pearson correlation matrix between variables used to impute the number of requests for Italian Superbonus

	Cumulative number of requests for Superbonus (Google Trend)	Number of requests for Italian Superbonus
Cumulative number of requests for Superbonus (Google Trend)	1.00	0.95
Number of requests for Italian Superbonus	0.95	1.00

Notes: Pearson correlation measures the strength of the linear relationship between two variables. It has a value between -1 to 1, with a value of -1 meaning a perfect negative linear correlation, 0 being no correlation, and + 1 meaning a perfect positive correlation.
Source: Authors' elaboration.

Where official time series are available at a higher frequency and exhibit high correlation both statistically and topically, they are preferred to sentiment and attention proxies. For example, the monthly interest rate data pertaining to deposits, bank borrowings and loans are transformed into weekly data using official ECB policy rates.

Figure 3 Monthly interest rates vs high frequency ECB rates, Italy, 1 January 2003-31 December 2022

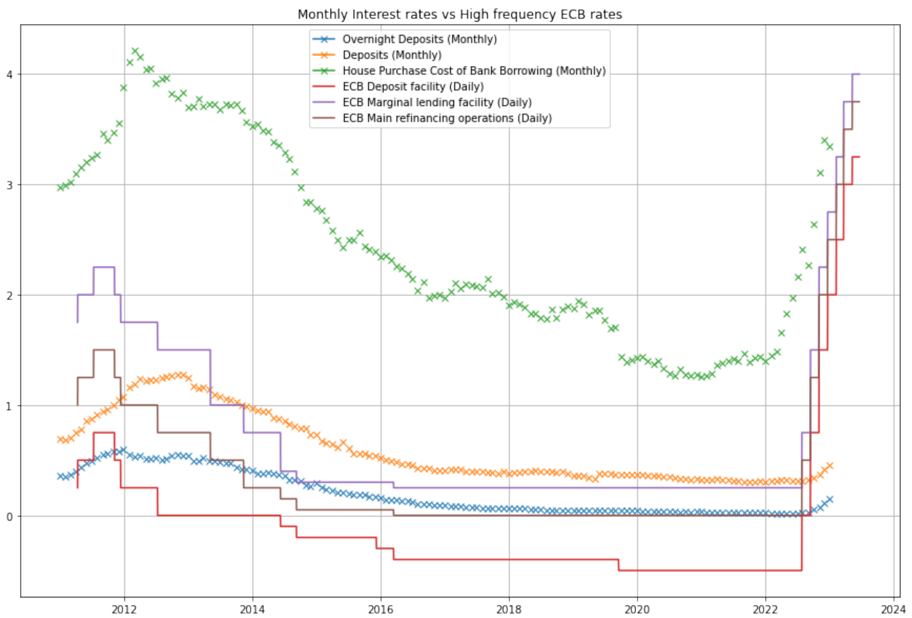

Notes: Overnight deposits, deposit and house purchase of bank borrowing for Italy, which are monthly interest rates (represented by cross) and official daily ECB rates of deposit facility, marginal lending facility and main refinancing operations (solid lines).
Source: Authors' elaboration based on Bank of Italy and ECB data.

The trends observed in the monthly interest rates used as control variables and the official ECB rates are notably similar, suggesting that the official daily ECB rates may be used as a reliable proxy for imputing higher frequency (monthly) interest rate data.

METHODOLOGY

The empirical analysis comprises a partially linear model (i.e. linear in the treatment effect and possibly non-linear in controls for confounding factors) and the DML method to estimate the average effect of inflation on the target variables.

DML AVERAGE TREATMENT ESTIMATES

The methodology is specifically designed to estimate the true relationship between an outcome variable, y_t, and a treatment variable, d_t, when a very large number of other factors (X_t) might affect this relationship through their relations with both outcome and treatment variables. The omission of these other factors, known as confounding variables, could bias the estimates of the treatment variable effects on the outcome variable. Machine learning algorithms provide an effective tool to select the most important from a large set of controls for the outcome and treatment variables. They also allow for both linear and more complex non-linear relationships with the outcome (function g_0) and treatment (function m_0) variables. Formally:

$$y_t = \theta_0 d_t + g_0(X_t) + \xi_t, \qquad \xi_t \sim \mathcal{N}\left(0, \sigma^2_{\xi_t}\right) \tag{1}$$

$$d_t = m_0(X_t) + v_t, \qquad v_t \sim \mathcal{N}\left(0, \sigma^2_{v_t}\right) \tag{2}$$

The model, in its structural form, thus consists of two main equations: the first equation focuses on the outcome of interest (i.e. real sector and labour market variables); and the second focuses on the treatment variable considered (i.e. HICP and energy price shocks). The key parameter of interest is θ_0. This indicates how much the treatment variable affects the outcome variable. In this model specification, the parameter is seen as an average causal effect of the treatment variable on the outcome variable, computed using all information in the sample. In practice, DML works on a reduced-form representation of the structural system exemplified above. Removing the treatment variable from the first equation, it becomes:

$$y_t = m_{y,0}(X_t) + v_{y,t}, \qquad v_{yt} \sim \mathcal{N}\left(0, \sigma^2_{\xi_t}\right) \tag{3}$$

$$d_t = m_0(X_t) + v_t, \qquad v_t \sim \mathcal{N}\left(0, \sigma^2_{v_t}\right) \tag{4}$$

where $m_{y,0}$ is a generic, possibly non-linear function relating the outcome variable to the confounding factors (controls) once the role of the treatment is not considered in the equation. This makes a reduced-form system, or an essentially predictive-purposes system, for which machine learning algorithms are specifically designed. Such algorithms, by searching for the best predictors among a high-dimensional set of controls, reduce the regularisation bias (i.e. they solve the problem of focusing solely on a set of user-defined predictors). The analysis is repeated multiple times on different datasets to avoid overfitting. This gives a 'clean' measure of how much the treatment actually affects the outcome, free of the influence of other variables. For example, consider the impact of energy price changes on the adoption of renewable energy policies: in order to estimate the impact, the effects of all other types of inflation must first be removed.

After a system of outcome and treatment variables reduced-form equations with high-dimensional controls is estimated, applying the machine learning algorithms, two residual components are obtained, v_{yt} and v_t, the first constructed by excluding the treatment variable d_t from equation (1) and the second by defining the unexpected component (i.e. the shock) of the treatment variable.

Based on these estimated residuals, a simple residual-on-residual Ordinary Least Squares (OLS) regression provides unbiased estimates of parameter θ_0:

$$\widehat{v_{y,t}} = \theta_0 \widehat{v_t} + u_t \tag{5}$$

In a time-series context, DML follows the same approach, with two key differences. The first is that, in a time series setting, time dependence (memory) must be addressed. This rules out the use of classical (time-independent) random sampling, as the presence of memory components calls for time-consistent sequences, such as excluding a number of observations before and after each test set, and ensuring independence between train and test sets. This is relevant in controlling regularisation bias (tuning of prediction equations) and overfitting bias. In both cases, time series-specific sampling is applied based on the 'neighbours-left-out'[11] cross-fitting procedure recently introduced by Semenova et al. (2023). This study considers 12 subsamples of equal sizes.

The second difference is the use of a dynamic model, which is made possible by the presence of time series information. The estimation of the reduced-form outcome equation (3) is simply repeated for different time leads, with the residual-on-residual equations (5) estimated as a local projection. More specifically, the different leads of the outcome equation's residual component are regressed on the present-time treatment equation's residual component.

[11] Neighbours-left-out method consists of dividing the sample in K blocks (subsamples) of equal size. K blocks are equivalent to K folds, i.e. in each fold, one block is used as a test, the neighbour blocks are left out, and the rest are used as train set.

RESULTS

This section presents the study findings, notably the average causal implications of two treatment variables: HICP inflation and energy price inflation. The controls are modified for different treatment variables. For example, when using energy price inflation as a treatment variable, HICP inflation is excluded from controls, as the former is a component of the latter. Similarly, when using the HICP as treatment, the different components of inflation are excluded from controls. In a macroeconomic setting, where aggregation issues are pervasive, accurate separation of treatments and transmission channels is vital for obtaining a clear evaluation of causal effects. Take, for example, an interest in determining the causal effect of aggregate HICP inflation on employment: if some of the aggregate price variation components are considered controls, the effect on unemployment estimated relates to the uncontrolled components, which is not the desired estimation.

This procedure enables the identification of inflationary shocks from other sources of macroeconomic variability. However, it cannot address the precise nature of the price shocks without making additional assumptions. In other words, an unexpected (exogenous) increase in prices can be identified, but its origin cannot be precisely determined. The positive inflation change might stem from an increase in demand or from a contraction in the supply-side. Simultaneously, the channels through which inflationary shocks affect the outcome variables considered cannot be addressed empirically. This analysis can only rely on theoretical speculations or hypotheses.

The effects being analysed should be considered independently of any time-specific aspects, including dynamic fluctuations. Commonly adopted in causal analysis, this approach interprets these impacts as an average treatment effect (ATE). Essentially, it explores the average causal influence of inflationary shocks on selected target variables of interest, where the average pertains to the time dimension. In incorporating dynamic patterns in the response of various variables, ATEs may lead to ambiguous results.

As the impact of inflationary shocks on outcome variables can fluctuate over time, impulse response functions obtained through linear projections are provided. Leads of structural outcome variable residuals are regressed on present-time structural residuals of the treatment equation, offering an estimate of impulse response functions. These dynamics give key insights into the time-sensitive reactions of the outcome variables to alterations in inflationary pressures, facilitating a more profound and nuanced understanding.

ITALY

Consistent with the broader pattern within the EU, Italy has experienced an intense escalation in energy costs, with the annual energy inflation rate surging to an unprecedented 38.3% in January 2023. This dramatic uptrend is principally attributable to the geopolitical tensions ignited by the Russia-

Ukraine crisis, a key consideration given Ukraine's historically significant role as a major supplier of Europe's natural gas. Concurrently, Italy is grappling with a peak in HICP inflation (January 2023), as the country with one of the highest inflation rates within the EU. During November 2022, the inflation rate increased markedly, achieving a record high of 12.6%. This primarily reflected escalating energy and electricity costs, underlining the profound macroeconomic effects of rising energy prices. Fluctuations in HICP inflation and energy inflation seem to influence labour market dynamics in opposite ways (see Table 2). An increase in HICP inflation does not appear to have a statistically significant impact on the aggregate employment rate, although a clearer picture could be obtained with sectoral-level data (out of scope of this analysis). The sole exception is the total employment rate for individuals aged 25-34, which demonstrates negative elasticity to an inflation shock. More specifically, a 1% rise in HICP inflation results in a 0.14% decrease in the employment rate for the period considered. Conversely, an increase in energy inflation appears to enhance employment rates across most demographic groups, potentially owing to sectors sensitive to energy price changes adapting in ways that encourage job creation.

Table 2 ATE – employment, Italy

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Employment rate – total, aged 25-34 years	-0.1416***	0.3038***
Employment rate – total, aged 35-49 years	-0.0341	0.0497
Employment rate – total, aged 50-64 years	-0.0005	0.0603
Employment rate – total, aged 15-64 years	-0.0394	0.0681**
Employment rate – women, aged 15-64 years	-0.0372	0.0854***
Employment rate – men, aged 15-64 years	-0.0369	0.0347

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

A unique reaction to inflation indicators is observed in relation to part-time and temporary employment rates. While an increase in HICP inflation does not show a statistically significant impact, energy inflation appears to increase part-time employment rates across various demographic groups. This may signify a transition in labour practices towards roles characterised by greater flexibility, possibly as a response to the change in energy prices.

Table 3 ATE – part-time employment, Italy

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Part-time employment rate - total, aged 15-64 years	0.0526	0.2616***

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Part-time employment rate - total, aged 25-49 years	0.0710	0.2873***
Part-time employment rate - total, aged 50-64 years	0.0247	0.1070*
Part-time employment rate - women, aged 15-64 years	0.0315	0.1734***
Part-time employment rate - women, aged 15-39 years	0.0849*	0.2843***
Part-time employment rate - women, aged 25-49 years	0.0369	0.2224***
Part-time employment rate - women, aged 50-64 years	-0.0001	0.0064
Part-time employment rate - women, aged 50-74 years	0.0073	0.0435
Part-time employment rate - men, aged 15-64 years	0.1015	0.2432***
Part-time employment rate - men, aged 15-39 years	-0.09	0.0337
Part-time employment rate - men, aged 25-49 years	0.1	0.3223***
Part-time employment rate - men, aged 50-64 years	0.0404	0.3748***
Temporary contract employment rate – total	0.01	0.2555***
Temporary contract employment rate – women	0.0458	0.3535***
Temporary contract employment rate – men	-0.0189	0.2276***

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

Looking at Italy's unemployment rates, a rise in HICP inflation is associated with a significant reduction in unemployment rates across all categories, implying that during inflationary periods, the labour market may experience a surge in job creation. This is broadly consistent with a standard Phillips relation (i.e. a demand shock-driven inflation-unemployment trade-off). By contrast, an increase in energy inflation tends to increase the unemployment rate across the different breakdowns. The emergence of such a positive relation between energy inflation and unemployment is consistent with (negative) supply shocks, confirming the results in Coibion et al., (2020). As the coefficients represent elasticities, and after controlling for participation rates, the overall effect of a HICP shock results in reduced unemployment. For an energy inflation shock, the overall elasticity of the unemployment rate is higher compared to employment, suggesting that the economy may experience an increase in unemployment proportionally higher than the positive stimulus to employment.

Table 4 ATE – unemployment, Italy

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Unemployment rate - total	-0.4399***	0.3706***
Unemployment rate - total, aged 15-64 years	-0.4750***	0.3661***
Unemployment rate - total, aged 25-34 years	-0.2607***	0.5199***

	Treatment variable:	
	HICP inflation	Energy inflation
Outcome variable:		
Unemployment rate - total, aged 35-49 years	-0.4698***	0.5651***
Unemployment rate - Total, aged 50-64 years	-0.5467***	0.1457
Unemployment rate - Total, aged 50-74 years	-0.3678***	0.1671
Unemployment rate - women	-0.6815***	0.3448***
Unemployment rate - men	-0.1909**	0.4592***

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

Finally, the study examines the effects of inflationary shocks on some key macroeconomic variables that indicate household economic behaviour. Italian households seem to respond to HICP inflation by markedly reducing their consumption and investment rates, suggesting a more conservative financial approach during inflationary periods. This result is in line with the theoretical hypothesis that inflationary shocks are detrimental to investment and consumption decisions, with savings prioritised. The results for consumer price inflation confirm those in the early empirical literature on the effects of inflation for savings and investment (Howard, 1978; Callen and Thimann, 1997). Interestingly, energy inflation is found to reduce saving but increase investment, possibly indicating that households may be reallocating resources from savings to investments, even through borrowing, in order to exploit the possible gains from participation in the green policies introduced by the government. This finding for the energy component of inflation is consistent with the results in Long (1993) and Opoku (2020).

Table 5 ATE – households, Italy

	Treatment variable:	
	HICP inflation	Energy inflation
Outcome variable:		
Household consumption expenditure	-0.1388***	-0.0526
Gross household saving rate	0.1265*	-1.0177***
Gross household investment rate	-0.3082***	0.1490***
Number of requests for Italian Superbonus	-12.5714***	5.4772***

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

Such a crowding-in result is confirmed by the direct analysis of participation in green transition policies in Italy, where an interesting difference emerges in the average treatment effect of HICP and energy price shocks: the number of applications for the Italian Superbonus responds negatively to the overall HICP inflation shock, whereas its response is positive and statistically significant for energy-specific shocks. More specifically, a percentage point (pp) increase in energy inflation causes a 5.5% increase (approx.) in the number of requests for the Italian Superbonus. This result calls for a distinction of the effects by specific source of the shock, i.e. for energy price shocks, households could be encouraged to leverage green transition policies to undertake infrastructural upgrades to minimise natural gas usage and shift to alternative forms of energy, minimising their exposure to future energy shocks.

FRANCE

In France, HICP inflation accelerated significantly during 2022. After limited disinflation [12]in summer 2022, driven by lower petrol prices, the inflation rate reached a new peak in October 2022. France is comparatively less affected by the energy crisis than other European countries, given that its economic growth relies less on industrial production and its dependence on gas is therefore lower. Nor is France as reliant on Russia for its gas supply, as its energy mix is dominated by nuclear power sources. It appears that HICP inflation does not have a substantial impact on employment across most demographic sections. The only significant result is in the employment of men, which has a negative elasticity to inflation. However, a rise in energy inflation appears to significantly stimulate employment, particularly among women. This implies that energy price changes might prompt certain sectors to adapt in ways that encourage increased hiring, with women benefitting most from this shift.

Table 6 ATE - employment, France

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Employment rate - total	-0.0210	0.0982***
Employment rate - total, aged 15-64 years	-0.0088	0.1048***
Employment rate - total, aged 25-49 years	-0.0031	0.0446*
Employment rate - total, aged 50-64 years	-0.0004	0.0430
Employment rate - women	0.0108	0.0943***
Employment rate - women, aged 15-64 years	-0.0043	0.0919***
Employment rate - women, aged 25-49 years	-0.0122	0.0193
Employment rate - women, aged 50-64 years	0.0245	0.0805***
Employment rate - men	-0.0602***	0.0530*
Employment rate - men, aged 15-64 years	-0.0565**	0.0298
Employment rate - men, aged 25-49 years	-0.0266	0.0595*
Employment rate - men, aged 50-64 years	-0.0080	0.0167

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

In France, HICP inflation shock positively affects the part-time employment rate of younger women. Energy inflation significantly increases part-time employment among women for all age groups. By contrast, both HICP and energy inflation shocks reduce part-time employment among men across all

12 Reduction in the rate of inflation

age groups. This could be an indicator of a gendered labour market response to the types of price fluctuations considered, where women might be taking on more part-time roles, while men's part-time employment declines.

Table 7 ATE – part-time employment, France

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Part-time employment rate - total, aged 15-64 years	-0.0128	0.0976*
Part-time employment rate - total, aged 25-49 years	-0.0409	0.0229
Part-time employment rate - total, aged 50-64 years	0.02018	0.0365
Part-time employment rate - women, aged 15-64 years	0.1217**	0.2375***
Part-time employment rate - women, aged 25-49 years	0.1593***	0.1264*
Part-time employment rate - women, aged 50-64 years	0.0780	0.2204***
Part-time employment rate - men, aged 15-64 years	-0.5044***	-0.2585***
Part-time employment rate - men, aged 25-49 years	-0.9165***	-0.4333***
Part-time employment rate - men, aged 50-64 years	-0.3002***	-0.1982***
Temporary contract employment rate - total	0.3040***	0.5707***

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

Exploring French unemployment rates, a remarkable reduction across all categories is found as HICP inflation increases. This phenomenon can be attributed to two potential mechanisms. Firstly, periods of higher inflation may stimulate an increased demand for labour, as firms seek to capitalise on economic activity. Secondly, the erosion of purchasing power due to inflation might prompt individuals to accept job offers they might otherwise decline. Both factors contribute to a tighter labour market, reducing unemployment rates. Interestingly, energy inflation has varied effects, increasing unemployment rates overall (Coibion et al., 2020) but decreasing among women aged <25 years. This implies a complex interplay between energy prices and job market stability. Similar to the Italian results, a HICP shock produces an overall reduction in unemployment that is more than proportional to the reduced employment rate. Energy inflation shock seems to produce an overall increase in unemployment, whose elasticity to this type of shock is proportionally higher than that of employment rate.

Table 8 ATE – unemployment, France

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Unemployment rate - total	-0.6351***	0.2100***
Unemployment rate - total, aged 25-49 years	-0.9308***	0.3915
Unemployment rate - total, aged 50 years or over	-0.6600***	0.4768***

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Unemployment rate - total, aged below25 years	0.2107***	-0.3944***
Unemployment rate - women	-0.9335***	0.0256
Unemployment rate - women, aged 25-49 years	-0.9895***	0.3131***
Unemployment rate - women, aged 50 years or over	-1.0229***	-0.0218
Unemployment rate - women, aged below 25 years	-0.3865***	-0.6742***
Unemployment rate - men	-0.4601***	0.3762***
Unemployment rate - men, aged 25-49 years	-0.8167***	0.5396***
Unemployment rate - men, aged 50 years or over	-0.3658***	0.8120***
Unemployment rate - men, aged below 25 years	0.5125***	-0.2981***

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

Finally, households tend to decrease their consumption expenditure and investment rates while increasing their saving rate in response to HICP inflationary shocks. This suggests a shift towards more conservative financial strategies. However, rising energy inflation appears to drive down gross household saving rate, possibly due to higher costs of living. The positive effect on household investment can be linked to energy-saving technologies, as households may seek to decrease the magnitude of future energy shocks (Long, 1993; Opoku, 2020).

Table 9 ATE - households, France

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Gross household investment rate	-0.1918***	0.6213***
Gross household saving rate	0.6143***	-1.6943***
Household consumption expenditure	-0.1719***	0.3562***

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

GERMANY

In Germany, the annual inflation rate reached 7.9% in 2022, its highest in over seven decades (albeit lower than that seen in other EU Member States).

Energy inflation shows a statistically significant negative effect on employment rates among men aged 15 to 74 (See Table 10). The effect on younger cohorts is not statistically significant. This suggests that energy price increases might be more detrimental to older men's employment in Germany, possibly due to higher employment in energy-sensitive sectors. Overall, HICP inflation shock does not have a statistically significant effect on employment in Germany.

Table 10 ATE - employment, Germany

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Employment rate - total, aged 15-64 years	0.0372	-0.0025
Employment rate - total, aged 15-74 years	0.0243	-0.0026
Employment rate - total, aged 25-49 years	0.0159	-0.0128
Employment rate - total, aged 50-64 years	0.0113	-0.0069
Employment rate - women, aged 15-64 years	0.0212	0.0147
Employment rate - women, aged 15-74 years	0.0156	0.0131
Employment rate - women, aged 25-49 years	-0.0112	0.0005
Employment rate - women, aged 50-64 years	0.0231	-0.0084
Employment rate - men, aged 15-64 years	0.0279	-0.0378*
Employment rate - men, aged 15-74 years	0.0140	-0.0457**
Employment rate - men, aged 25-49 years	0.0295	-0.0194
Employment rate - men, aged 50-64 years	0.0429	-0.0098

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

An increase in HICP inflation leads to a modest increase in part-time employment, particularly for men, and a statistically significant positive effect on temporary contracts for the different groups considered. By contrast, energy inflation appears to decrease part-time employment rates significantly, particularly for men and more intensely among those aged 20-64 years. It may be that sectors employing part-time workers are more susceptible to energy price fluctuations, leading to a reduction in part-time work opportunities when energy prices escalate.

Table 11 ATE – part-time employment, Germany

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Part-time employment rate - total, aged 15-64 years	0.0344	-0.0323
Part-time employment rate - total, aged 15-74 years	0.019	-0.0537
Part-time employment rate - total, aged 20-64 years	0.0095	-0.0958**
Part-time employment rate - women, aged 15-64 years	-0.01	0.0258
Part-time employment rate - women, aged 15-74 years	0.0049	0.0331
Part-time employment rate - women, aged 20-64 years	0.0124	0.0066
Part-time employment rate - men, aged 15-64 years	0.06	-0.1719***
Part-time employment rate - men, aged 15-74 years	0.1032	-0.1959***

	Treatment variable:	
	HICP inflation	Energy inflation
Outcome variable:		
Part-time employment rate - men, aged 20-64 years	0.1897*	-0.3059***
Temporary contracts employment rate - total	0.2378***	-0.0634*
Temporary contracts employment rate - women	0.3129***	-0.0657*
Temporary contracts employment rate - men	0.131*	-0.1321***

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

In examining the impacts on the unemployment rate, a generally significant negative relationship with HICP inflation is observed across all groups, indicating that as consumer price inflation increases, unemployment rates decline, potentially due to increased economic activity in the short term, which facilitates job creation. Conversely, energy inflation seems to increase unemployment rates for all groups considered (Coibion et al., 2020).

Table 12 ATE – unemployment, Germany

	Treatment variable:	
	HICP inflation	Energy inflation
Outcome variable:		
Unemployment rate - total	-0.3653***	0.1620***
Unemployment rate - total, aged 15-25 years	-1.0716***	0.1386*
Unemployment rate - total, aged 25-74 years	-1.0162***	0.3341***
Unemployment rate - women	-0.9802***	0.7910***
Unemployment rate - women, aged 15-25 years	-1.3905***	0.5182***
Unemployment rate - women, aged 25-74 years	0.0937	0.6082***
Unemployment rate - men	-0.9063***	0.0190
Unemployment rate - men, aged 15-25 years	-0.5440***	0.1180*
Unemployment rate - men, aged 25-74 years	-0.3099***	0.3066***

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

Finally, in Germany, a rise in HICP inflation slightly decreases gross household saving rate and increases consumption expenditure, while energy inflation substantially increases both gross household saving and investment rates. This suggests that German households may react to energy price increases by reallocating more resources to investments and savings, possibly as a precaution against prospective energy price inflation.

Table 13 ATE - households, Germany

	Treatment variable:	
	HICP inflation	**Energy inflation**
Outcome variable:		
Gross household saving rate	-0.1055*	0.1918***
Gross household investment rate	0.0572	0.3356***
Household consumption expenditure	0.6432***	-0.5286***

*Notes: *statistically significant at 90%; **statistically significant at 95%; ***statistically significant at 99%.*

COMPARATIVE ANALYSIS

Across all three countries, HICP inflation shows only minimal impact on employment rates across most demographic sections. However, the influence of energy inflation on employment rates varies: in Italy and France, an increase in energy inflation stimulates employment rates (in line with existing literature (Coibion et al., 2020), with the effect spread across most demographic groups, particularly women (notably in France). In Germany, by contrast, energy inflation negatively affects employment rates among men aged 15-74, indicating potential harm to employment in energy-sensitive sectors. These findings highlight the nuanced interplay between inflation types and employment across different demographic groups and countries.

When looking at part-time and temporary employment rates in Italy, France, and Germany, the impact of HICP and energy inflation is diverse. Italy and Germany share a non-statistically significant impact of HICP inflation on part-time employment, while France exhibits a general positive effect for women and a negative effect for men. In both Germany and France, there is a positive effect on temporary contracts. The influence of energy inflation varies significantly: in Italy, energy inflation appears to drive up part-time employment rates across several demographic groups; in Germany, it significantly decreases part-time employment rates, particularly for men and more intensely among those aged 20-64 years; and in France, there is a general positive effect for women and a negative effect for men. The effect on temporary contracts is positive for Italy and France, but negative for Germany.

Similarly, the impact of HICP and energy inflation on unemployment rates presents a diverse picture across Italy, France and Germany. All three show a significant reduction in unemployment rates associated with a rise in HICP inflation, implying that inflationary periods might spur job creation. However, the influence of energy inflation varies significantly: in Italy and Germany, a surge in energy inflation appears to increase unemployment rates among the different groups; while, in France, energy inflation has varied effects, increasing unemployment rates overall but decreasing them among women aged <25 years.

On the key macroeconomic variables indicating household economic behaviour, the impact of HICP seems to reduce consumption and investment in France and Italy, while having a positive effect on consumption expenditure (Lieb and Schuffels, 2022) and a negative effect on savings in Germany (Opoku, 2020). Across all three countries, a common effect of energy inflation is increased investment – savings seem to increase in Germany and decrease in Italy and France, with the opposite holding true for consumption.

The observed differences in the impacts of HICP and energy inflation in the three countries considered reflect a constellation of factors. Their distinct labour market structures play a pivotal role, as the composition, flexibility, and regulatory environment in each case can significantly influence how inflation impacts employment rates, part-time and temporary employment, and unemployment rates. Additionally, variations in household economic behaviour also contribute to these differences. The manner in which households respond to inflation, i.e., their consumption, saving and investment patterns, all shape the macroeconomic response to inflationary shocks. Finally, the energy mix and energy dependence of each country also play a crucial role. Countries with different levels of reliance on energy-sensitive industries or those with diverse energy policies and green incentives respond differently to energy inflation. Accordingly, the nuanced interplay between labour market structure, household economic behaviour, and energy dynamics underpins the observed differences in the impacts of HICP and energy inflation in each country. These findings underline the importance of tailored policy interventions that consider the different reactions to inflation shocks across various demographic groups and household types.

DYNAMIC CAUSAL IMPACT

Local projections (a method first developed by Jordà (2005)) are used here to interrogate the time-bound nature of causal effects. This approach is particularly effective at describing the temporal dynamics of the responses of the outcome variables to inflationary shocks.

In its original form, the local projection technique entails regression of leads of outcome variables on identified shocks. Here, the contemporaneous and future values of the outcome variables' prediction errors (including the effect of the treatment variable by construction) are regressed on the prediction errors of the treatment variables (both obtained through DML, see equation (5)). Even in this case, no further assumption of shock identification is needed, as the procedure models the direct impact of a shock on future values of the prediction error of the variable of interest. As such, it is possible to understand the propagation on $t + h$ periods ahead of a shock hitting a variable of interest at time t.

To estimate the impulse response functions, the explanatory variables used are the Neyman orthogonalised residuals obtained from the target variable regressions described earlier. The shock variables are the orthogonalised residual of the treatment variables obtained from the DML model. The analysis considers the effect of inflationary shocks until 12 periods (months) ahead.[13] The average causal effect

[13] Variables are in monthly (weekly) frequency, thus estimating an impulse response function of 12 periods means analysing the effect on one year (3 months) of a shock hitting the variable of interest at time t.

should be quite similar to the mean of the corresponding impulse response functions. However, there are cases where the results might not match because the two models are not the same. At this stage, residual-on-residual equations are estimated using a local projection method, and these equations come from different setups in the DML model. The variables chosen by machine learning algorithms can be different for each lead of the outcome variables, which may result in larger standard errors.

Italy

Figure 4 presents the estimated impulse responses for a selected number of variables for Italy in the context of both a HICP shock and an energy inflation shock.

Starting from the effect of a HICP inflation shock on employment, a statistically significant negative effect is observed for the first month only. That shock is subsequently absorbed, with the employment rate returning to its steady state from the following two months onwards. Similar behaviour is observed for an energy inflation shock, when the impact on employment is initially positive but vanishes after a few periods.

For the unemployment rate, the impulse response to a HICP inflation shock aligns with the average effect presented earlier. It is, in fact, negative on impact and recovers in the next period. The effect of an energy inflation shock is interesting and highlights the importance of undertaking a dynamic causal effect analysis: initially, the observed response on impact on unemployment rate is positive but not statistically significant, but a negative and significant effect becomes evident after six months. This behaviour cannot be grasped simply by analysing an ATE and informs how the negative impact of energy inflation on the unemployment rate is delayed.

As for the effect of a HICP inflation shock on consumption and investment rate, the negative responses last four and three months, respectively. Consumption decreases after an energy price shock and the effect is delayed by three months.

Looking at the responses of saving and investment rates to an energy inflation shock, it is evident that savings exhibits a negative trend and becomes significant and severe after many periods, demonstrating a delayed effect. By contrast, the investment rate response is positive and returns to a steady state after four months.

The dynamic causal analysis reveals that the number of requests for the Italian Superbonus shows a negative initial response to a HICP inflation shock, which intensifies over a span of six weeks. Remarkably, for an energy inflation shock, the effect turns positive and statistically significant after a six-week period. These findings underscore the value of employing dynamic causal analysis in identifying the delayed effects of shocks on key variables.

Figure 4 Local projections – Italy

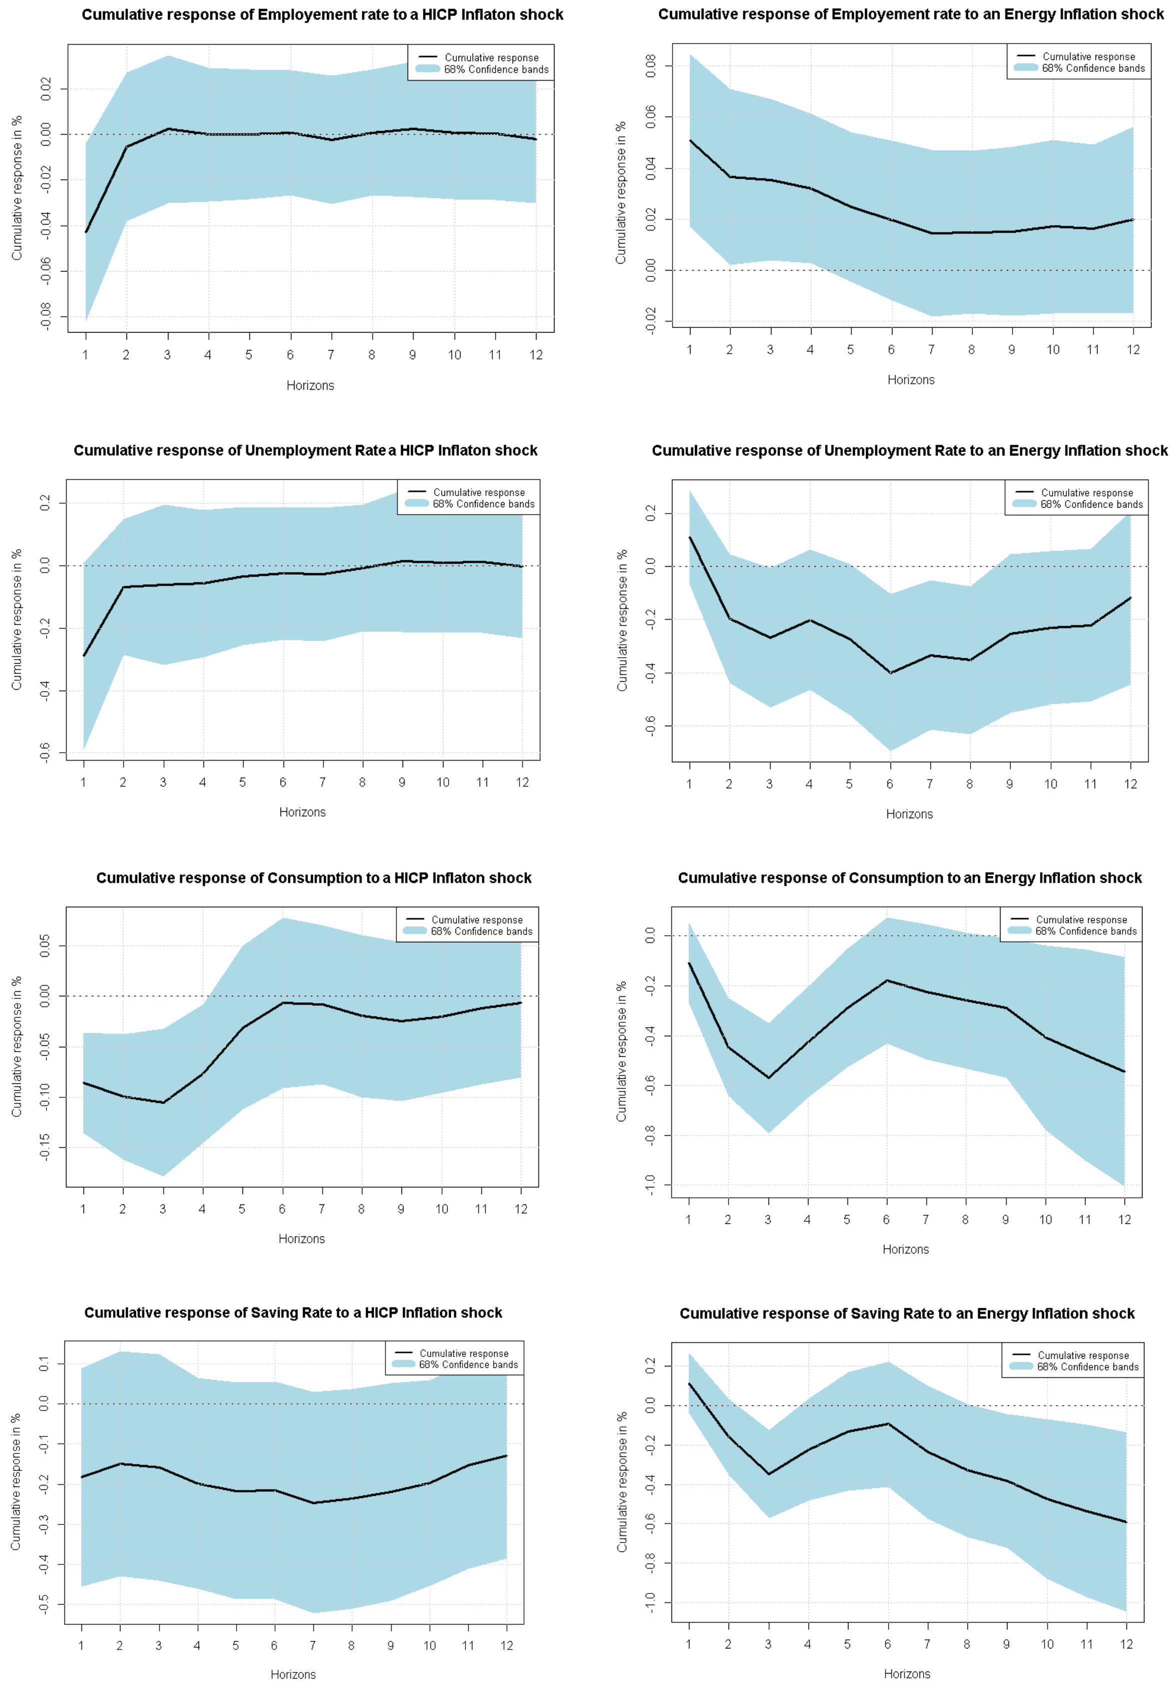

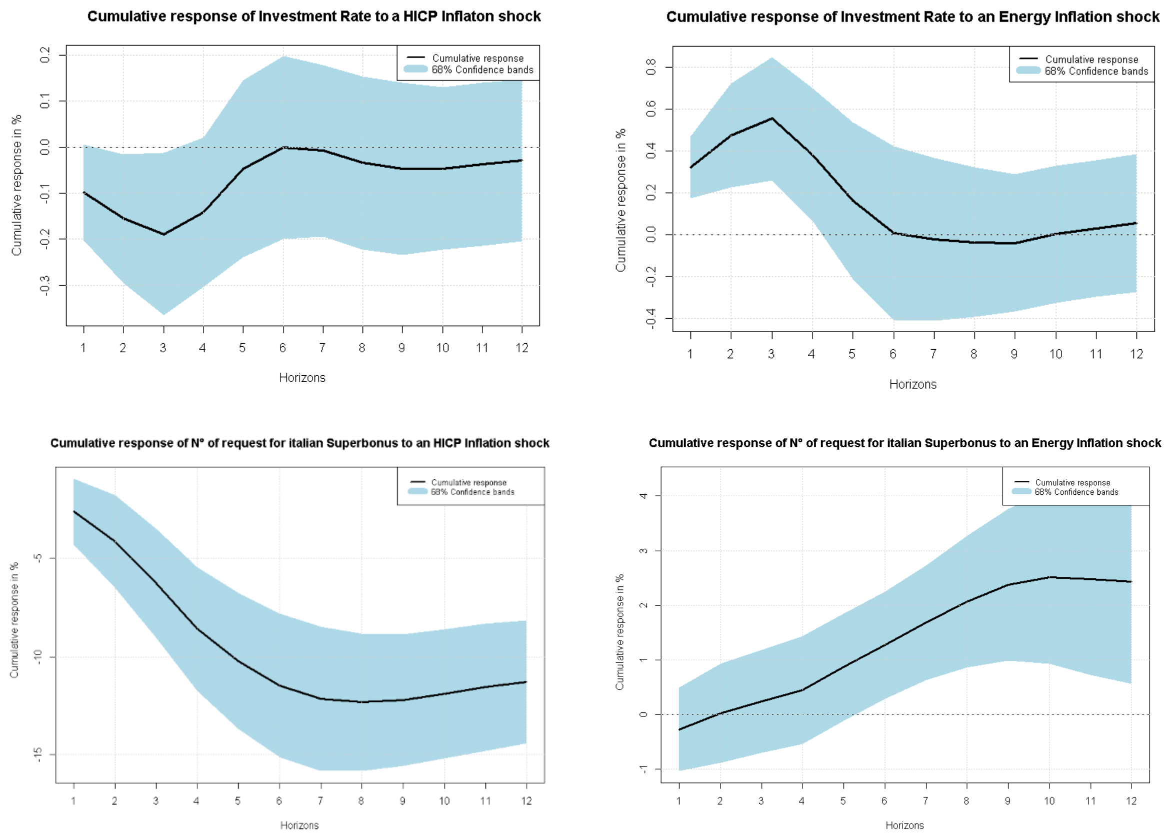

Notes: Impulse response functions for Italy; confidence bands computed at 68% confidence level; horizons represents months for all graphs, except impulse response functions of Superbonus, where horizons represent weeks. Source: Authors' elaboration.

France

The impulse responses of variables specific to the French context are presented in Figure 5. The response in respect of impact of the employment rate is coherent with the ATE results described earlier, i.e. negative to a HICP inflation shock, recovering immediately after, and positive to an energy inflation shock, returning to a steady state after four months.

The response of the unemployment rate shows a downward trend for both inflation shocks, but is not statistically significant.

The response of the household investment rate shows statistically significant and positive responses extending over six periods, but only for an energy shock. This is consistent with the results obtained for Italy. This observation can be interpreted through the prism of household economic behaviour, which appears to corroborate the theory that during periods of energy inflation, households may choose to reallocate their investments towards energy-saving assets, providing a potential hedge against this type of shock. The household saving rate is similar for both HICP and energy inflation shocks. The response to the energy shock is higher in magnitude, while both show statistically significant negative responses, phasing out after five months.

A closer scrutiny of consumption patterns reveals a decrease on impact of a HICP shock, which converges to a neutral position following one period. The opposite holds true for an energy inflation shock, where the response is positive on impact and returns to a steady state after four periods.

Figure 5 Local projections – France

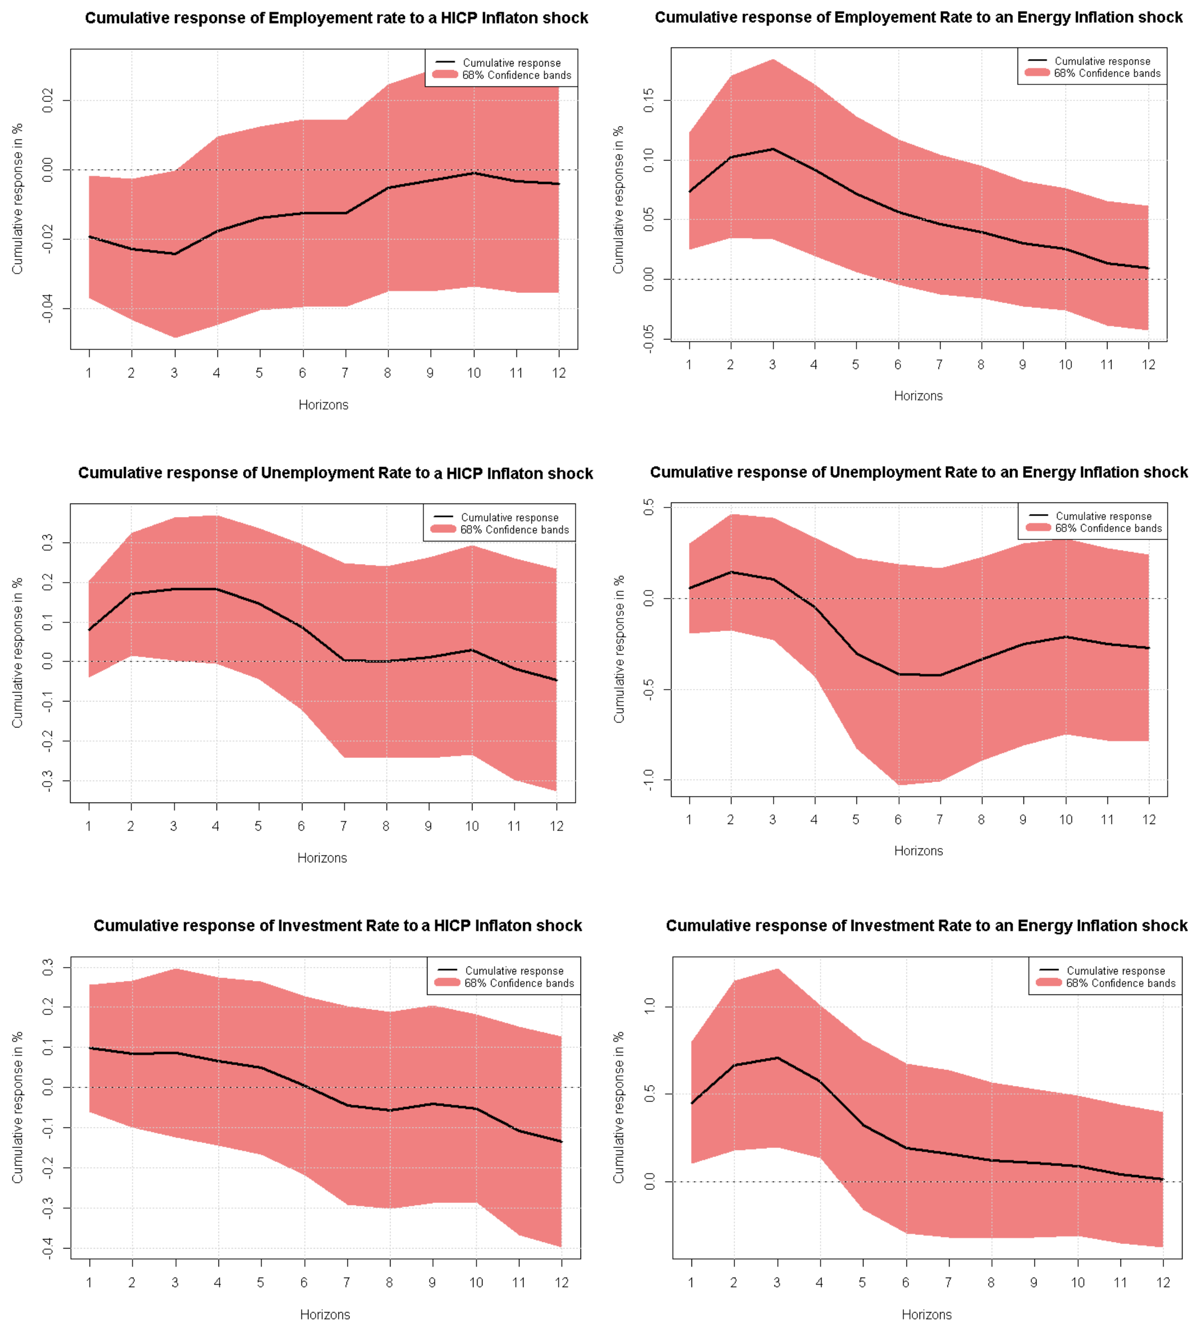

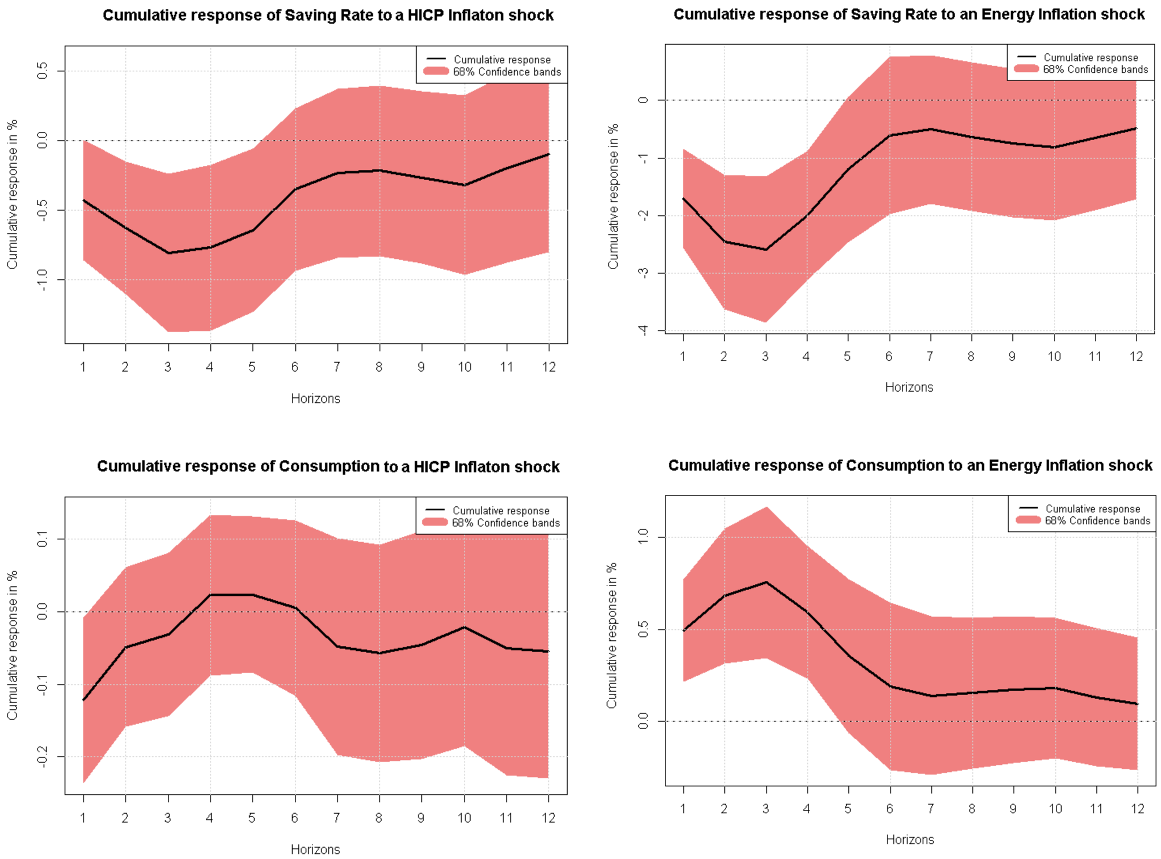

Notes: Impulse response functions for France; confidence bands are computed at 68% confidence level; horizons represent months.
Source: Authors' elaboration.

Germany

Figure 6 presents an in-depth analysis of the dynamic outcomes for Germany. Interestingly, the employment rate responds negatively to impact of an energy shock only, and the response reduces more markedly after three months. The unemployment rate responds significantly and increases with an energy inflation shock, recovering to a neutral level after three months. The opposite is observed for response of the household investment rate. The response to a HICP inflation shock is negative, recovering after four periods. However, similar to the responses of Italy and Germany, the energy inflation shock raises the investment rate, suggesting that households may shift their investments towards energy-saving assets. This is particularly important, given that a general inflation shock will reduce the investment rate overall.

In the case of a HICP inflation shock, the impact yields a positive effect; however, over a period of six months, a delayed negative effect is observed, which aligns with ATE results. Conversely, a positive effect is observed in response to an energy inflation shock, which may plausibly be ascribed to the comprehensive energy-saving campaign by the German government. Policy initiatives targeting energy efficiency may stimulate households to increase their saving rate. The rationale here is twofold:

firstly, such initiatives often entail incentives or subsidies for energy-efficient purchases or renovations, which may reduce household expenditure in the short to medium term, thereby allowing for an increase in savings; secondly, a cultural emphasis on frugality and efficiency might engender a behavioural response where households opt to save in order to invest in energy-efficient technologies or practices, especially in a context of increasing energy prices.

This effect is temporary, however, as the household saving rate response to energy shock gradually converges to neutrality over time. This suggests that the increased savings represent an adaptive but transient response to the energy shock, likely reflective of households adjusting to the new energy landscape before returning to their normal saving behaviour. Further research is required to ascertain the exact mechanisms and persistence of this phenomenon.

The reaction of consumption to a HICP inflation shock is delayed over time. More specifically, it turns positive from the second until the eighth months after the shock occurs, returning to a steady state thereafter. By contrast, consumption responds negatively and remains under the initial level when affected by an energy inflation shock.

Figure 6 Local projections – Germany

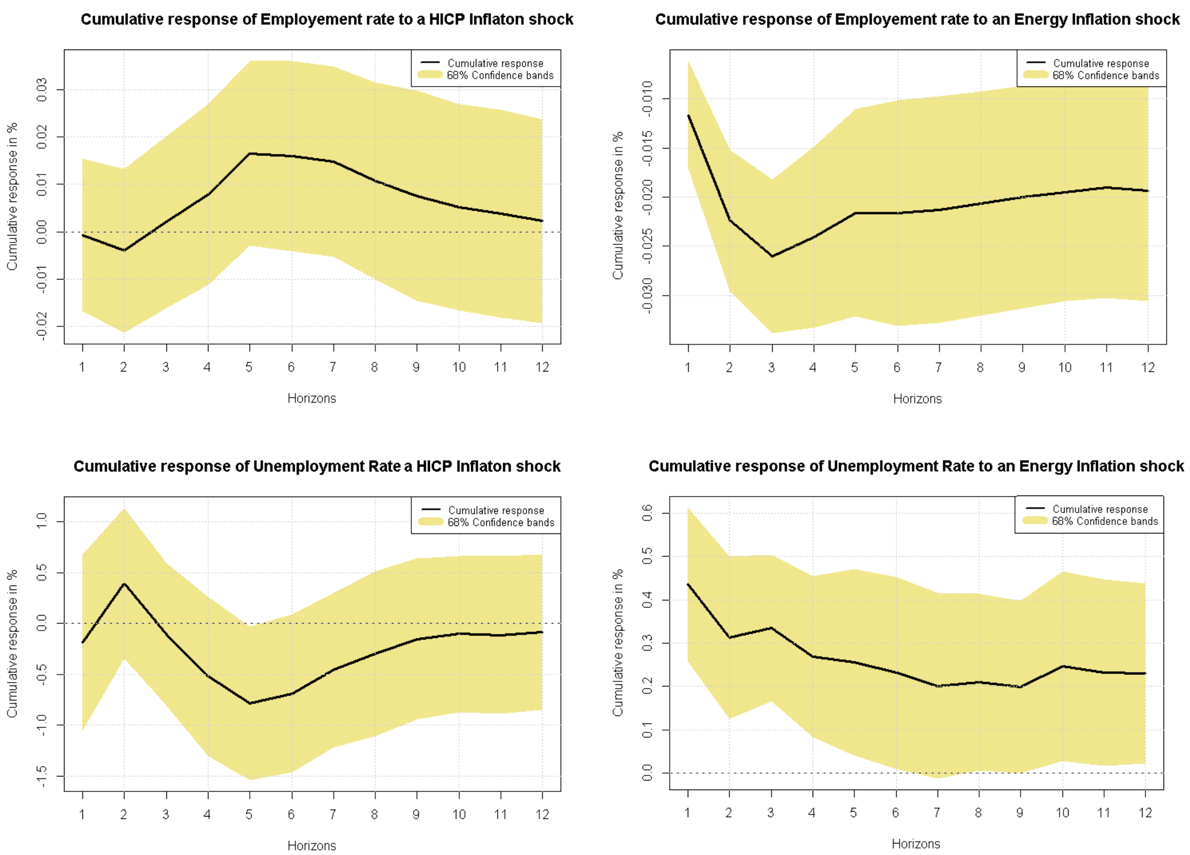

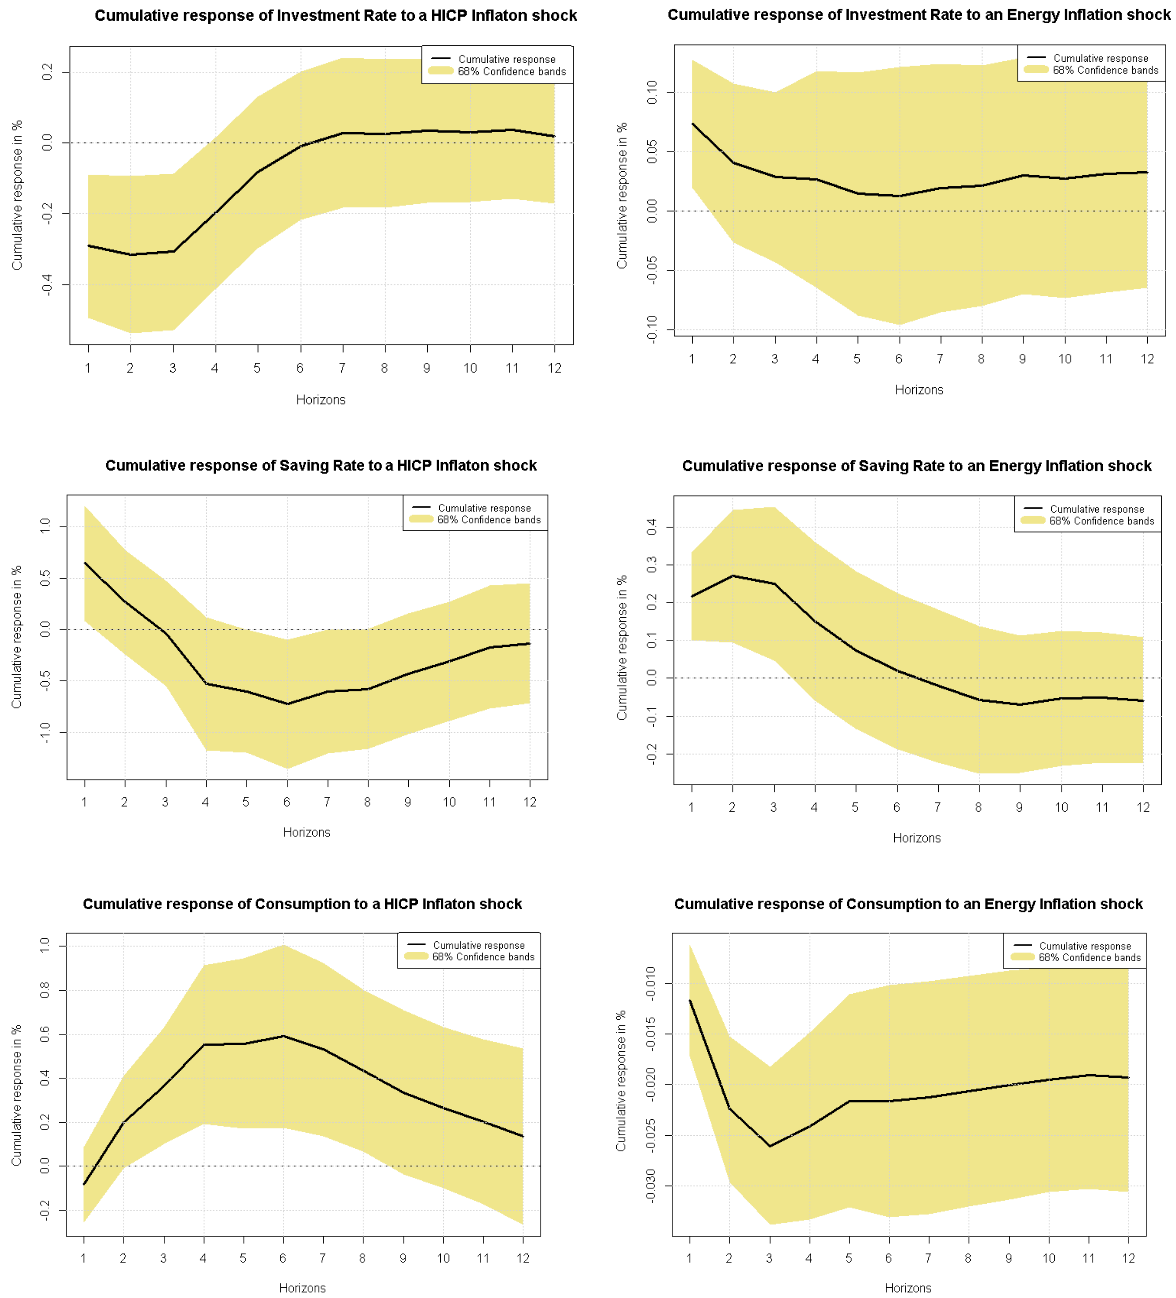

Notes: Impulse response functions for Germany; confidence bands are computed at 68% confidence level; horizons represent months.

Source: Authors' elaboration.

CONCLUSIONS AND FURTHER RESEARCH

This research note empirically evaluates the effects of inflationary shocks on macroeconomic indicators, labour market trends, and, for Italy, participation in one of the most-financed green transition policies, the Superbonus 100% initiative. The analysis relies primarily on a novel empirical methodology specifically developed for causal inference in high-dimensional data settings. Applied to a 'pure' time-series setting, it is based on the DML approach.

The causal machine learning approach is applied to an extensive dataset for three major European economies, Germany, France and Italy. Data are characterised by mixed frequencies, i.e. events observed at different points in time, giving rise to high-frequency missing values for low-frequency observations. This problematic aspect of the analysis is addressed by means of well-established techniques based on KF and the EM algorithm.

Identifying the effects of inflationary shocks is somewhat problematic in a macro environment, due to aggregation biases and strong dependency among shocks (treatments), transmission channels, and controls. The methodology used here is particularly suitable for this setting, provided that it is specifically designed to 'learn' from highly complex and correlated information sets.

The analysis first considers average treatment effects (i.e. not detailing the dynamic dimension in the causal effects). In order to better understand the causal relations, the dynamic responses of the target variables to the identified shocks are then analysed using local projections. These impulse responses help to detail and justify some unclear results of the average treatment effect analysis.

Two inflationary shocks are considered, one aggregate (HICP) and one for the energy price sector. Outcome variables at household level are consumption, saving and investment. For the labour market segment, the effects on employment rates, including part-time and temporary employment, and unemployment rates are depicted, considering aggregate measures, as well as age cohorts and gender decompositions. Despite the difficulties in precise understandings of an inflationary shock (demand-side or supply-side), results point to significant heterogeneous effects of inflationary shocks for the macroeconomy and the labour market. An aggregate inflation shock tends to trigger a reduction in consumption and household saving rates in all three countries, while the effects on household investment are negative and significant only for Italy and Germany. Interestingly, the shock hitting the energy prices does not necessarily lead to a macroeconomic contraction. Investments respond positively in all countries, while more heterogeneous effects are observed for the other macroeconomic variables related to household economic behaviour.

In Italy, the effects of inflationary shocks on participation in the foremost green transition policy are of particular interest: a generalised inflationary shock significantly dampens policy participation, while a price shock in the energy sector has a positive effect.

The labour market is heterogeneously affected by the two inflationary shocks. Aggregate inflation shock reduces the employment rate, while an energy inflation shock tends to increase the employment rate. Effects on the unemployment rate are more ambiguous, possibly due to the counteracting response of the participation rate. Results on the effects on the different labour market measures highlight the role of structural and labour market differences in the three countries. As the decompositions considered here reflect individual differences, they suggest that inflationary shocks might entail relevant distributional differences – future research could investigate further.

The dynamic analyses for Italy, France and Germany show that the economic variables have varying sensitivities over time to overall HICP and energy inflation shocks. The investment rate exhibits prolonged and statistically significant positive responses to energy inflation shocks, lasting up to four periods in France and Italy. This effect suggests a reallocation of household investment towards energy-saving assets, potentially as a hedging strategy against rising energy prices. In general, the effects linked to energy inflation shock tend to be more persistent than those of HICP inflation.

The dynamic analyses suggest the need for a multifaceted policy approach to the ongoing inflationary shock. The marked impact of energy inflation shocks on household investment rate, eventually coming back to equilibrium in nearly five months, suggests that policymakers should consider incentives to encourage investment in energy-efficient technologies or renewables with longer horizons. This could amplify the evident short-term propensity of households to shift their investment portfolios as a hedge against energy price volatility. In Germany, the long-lasting impact on the unemployment rate indicates a need for labour market interventions that can address the persistent effects of such shocks. Policymakers should also be aware that household behaviour, especially saving and investment, may be a temporary adaptation and revert to a baseline after the initial shock dissipates. Policy measures thus need to be dynamic, country-specific and adjustable if they are to effectively mitigate the varied and often delayed impacts of economic shocks.

REFERENCES

Aliaj, T., Ciganovic, M. and Tancioni, M. (2023). ‘Nowcasting inflation with Lasso-regularised vector autoregressions and mixed-frequency data’. *Journal of Forecasting*, 42(3), 464-480.

Bańbura, M., Giannone, D. and Reichlin, L. (2010). ‘Large Bayesian vector auto regressions’. *Journal of Applied Econometrics,* 25(1), 71-92.

Belloni, A., Chernozhukov, V., Fernández-Val, I. and Hansen, C. (2017). ‘Programme evaluation and causal inference with high-dimensional data’. *Econometrica,* 85(1), 233-98.

Belloni, A., Chernozhukov, V. and Hansen, C. (2014). ‘High-dimensional methods and inference on structural and treatment effects’. *Journal of Economic Perspectives,* 28(2), 29-50.

Blanchard, O.J. and Bernanke, B.S. (2023). 'What caused the US pandemic-era inflation?' *NBER working paper*, 31417.

Callen, T. and Thimann, C. (1997). *Empirical determinants of household saving: evidence from OECD countries*. International Monetary Fund (IMF).

Cascaldi-Garcia, D., Londono, J.M. and Wilson, B.A. (2023). SNB-FRB-BIS High-Level Conference on Inflation Risk and Uncertainty.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W. and Robins, J. (2018). *Double/Debiased Machine Learning for Treatment and Structural Parameters*. Oxford University Press Oxford, UK.

Cimadomo, J., Giannone, D., Lenza, M., Monti, F. and Sokol, A. (2022). ‘Nowcasting with large Bayesian vector autoregressions’. *Journal of Econometrics,* 231(2), 500-519.

Coibion, O., Gorodnichenko, Y. and Ropele, T. (2020). ‘Inflation expectations and firm decisions: new causal evidence’. *Quarterly Journal of Economics,* 135(1), 165-219.

Farrell, M.H. (2015). ‘Robust inference on average treatment effects with possibly more covariates than observations’. *Journal of Econometrics,* 189(1), 1-23.

Gagliardone, L. and Gertler, M. (2023). 'Oil prices, monetary policy and inflation surges'. *NBER working paper*, 31263.

Hartford, J., Lewis, G., Leyton-Brown, K. and Taddy, M. (2017). ‘Deep IV: A Flexible approach for counterfactual prediction’. In *International Conference on Machine Learning*, 1414–23. PMLR.

Hille, E. and Möbius, P. (2019). ‘Environmental policy, innovation, and productivity growth: controlling the effects of regulation and endogeneity’. *Environmental and Resource Economics,* 73, 1315-55.

Howard, D.H. (1978). ‘Personal saving behaviour and the rate of inflation’. *Review of Economics and Statistics*, 547-54.

Jordà, Ò. (2005). ‘Estimation and inference of impulse responses by local projections’. *American Economic Review,* 95(1), 161-82. https://doi.org/10.1257/0002828053828518.

Lieb, L. and Schuffels, J. (2022). ‘Inflation expectations and consumer spending: The role of household balance sheets’. *Empirical Economics,* 63(5), 2479-2512.

Long, J.E. (1993). ‘An econometric analysis of residential expenditures on energy conservation and renewable energy sources’. *Energy Economics,* 15(4), 232-238.

Opoku, P.K. (2020). ‘The short-run and long-run determinants of household saving: evidence from OECD economies’. *Comparative Economic Studies,* 62(3), 430-64.

Semenova, V., Goldman, M., Chernozhukov, V. and Taddy, M. (2023). ‘Inference on heterogeneous treatment effects in high-dimensional dynamic panels under weak dependence’. *Quantitative Economics,* 14(2), 471-510.

APPENDIX A:

Table A1 Superbonus dataset description and weekly proxy used for imputation, Italy

Variable name	Time frequency	Category	Weekly proxy
Number of requests for Italian Superbonus	Monthly	Target variable	Cumulative sum of Google Trends: Superbonus
Energy price inflation	Monthly	Treatment variable	Gas price inflation (GMELNOBS Index)
HICP inflation	Monthly	Treatment variable	Time interpolation + gas price inflation (GMELNOBS Index)
Interest rate - overnight deposits - households and non-financial corporations	Monthly	Control variable	ECB deposit facility - date of changes (raw data) – Level ECB marginal lending facility - date of changes (raw data) – Level ECB main refinancing operations – fixed-rate tenders (fixed rate) (date of changes) - level
Interest rate - loans for house purchases	Monthly	Control variable	ECB deposit facility - date of changes (raw data) – level ECB marginal lending facility - date of changes (raw data) – Level ECB main refinancing operations – fixed-rate tenders (fixed rate) (date of changes) - level
Interest rate - loans other than bank overdrafts to non-financial corporations	Monthly	Control variable	ECB deposit facility - date of changes (raw data) – level ECB marginal lending facility - date of changes (raw data) – level ECB main refinancing operations – fixed-rate tenders (fixed rate) (date of changes) - level
Interest rate - deposits - households and non-financial corporations - outstanding amounts	Monthly	Control variable	ECB deposit facility - date of changes (raw data) – level ECB marginal lending facility - date of changes (raw data) – level ECB main refinancing operations – fixed-rate tenders (fixed rate) (date of changes) - level

Variable name	Time frequency	Category	Weekly proxy
Interest rate - composite cost of bank borrowing indicator for households and non-financial corporations - short-term loans	Monthly	Control variable	ECB deposit facility - date of changes (raw data) – level ECB marginal lending facility - date of changes (raw data) – level ECB main refinancing operations – fixed-rate tenders (fixed rate) (date of changes) - level
Interest rate - composite cost of bank borrowing indicator for households for house purchase	Monthly	Control variable	ECB deposit facility - date of changes (raw data) – level ECB marginal lending facility - date of changes (raw data) – level ECB main refinancing operations – fixed-rate tenders (fixed rate) (date of changes) - level
Interest rate - composite cost of bank borrowing indicator for non-financial corporations - total loans	Monthly	Control variable	ECB deposit facility - date of changes (raw data) – level ECB marginal lending facility - date of changes (raw data) – level ECB main refinancing operations – fixed-rate tenders (fixed rate) (date of changes) - level
Employment rate 15-64Y	Monthly	Control variable	Google Trend: *Offerta di Lavoro* *Cerco Lavoro*
Unemployment rate 15-64Y	Monthly	Control variable	Google Trends: *Offerta di Lavoro* *Cerco Lavoro*
Gross household saving rate	Quarterly	Control variable	Google Trends: *Come Investire*
Gross household investment rate	Quarterly	Control variable	Google Trends: *Come Investire*
Cost of labour	Monthly	Control variable	Time interpolation
Construction companies' trust in economy	Monthly	Control variable	Time interpolation
HICP (Codes 01-03, 06-12)	Monthly	Control variables	Time interpolation

Notes: Variables used to estimate the causal effect of inflation on the number of requests for the EIP Superbonus.

APPENDIX B:

The selected estimation methodologies are primarily rooted in the latest advancements in the field of causal machine learning. This discipline was significantly influenced by a paper published by Belloni et al., (2014), outlining techniques with authentic predictive potential that could alleviate the hazards of overfitting and false discovery. The authors underlined the need to differentiate between in-sample fitting and out-of-sample prediction capabilities. The paper stressed the requirement for techniques that accurately handle multiple hypothesis or model testing on the same dataset.

Building on this foundational research, Farrell (2015) devised a new method for constructing confidence intervals with a doubly robust estimator. The intervals produced retained uniform validity across a wide range of models, marking significant progress in the field. This method is useful in probing multivalued treatments with diverse effects and selection from a higher number of covariates than observations (previously considered impossible).

Belloni et al. (2017) introduced a new technique to determine treatment effects in environments abundant in data. This innovative solution merged the latest progress in machine learning and causal inference, enabling a more precise evaluation of different treatment effects. The authors supplemented their techniques with a thorough theoretical examination, further endorsed by simulations and real-world examples that proved their empirical robustness.

Chernozhukov et al. (2018) proposed a novel way to employ deep learning to examine relationships between variables amid confounding factors. Their 'Deep IV' framework involves two distinct prediction tasks, each tackled with deep neural networks. The first predicts the treatment, with the subsequent network using this forecast to estimate the causal effect by integrating it into the loss function. The authors validated this method as a significant upgrade from traditional machine learning techniques, with wide-reaching applicability in fields requiring causal inference.

Finally, Chernozhukov et al. (2018) offered a ground-breaking approach, DML, for estimating treatment and structural parameters using diverse machine-learning algorithms. They introduced a theorem to obtain a non-asymptotic debiased machine learning result applicable to any function of any machine learning algorithm, provided certain fundamental interpretable constraints are met. This method is particularly effective in complicated settings with numerous variables, proposing a straightforward and universally applicable approach for estimating and inferring low-dimensional parameters. It is highly acclaimed as one of the most substantial recent breakthroughs in causal inference.

APPENDIX C:

The estimator of θ_0 that overcomes the regularisation and overfitting bias is built as follows.

Split the data into two samples: a main sample of size T_1, with observation numbers indexed by t ∈ I, and an auxiliary sample of size $T - T_1$, with observations indexed by t ∈ I^c. For simplicity, in this research note $t = T/2$.

The orthogonalised formulation is achieved through the process of directly partialling out the effect of X_t from d_t, resulting in the creation of the orthogonalised regressor $v_t = d_t - m_0(X_t)$. This involves obtaining an estimator of m_0, denoted by $\widehat{m}_0$, through a machine learning (ML) estimator using the auxiliary sample. The resulting estimator for the orthogonalised regressor is denoted by $\hat{v}_t = d_t - \widehat{m}_0(X_t)$, which is estimated using the main sample. After partialling the effect of X_t out from d_t and obtaining a preliminary estimate of g_0 from the auxiliary sample as before, the DML estimator for θ_0 is estimated using the main sample:

$$\hat{\theta}_0 = \left(\frac{1}{T_1}\sum_{t \in I} \hat{v}_t d_t\right)^{-1} \left(\frac{1}{T_1}\sum_{t \in I} \hat{v}_t(y_t - \hat{g}_t(X_t))\right) \tag{6}$$

To remove overfitting bias and slow rate of convergence, the samples above are inverted and a simple average the results of $\hat{\theta}_0$ is taken. (Chernozhukov et al. 2018) showed the under-regularity conditions obeyed by the estimate of theta using this DML estimator.

$$\sigma^{-1}\sqrt{T}\left(\hat{\theta}_0 - \theta_0\right) \sim \mathcal{N}(0,1), \tag{7}$$

where $\sigma^2 = (E[v_t 2])^{-1} E[v_t^2 \xi_t^2](E[v_t 2])^{-1}$. This means that standard procedures can be used to compute standard errors.[14]

The ML algorithm used to estimate nuisance parameters is Random Forest. It was finetuned through the utilisation of the Python 'scikit-learn' library, in combination with time-series cross-validation, using 50 iterations. By exploring the hyperparameters of maximum tree depth, number of randomly selected features, and number of trees, the optimal combination of parameters that minimises mean squared error (MSE) was determined.

[14] Standard error of a parameter $\hat{\theta} = \sqrt{\frac{\sigma_{\hat{\theta}}^2}{T}}$.

GETTING IN TOUCH WITH THE EU

In person

All over the European Union there are hundreds of Europe Direct information centres. You can find the address of the centre nearest you at: https://europa.eu/european-union/contact_en

On the phone or by email

Europe Direct is a service that answers your questions about the European Union. You can contact this service:

– by freephone: 00 800 6 7 8 9 10 11 (certain operators may charge for these calls),

– at the following standard number: +32 22999696 or

– by email via: https://europa.eu/european-union/contact_en

FINDING INFORMATION ABOUT THE EU

Online

Information about the European Union in all the official languages of the EU is available on the Europa website at: https://europa.eu/european-union/index_en

EU publications

You can download or order free and priced EU publications at: https://publications.europa.eu/en/publications. Multiple copies of free publications may be obtained by contacting Europe Direct or your local information centre (see https://europa.eu/european-union/contact_en).

EU law and related documents

For access to legal information from the EU, including all EU law since 1952 in all the official language versions, go to EUR-Lex at: http://eur-lex.europa.eu

Open data from the EU

The EU Open Data Portal (http://data.europa.eu/euodp/en) provides access to datasets from the EU. Data can be downloaded and reused for free, for both commercial and non-commercial purposes.

Publications Office
of the European Union